



BANK OF ENGLAND

Speech

The Phillips curve: lower, flatter or in hiding?

Speech given by

Sir Jon Cunliffe, Deputy Governor Financial Stability, Member of the Monetary Policy Committee, Member of the Financial Policy Committee and Member of the Prudential Regulation Committee

Oxford Economics Society

Tuesday 14 November 2017

I am grateful to David Bradnum, Pavandeep Dhani, Lien Laureys, Katie Low, Brad Speigner and Matthew Waldron for their help in preparing these remarks.

The unemployment rate in the UK today is 4.3%. The last time it was that low was 1975 – the year I graduated from university.

That year, average wages grew by 24%. 42 years later, with unemployment at the same level, whole economy average weekly earnings grew by 2.2% (Chart 1).

Equally strikingly, that 2.2% is about the same rate of wage growth as in 2011 when unemployment rose above 8% for the first time since the mid-1990s. Over the following 6 years unemployment has fallen quickly and continuously but nominal pay growth has largely remained bound between 1 and 3%.¹

Real pay growth, of course, has fluctuated much more widely. Over the past 6 years the UK has been buffeted by externally generated inflationary and disinflationary winds: by energy price inflation between 2010 to 2013; by energy price disinflation between 2014 and 2015; and, since the EU referendum, by externally generated inflationary pressure from the depreciation of sterling and some recovery in oil prices.

As a result, real pay growth has ranged from minus 2% to plus 3%. But, due in no small part to the stickiness of nominal pay and its apparent insensitivity either to inflation or a very rapid fall in unemployment, domestically generated inflation pressure has remained low.

A similar evolution of unemployment and pay has been seen in other advanced economies – particularly the US.² Indeed, the apparent disappearance of the relationship between pay and unemployment is for me the second big puzzle in advanced economies post-crisis. The first is the related puzzle of the disappearance of productivity growth which has almost certainly been responsible for some of the weakness in pay growth but that does not explain the full weakness (Chart 2).³

Why does this matter? The Bank of England does not target unemployment, let alone pay growth. Monetary policy makers learned, painfully, around the time unemployment was last at 4.3% that monetary policy cannot in the longer run determine real things like unemployment. Attempts to do so can result in very perverse effects. Rather, they should direct their efforts to controlling nominal things like inflation.⁴

¹ Of the 22 quarters over 2012-2017, whole economy total (regular) wage growth fell below 1% in five (four) of them and breached 3% in one (none).

² While some indicators of US pay growth have partially recovered (Yellen, 2017), the general picture remains one of relatively modest pay growth. Growth in average hourly earnings remains around 2½%, notably below the pre-recession norm of above 3%. A similar picture holds for other measures of pay, such as the employment cost index. This is despite the unemployment rate falling to its lowest level since late 2000. Constâncio (2017) surveys the possible causes of this “missing reflation” in advanced economies during the recovery from the 2008 recession.

³ As shown in Charts 3 and 4, both wage growth and unemployment have persistently surprised the MPC. See ‘Why has wage growth been subdued?’ in February 2017 Inflation Report for fuller discussion of the errors in our wage forecasts.

⁴ Monetary authorities can use their control over nominal quantities to either peg a nominal quantity or the rate of change in such a nominal quantity. They can, however, not peg a real quantity. (See Friedman (1968)).

That is why the Bank of England has a clear primary objective of price stability and a forward-looking inflation targeting remit. We have an objective to support the government's economic policy but it is a secondary objective and subject to the first.

However while neither unemployment nor pay is part of our target, the relationship between them is important to our understanding of domestically generated inflationary pressure in the economy and how it is likely to evolve over the 2-3 year period over which monetary policy can have influence.

And arguably, since the crisis, that relationship has become more important to our understanding of inflationary pressure in the economy.

Before the crisis the supply side of the economy posed far fewer puzzles. Productivity grew at a relatively steady rate around 2% a year and we had, or thought we had, a relatively clear idea of labour supply. The MPC certainly did not need to give as much attention to the evolution of the supply side of the economy as its present day counterpart.

That is not just my recollection as the Treasury representative on the Committee from 2002 to 2007. It is borne out by the minutes – the rate of mentions of 'supply' and 'productivity' in the minutes of MPC meetings in the past decade has been about twice that in the decade prior to the crisis.

As my colleague, Ben Broadbent has pointed out, pre-crisis the MPC could navigate primarily by demand and output. Changes in unemployment lagged changes in demand and output and, in the UK at least, the relationship between pay and unemployment appeared relatively stable.⁵

The post-crisis uncertainty over productivity and the dependence of economic growth on labour supply has, as Ben suggests, led the MPC to depend more on the labour market as a guide to spare capacity in the economy and on the relationship between pay and unemployment as a guide to domestically generated inflationary pressure in the pipeline.⁶

There is of course a vast academic literature on these issues which generally goes under the heading of the "Phillips Curve", the relatively constant, negative and non-linear relationship between wages and unemployment in 100 years of UK data that A.W. Phillips identified in 1958 (Chart 5).⁷

⁵ Broadbent 2014

⁶ To illustrate this dependence, growth in hours worked has accounted for 80% of growth in output in the UK since 2013, where it accounted for just 27% of growth in the decade prior to the crisis.

⁷ Phillips found that the data for 1948-1957 fit the same relationship as seen over 1861-1913. There was greater variation in the war and inter-war years.

That relationship has not remained constant since then – simple lines of best fit suggest that the relationship between the unemployment rate and nominal wage growth has moved lower and flatter over time (Chart 6).

Phillips's work has been extensively revised and adapted over the past 60 years, most importantly by adding some influence of inflation expectations to explain the apparent breakdown of the relationship in the 1970s when unemployment and pay appeared positively correlated.

The framework has been extended to prices rather than wages.⁸ And there has been extensive work on labour market dynamics and frictions to establish the point at which the level of unemployment accelerates the rate of inflation which, for the sake of simplicity, I will refer to here as the natural rate.⁹

However for nearly 5 decades the essential framework of a labour market impacted by inflation expectations that clears with lags around a natural rate of unemployment over the business cycle (or trade cycle in Phillips' original work) has played a fundamental role in modern macroeconomic models, including those used as part of their forecasting machinery by monetary policy makers, like the Bank of England.

So for monetary policymakers the apparent disappearance of the link between pay growth and employment is, as I have said, one of the key puzzles of the post recovery economy in the UK and in advanced economies. It has of course attracted a great deal of attention.

I am afraid I do not have the answer to the puzzle to give you here tonight. I suspect that for many years to come, economists in academe and elsewhere will be occupied - and arguing - about this period and revisiting (again) the established Phillips curve framework.

But policy makers cannot just 'down tools' until we have a better understanding. So I do want tonight to look at some of the main candidate explanations and what they might mean for the policy maker. And I will group them under four headings in the terminology of the Phillips curve: we have under-measured slack in the labour market; the curve has shifted downwards; the curve has become flatter; and the curve is not there.

⁸ The premise of the Phillips curve – that the rate of change in prices is related to the gap between supply and demand – extends beyond the labour market. Indeed, although Phillips' paper looked at wage growth the term Phillips curve is now often used to describe the relationship between general prices – inflation – and the unemployment gap (or other measures of supply/demand imbalance). In a stylised world, price and wage Phillips curves would tell us the same story – in a closed economy (where there is no impact from changes in import prices) in which the labour share is constant and productivity is growing in line with trend, there should be a steady relationship between wage inflation and price inflation. The labour share in the UK has been relatively flat since the turn of the century.

⁹ The natural rate of unemployment, sometimes referred to as U^* , typically refers to equilibrium unemployment in the long-run. However, there may be temporary factors that affect the level of unemployment consistent with stable inflation, and therefore the medium-term equilibrium unemployment rate, sometimes called the NAIRU – the 'non-accelerating inflation rate of unemployment'. For example, persistent weakness in demand can mean that more people will remain in unemployment for some time, which makes it difficult to re-enter employment. Once those temporary factors have dissipated, the equilibrium rate will tend back towards the natural rate of unemployment – the level to which unemployment would return eventually when supply and demand for labour are in balance in the long run.

Under-measurement of labour market slack

A straightforward explanation of why pay growth is subdued at very low levels of unemployment is that we are under-measuring the amount of spare capacity – or ‘slack’ – in the labour market. Recent trends in the world of work have meant greater (voluntary and involuntary) self-employment and part-time employment. Measures incorporating under-employment as well as unemployment – i.e. how much more people who are in work would *like* to work – may give a better indication of the amount of spare capacity in the labour market.

In such a world, low pay is simply telling the policy maker that there is more labour market slack than the unemployment indicators are registering, that the output gap is larger than thought and that the economy can grow at a faster rate without generating domestic inflation pressure.

I am reasonably sure that the under-measurement of labour resource utilisation accounts for some of our failure to forecast pay growth more accurately in the early years of the recovery. But I am less convinced it is now the main driver. The underemployment rate, which measures the fraction of the workforce that want to work more hours, still remains more than half of a percentage point above its pre-recession level in 2007. However, although some potential adjustment could possibly remain, we have nevertheless seen a substantial unwinding of spare capacity even in these broader indicators of underemployment.¹⁰

The Curve has shifted downward

The Phillips curve traces the relationship between pay growth on the one hand and the balance of labour market supply and demand, represented by unemployment, on the other. The same *shape* of curve can be drawn at lower or higher levels of pay growth and unemployment. A shift up or down in the level of the curve does not affect the coefficient of the relationship between pay and unemployment. But any given level of unemployment will generate a higher or lower rate of growth of pay.

There are a number of very different reasons why the Phillips curve may have shifted downwards so that current very low levels of unemployment are generating very weak upward pressure on pay. These include a lower natural rate of unemployment, lower or more anchored inflation expectations and changes in expectations of real pay growth.

¹⁰ The same picture holds even if we look at *net* desired hours, controlling for the number of people who want to work less hours than they currently are. On this measure, the workforce is basically content with the quantity of hours worked currently. However, the workforce wanted to work *shorter* hours in the run up to the crisis (2002—2007 period). So relative to that previous norm, there appears to be greater capacity for work currently. This is due to a combination of a surge in involuntary part-time work since the crisis as well as a reduction in the extent to which full-time workers want to reduce their work time. Both of these effects are unwinding with the tightening in the labour market but there is still potentially some adjustment left to go.

Lower natural rate

The Phillips curve framework depends on the concept of a natural, equilibrium, level of unemployment – the level to which unemployment would return eventually when supply and demand for labour are in balance. Within the framework this level is the anchor - unemployment below this level leads to upward pressure on pay (and hence inflation) and vice versa.

The natural rate cannot of course be directly observed; it can only be estimated. It depends on the structure and the dynamics of the labour market – on the fact that some in the labour force will find it much harder to find employment even when the labour market is tight and on the fact that churn in the labour market will lead to frictional unemployment, even if there are unfilled vacancies, as it takes time for workers who have left jobs to find new ones.

If the level of structural or frictional unemployment is reduced the natural rate will decline. Both are influenced by many factors. Structural unemployment, for example, can be (and in the UK almost certainly has been) reduced by changes to social benefits.

Frictional unemployment can be reduced if the labour market becomes more efficient at matching unemployed workers to unfilled vacancies.¹¹ That would also push the natural rate down. Indeed, in its review of the supply side of the economy in February of this year, the MPC appealed to evidence that secular improvements in the educational attainment of the workforce are likely to have improved the efficiency of the labour market and lowered trend unemployment and, as a result lowered the natural rate assumed in our forecast (Chart 7).

Inflation expectations

The same effect of lowering the curve might be occurring, albeit via a very different mechanism, by the impact of inflation expectations on pay growth.

As I noted earlier, one of the major developments of Phillips' original work was the incorporation of inflation expectations into the curve plotting the relationship between pay growth and unemployment. This enabled the Phillips curve framework to explain the experience of the 1970s in which unemployment rose but pay did not fall.¹²

¹¹ That can happen if the efficiency of matching changes – perhaps due to technology – or because of compositional changes to the workforce – if different cohorts of worker have different rates of frictional unemployment. Somewhat counterintuitively, the data do not tend to support the idea that matching efficiency is improving over time. See Pizzinelli and Speigner (2017) for a discussion.

¹² Phillips (1958) does investigate the role of inflation – noting that the relationship was absent in years when a rapid rise in import prices had initiated a wage-price spiral – but does not incorporate inflation or inflation expectations into the mathematical relationship.

The inflation adjusted Phillips curve adds inflation expectations as a driver of pay growth alongside excess supply (or demand) for labour relative to the natural rate. In order to preserve the value of real wages workers are assumed to require compensation for expected inflation as well as whatever pay increase or decrease the tightness or looseness of the labour market would warrant. If workers are surprised on the upside by inflation they will build that higher expected inflation into their bid for the next year and vice versa. It is this mechanism that drove the inflationary wage price spirals of the 1970s and disinflationary spirals of the 1980s.

A drop in inflation expectations, in the inflation expectations adjusted Phillips curve framework, will therefore have the effect of shifting the curve downwards and reducing pay growth at any given level of unemployment (Chart 8). Some in the US have argued that a reduction in inflation expectations could lie behind much of the weakness in US pay growth at current very low, by historic standards, levels of unemployment.¹³

In the UK, as I have noted, the current rate of inflation does not seem to have affected pay growth very much over the last 10 years. Indeed we have seen relatively large moves both up and down in real pay that do not seem to have been reflected much in nominal pay growth.

One reason for this might be the stability of inflation expectations around the Bank of England's 2% target. This is not a self-congratulatory point. Central bank credibility is crucial to anchoring inflation expectations. It was the de-anchoring and subsequent re-anchoring, through the establishment of central bank credibility – of expectations – that led to the inflationary and disinflationary spirals of the 1970s and 1980s.¹⁴

The impact of central bank credibility on inflation expectations in the UK would over time have lowered the inflation expectations adjusted Phillips curve in much the same way, though by a very different mechanism, as a drop in the natural rate. Arguably we are only now discovering the impact at very low levels of unemployment of the Bank of England's credibility as an inflation anchor.

Change in pay norms

One way of looking at the Phillips curve relationship between pay and employment is that it incorporates not only inflation expectations but also some 'norm' for real pay increases – some 'going rate' for an expected increase in real disposable income. Some of this will reflect increasing productivity – as output per worker rises, so their real pay ought to rise. But this norm may be demanded regardless of tightness in the labour market and regardless of whether increases in real pay are justified by increases in productivity.

¹³ See for example Brainard (2017)

¹⁴ Since the 1970s several monetary frameworks have been in place with the goal to provide an anchor for the price level and credibility for the government's commitment to low inflation. In the 1970s, targets for monetary aggregates were introduced, first for broad money and then for narrow money. With membership of the ERM in 1990, the exchange rate became the explicit target. Finally, following the departure of the ERM in 1992, an explicit inflation target was adopted. Operational independence was granted in 1997.

My colleague Ben Broadbent has suggested that one effect of the crisis and perhaps of other changes in the world of work and in workers' perceived pricing power has been to lower or remove that 'norm'.¹⁵

I suspect changes in workers' perceived pricing power and perhaps their appetite for risk have been responsible for some of the behaviour of pay and employment post crisis. If this has happened, and I will come onto why it might have done in a moment, the effect could indeed have been to lower the Phillips curve.

With one caveat, the policy implications of a lower Phillips curve are the same whether the driver is a lower natural rate, lower/more anchored inflation expectations, or some change in real pay increase norms net of changes in productivity growth. The evolution of unemployment and pay growth will trace the same shape of curve, but pay growth, and hence domestic inflation pressure, will be lower for any given level of unemployment.

All else equal the policy maker can therefore tolerate a lower level of unemployment before worrying about domestic inflation pressure in the pipeline.

The caveat, and it is a very important one, is that if the Phillips curve has been lowered by the anchoring effect of central bank credibility on inflation expectations, attempting to exploit that change might damage the central bank's credibility and increase expected inflation, shifting the curve back up again.¹⁶

The curve has flattened

A *lowering* of the curve does not change the sensitivity of pay growth to unemployment, and changes in unemployment, will put pressure on pay growth in the same way. It will all just happen at a lower level.

A *flattening* of the curve, however, reduces the sensitivity of pay growth to unemployment at all levels and to all rates of change of unemployment. This appears to have happened in some estimations of the post crisis estimations of the curve.¹⁷

¹⁵ See Broadbent (2014)

¹⁶ This could result in a 'double whammy': pay growth would be accelerated higher demand for labour while at the same time the inflation adjusted curve would shift up, accelerating pay growth further. This can be seen in Chart 8: a fall in unemployment would result in moving from point A to point B (at higher inflation expectations) or point C to point D (at lower expectations). If those inflation expectations were to change at the same time that unemployment falls, the move would be from C to point B. See also Mishkin (2007).

¹⁷ See for example BIS (2017) and IMF (2013). Blanchard et. al. (2015) argue that since 1990 there is no statistically significant slope to a price Phillips curve in many countries, including the UK.

I noted above that changes in the world of work have very possibly changed the pricing power of labour and workers' appetite for risk (i.e. job insecurity). This is in itself a large area of current debate and I do not want here to go into these in great detail.¹⁸

Some of these important changes in the structure of the labour market, such as the rise in self-employment and decline in union membership, predated the financial crisis. Others, like the rise in temporary work and zero hours contracts, are more recent. Technology – and the rise of the gig economy – has further increased what my colleague Andy Haldane has called the 'divisibility' of labour.

The impact of the financial crisis and subsequent recession itself may also have changed workers appetite for risk. Anecdotally this is certainly true; I have lost count of the number times in the last four years, when the discussion on regional visits has turned to pressure for pay increases that I have heard the explanation that workers were "happy just having a job".

This may well be more than anecdote. There is literature on the Great Depression that identifies the lifetime effects on the so called 'depression babies' and their attitude to risk.¹⁹

Changes to the world of work, technology and a greater sense of insecurity, therefore may well have changed workers perception of their bargaining power, their appetite for risk and weakened their response, in terms of seeking higher pay, to signals of a tightening labour market. This would have flattened the Phillips curve as well as lowered it.

To the extent that greater risk aversion has played a role, one might expect the impact to be asymmetric so that the effect on pay from falling unemployment is less than the effect of rising unemployment.²⁰

Globalisation and technology have almost certainly contributed to a greater sense of job insecurity. But they may also have a more direct effect on the Phillips curve relationship between unemployment and pay by enabling global supply and production chains that have resulted in a more global labour market.

To put this in its starkest terms, the global Phillips curve matters much more now. Claudio Borio and Andrew Filardo at the BIS argue that traditional models are too 'country-centric' and that a more 'globe-centric' view of the inflation process is now not only warranted but has been growing in importance over time.²¹

¹⁸ My colleagues on the MPC, Michael Saunders and Andy Haldane have explored these issues in great detail. See Saunders (2017) and Haldane (2017) for a discussion of how labour market flexibility may have affected pay determination.

¹⁹ For example, Nagel, S and Malmendier, U (2011)

²⁰ Stiglitz (1997) has described a similar process for the price Phillips curve based on asymmetric price adjustment.

²¹ See Borio and Filardo (2007).

Mark Carney has pointed out how global influences will shift the Phillips curve both lower (due to the positive supply shock from increased product and labour market integration) and flatter (due to integration of global value chains, increased competition and contestability of product and labour markets). The downward shifts, which will persist as long as integration continues, are disinflationary, pointing to running the domestic economy at tighter spare capacity to meet the inflation target.²²

For the policymaker a flattening of the Phillips curve poses a more difficult challenge than under-measurement of labour utilisation or a lowering of the natural rate or of inflation or real pay expectations (Chart 9).

The latter pose the challenge of judging the timing of policy action right; excess (deficient) demand for labour will drive pay up (down) as before as unemployment falls (rises).

But the point – i.e. the level of unemployment and the rate of growth of pay – at which that will happen is lower. Getting the timing of monetary policy right is a formidable challenge for the monetary policy maker. But it is a familiar one.

A flatter curve on the other hand, not only affects the timing of the build-up of pay pressure as unemployment falls and hence the timing of policy action. It affects also the sensitivity of pay pressure to unemployment and hence the degree of policy action necessary to bring inflation to target.

If the curve is pretty close to flat, monetary policy will have to work very hard to affect domestically generated inflation. As Borio, and other, have noted this may bring into question the very ability of central banks to meet their inflation targets.

There is no curve

The Phillips curve framework has been a central part of our thinking about the economy for over 50 years. But it has not always been that way.

Keynes, after the Great Depression, posited a world where absent government intervention labour markets did not clear even with lags. In this world unemployment could get stuck at high levels even though wages fell and fiscal policy was therefore required to bring unemployment down.

And it is true also that the Phillips curve framework has had to be adapted and extended substantially over the past 50 years to maintain its explanatory power in the face of the actual evolution of pay growth and unemployment over the period in a range of countries.

²² Carney (2017b)

We cannot exclude other possibilities. Roger Farmer for example argues that the statistical evidence for a natural rate of unemployment, an equilibrium to which the labour market returns is unconvincing. Rather, he argues, there are multiple equilibria for the labour market. Which unemployment rate prevails depend on investors' animal spirits, captured by a 'belief function' comprising expectations of future economic prospects.²³

The policy implications of a very different framework for the relationship between demand and supply would be fundamental. Given the success, in my view, of inflation targeting frameworks over the past few decades, policy makers should be rather cautious here and set a pretty high evidential bar – in both theoretical and empirical terms – for assessing radically different frameworks.

Current UK monetary policy

These issues are particularly relevant now given the current conjuncture in the UK.

Inflation is currently above target as a result of the post referendum depreciation of sterling and forecast, for that reason, to remain so over the next three years.

The Monetary Policy Committee's (MPC's) remit requires the MPC to meet the inflation target, which is forward looking, at all times. The remit recognises that in exceptional circumstances a large and persistent shock to the economy could present the MPC with a significant trade-off between the speed at which it brings inflation back to target and the impact on jobs and economic activity. The post referendum depreciation of sterling is precisely such a shock.

Given the post-crisis fall in the economy's productive capacity we appear to be operating close to full capacity, even at a growth rate of less than 2%. Unemployment is at a 42 year low. Labour market indicators such as job to job flows, quit rates, and recruitment difficulties are signalling tightness in the market.

Domestically generated inflation pressure, however, appears low. Measuring domestically generated inflation when externally generated inflation pressure is high, as at present, is not straightforward. Higher import prices work their way into the domestic production chain and can appear as domestic inflation.

Bank staff calculations suggest that adjusted for this effect indicators of domestic inflation pressure are below levels consistent with the 2% target. And that is particularly true for the most domestic of the domestic measures, pay and unit wage costs.

²³ See Farmer (2013) and Farmer and Nicolo (2017)

Whole economy average weekly earnings grew in the year to August at a rate of 2.2%. Higher frequency indicators suggest a faster rate of growth: 3 month on 3 month annualised AWE growth for *regular* pay is 2.9%. But wage data are volatile and can mislead. Between 2014 and 2016, the 3 month on 3 month measure annualised AWE gave readings well above 3% which did not translate into sustained higher rates of pay growth.²⁴

And, perhaps in keeping with changes in workers' risk appetite, labour market indicators like the unemployment rate have signalled tightness for most of my time on the MPC without translating into pay pressure.²⁵

The MPC's November Inflation Report forecast, which is the best collective judgement of the Committee, is that inflation will peak around 3% this year. In the later part of the forecast period, as the impact of externally generated inflation pressure wanes, domestic inflation pressures, particularly pay, rise so that inflation is slightly above the target at the 3 year forecast horizon. The forecast is based on the increase in Bank rate to 0.5% which the Committee agreed at its last meeting and on the market path for interest rates, which implies Bank rate rising to 1% over the next three years.

At the November meeting I voted against increasing Bank rate. I did so not because I had a markedly different view to the majority of the right trade-off between output/inflation over the forecast period – or to put it another way, on the limits to tolerance of above-target inflation. In this respect, I share the overall framework of the Committee.

Our framework, however, depends crucially on our estimation of the Phillips curve – both the slope and level.²⁶ Given the uncertainties I have described and the serial disappointments we have had in recent years in forecasting the impact of unemployment on pay growth, there is in my view a not immaterial risk that the trade-off is not as it currently appears and that domestic inflation pressure will undershoot the Committee's collective expectation.

In my view the low level of domestic pressure on inflation now, the absence of second round effects from the depreciation of sterling, and inflation expectations around their historical averages, make it possible to wait

²⁴ I have quoted the shorter-term measure for *regular* pay, rather than total that I use elsewhere, as the inclusion of bonuses in *total* pay leads to more short-term volatility in that series. 3-month-on-3-month annualized growth for *total* pay was 2.1% in August but had been above 3% earlier in the year, and has risen well above 4% at two points in the preceding three years without signaling sustained higher wage growth.

²⁵ Job to job flows within a quarter have been at about 2.5% - which is what they averaged over 2005-07 before falling down below 2% in 2009 – since early 2016. Similarly, the MPC's swathe of surveys have been pointing to recruitment difficulties for some time, for example the REC survey has indicated below-average staff availability since 2013 – and significantly so since mid-2014.

²⁶ As is set out in Carney (2017a): The actual trade-off struck will be influenced by individual MPC members' views of both the nature of the shocks hitting the economy and the transmission mechanism of monetary policy – in other words, why the economy is behaving as it is, and how interest rates affect output and inflation. Put another way, if members have, for example, different views of κ – the slope of the Phillips curve – their λ s will appear to be different. Indeed, it is quite possible for the Phillips curve slope to vary over time, and some have argued that it may have flattened in recent years.

before tightening policy until there is clear evidence that pay growth is responding to the level of unemployment in line with our forecast.

Finally, I have managed to get almost to the end of this speech without mentioning Brexit. Almost, but not quite. While Brexit is a further source of uncertainty, my policy decision does not anticipate any particular Brexit outcome.

The questions I have described in relation to the Phillips curve framework are independent of Brexit. They arise in other advanced economies. But Brexit is relevant to monetary policy and to the MPC.

Expectations of Brexit - the full and diverse range from the most optimistic to the most pessimistic – are in the UK economy now. These expectations have clearly had an impact on the exchange rate with all that implies for output and inflation and for the exceptional circumstances trade off faced by the MPC that I have described earlier.

And though it is more difficult to measure, these expectations are affecting to some degree many of the economic decisions of businesses and households.

The full and diverse range of expectations of Brexit and of the path to it cannot by definition all come to pass. Over our forecast period, a specific Brexit and a specific path to it will emerge. The MPC cannot forecast what that outcome will be or how it will measure against the diverse expectations in and between households, businesses and financial markets.

Nor can the MPC forecast how that outcome will affect the adjustment of the paths of demand, supply and the exchange rate, relative to each other. It is that relative adjustment of these variables that will determine the path of inflation and hence whether – and in which direction – policy may need to respond to meet the Committee's primary objective. We can, in my view, only work on what we see in the economy now and stand ready to respond as necessary as Brexit emerges.

Conclusion

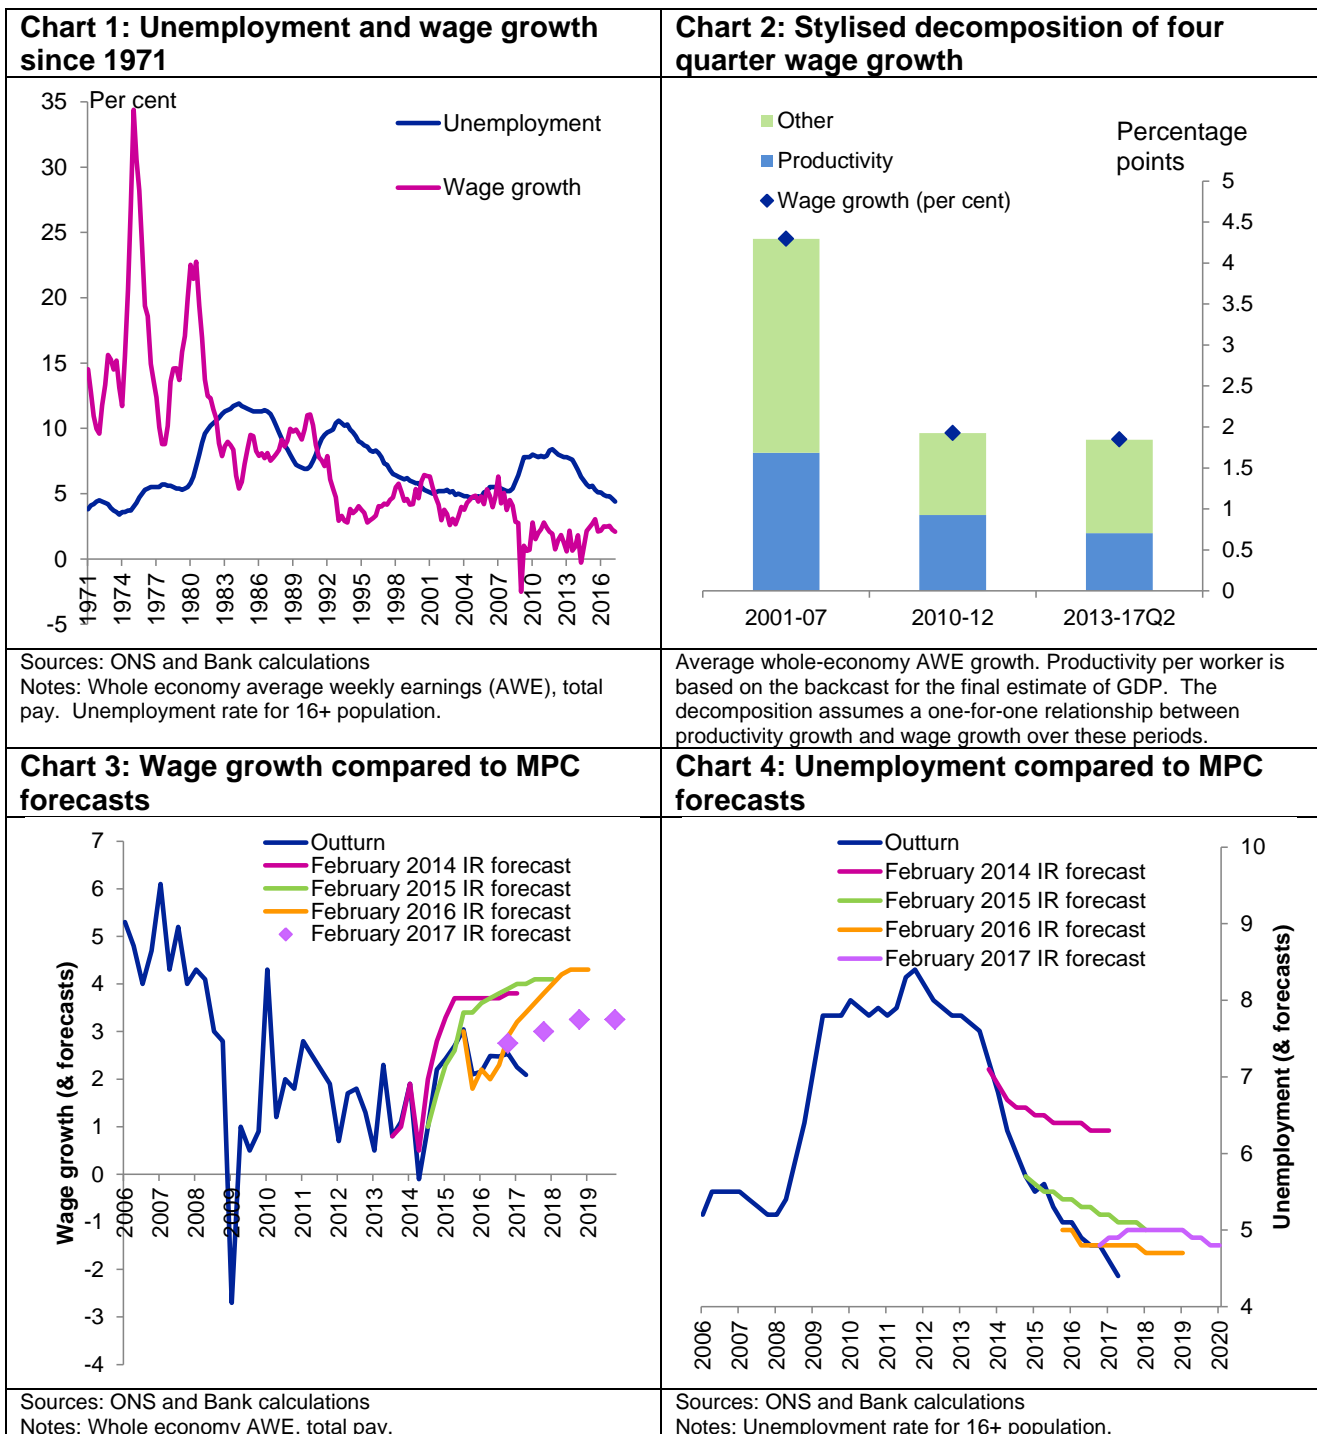
Monetary policy makers complaining about uncertainty is a bit like sailors complaining about the sea. Uncertainty is not a reason for inaction. Economic data are never clear. Economies are adaptable organisms that are always evolving.

Monetary policy cannot simply wait until we have a clearer understanding of whether we are correctly measuring employment, whether the Phillips curve has shifted down, tilted or both – or whether there is a better framework for understanding the relationship between supply, demand and inflation.

But we are almost certainly going through a period of heightened change at present. The current level of uncertainty around the key assumptions on which we model and understand the economy is, in my view, different in kind to uncertainties about the data or signals we receive which are always mixed. Some of the relationships between economic variables that we depended on in the past appear to have gone on a longish leave of absence, but we are not sure why, or whether this is a temporary or more persistent departure.

Against that background – and although it makes forward looking monetary policy more difficult – I tend to put more weight on the evidence we can or cannot see in the data and a little less on the un-observables and on how we think the economy works.

In any event, monetary policy makers have a personal reason to be careful of complaining too much about the burden of uncertainty. It is not just that it is one of the things that makes the job challenging and worth doing. It is also why it cannot be done by a rule – or, at least as yet, by a robot.



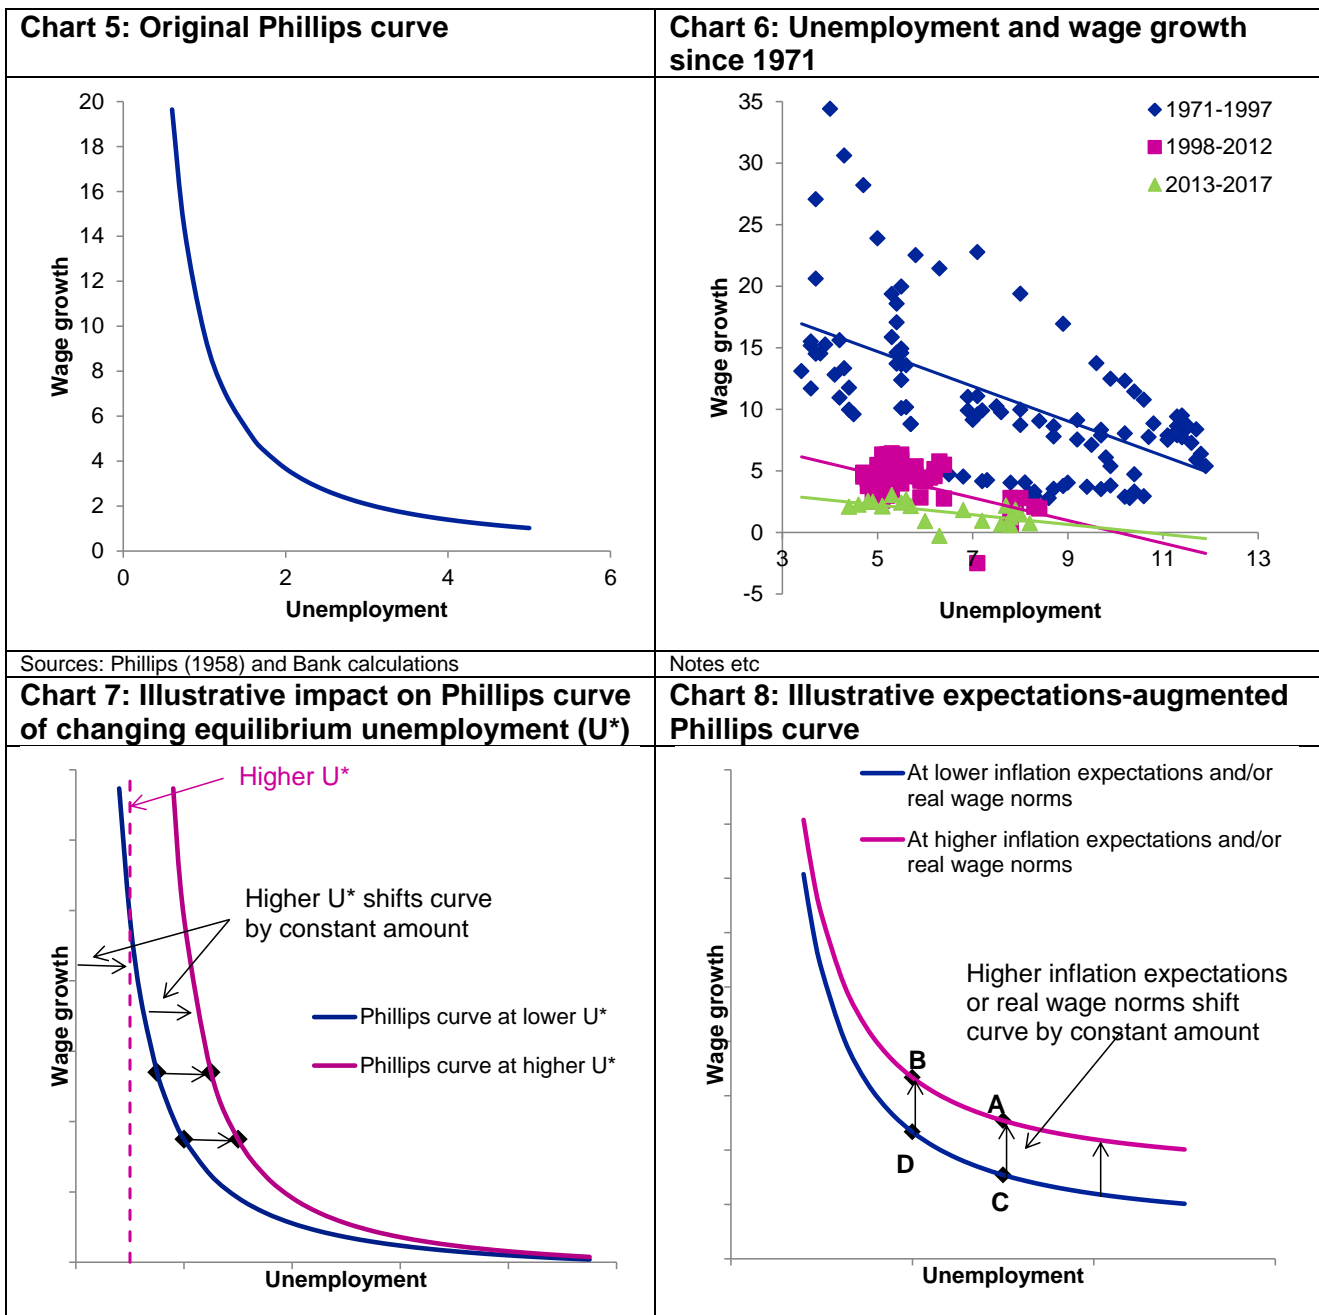
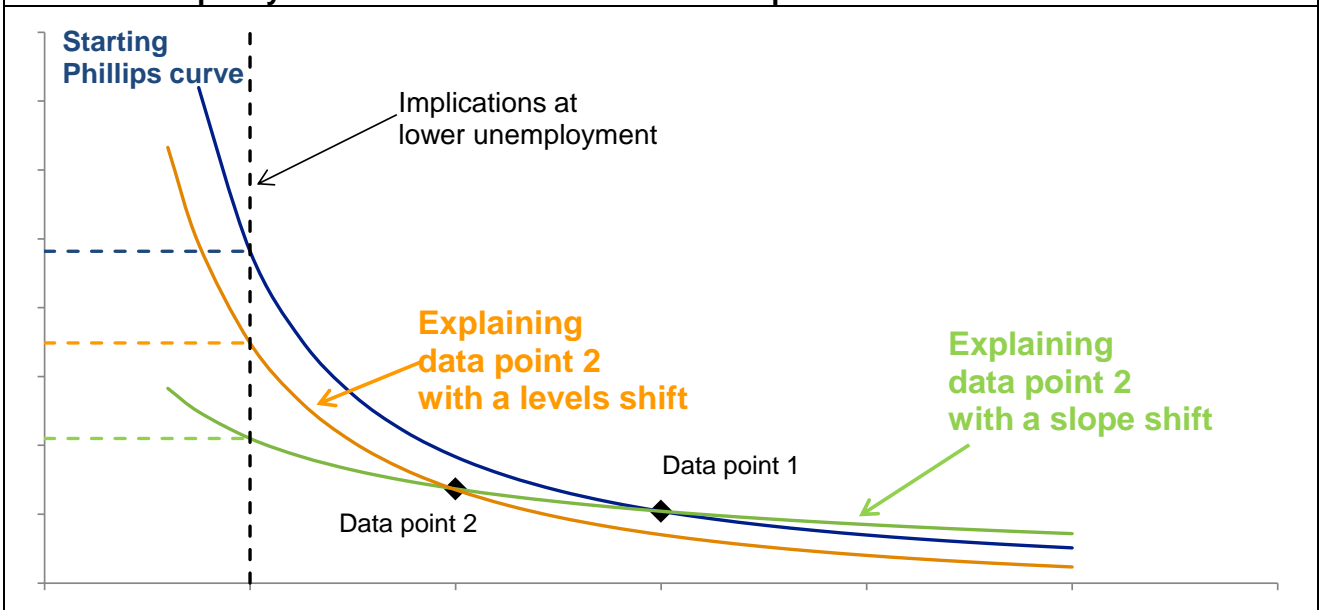


Chart 9: The policy-maker's dilemma: level shift or slope?



References

- Borio, Claudio and Andrew Filardo (2007)** 'Globalisation and Inflation: New cross-country evidence on the global determinants of domestic inflation' *BIS Working Papers, No. 227*.
- Bank for International Settlements (BIS) (2017)** '87th Annual Report, Basel'
- Blanchard, Cerutti, Summers (2015)** 'Inflation and Activity – Two Explorations and their Monetary Policy Implications' *IMF working paper*
- Brainard, Lael (2017)** 'Understanding the Disconnect between Employment and Inflation with a Low Neutral Rate' *speech given at the Economic Club of New York*
- Broadbent, Ben (2014)** 'Unemployment and the conduct of monetary policy in the UK' *speech given at the Federal Reserve Bank of Kansas City 38th Economic Symposium, Jackson Hole*
- Carney, Mark (2017a)** 'Lambda' *speech given at the London School of Economics*
- Carney, Mark (2017b)** '[De]Globalisation and Inflation' *2017 IMF Michel Camdessus Central Banking Lecture*
- Constâncio, Vítor (2017)** 'Understanding and overcoming low inflation' *remarks at the Conference on "Understanding inflation: lessons from the past, lessons for the future?", Frankfurt*
- Farmer, Roger (2013)** 'The Natural Rate Hypothesis: an idea past its sell-by date' *Bank of England Quarterly Bulletin Q3 2013*
- Farmer, Roger E.A. and Giovanni Nicolo (2017)** 'Keynesian Economics without the Phillips Curve' *NBER Working Paper No. 23837*
- Friedman, Milton (1968)** 'The Role of Monetary Policy' *The American Economic Review, Vol. 58, No. 1, pp. 1-17*
- Haldane, Andrew G. (2017)** 'Work, Wages and Monetary Policy' *speech given at the National Science and Media Museum, Bradford*
- IMF (2013)** 'The dog that didn't bark: has inflation been muzzled or was it just sleeping?' *April World Economic Outlook, Chapter 3*
- Mishkin, Frederic (2007)** 'Inflation Dynamics' *NBER working paper 13147*
- Nagel, S and Malmendier, U (2011)** 'Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?' *Quarterly Journal of Economics, Vol126(1), pp373-416*
- Phillips, A W (1958)** 'The relation between unemployment and the rate of change of money wage rates in the United Kingdom, 1861–1957' *Economica, Vol. 25, No. 100, p. 283–99*
- Pizzinelli, Carlo and Bradley Speigner (2017)** 'Matching Efficiency and Labour Market Heterogeneity in the United Kingdom' *Bank of England Working Paper No. 667*
- Saunders, Michael (2017)** 'The Labour Market' *speech given at the Resolution Foundation, London*
- Stiglitz, Joseph (1997)** 'Reflections on the natural rate hypothesis' *Journal of Economic Perspectives, 11, 3–10*
- Yellen, Janet (2017)** 'Inflation, Uncertainty and Monetary Policy' *Speech at the "Prospects for Growth: Reassessing the Fundamentals" 59th Annual Meeting of the National Association for Business Economics, Cleveland, Ohio*