



BANK OF ENGLAND

Speech

Monetary policy and the damaged economy

Speech given by

David Miles, External Member of the Monetary Policy Committee, Bank of England

At the Society of Business Economists Annual Conference, London

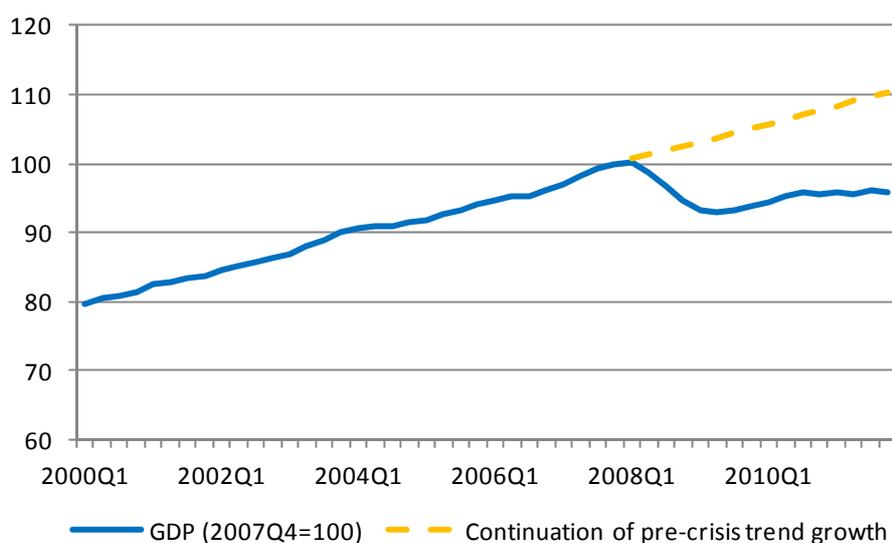
24 May 2012

I would like to thank Jochen Schanz and Arpad Morotz for research assistance, and others for helpful comments. The views expressed are my own and do not necessarily reflect those of the Bank of England or other members of the Monetary Policy Committee.

Introduction

In the first quarter of this year, the level of GDP in the UK was about 4% lower than it was just before the banking crisis which came in the second half of 2008. A wide gap has emerged between GDP and where it might have been in the absence of the crisis: had growth followed its pre-crisis trend of a bit above 2.5% per year, GDP now would have been about 10% higher than in 2008, and not 4% lower. The size of that gap is about 13% of trend GDP (Figure 11).

Figure 1: GDP, and how GDP might have evolved in the absence of the banking crisis



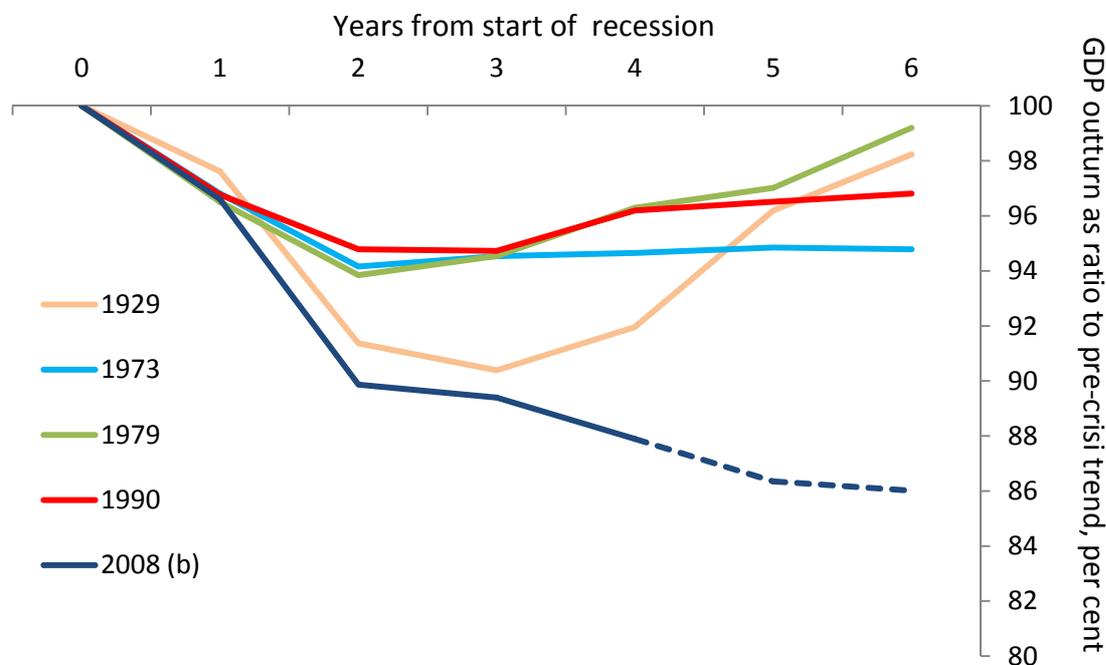
Source of data: ONS, Bank of England calculations. Notes: The trend is assumed to be linear in log GDP and is estimated over 1971Q1 – 2007Q4. The 'continuation of pre-crisis trend growth' line has been constructed by shifting the trend upwards so that it coincides with GDP in 2007Q4. Last observation is 2011Q4.

This is considerably larger than the gap that had opened in the UK four years into the Depression of the 1930s. Unless growth picks up very sharply to well above trend over the next couple of years the size of the output gap – at least based on this sort of fairly crude trend-fitting – will become even larger relative to the other recessions of modern history (see Figure 2).

But how much of that apparently huge gap actually reflects a fall in the amount the UK *could* produce rather than a chasm between actual production and the economy's potential output? That question is central to monetary policy. Potential output is also directly relevant to the welfare of the country – it is the income it can earn without resources being under-utilised or squeezed too hard. By "squeezed too hard" I mean that if output exceeds potential there will be triggered some adjustments – e.g. workers might require higher wages to continue providing a higher level of hours which might for a short while be provided at unchanged rates of pay. If actual output nears potential output, inflationary pressures generated domestically are likely to rise. But when a wide gap exists between potential and actual output, inflation pressures will tend to be weak.

While a number of other factors affect inflation, the distance between potential and actual output, sometimes referred to as slack, is a crucial determinant of inflation.

Figure 2: Evolution of GDP from the start of the 2008 and previous recessions relative to the respective trends^(a)



Source of data: ONS, Bank of England calculations. Notes:

- (a) This chart shows the deviation of GDP from a time-varying trend path. That trend is estimated using a Hodrick-Prescott (HP) filter on log GDP from 1700 to 2007. The lines show the difference between trend GDP and actual GDP, where trend GDP during the crisis is assumed to continue to grow at its estimated rate immediately prior to the crisis. This growth rate is 1.5% for the 1929 recession, 2.3% for 1973, 1.6% for 1979, 2.4% for 1990. As HP trends are sensitive to end-sample observations, it has been assumed that trend GDP would have grown by 2.5% since 2008, even though the unadjusted HP filter gives a higher growth rate. Higher estimated trend growth rates tend to lead to larger estimates of the gap.
- (b) Forecasts for 2012 and 2013 are the central projections taken from the Bank's May Inflation Report.

Estimating potential output is therefore key for monetary policy makers. But it is also difficult. And that difficulty is acute when, as now, the economy has been buffeted by huge shocks that have taken it a long way from what seemed to be its steady trajectory. One indication of this difficulty is that different measures of slack in the UK are now giving very different signals about the level of potential output.

Weak GDP growth and fairly stubborn inflation, along with evidence from surveys of firms, suggest that the amount of slack is considerably smaller than the more than 10% gap between the level of GDP and its pre-crisis trend. That evidence should make one doubt estimates from simple extrapolation of past trends; there is likely to be less slack than simply focusing on GDP relative to trend suggests. But while the amount of slack is probably lower than just fitting trend lines to past activity suggests, it is not likely that renewed growth in GDP would immediately create substantial inflationary pressures. That is not just because slack is still likely to be substantial today; it also reflects the likelihood that effective capacity is itself likely to expand

with demand and activity. Just as weak growth in GDP drags down potential output, strong growth tends to increase it. This considerably complicates the task for monetary policy: inflation, output, and potential output are interdependent, and all three are affected by monetary policy.

I want to explore this interdependence in more detail, and consider what it implies for how monetary policy should be set now. I'll first discuss how potential output may have evolved since the crisis started, and how it might evolve in the future under different assumptions for GDP growth. I'll then turn to what that means for monetary policy. I want to explain why there is a case for making monetary policy more expansionary, even when inflation has surprised us repeatedly on the upside.

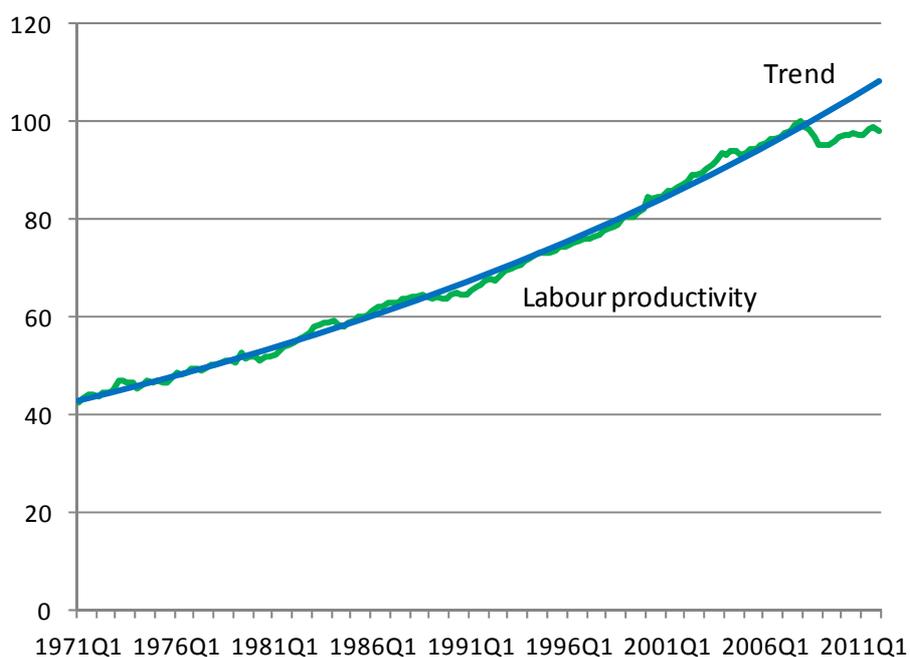
Labour productivity and labour supply

One way to estimate potential output is to decompose it into two factors, labour supply and labour productivity. Labour productivity measures what a given amount of labour can produce on average. Multiplying this by the supply of labour yields potential output.

Labour productivity

Figure 3 shows the level of labour productivity, measured as GDP over total hours worked, together with a simple estimate of its trend since 1971.

Figure 3: Labour productivity and an estimate of its pre-crisis trend



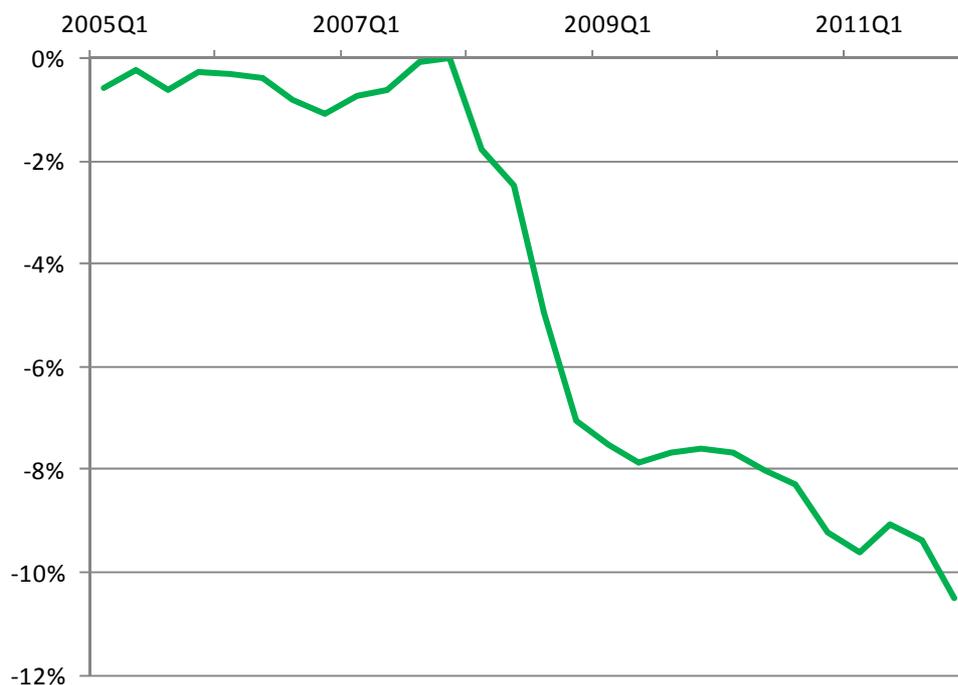
Source of data: ONS, Bank of England calculations. Notes: Labour productivity is an index of GDP (Source: ONS) over total hours worked (ONS Workforce Jobs series), where 2007Q4=100. The trend is the estimated long-run relationship in a VECM between GDP and hours worked, which is estimated over 1971Q4 – 2007Q4 (see Annex A1 for details). The final observation is 2011Q4.

The trend – which is assumed to be linear in the log of labour productivity – describes the long-run evolution of labour productivity reasonably well; but only up to the start of the banking crisis. Since then, productivity has fallen massively below trend (Figure 4).

If labour productivity had followed its pre-crisis trend, it would now be about 10% higher. This is almost as much as GDP has fallen behind its pre-crisis trend. But it seems unlikely that underlying labour productivity, and with it potential GDP, has declined permanently by almost as much as actual GDP. Labour productivity in the past has shown considerable cyclical – by definition non-permanent – variability. A simple statistical (time-series) model of labour productivity suggests that in the past about 10% of the gap between labour productivity and its trend is closed in each quarter (see Annex A1 for details). This suggests that when productivity departs from its trend, almost 60% of the gap to trend is typically closed after two years¹.

¹ If 10% of a gap is closed each quarter the amount of a gap closed after 8 quarters is $(100\% - (100\% - 10\%)^8)$.

Figure 4: Deviation of labour productivity from its pre-crisis trend



Source of data: ONS, Bank of England calculations.

On the basis of these results, we should expect labour productivity growth to pick up strongly and very soon. Indeed, on this view it is puzzling why it has not done so already. One answer is that we had not experienced a banking crisis of similar severity during the past forty years, the period over which the typical catch up of labour productivity to its smooth trend has been estimated. A related reason is that GDP has not yet recovered: labour productivity growth and GDP growth are strongly positively correlated.²

It is very hard to untangle the causal mechanisms at work behind the correlation between GDP growth and the growth in labour productivity – is labour productivity very weak because GDP growth has been very weak, or is weak labour productivity the cause of weak GDP? One way to try to figure out what is going is by estimating vector autoregressions – that is time series models that allow for the interdependence between GDP and productivity. I have used that type of model to explore the relationship between GDP and labour productivity (see Annex A1). The model suggests that GDP growth will affect the path for labour productivity.

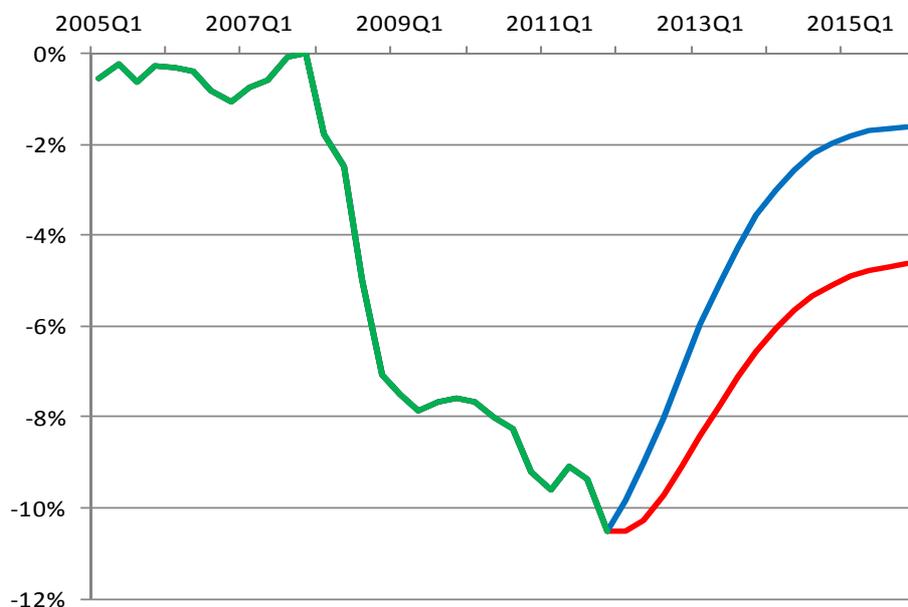
The red line in Figure 5 shows the evolution of labour productivity that the regressions predicts if GDP growth from now to 2015 was zero.³ Because labour productivity has a tendency to close the gap between its current level and its trend, this gap tends to close even when GDP growth is zero. By the end of 2015, the shortfall – currently about 10% – would have declined to about 4.5% of the estimated trend for labour

² The observation that faster GDP growth can increase productivity is sometimes referred to as 'Verdoorn's Law'.

³ The chart shows the log of labour productivity data until 2011Q4, and then uses the regression equation reported in Annex A2 to forecast how log productivity will evolve under the two scenarios for GDP growth.

productivity. But the adjustment would be considerably faster if annual GDP growth was positive. The blue line illustrates this under the assumption that annual GDP growth was 3% for the next three years. Here, the gap would have declined to about 1.5% by the end of 2015.

Figure 5: Estimated evolution of labour productivity for different scenarios for GDP growth



Source of data: ONS, Bank of England calculations. Notes: Red line: 0% GDP growth p.a. during 2012-15; blue line: 3% annual GDP growth.

This simple exercise suggests that while labour productivity is likely to pick up over time, even if GDP growth remains weak, faster GDP growth is likely to increase the speed of the recovery of labour productivity. If these results are a guide to where we are now in the UK it has implications for inflation and for monetary policy. Stronger labour productivity would pull down unit labour costs unless it is fully matched by higher wages. This means that stimulating GDP growth with monetary policy does not have to lead to strong inflationary pressures.

How reliable is this evidence? There are certainly plausible reasons why measured labour productivity may be temporarily depressed in a recession and those go beyond the apparent empirical regularity that periods when labour productivity seems to fall far below trend are followed by productivity bouncing back towards trend. It is not simply a case of blindly believing past correlation for which there is little plausible interpretation based on business behaviour. The costs of closing and re-opening capacity can be high. Rather than firing staff immediately, firms may accept that their employees work less efficiently in a downturn. For example, keeping a shop open may require the presence of a minimum number of cashiers and shop assistants, even if they serve fewer customers. But when business picks up the productivity of those staff rises fast.

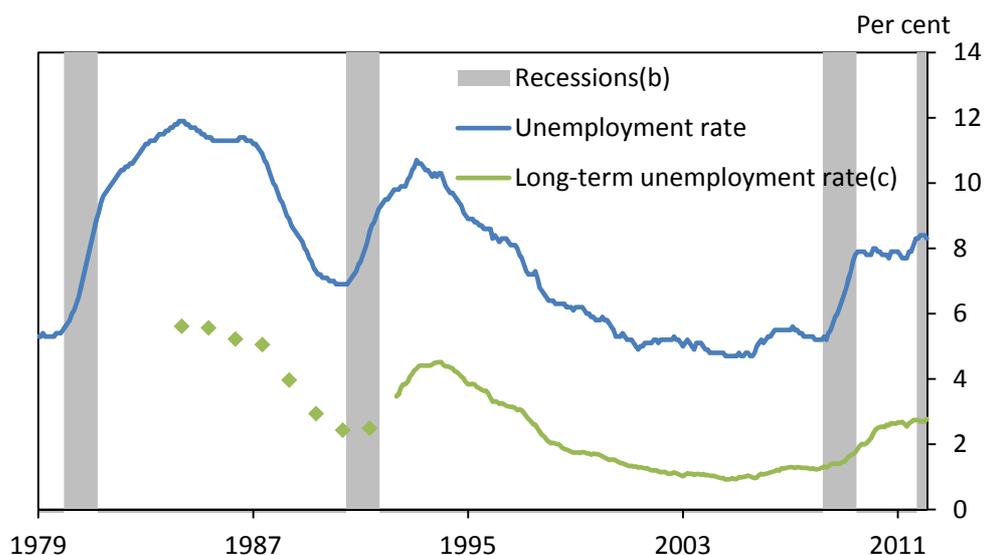
But all the quantitative results described above are based on historical correlations. They only inform us about the future if we assume that the current recession is structurally similar to previous ones, even though it is larger. But the structure of this recession *is* different: it was triggered by what was probably the largest banking crisis in the UK's history. The results therefore have to be treated with care.

Labour supply

If labour supply were to be much lower because of the crisis then even if labour productivity recovers completely the hit to productive capacity could still be very large. That would mean that effective slack in the economy could be very much lower than the time series evidence on productivity seems to suggest. Employment has fallen since the start of the crisis, and long-term unemployment has picked up (Figure 6). The rise in the jobless is much less than would have been expected based on the fall in output – which is another way of posing the puzzle of the collapse in productivity relative to past trends. And indeed jobs in the private sector have increased fairly substantially over the past year – which is one sign that spare capacity may be less than the trend lines described above suggest. Nonetheless long term unemployment may increase further if short-term unemployed are not able to join the labour market soon. The increase in long-term unemployment means that there is a risk that a part of the labour force has lost skills, and may struggle to find employment even when growth picks up. In some sense the effective labour supply has shrunk.

But other factors work in the opposite direction. There is a large pool of short-term unemployed, most of whom are likely to be willing to work at existing wages if they were able to find a job, and many of whom may be *more* willing to work as a result of being unemployed. Stronger net immigration might also help to contain wage growth once prospects improve here.

Figure 6: Unemployment and long-term unemployment rates^(a)



Source: ONS (including the Labour Force Survey). Notes:

(a) Rolling three-month measures unless otherwise stated.

(b) Recessions are defined as at least two consecutive quarters of falling output (at constant market prices) estimated using the latest data. The recessions are assumed to have ended once output began to rise.

(c) Defined as those people who have been unemployed for more than twelve months divided by the economically active population. Data prior to 1992 are based on seasonally unadjusted, annual LFS microdata. These annual observations correspond to the March-May quarter.

And there are other factors which are likely to lead to more labour supply. The hit to wealth that many households have felt means that people with a target level of savings (for example to finance retirement spending) may both aim to work longer and possibly consume less. Currently both real share prices and real house prices are about a quarter below their pre-recession peaks (Figures 7 and 8).

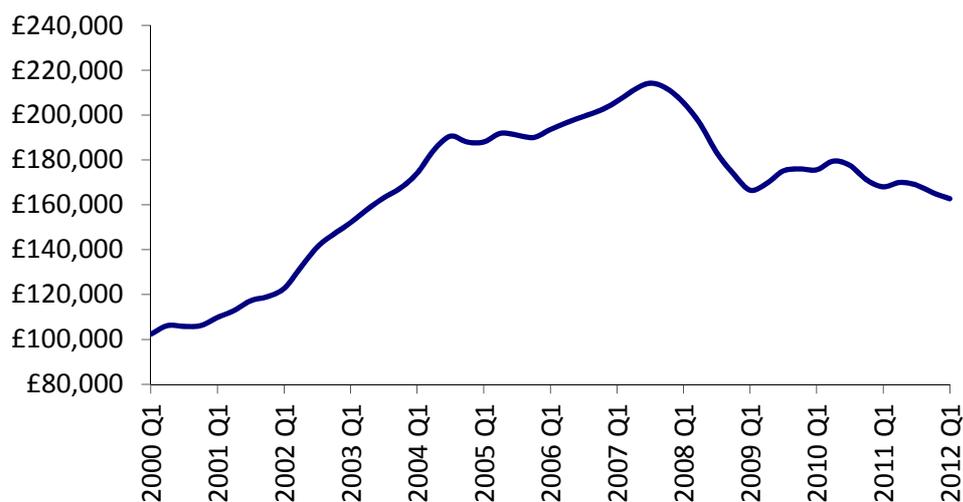
The factors that increase labour supply offset – and maybe more than offset – the negative impact of higher actual unemployment upon future effective labour supply.

Figure 7: UK equity prices – FTSE All-share index adjusted for inflation



Source: Thomson Reuters Datastream, Bank of England calculations. Notes: Monthly averages of the FTSE All-share index, converted to current prices using the Consumer Price Index (CPI).

Figure 8: UK house prices adjusted for inflation



Source: Nationwide, Bank of England calculations. Nominal house prices, converted to current prices using the Consumer Price Index (CPI).

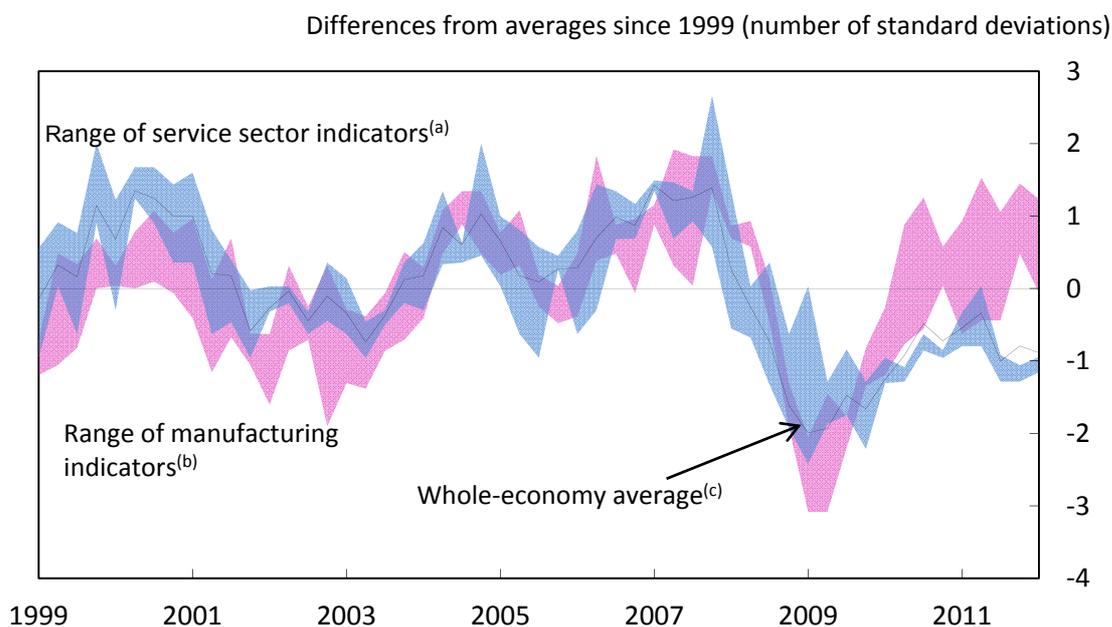
Looking at the overall forces that are affecting labour supply and labour productivity I think it is plausible that potential output may have fallen in the aftermath of the crisis, but not to the same extent as GDP. The decline in labour productivity could reverse quickly when GDP growth picks up, and any additional demand for labour could probably be satisfied without substantial increases in wages. In short, there is probably substantial slack. Yet some people believe that spare capacity is, despite the big falls in GDP and low growth of recent years, very small. Surveys of spare capacity, and the path of inflation itself, lend some weight to that. But how should we interpret the survey evidence and what light does CPI inflation throw on spare capacity within the UK?

Arguments for little spare capacity: surveys and persistent above-target inflation

Surveys

Firms are regularly asked about whether their spare capacity has changed. At face value these surveys suggest that on the whole firms in the UK now have little spare capacity (Figure 9), though there is some variation across sectors.

Figure 9: Survey measures of capacity utilisation by sector



Sources: Bank of England, BCC, CBI, CBI/PwC and ONS. Notes:

- (a) Includes measures of services capacity utilisation from the Bank's Agents, BCC and CBI. The Agents data are end-quarter observations. The CBI measure weights together financial services, business/consumer services and distributive trades surveys using shares in nominal value added. The BCC data are non seasonally adjusted.
- (b) Includes measures of manufacturing capacity utilisation from the Bank's Agents CBI, and a measure of non-services capacity utilisation from the BCC. The Agents data are end-quarter observations. The BCC data are seasonally unadjusted.
- (c) Average of the number of standard deviations from averages since 1999 of three survey measures of capacity utilisation for services and manufacturing calculated by weighting together those measures described in (a) and (b) by their shares in nominal value added.

In principle this is powerful evidence – it comes directly from companies. But the survey answers are somewhat ambiguous; they can be interpreted in different ways. It is not entirely clear what respondents to the surveys understand by 'full capacity'. Some might believe that it refers to their scope to expand production in the very short term. That immediate scope might be limited. Some manufacturing companies, for example, might have temporarily set aside some of their productive capacity and laid off some of their employees. They may not be able to re-activate mothballed capital quickly. Hiring labour also takes time.

But those firms may still be able to expand production very substantially within the longer time frame of several quarters over which monetary policy typically affects the economy.

And many firms might see their staff busy trying to do business and win contracts but with less success than in more normal economic times. Estate agents, advertising executives and car salesmen may have to work more now on the average sale. That shows up as poor productivity and would not typically show up as spare capacity in a survey. But if demand were stronger the productivity of those workers would improve. This phenomenon may be more relevant in services than in manufacturing – though even within manufacturing firms there will be people working on sales and marketing.

Finally, the survey results are difficult to translate into quantitative estimates of spare capacity. The reason is that the survey questions usually only generate qualitative responses: firms typically indicate whether output is below capacity or not. The aggregate responses only tell us how many more firms work below capacity than at or above capacity. So a zero balance could be the result of half of all respondent working just below capacity whereas the other half operates considerably above capacity. Or, a zero balance could occur when half of all respondent work considerably below capacity, whereas the other half work just above capacity. Clearly, there is much more spare capacity in the second case; but that will not show up in the survey balances.

Persistent inflation

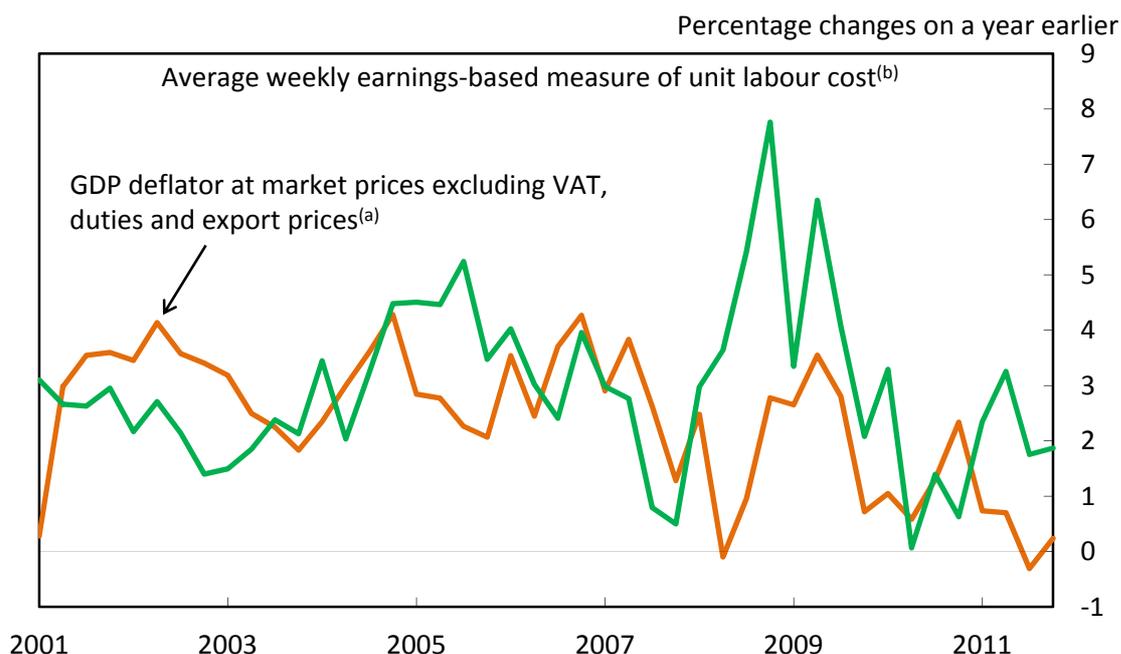
Besides survey responses there is another, less direct, bit of evidence that seems to support the view that there is much less spare capacity in the UK than the deviation in GDP or in labour productivity from trend suggests: inflation has not fallen back as quickly as many people – including the Monetary Policy Committee – forecast.

Spare capacity is likely to affect what one might call domestically generated inflation pressures; cost pressures arising from outside the UK are not at all likely to be affected by slack within the UK economy (though the extent to which they are passed on into domestic prices might be). The main reason why inflation over the past few years has fairly consistently been above the forecasts made by the MPC is that externally generated cost pressures turned out higher than had seemed most likely: sharp increases in energy prices, rising import prices, in part due to the depreciation of sterling, and the increase in VAT all played a role. Domestically generated inflationary pressures are likely to have been much lower, and may even have been negative, dragging down headline inflation (Figure 10). Yet those forces restraining inflation have been drowned by the forces driving up inflation. Why has the impact of slack not been greater?

One factor that could explain why slack has not held inflation lower is that impairment to credit makes it harder for some firms to respond to spare capacity by cutting prices (or raising them less than costs) to gain

market share. Take a firm with spare capacity that considers reducing the prices of its products, attempting to attract a competitor's clients. If the firm reduces its prices, initially its revenue and profit are likely to fall: its existing customers take advantage of the price reduction while the firm is still waiting for new clients to arrive. When bank credit is ample, the firm may be willing and able to accept a likely initial decline in revenue and profits which is compensated for by higher revenues later. That strategy may not be feasible when bank funding is harder to get.

Figure 10: Measures of domestically generated inflation



Sources: ONS (including the Labour Force Survey) and Bank calculations. Notes:

- (a) Export prices include a contribution from imported components, so this measure uses an estimate of export prices excluding the contribution from import prices.
- (b) Calculated using average weekly earnings data, adjusted using the ratio of National Accounts compensation and wages and salaries data, and divided by output per worker.

Let me summarise these thoughts on the evidence of the degree of spare capacity and on the impact of a given degree of spare capacity upon inflation pressures. First, both the degree of slack and the impact slack has on inflation are not possible to measure directly – we have to rely on indirect evidence. Second, one powerful set of indirect measures – the path of GDP and of labour productivity relative to trend – suggest there is lots of spare capacity. But, third, surveys suggest slack may not be so large. Fourth, while the inflation evidence is to a significant extent a reflection of cost pressures not really affected by UK spare capacity, I think it does suggest that the degree of slack may be having less of an impact than usual, a phenomenon that could be linked to the disruption to credit.

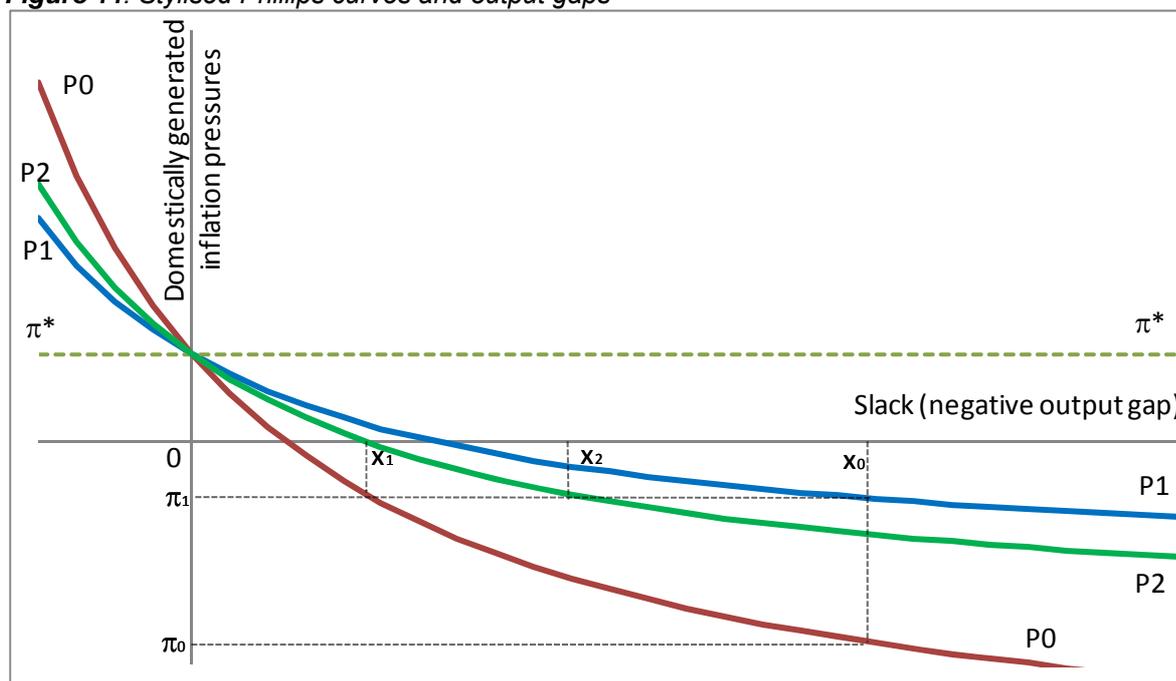
What to make of all this? My own (Bayesian) view is that rather than interpret the recent slightly greater inertia in inflation as simply reflecting much less slack in the economy I would ascribe a significant part to a

smaller downward impact of a given amount of slack on inflation pressures. It might look like splitting hairs to apportion more of an explanation to a weaker impact of spare capacity and less to an erosion of capacity. But I think this makes a difference to monetary policy. If spare capacity does have a weaker impact on domestically generated inflation pressures it changes the costs and benefits of bringing inflation back to target faster or slower. In the next section I want to consider this in more detail.

Policy implications

To draw out the policy implications of changing assessments of the degree of spare capacity and of its impact on inflation let me use a simple graph (Figure 11), whose shape is familiar to all economists. This graph shows a relation between a measure of inflation pressures (on the vertical axis) and a measure of slack. In its original form, as developed by Phillips, the measure of slack was unemployment. In my version (and in many current economic models) slack is the output gap – here defined as the difference between potential output and actual output. The relation between an amount of slack and a degree of inflation pressures is reflected in the curves in the diagram – let me call these (abusing slightly his original meaning) Phillips curves.

Figure 11: Stylised Phillips curves and output gaps



Note: π^* is the inflation target. Inflation is equal to target if there is a zero output gap, all agents expect inflation to coincide with the target, and there are no external price shocks so that inflation equals domestically generated inflation pressures. We assume that changes in the slope of the Phillips curve do not affect long-term inflation expectations.

Assume that one initially believed that the relationship between slack in the economy – that is potential output minus actual output – and domestically generated inflation is described by the red curve in Figure 11 (labelled P0). One would then believe that a big output gap of x_0 would be associated with very low

domestically generated inflation of π_0 . But suppose our best measure of inflation pressures turned out to be higher than that, say at π_1 .

This could be rationalised either by assuming that the red curve is correct but that there is substantially less slack (of x_1). Or, one might think that the red curve is wrong, and that the relationship between slack and inflation pressure is better described by the flatter, blue curve (P1). In this case, observed inflation pressure of π_1 would be consistent with an unchanged assessment of the output gap.

Discriminating between these rationalisations of higher than anticipated domestically generated inflation requires judgement, because we can't observe directly either the output gap or the slope of the Phillips curve; and indicators of slack may be giving us conflicting signals. A Bayesian would make some adjustment to the estimate both of slack and the slope of the curve. That is what I have done over the past year or so. I conclude that there is less slack than I would have thought – given the big falls in GDP, weak growth and very poor labour productivity – but that the Phillips curve is also somewhat flatter. This is illustrated by the shift from the red to the green curve (labelled P2), where surprisingly high inflation of π_1 is associated with a smaller impact of slack on inflation and also a lower output gap, x_2 .

All else equal, monetary policy should be tightened when inflation surprises on the upside and the output gap is estimated to be narrower than thought earlier. But all else is not equal here if the Phillips curve has indeed flattened. A flat curve has two implications that are essentially just two sides of the same coin. On one side, tighter monetary policy will have less of an impact on bringing inflation back to target. A lot of spare capacity is needed to reduce inflation quickly – and a lot of spare capacity means underutilisation of capital and more unemployment, which is costly in terms of welfare. On the other side, stimulating demand will put less pressure on inflation. So the price of bringing above-target inflation back down quickly is high, while the cost of more expansionary policy – which means a slower trajectory of inflation coming back to target – has fallen. As a result, the optimal path of monetary policy is likely to entail a slower return to the inflation target. It may even be optimal to make policy more expansionary even as one's estimate of spare capacity is lowered. I discuss this more formally in Annex A2 within the framework of a standard model of monetary policy.

There is another powerful reason for accepting that inflation should return to target gradually if we start from an unhappy position of above target inflation and significant slack, as I believe we do today. The longer output lies below potential, the more likely potential output itself declines – perhaps permanently. Hysteresis means that the assumption of an unchanging path for underlying capacity is a poor one. Workers lose some of their skills during long spells in unemployment. Long-term unemployment has already picked up during the crisis. The pool of qualified labour from which firms can hire then risks shrinking. The more quickly economic activity recovers, the smaller the risk of substantial declines in potential output.

Conclusion

Accurate assessment of the degree of spare capacity is difficult. There are good arguments for why potential output has fallen significantly since the financial crisis. But there are also reasons for believing that the impact of slack on inflation is lower – the Phillips curve is flatter. Some of the reasons for the decline in demand and activity and the flattening of the Phillips curve have the same source in impaired supply of credit which makes it harder for firms to respond to slack by cutting prices to slowly gain market share.

None of this is certain. The existence of substantial slack, of a flatter Phillips curve and of a high degree of dependence of productive potential upon demand (hysteresis) are all uncertain. But I believe they are consistent with the evidence and that they make an exceptionally expansionary monetary policy appropriate. No one on the MPC feels comfortable with the prolonged and substantial overshoot of inflation above its target level. But that does not mean bringing inflation back to target very rapidly is the best thing to do. In a situation where weak demand is likely to be having a negative impact upon productive capacity the cost of having a tighter monetary policy to bring inflation back to target fast will be some long lasting damage to incomes.

Annex A1: Time-series properties of labour productivity

I estimate a 2-equation Vector Error Correction Model (VECM) between log GDP (ln y) and log hours (ln h).

Formally,

$$\begin{aligned}\Delta \ln y_t &= \gamma_y (\ln y_{t-1} - (\delta_0 + \delta_1 t + \delta_2 \ln h_{t-1})) + \alpha_1 \Delta \ln y_{t-1} + \alpha_2 \Delta \ln y_{t-2} + \alpha_3 \Delta \ln h_{t-1} + \alpha_4 \Delta \ln h_{t-2} + u_{y,t} \\ \Delta \ln h_t &= \gamma_h (\ln y_{t-1} - (\delta_0 + \delta_1 t + \delta_2 \ln h_{t-1})) + \beta_1 \Delta \ln y_{t-1} + \beta_2 \Delta \ln y_{t-2} + \beta_3 \Delta \ln h_{t-1} + \beta_4 \Delta \ln h_{t-2} + u_{h,t}\end{aligned}$$

where

- the long-run relationship between GDP and hours is assumed to be log linear with a trend and a constant. δ are the coefficients in the long-run equations.
- γ are the coefficients in the adjustment equations for hours and output on the deviation of from the long-run relation between log GDP and log hours.
- α are the coefficients of the short-run effects on GDP growth
- β are the coefficients of the short-run effects on hours growth

The difference between the two equations is the growth in labour productivity:

$$\Delta \ln y_t - \Delta \ln h_t = \Delta (\ln y_t - \ln h_t) = \Delta \ln (y_t / h_t)$$

Ignoring the short-run parts of the VECM, the difference between the equations for $\Delta \ln y_t$ and $\Delta \ln h_t$ is given by

$$\Delta \ln (y_t / h_t) = (\gamma_y - \gamma_h) (\ln y_{t-1} - (\delta_0 + \delta_1 t + \delta_2 \ln h_{t-1}))$$

If we restrict $\delta_2 = 1$, then we can also interpret the long-run residual as the deviation of labour productivity from its trend, and $\gamma_y - \gamma_h$ as the speed of adjustment to this trend.

For the VECM specification to be valid,

1. both log GDP and log hours need to have a unit root (i.e., the first difference of log GDP and the first difference of hours need to be stationary). An ADF test (not reported) does not reject the hypothesis that the log levels of these variables have a unit root.
2. In addition, log GDP and log hours need to be cointegrated. A Johansson test for cointegration (not reported) does not reject the hypothesis that they are.

The estimation results are in Table 1.

Table 1: Estimation results

Long-run residual:

Change in log GDP	1
Constant	-5.12
Time trend	-0.005**
Change in log hours: restricted to	-1

Error correction part:

	Change in log GDP	Change in log hours
Long-run residual, lagged once	-0.014	0.086*
Change in log GDP, lagged once	-0.066	0.06
..., lagged twice	-0.035	0.038
Change in log hours, lagged once	0.505*	0.27*
..., lagged twice	0.087	0.28*
Constant	0.006*	-0.0003

Sample: 1971Q4 2007Q4. *denotes significance at the 5% level. An LR test on the restriction just rejects it at the 5% level. We nevertheless proceed with the restricted model as the unrestricted coefficient is 1.28, which appears unrealistically high. R-squared is 9.4%.

The difference coefficients on the error correction term is 0.10; this is our estimate of the speed with which labour productivity closes a gap from its long run trajectory. The forecasts for labour productivity 2012-15 are computed by using actual data until 2011Q4, imposing a path for GDP on the model for 2012-15, and allowing hours worked to vary assuming that their behaviour is correctly described by the second equation of the VECM. (That is, we set the residuals to that equation to zero throughout the forecast horizon.)

A2: Optimal monetary policy when the estimated Phillips curve and output gap change

The idea that the speed with which inflation should be brought back to target through setting monetary policy depends upon both size of the output gap and the sensitivity of inflation pressures to the gap can be formalised using results from the literature on optimal monetary policy. Woodford (2004) provide one such formalisation, based on earlier results in Woodford (2003) and Woodford and Giannoni (2003). I briefly describe the model from Woodford (2004) and the optimal monetary policy that it implies.

Suppose that prices in the economy are fixed one quarter in advance and that prices are only re-optimized at random intervals (following Calvo's original idea). Between occasions on which prices are optimally set they are just automatically indexed to the aggregate price level.

With fixed capital and no labour market frictions this generates an aggregate supply relation of the form:

$$\pi_t - \pi_{t-1} = kE_{t-1}(x_t) + \beta E_{t-1}(\pi_{t+1} - \pi_t) + u_{t-1} \quad (1)$$

π_t is the inflation rate in quarter t

x_t is the output gap (output – potential output)

E_{t-1} is the expectations operator

u_{t-1} is an exogenous shock to inflation

Let the objective of monetary policy be to minimise the weighted sum of all future expected inflation distortions and squared deviations of output from its neutral (or natural) level – that is the square of output gaps (x_{t+i}). In this particular model where we assume indexation of prices as a default for those prices not optimally re-set the inflation distortions come from changes in the rate of inflation. Let the weight upon output deviations relative to that on inflation distortions be λ .

The size of the output gap depends on real expenditure which we assume is sensitive to interest rates (or more generally to the setting of monetary policy). It is through influencing the size of the output gap that monetary policy affects the rate of inflation; parameter k in equation (1) reflects the power of that effect and it ties down the slope of the Phillips curve.

Woodford shows that with these assumptions optimal policy is set such that the following condition holds:

$$E_t(\pi_{t+1} - \pi^*) = -\frac{\lambda}{k} E_t(x_{t+1}) \quad (2)$$

π^* can be thought of as the optimal long-run inflation target. This condition says that monetary policy should be set so that inflation should be expected to fall back towards the target rate in a way that depends on the likely size of the output gap next period ($E_t(x_{t+1})$) multiplied by the ratio of the relative importance of the cost of non-zero output gaps to inflation distortions (λ) to the sensitivity of inflation to the output gap (k).

For illustration let us assume that $\lambda = 0.04$ and $k = 0.2$. The value of λ means that much more weight is placed on inflation distortions than the loss from non-zero output gaps in setting monetary policy. (In Woodford and Giannini (2003) a value of $\lambda = 0.048$ is used). The value of k means that the effect of a 5% output gap (a very large amount of slack) would be making the quarterly inflation rate fall by about 1 percentage point ($= (0.04/0.2) * 5\%$). This is a rapid rate of fall in inflation since it means the annualised quarterly rate would be falling by 4%.

Suppose initially we start with inflation significantly above target and we have set policy optimally such that the expected output gap next period is 5% ($E_t(x_{t+1}) = -0.05$). Given these assumptions the optimality condition implies that monetary policy is being set so that in the following period the likely quarterly inflation rate will be 1 percentage point above the target level. If the long run target for the annual inflation rate were to be 2% then the quarterly target rate, π^* , would be 0.5%. Under the optimal monetary policy, the expected quarterly rate of inflation next period would be 1.5% – an annual rate of 6% inflation. What is happening here is that the output gap is so large that the optimal strategy is not to tighten policy so much as to get inflation quickly back to 2%. Rather we allow the quarterly rate next period to be much higher than the target – knowing that the force of the output gap on inflation will be bringing it down sharply further ahead.

Now suppose we are on this optimal policy path and then obtain new information, in one scenario, the new information is only about the output gap; and in another scenario, about the output gap *and* the impact of the output gap on inflation. How might the monetary policy response be different in each scenario?

If we thought the likely size of the output gap was much smaller – in fact halved so that with the same monetary policy setting as before the level of $E_t(x_{t+1})$ would be -0.025 rather than -0.05 – the optimality condition implies that the expected overshoot of the inflation target also be half the size; that is an annualised rate of 4% ($=4*(0.5\% + 0.04/0.2*0.025)$) rather than 6%. In order to reduce inflation from the previously optimal rate of 6% to 4%, monetary policy would have to be tightened: if we left policy unchanged, expected inflation next period would not be 4%, but above 6% because we have reduced our estimate of the output gap. This makes intuitive sense: with less slack in the economy the forces driving down inflation are weaker and also the loss from having unused capacity is smaller. Both those factors mean that monetary policy should be tighter and the second of them means that we should be willing to accelerate the return of inflation to target.

Now consider the second scenario. Here we reduce our estimate of the output gap as above (so that at unchanged policy the estimated size of the output gap next period, $E_t(x_{t+1})$, is 2.5% and not 5%) and simultaneously reduce our estimate of the impact of an output gap on inflation, k . (In terms of Figure 11, we believe that the Phillips curve has flattened (green line), and that the output gap is somewhat lower at x_2 .) Suppose we sharply cut our estimate of k – reducing it by more than half. At unchanged policy this would mean the right hand side of equation (2) rises, since the proportionate decline in the absolute size of the output gap is more than offset by the increase in the ratio λ/k . That would mean that unchanged monetary policy would need to generate a higher rate of inflation than 6% were it to remain optimal.

In other words the combination of less of an output gap – which in itself requires less of an inflation overshoot to the target and a clear tightening in monetary policy – is offset by the extra price that has to be paid in terms of more output losses from bringing inflation down faster. The net result is likely to be that a slower path of returning inflation to target is optimal. And it is possible that the optimal monetary policy is

looser when there are simultaneous falls in both the estimated size of the output gap and in the likely responsiveness of inflation to that gap. A sufficiently large decline in k will always make a loosening of monetary policy optimal, even if the reassessment of the size of k comes with a simultaneous decline in the estimated size of the output gap.

References

Woodford, M (2004) "Inflation Targeting and Optimal Monetary Policy", *Federal Reserve Bank of St Louis Review*, July/August, pages 15-41.

Woodford, M and Giannoni, M (2003) "How Forward looking is Optimal Monetary Policy?", *Journal of Money, Credit and Banking*, vol. 35, no 6, pages 1425-1469.

Woodford, M (2003) Interest and Prices: Foundations of a Theory of Monetary Policy, *Princeton University Press*.