RECOMMENDED READING

MPC MEMBERS

Copies to: James Bell
Simon Hayes
Forecast team
Costs and Prices Team

From:

HOW LARGE IS THE UNEMPLOYMENT GAP?

Key Message

Macro-econometric evidence suggests that the risks to our current estimate of 5% for the long-run equilibrium unemployment rate may be more balanced than suggested by the persistent weakness of wages and other downside factors such as compositional changes in the labour market and changes to work incentives. Indeed there is some evidence of a deterioration in labour market efficiency in recent years.

Main points

- The unemployment rate has fallen sharply since mid-2013 but has since flattened off somewhat above Bank staff estimates of the medium-run equilibrium rate of 5.3% (for 2015Q2) and the long-run equilibrium rate of about 5%. This note assesses a range of evidence on equilibrium unemployment, particularly given that this was not a focus of the supply stocktake earlier this year.
- The persistent weakness of wages in 2013 and 2014 pointed to downside risks to the equilibrium unemployment rate. Changes in the composition of the workforce towards groups that typically have low unemployment rates, such as the old and highly educated, could reduce aggregate unemployment, as could the increasing stringency of unemployment benefits. But other changes that incentivise greater participation could lead to higher unemployment if those encouraged out of inactivity find it difficult to obtain jobs. The net effect of these specific factors is difficult to quantify.
- The efficiency with which unemployed people are being matched to job openings deteriorated just prior to the crisis, and has remained low as vacancies (a measure of labour demand) recovered to a robust level. This apparent outward shift in the Beveridge curve could indicate an increase in the long-run equilibrium unemployment rate from 5% to 5½%, although it could also reflect transitory factors such as increases in participation and inward migration.
- An alternative approach based on statistical filters that make use of a Phillips curve relationship between wages and the unemployment gap are more likely to capture the effective slack in unemployment for wage pressure, rather than the long-run equilibrium, so they are more akin to our measure of medium-term equilibrium (U*). Different specifications are all close to our current estimate of U*.
- Overall, the macro-econometric evidence does not point to a large unemployment gap remaining, and there is some evidence of upside risks. That potentially leaves the risks around our estimates of equilibrium unemployment more balanced than previously thought.

RECOMMENDED READING

Introduction

After remaining elevated for a number of years following the recession, the unemployment rate has fallen sharply since mid-2013 but has since flattened off. Currently, unemployment remains somewhat above Bank staff estimates of the medium-run equilibrium rate of 5.3% (for 2015Q2) and the long-run equilibrium rate of about 5% (Figure 1). In our supply stocktake earlier this year, we focused mainly on average hours and the participation rate, which reinforces the benefits of re-visiting our judgement on the long-run sustainable unemployment rate.

Our estimate of the long-run equilibrium unemployment rate is consistent with the average rate seen in the 2002-07 period, when the unemployment rate was relatively stable. Our assumptions about the medium-term equilibrium (U*) attempt to capture the potential for the long-term unemployed to exert less downward pressure on wages than the short-term unemployed. That implies a positive wedge between U* and the long-run equilibrium when long-term unemployment is elevated, but one which varies alongside the long-term unemployment rate.

The persistent weakness of wages in 2013 and 2014, as the labour market recovered, pointed to potential downside risks to our estimate of the equilibrium unemployment rate. In August 2014, the Committee reduced the assumed impact of elevated long-term unemployment on U*, narrowing the wedge over the long-term equilibrium. Despite that, wages remained weak relative to their determinants. While other factors, such as compositional shifts within the workforce could explain the weakness of wages, the risks to our estimates of the equilibrium unemployment rate appeared to be weighted to the downside. This note considers some further evidence, which makes the risks look more balanced.

Specific factors influencing equilibrium unemployment

There are a range of specific factors that have potentially shifted the equilibrium unemployment rate over the recession or could do so in the future. A year ago, Bank staff had previously outlined a number of these in a briefing note concluding that it was not clear from the evidence that the natural rate had changed significantly in recent years, because different factors—including a complex web of labour market policy changes—have pushed in opposite directions, and it is difficult to quantify the impact of many of these factors. While this is still our view, it is nevertheless useful to recap on some of the individual demographic and policy factors that could be influencing the path of equilibrium unemployment. Following that, we report results from our updated range of quantitative macro models that we employ to estimate the unemployment gap.

Trends in **educational attainment and demographics** are likely to have decreased unemployment over the recent past; older people and those with higher education levels tend to have lower unemployment rates. Figure 2 illustrates the downward pressure on unemployment from labour force ageing and increasing education levels. In principle, this could reduce equilibrium unemployment by as much as 1pp. While such changes in labour force composition could be expected to exert some downward pressure on long-run equilibrium unemployment, assuming that the simple batting average effect will apply in full probably overstates the effect of compositional changes on the unemployment rate. For example, as the workforce becomes better educated, the job search process will become more congested for high skilled workers unless labour demand changes as well. And within-group unemployment rates are liable to change for the lower educated groups as well if only those with relatively poor job prospects are left once those who are sufficiently capable move on to obtain degrees.[1]

[1] Although the unemployment rate for degree qualified people has normalised following the recession, low educated unemployment rates are currently persistently above their pre-recession norm.

Figure 1: Unemployment August IR estimate of equilibrium

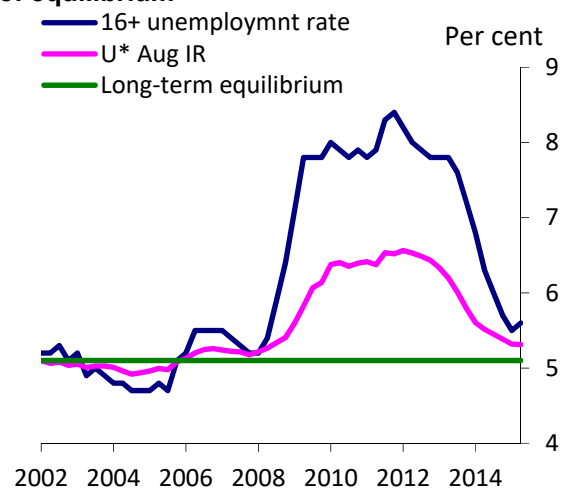

Figure 2: Age and education effects on unemployment

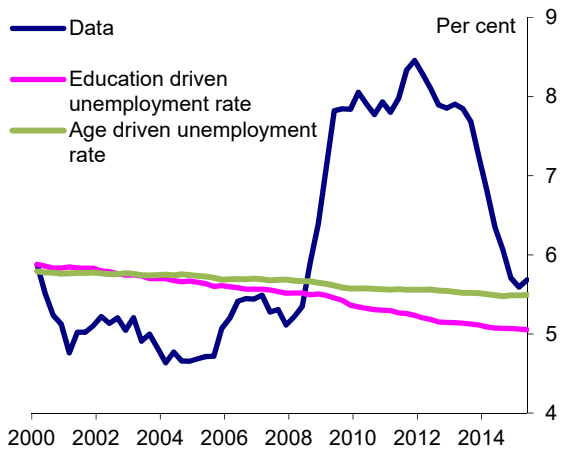

Inward migration might counteract the downward pressure on unemployment from age and education to some extent given the higher unemployment rates that new migrants typically face in the near term. But the share of newly arrived migrants in the aggregate labour force is negligible and the unemployment rate for longer-term migrants is not that different to UK nationals, so the impact is not likely to be large. There may be other effects of migration on unemployment but there is little evidence in the literature of substantial effects.

Overview of policy changes

Since the recession, wider changes in welfare and tax policies (aside from Jobseekers' Allowance) will have broadly increased the financial incentives to participate in the labour market, drawing inactive workers into the labour force. The IFS estimates that the participation tax rate, which measures the proportion of total gross earnings lost in the form of tax and withdrawn benefits by moving into work, is expected to decrease by 1.5pp from 51.3% in 2010 to 49.8% in 2015 although this varies considerably across different types of household.[2] Table 1 decomposes the expected changes in the participation tax rate due to changes in taxes, benefits and Universal Credit.[3] The IFS estimates that the number of people in the labour market could increase by 125,000 to 540,000 in response to financial work incentives although this is subject to considerable uncertainty.[4]

The impact of this increase in labour supply on the equilibrium unemployment rate is unclear. On one hand, previously inactive workers might just end up unemployed for longer periods of time than previous cohorts of unemployed workers. This could be because they are less attached to the labour market and might not have the relevant skills demanded by employers. But, on the other hand, benefit eligibility and generosity have decreased and conditionality has increased. This might incentivise workers to find jobs more quickly than before.

[2] S. Adam and J. Browne (2013), "Do the UK Government's reforms make work pay?", *IFS Working Paper Series*, W13/36.
[3] Replacement rates, another measure of work incentives, show a similar picture to participation tax rates.
[4] Phillips et al (2013), "How might labour supply respond to the changes in financial work incentives", Presentation at the IFS Briefing on Welfare reform, universal credit and supply.

Table 1 – Changes in work incentives due to planned tax and welfare reforms

	2010 level	Change in the PTR due to: Tax changes	Benefit changes	UC	2015
Single, no children	53.8	-0.8	-2	-1.5	51.2
Lone parent	53.6	-0.4	1.3	3.6	56.4
Partner not working, no children	60.4	-0.1	-2.1	-3.4	55.2
Partner not working, children	70.8	-0.5	4.8	-10.7	64.9
Partner working, no children	42.4	-0.6	-1.5	0.1	41.2
Partner working, children	48.4	-0.8	-1.8	2.5	49.4
Without children	50	-0.6	-1.8	-1.1	47.7
With children	53.5	-0.7	0	0	53.6
Non-workers	53.2	-0.5	-1.5	-1	51.5
Workers	50.6	-0.7	-1	-0.6	49.2
Total	51.3	-0.7	-1.2	-0.7	49.8

Source: Adam and Browne (2013)

Figure 3 – JSA benefit sanctions

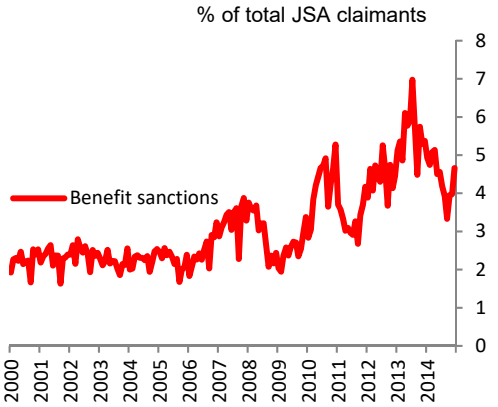

Changes to Jobseeker's Allowance (JSA) are an example of a welfare policy change that might result in an improvement in job search and lower equilibrium unemployment. Conditionality associated with claiming JSA has increased as demonstrated by the increased use of sanctions (Figure 3) as has the substance and availability of training programmes. This might result in quicker and more effective job search. But other reforms have had more limited impact. For example, the benefit cap or the housing benefit reform are not likely to have a significant impact on the unemployment rate in the long-term. And some reforms such as the Employment Support Allowance or the Personal Independence Payment have been delayed. So far, there is limited evidence that points to a significant net effect on unemployment in either direction. The box at the end of the note goes into more detail of specific welfare reforms.

The reforms outlined above—which had been announced prior to the Summer Budget 2015—are planned to be fully carried out by the end of the government's mandate in 2020. Looking forward, there are two important upcoming labour market policy changes: Universal Credit[5] and the "national living wage".

Universal Credit has currently enrolled 3% of all claimants. But once it is fully rolled out by 2020, it is also expected to also increase participation (see Table 1 above). The existing and potential claimants joining the labour market, perhaps due to different skills or less experience, might not be as successful at moving into employment than previous unemployed cohorts. This could reduce the aggregate job finding rate and raise the equilibrium unemployment rate.

The "national living wage", introduced in the Summer Budget 2015, represents an unprecedentedly large increase in the wage floor. The OBR have raised their estimate of the natural rate of unemployment by 0.2pp in response to this. The individual assumptions underpinning this calculation appear to be reasonable, but the fact that firms can respond to the increased cost along several margins (including by adjusting prices or seeking productivity improvements) creates uncertainty about the impact.

But it is possible that welfare policy changes over the next five years have more of a downward impact on unemployment in the long-term. New cohorts of potential claimants might increase their job search effort or might be more effective at finding work particularly if they respond to welfare changes by improving their skills.

In practice, our approach is not to refine our quantitative estimate of the unemployment gap to reflect individual policy changes, although the macro models we use will implicitly reflect changes to the policy environment and other slow moving trends. These approaches are outlined in the next sections.

[5] Universal Credit will subsume the main six benefits (income support, income-based JSA and ESA, housing benefit, working tax credit and child tax credit) into one payment. See reference in footnote 20 for more detail.

Macro-econometric approach 1: Equilibrium matching model

This section details a way of analysing labour market outcomes based on basic equilibrium unemployment theory. Equilibrium is characterised as a steady state in which two conditions are jointly satisfied; (i) the Beveridge curve—that is, a downward sloping relation between steady state vacancies and unemployment, and (ii) a job creation condition that is derived under the assumption of profit maximisation by firms, which effectively pins down the labour market's location on the set of feasible points given by the Beveridge curve. The long-run equilibrium is determined by the intersection of the two curves.

Central to the model is the idea of an aggregate matching function which pairs vacancies with job seekers, standing as a metaphor for the time consuming and costly process characterising job search. The aggregate matching function, M_t, takes as its arguments the measure of job vacancies, v_t, and the measure of job seekers, s_t, and returns the flow of newly created job matches. We extend the standard two-state equilibrium unemployment model—with the two states being employment and unemployment—to allow for unemployment by duration as well as passive job search by non-participants. Accordingly, our definition of job seekers that is relevant for measuring tightness in the labour market is therefore much broader than the conventional unemployment rate. Altogether we then have;

$$s_t = u_t^{<6} + \beta^{<12} u_t^{<12} + \beta^{12+} u_t^{12+} + \beta^i i_t$$

where $u_t^{<x}$ is the unemployment rate of less than *x* months' duration and i_t is the inactivity rate. The β^i parameters are the respective efficiency weights, calculated as the transition probability from the respective base state to employment, relative to the employment transition probability of the short-term unemployed. Therefore, s_t is a measure of effective short-term unemployment (whose search intensity is normalised to unity). We make the conventional assumption of a Cobb-Douglas functional form with constant returns to scale given the considerable supporting evidence in the applied literature;

$$M_t = m v_t^{\gamma} s_t^{1-\gamma}$$

In our current assessment, we will pay particular attention to the parameter *m*, referred to as matching efficiency, in the matching equation which is crucially important to our outlook for unemployment. Persistent changes in *m* can have long lasting effects on the equilibrium unemployment rate (by affecting the unemployment outflow) and can be caused by occupational mismatch, skills erosion, or more generally any factor that is synonymous with 'structural' unemployment.

Estimating the matching function

The applied literature on matching function estimation is immense and there are a number of technical and data issues which are omitted here. The key to our approach is to allow for a reasonably flexible stochastic process for the matching efficiency parameter, *m*, which we assume is time-varying. Given that matching efficiency is unobservable, the matching equation is cast in a state-space setup with matching efficiency as a latent state.[6] We assume that m_t follows a random walk, adopting a stochastic specification that affords us greater flexibility than the assumption of a deterministic time trend that is commonly applied in the literature.

[6] In practice, vacancies are also treated as an unobserved state. Prior to the current ONS monthly vacancies survey, which commenced in 2001, the only available data on job openings was administrative data collected from government employment services which has only imperfect coverage of the aggregate stock of vacancies. This measurement error gets confounded with matching efficiency shocks if not controlled for, and so we explicitly take this into account by treating the administrative vacancies data as a noisy signal about the true supply of vacancies in the period prior to the establishment of the ONS's current survey.

The Kalman smoothed path for matching efficiency is shown in Figure 4. The point estimates towards the start of the sample are subject to considerable statistical uncertainty—this is attributable to the less reliable job openings data that we have prior to 2001 (see footnote 6). We therefore focus on events from the mid-2000s onwards.

The model detects a notable decline in the performance of the labour market just before the latest recession, around 2004 or 2005. As discussed in our *Inflation Reports* at the time, the economy was entering a phase of unusually rapid labour force growth as the unemployment rate edged up from sub-5% to about 5.5% in 2006. This was also the time when migration to the UK from EU accession countries was particularly strong following the 2004 EU enlargement. As the pool of available workers expands, it is natural to expect some—at least temporary—deterioration in the matching process at the margin as heterogeneity increases across labour force participants, particularly if new entrants are less productive and less attractive to hire at the margin.

Figure 4: Matching efficiency (with 95% confidence bands)

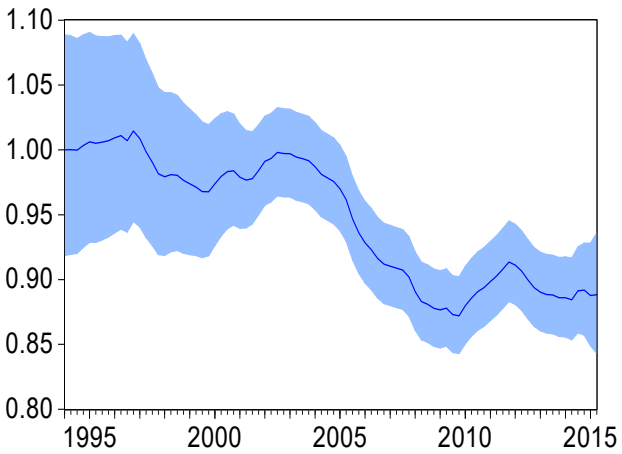

Figure 5: Employment exit rates (% of employment in previous quarter) across selected job types

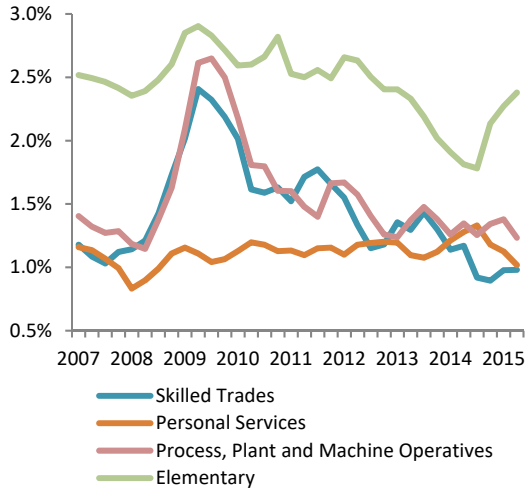

Note: normalised to 1 at the start of the sample

The crucial point is that this apparent dislocation in the labour market occurred prior to the 2008 shock. There is another visible dip in matching efficiency right as the 2008 shock occurred, which reverses by the end of 2011 only to subsequently fall again thereafter. The decline in 2008 is likely to be related to dispersion in job destruction rates across occupational types in the immediate aftermath of the recession (Figure 5), causing a discrepancy between job losers and job openings, with the people leaving their jobs not necessarily being well suited to the jobs available. The double dip in 2012 is again coincident with a rebound in labour force growth. Interestingly—and somewhat reassuringly—the path for matching efficiency that we estimate using aggregate data is fairly consistent with other work on occupational mismatch for the UK based on panel regressions (Patterson et al., 2013).[7]

Overall, the lower level of matching efficiency by 2015q2 implies that job finding rates are currently relatively low conditional on the degree of tightness in the labour market that is observed. This is also clearly evident without a model by simply noting that the job finding rate is about 10% lower than its pre-crisis norm even though labour market tightness has made a full recovery (driven by very robust vacancy supply).

[7] Patterson, Sahin, Topa and Violante (2013), "Mismatch Unemploymetn in the UK", Unpublished.

The Beveridge curve

We are interested in whether there has been a structural shift in the Beveridge curve recently. The Beveridge curve traced out by the data over the past couple of years currently remains somewhat to the right of its pre-2008 location, signalling a potential increase in the long-run equilibrium. Many factors can be responsible for such a shift. For example, unusually high job destruction rates or inflows into the labour force from non-participation raise the unemployment rate for given labour demand, shifting the Beveridge curve outward temporarily. But job destruction rates and flows between unemployment and inactivity have now normalised to levels close to those observed prior to the recession, a period which we judge to be reasonably consistent with long-term labour market stability.

The primary concern is the possibility that the decline in job matching performance highlighted previously has persistently shifted the Beveridge curve outward. Given the current near normalisation of most labour market flows, the only parameter which could displace our estimate of the latest Beveridge curve from its pre-recession location is therefore *m*. Figure 6 shows the extent to which the decline in job matching efficiency as at 2015q2 compared to the 2002-2007 average shifts the Beveridge curve out, effectively bringing the model in line with the latest cluster of data points. The solid Beveridge curve is calibrated using the average value of the matching efficiency parameter over the 2002-2007 period. The dashed version sets matching efficiency at its latest point estimate. The counter-clockwise loop that the data points in Figure 6 trace out during the recession is typical of most business cycles, so some temporary rightward shift is to be expected. The question is whether the Beveridge curve will continue to shift further in from its current position, or whether there has been a structural outward shift relative to pre-recession data.

Figure 6: Matching model equilibrium

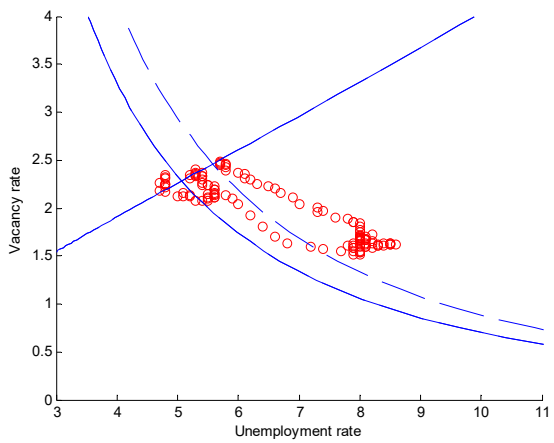

Solid lines are the pre-recession equilibrium (~5% unemployment). Dashed line is the low matching efficiency Beveridge curve (~5.5% unemployment).The job creation curve slopes upwards. Red dots are actual data.

In Figure 6, we have also shown our estimate of the job creation curve, the model's other remaining component, which is briefly outlined next prior to discussing the implications for equilibrium unemployment.

The job creation curve

The second block of our equilibrium unemployment model consists of a forward-looking job creation condition that is derived under the assumption of profit maximisation by firms. We omit the details here, but the key equation is an optimality restriction on the supply of vacancies which relates the cost of creating a vacancy to the expected returns accruing to the firm from employing a worker. In this model, jobs are asset values to the firm and hiring is an investment decision, thereby allowing for a potentially rich set of determinants of aggregate job creation, such as (but not limited to); the firm's discount factor, the survival rate of matches and productivity growth. The end product is an upward sloping line which intersects the Beveridge curve at a unique point, pinning down the model's equilibrium.

In order to estimate the parameters of the job creation condition, we make use of suitable econometric techniques to handle the expectational dynamics arising from the forward-looking behaviour of firms. This requires the use of instruments, which complicates the estimation procedure. Data-based information criteria can help with instrument selection and reduce arbitrariness, but in practice we found that the position of the job creation curve is sensitive to even small changes to the instrument set. Here we report results from a specification which we think is economically and statistically sensible and note that this continues to be an ongoing area of research in SEAD.

Our best estimate of the job creation curve, which is estimated over our full sample of available data spanning 1993-2015, is plotted alongside the Beveridge curve in Figure 6. Given the difficulty of precisely pinning down this relationship, we treat this as our best estimate of long-run structural job supply and ignore potential shifts in the flow cost of bringing a vacancy to market over the estimation sample. As with the calibration of the Beveridge curve, we also have to make assumptions regarding the steady state values of some of the variables which enter the job creation condition. The main assumptions which affect the location of the estimated job creation curve pertain to productivity growth and the firm's discount factor. Higher productivity and a lower interest rate both raise the asset value of a job to the firm, spurring job creation. These variables are calibrated at 2002-2007 averages as a reasonable approximation to long-run stability. In the rest of the discussion, we will abstract from shifts in the job creation curve on the grounds that the data do not appear to suggest that it has shifted materially.[8]

Discussion of equilibrium in the matching model

The model depicted in Figure 6 implies a pre-crisis equilibrium unemployment rate of about 5% for the 16+ population.[9] The decline in matching efficiency that we detect raises the equilibrium rate to about 5.5%. Is this increase likely to persist? The fact that growth in vacancies has stalled in the recent data indicates that the hiring recovery is likely to be nearing its end. Consequently, it seems unlikely that a surge in labour demand presents a material downside risk to unemployment.

The real downside risk is that the shocks to matching efficiency are in fact transitory. As Figure 4 showed, matching efficiency is fairly volatile, and the shocks that we have identified it with may not be permanent. Factors such as occupational reallocation or fluctuations in the size of the labour force do appear to have a significant temporary effect on job finding rates, but are not the type of secular forces that would cause us to change our view of long-run equilibrium in the *steady state* sense. Furthermore, the estimation exercise is sensitive to outliers. There is a marked downward correction in the job finding rate in the latest data (2015q2). Omitting the latest data point reduces the extent to which job matching performance is estimated to have declined, tempering the outward shift in the Beveridge curve.

We should not be too optimistic, however, because matching performance has been weak for several years now and currently shows no signs of recovering, indicating that it may be a factor keeping unemployment above 5%. One fairly robust (and somewhat counterintuitive) finding in the applied literature is worth mentioning here: matching efficiency tends to display a *negative* secular trend (Petrongolo and Pissarides, 2001)[10], contradicting the notion that better search technology facilitates faster job finding rates that some commentators have proposed (O'Kelly and Saunders, 2015).[11] Rather, a gradual move towards increasingly specialised labour markets could actually be lowering average vacancy yields.[12] Unfortunately, the applied literature on matching function estimation has not engaged with reconciling the apparent negative trend in matching efficiency with the widely held belief that the lowest sustainable unemployment rate has fallen over recent decades.

[8] Some of the variables, notably productivity and interest rates, will have shifted in ways that have offsetting effects on the location of the job creation curve.
[9] The model is estimated on data for the 16-64 population due to the availability of flows. We adjust the equilibrium very slightly to obtain a 16+ estimate.
[10] Petrongolo and Pissarides (2001), "Looking Into the Black Box: A Survey of the Matching Function", *Journal of Economic Literature*, Vol. XXXIX, pp. 390-431.
[11] O'Kelly and Saunders (2015), "Where is the Equilibrium Jobless rate?", UK Economics Weekly, Citi Research, 20th March 2015.
[12] In support of this hypothesis, Patterson et al. (2012) find some evidence that labour markets for more advanced skill sets appear to have lower matching efficiency parameters.

Macro-econometric approach 2: Time-varying Phillips curves and the Kalman filter

This section reports estimation results from a filtering exercise based on a state-space model of inflation-unemployment dynamics that conveniently captures time-variation in the NAIRU. The NAIRU defines an unemployment gap that is relevant for determining wage pressure. It is crucial at the outset to stress that the NAIRU is buffeted around by shocks in the short-run and we therefore do not view it as the unemployment rate that the economy is likely to settle at in the long-run. The usefulness of the NAIRU for monetary policy making is that it yields an estimate of the extent to which labour market slack is acting on wage growth, similar in spirit to our conventional measure of *U** which tries to account for the fact that the long-term unemployed may exert less downward pressure on wages than the short-term unemployed. Such models have the advantage that they can take into account wage developments.

We take the simple 'triangle model' of inflation dynamics (Gordon, 1997)[13] as a benchmark specification applied to wage inflation and consider some extensions from the literature which have quantitatively important implications for both the estimated path of the NAIRU and the measured degree of uncertainty around it.

The benchmark triangle model

The triangle model of inflation refers to a simple relation linking (wage) inflation to its three main determinants of inertia, demand and supply. A general representation of this basic structure is

$$\pi_t = a(L)\pi_{t-1} + b(L)(u_t - \bar{u}_t) + c(L)z_t + \varepsilon_t$$

where π_t is the wage inflation rate, $u_t - \bar{u}_t$ is the unemployment gap, z_t is a vector of supply shocks and L denotes a lag polynomial.

The measure of wage inflation that we choose to work with is the rate of change in whole economy average weekly earnings (AWE) due to a relatively long sample of data. The coefficients on lagged wage inflation are constrained to sum to one such that there is a value of the unemployment gap that is consistent with a constant rate of inflation in the absence of supply shocks. As in the original Phillips (1958)[14] paper, we allow the unemployment gap to enter the Phillips curve contemporaneously and with a lag in order to capture the idea that both the level and change in excess aggregate demand affects inflationary pressure in the economy. The vector of supply shocks z_t includes deviations of productivity from a slow moving trend.

In accordance with standard practice, the benchmark model assumes that the law of motion of the NAIRU is simply a random walk;

$$\bar{u}_t = \bar{u}_{t-1} + \epsilon_t$$

where ϵ_t is a normally distributed error term that can, in principle, have a general covariance structure with the error term in the Phillips curve. Free estimation of the variance parameter of ϵ_t is widely known to be subject to technical complications, however, typically resulting in a path for $\bar{u}_t$ that is implausibly smooth. We therefore follow established convention by constraining the signal-to-noise ratio to a value which yields sensible results that can form a suitable basis for policy decisions.

Denote the cyclical component of unemployment by $u_t^c = u_t - \bar{u}_t$. The final equation that completes the benchmark setup is a second-order autoregressive process for u_t^c;

$$u_t^c = \rho_1 u_{t-1}^c + \rho_2 u_{t-2}^c + \eta_t$$

[13] R. Gordon (1997), "The Time-Varying NAIRU and its Implications for Economic Policy", *The Journal of Economic Perspectives*, 11(1), pp. 11-32.
[14] A. W. Phillips (1958), "The Relation Between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861-1957", *Economica.*

The model is estimated by maximum likelihood implemented by the Kalman filter. Prior to discussing the estimation results, we first outline two variants of the benchmark specification that serve as robustness checks on the conclusions we draw about the current location of the NAIRU.

Modifying the benchmark filter

Extension 1: A multiple indicator approach

Estimating the NAIRU with a reasonable degree of statistical precision is a significant challenge which has prompted efforts to find ways of tightening the confidence bands to allow firmer conclusions to be drawn about the cyclical state of the economy. We make use of recent technical advances in the literature that have been shown to reduce the measured degree of uncertainty based on the idea that there is a single forcing variable which drives fluctuations across sectors and macroeconomic variables.[15]

We construct a version of the model by Basistha and Startz (2008)[16] which incorporates fluctuations in the output market as well as the labour market into the otherwise standard model of wage inflation dynamics set out in the benchmark model above. In practice, this amounts to jointly estimating an unemployment gap and an output gap, postulating an Okun's Law relation that links the joint dynamics of the two gaps. This richer time series structure provides better information about equilibrium unemployment compared to the more standard benchmark specification, in the sense that the confidence bands around the point estimates are smaller, and is similar in spirit to other top-down filters that we have used to measure the output gap.

The labour market block remains the same as before. To this, we add an output market block consisting of the following three simple equations;

$$\pi_t^y = a(L)\pi_{t-1}^y + b(L)y_t^c + v_t^1$$

$$\bar{y}_t = d + \bar{y}_{t-1} + v_t^2$$

$$y_t^c = \phi_1 y_{t-1}^c + \phi_2 y_{t-2}^c + v_t^3$$

where π_t^y, $\bar{y}_t$ and y_t^c denote the rate of change in the GDP deflator, trend (or 'potential' output) and the cyclical output gap, respectively. This setup is a parsimonious decomposition of actual real output into a permanent and cyclical component, $y_t = \bar{y}_t + y_t^c$, and it is the cyclical component which is assumed to drive fluctuations in GDP inflation. Potential output is assumed to evolve according to a random walk with constant drift.

The extended model is closed with a restriction that links the joint dynamics of the output and labour market gaps with a dynamic version of Okun's Law;

$$u_t^c = \sum_{k=1}^{K} \theta_k y_{t-k}^c$$

which stipulates that there is a single gap that drives business cycle fluctuations in output and unemployment—that is, shocks do not originate in the labour market, which merely serves to propagate them. Basistha and Startz (2008) demonstrated that this assumption tends to outperform the alternative strategy of incorporating separate, though possibly correlated, stochastic shocks to the unemployment and output gaps individually.

[15] The Fed have also demonstrated an interest in the multiple indicator approach—see Fleischman and Roberts (2011), "From Many Series, One Cycle: Improved Estimates of the Business Cycle from a Multivariate Unobserved Components Model", Federal Reserve Board Finance and Economics Discussion Series.

[16] Basistha and Startz (2008), "Measuring the NAIRU with Reduced Uncertainty: A Multiple-Indicator Common-Cycle Approach", The Review of Economics and Statistics, 90(4), pp. 805-11.

Extension 2: A nonlinear inflation-unemployment trade-off

Linearity is a common assumption in the applied literature on Phillips curve estimation, but one which is also subject to debate. We investigate how a simple piecewise linear Phillips curve affects our estimates of the NAIRU. While it is, of course, possible to allow for more complex nonlinear functional forms, this simple model suffices to illustrate our basic point, which is that a nonlinear Phillips curve alters the path of the NAIRU relative to a linear model over the available sample period in a way that has meaningful policy implications (as will be elaborated on further below). Another advantage of this simple nonlinear structure is that it nests the linear benchmark as an alternative that can be statistically discriminated against straightforwardly.

Estimation results

Time path of the NAIRU

The time paths of the NAIRU implied by our estimated models are shown in Figures 7 and 8. All of the models are consistent with a gradual rise in the NAIRU over the 70s and 80s when unemployment exhibited large, persistent fluctuations. Over the course of the 90s and 2000s, the NAIRU fell back to between 5% and 6%, depending on the model specification. Monetary policy reforms during the 90s and the decline of trade unionisation served to anchor inflation expectations and subdue wage inflation in a manner which would be consistent with a falling non-inflationary unemployment rate. The path of the NAIRU from the models is more or less in line with previous Bank research (Greenslade et. al., 2003).[17] Despite the somewhat different historical paths of the NAIRU from the different models, they all converge to roughly the same point in the latest data (2015q2), around 5.2% to 5.4%, suggesting only a very small unemployment gap remaining. This is consistent with our current estimate of U*, which is 5.3% for 2015q2 (Figure 8).

The conclusion that we take away from this analysis is that it does not support the idea that there has been a material decline in the NAIRU beyond that already incorporated in our estimate of *U**. Figure 9 adds 95% confidence bands to the Kalman smoothed time paths of the natural rate for our three models discussed above. The extended approach which jointly models the output and labour market gaps significantly reduces the degree of statistical uncertainty surrounding the point estimates (by about 30%). Interestingly, the nonlinear model also generates much tighter confidence bands than the baseline.

Figure 7: NAIRU time paths (full sample)

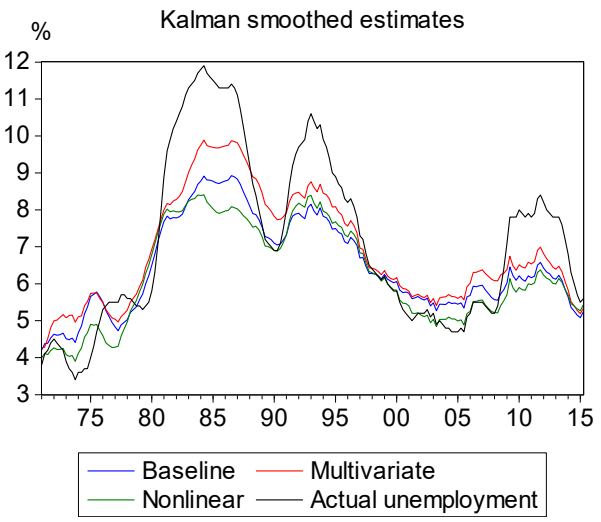

Figure 8: NAIRU time paths and U* (2000-2015)

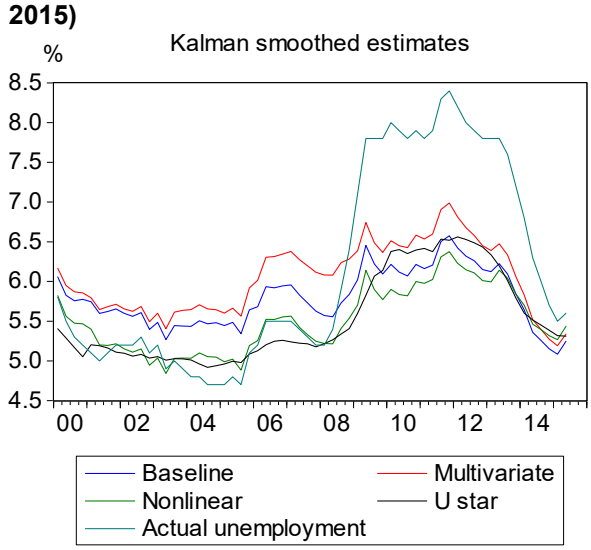

[17] (2003), "A Kalman Filter Approach to Estimating the UK NAIRU", BoE WP. No. 179.

Figure 9: Uncertainty bands around the NAIRU (actual unemployment data in red)

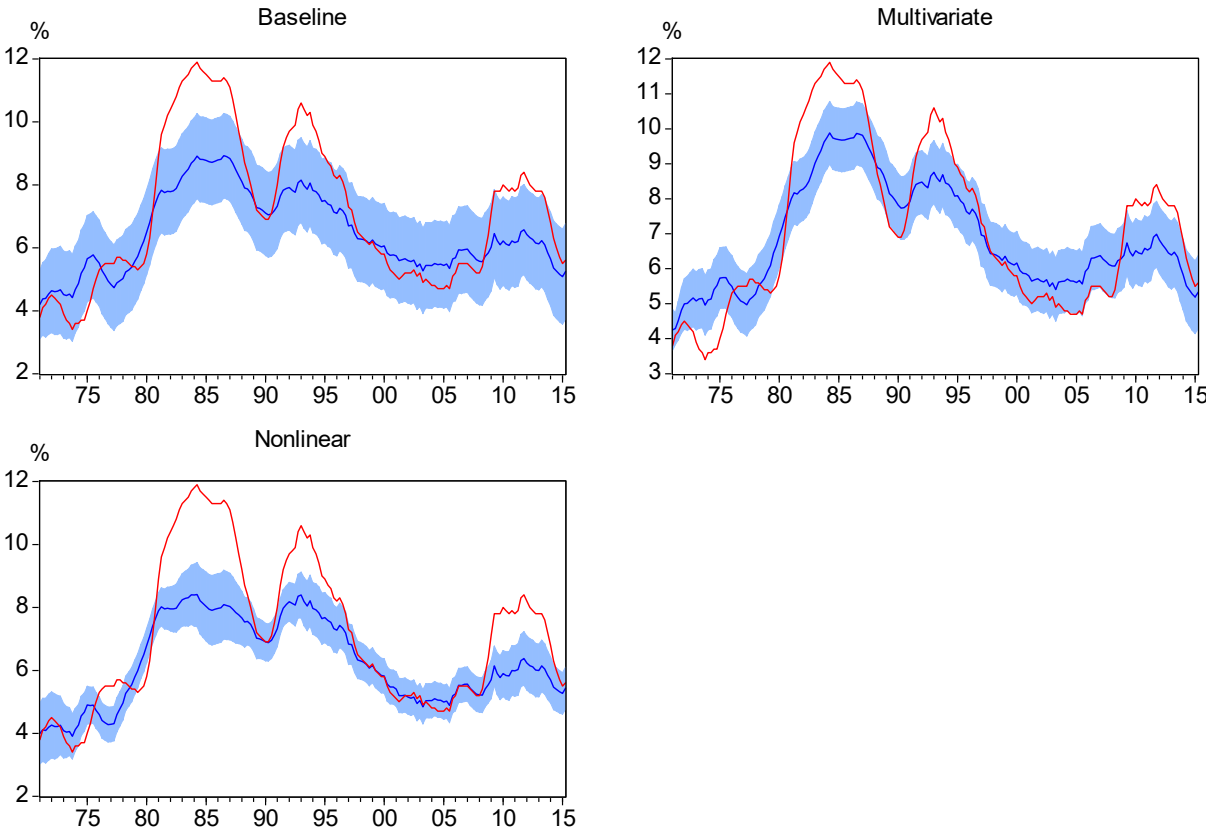

Concluding remarks

In this note, after suggesting that changes in the policy and institutional environment probably had a small impact on the equilibrium rate, we have reported estimation results from two broad classes of macro-econometric models that we use to quantify the unemployment gap. The first, based on a quantitatively implemented version of the workhorse model of modern macro-labour theory, and the second, based on statistical filters, both yield the conclusion that a substantial unemployment gap—say, on the order of a percentage point or so—seems unlikely.

While it is always difficult to disentangle a persistent shock from a structural shift, particularly on the small sample sizes we have available, gathering all of our evidence together we do not find suggestive macro-econometric evidence that there is significant scope for unemployment to fall below 5%. In the medium-run, it is likely that any further potential falls in unemployment will be slow and protracted, relying solely on falls in long-term unemployment.

These estimates are uncertain and there are risks in both directions. The full implementation of welfare policies such as Universal Credit as well the introduction of a national living wage could potentially exert upward pressure the equilibrium unemployment rate. But increased benefit conditionality, lower generosity and possible behavioural responses such as higher or more effective job search efforts might cause downward pressure on the equilibrium rate.

Box 1 – Overview of specific policy changes

This box goes into more detail on specific welfare policy changes and their possible impact on the equilibrium unemployment rate.

Welfare design has an impact on structural unemployment by determining job search incentives. Recent reforms to **Jobseekers' Allowance** (JSA) that have increased the stringency of these entitlements could potentially lower the equilibrium unemployment rate due to increased job search effort and improved training. Conditionality has increased markedly, with job seekers now having to sign a 'claimant commitment' that sets out what they have to do while claiming JSA. The sanctions for failing to comply with the requirements can make claimants ineligible for up to 3 years. This is stricter than in the past, as illustrated by the increased number of sanctions (Figure 3). New training schemes have also been introduced for JSA claimants. From 2011, the Work Programme replaced existing programmes with the purpose of improving the likelihood of long-term unemployed moving into and staying in work. But the evidence so far has not suggested that the new programmes have been more effective than the previous ones.[18] [19] It is worth noting that the coverage of JSA as a benefit has increased because those with a potential disability but assessed as fit-for-work—as well as some lone parents—now claim JSA instead of other benefits. This means that they are now required to look for work and also have the training programmes at their disposal. It is likely that JSA, through tighter conditionality, has slightly increased job search effort, potentially acting to reduce the equilibrium unemployment rate.

The effectiveness of the reforms in encouraging labour force participation is likely to differ across specific policies. In some cases, reforms which affect a large number of people may end up having only a small impact because of small behavioural responses. For example, the cap on the maximum amount a family can receive for **housing benefit**, introduced in 2011, is estimated to have affected around 600,000 social-rent households in 2013-2014. But only 14% of those affected started looking for work as a response to the cut.[20] The most common response has been to decrease spending. Since around two-thirds of housing benefit claimants also claim other benefits, it is possible they are less attached to the labour market and less likely to actively search for work.

Figure 10 – Outcomes of new ESA claims

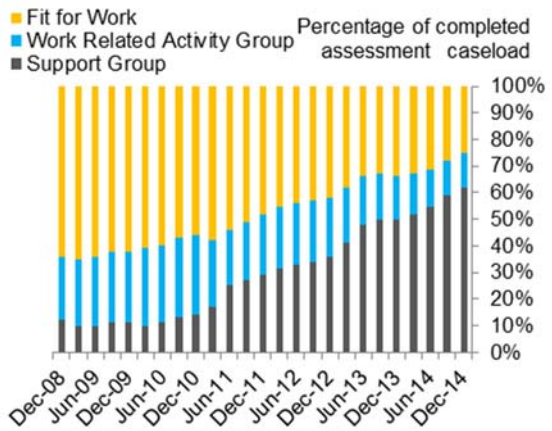

And in other cases, implementation of benefit reform has been hampered by administrative delays. **Incapacity Benefit** (IB) was replaced by the **Employment Support Allowance (ESA)** in 2008, but the roll-out to existing claimants was delayed until 2011 and progress has been slow since. Around 1.5 million people were expected to be migrated from IB on to the newer ESA benefit between 2010 and 2014. The latest available figures suggest that around a quarter of those reassessed were found fit for work which was lower than expected. Eligibility criteria for new claimants were also loosened in 2011, relative to initial plans, resulting in a decline in those assessed as fit-for-work or capable of taking steps towards moving into work ("work-related activity" group) (Figure 10). But those who are in the work-related activity group are effectively moved onto JSA where job search requirements have been intensified.

[18] The previous programmes include the New Deal, Employment Zones and the Flexible New Deal.
[19] Report by the Comptroller and Auditor General (2014), "The Work Programme", *National Audit Office*.
[20] A. Hood and D. Phillips (2015), "Benefit Spending and Reforms: The Coalition Government's Record", *IFS Briefing Note, BN160.*

Disability Living Allowance (DLA), the single largest disability benefit, provides another example of slow implementation. It has been growing rapidly from 2.0 million claimants in 1997-98 to 3.3 million in 2013-14, and the planned reform is to replace DLA with a **Personal Independence Payment** (PIP) that contains stricter eligibility criteria, time limits and automatic reassessment. But the transfer of existing claimants from DLA to PIP, which was due to start in late 2013, has been put on hold for most people until late 2015 and is expected to take three years. There is also some preliminary evidence that tighter eligibility criteria might not have been as strict as expected.

Policies that are likely to achieve the most success are those which target groups that have particularly elastic labour supply decisions. Lone parents, for example, are one of the groups most sensitive to financial incentives to work.[21] From 2008 to May 2012, the age of the youngest child enabling receipt of **Lone Parent Benefit** had been lowered to 5 years, so that lone parents with older children have to claim another benefit.[22] Impact assessments have found the reform led to a 7pp increase in employment, or a rise of 30,000, amongst those affected. Those who moved into work either moved directly or from a short spell on JSA. Those affected who moved onto **Employment Support Allowance** (ESA), a general benefit for people unable to work or with a disability, were less successful in finding employment. While successful, these numbers are small, and so this reform is likely to have had a negligible impact on equilibrium unemployment.

[21] R. Blundell, C. Crawford and W. Jin (2014), "What can wages and employment tell us about the UK's productivity puzzle?", *IFS Working Paper*, W3/11.
[22] S. Avram, M. Brewer and A.Salvatori (2013), "Lone Parent Obligations: An impact assessment", *Department for Work and Pensions*.