



BANK OF ENGLAND

Speech

Conditional guidance as a response to supply uncertainty

Speech given by

Ben Broadbent, External Member of the Monetary Policy Committee, Bank of England

At the London Business School, Regents Park

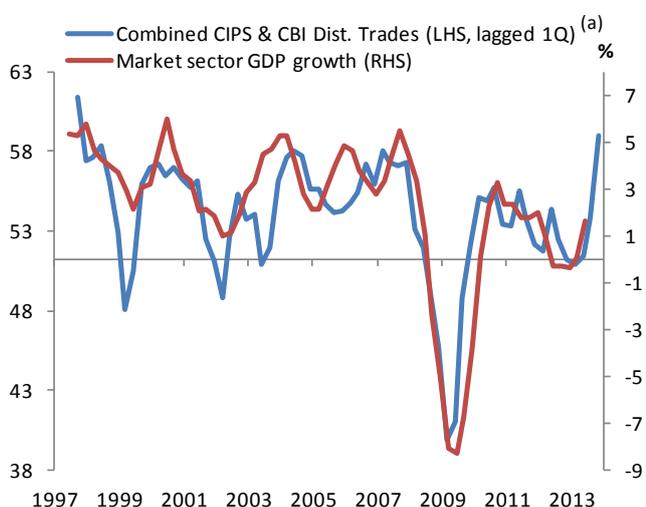
Monday 23 September 2013

I would like to thank Alina Barnett, Adrian Chiu and David Barkshire for research assistance, and I am also grateful for helpful comments from other colleagues, especially Spencer Dale, Gabor Pinter, Matt Waldron and Martin Weale. 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.

Conditional guidance as a response to supply uncertainty

The economy is growing. It's hard to say exactly how fast it's growing. Official third-quarter GDP estimates aren't published for another month and will be subject to significant revision for some time afterwards. But what output data we do have for the quarter are strong¹. And if you take survey readings as a guide – the red line in Chart 1 is a weighted average of the sectoral CIPS surveys – they suggest output in the market sector of the economy is expanding at an annualised rate of 5% or so. This may be an over-estimate. But the economy has clearly picked up significantly faster than the majority of forecasts (including those of the MPC) made last year².

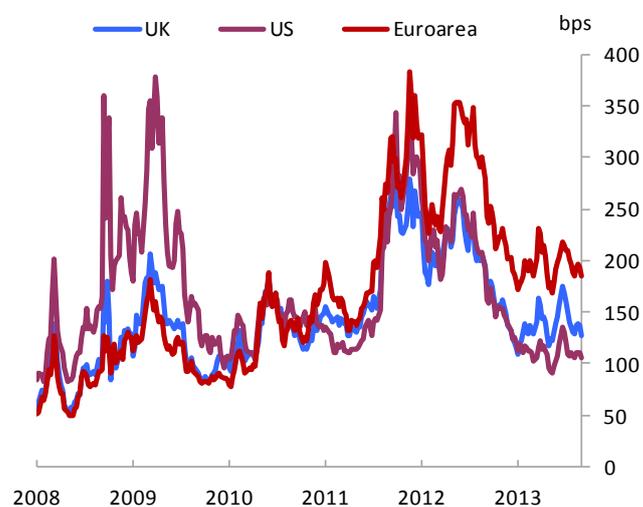
Chart 1: Surveys suggest significant pick-up in growth



Source: ONS, CBI, Markit Economics, Bank calculations

(a) Weighted average of the Markit/CIPS PMIs and the CBI Distribution Trades surveys.

Chart 2: Follows large drop in banks' funding costs through second half of 2012^(a)



Source: SNL Financial, Datastream and Bank calculations.

(a) CDS spreads of banks weighted by assets as of end 2012.

Some commentators have wondered how it can have done so. Citing high debt levels in the private sector and ongoing fiscal contraction in the public sector, they have also expressed doubts it can continue. But these have not proved insuperable barriers to growth in the past. The economy has grown through similar-sized reductions in government borrowing before now³. And the clear pattern after other financial crises, here and elsewhere, is that economies turn up while private-sector debt ratios are still falling⁴.

¹ Taken together, early estimates have output in manufacturing, construction and retail distribution – 30% of the non-oil private-sector economy – 1.5% higher in July than (the monthly average) in 2013Q2.

² At the end of last year the central consensus and MPC forecasts for growth over the following year (to 2013Q4) were 1.2% and 1.6% respectively. The latest forecasts are 2.1% and 2.4%. My guess is they'll both have to be revised up further in the next few months.

³ The ratio of the cyclically-adjusted government deficit to GDP has fallen by 4.3% pts over the past four years. Since 1965, when the OBR estimates begin, there have been three other (non-overlapping) episodes with a similar degree of fiscal tightening: financial years 1974-78, when the deficit:GDP ratio fell by 4.0% pts and GDP growth averaged 1.9%, 1978-1982 (4.2% pts and 0.5%) and 1993-97 (4.6% pts and 3.7%). We haven't quite reached the end of the four-year period to 2013-14 but it looks as if GDP growth will have averaged around 1.3% during that time.

⁴ Broadbent (2012a), especially Charts 20a and 20b.

Some acceleration in activity was always likely in 2013. Underlying growth in 2012 was masked by exceptional contractions in construction and oil output that were unlikely to be repeated. More importantly, the various central bank interventions through the course of last year were followed by a big decline in risk premia and funding costs across the advanced economies, not least for their banks (Chart 2). This wasn't exactly the same thing as a cut in official interest rates. But it nonetheless amounted to a sharp easing in financial conditions, one that might have been expected – after the usual lag of around a year – to have boosted real activity, here and in other parts of the developed world.

It is easy to be wise after the event. As I explained in a speech earlier this year, forecasters generally anticipate only around a third of changes in economic growth one year ahead, and we can be no more confident that the recovery can continue smoothly at this rate than in the view that it would never arrive. The survey readings in Chart 1 are near their series highs and, even if only from a statistical perspective, it would be reasonable to expect a degree of moderation from here. Domestic deleveraging pressures may not be enough to keep the economy in permanent stagnation but they nonetheless provide a headwind to growth. While they may have diminished, the underlying imbalances in the euro area, and the risks they entail, have not gone away.

And if the path of actual output is inevitably unpredictable, we are also facing an unusually uncertain outlook for potential output – the capacity of the economy to meet demand without generating inflationary pressure. That capacity depends centrally on how productivity evolves. And because productivity growth has been unusually poor, even allowing for the weakness of economic activity, it is hard to be sure how it will behave as the economy recovers. (You can see the scale of the puzzle in Chart 3, which plots five-year average growth rates in GDP and employment. I included the same graph in a speech on productivity last year.) This is a central feature – perhaps the central feature – of the economic landscape in the UK and it is in this light, I will argue, that we should see the policy of “conditional guidance” recently introduced by the MPC.

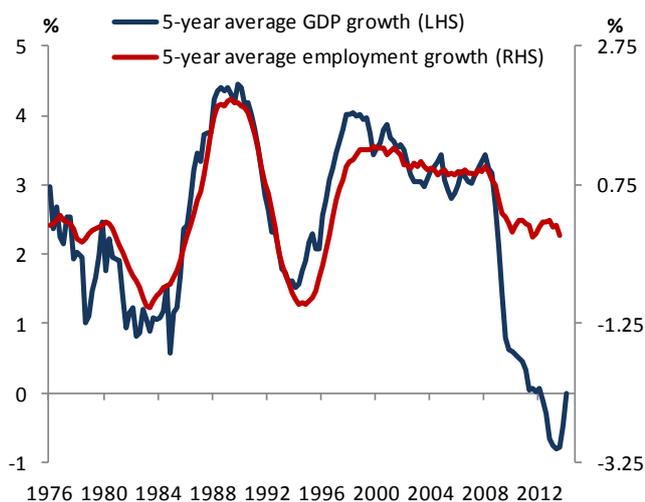
This policy has a number of virtues, not least the reassurance to investing businesses that monetary policy will not rise until economic recovery has become firmly entrenched. What I want to do in this speech, specifically, is to explain why it is appropriate to measure that progress with the rate of unemployment: conditioning policy on the labour market is the natural response to higher uncertainty about the path of potential output.

Productivity revisited: explanations for the recent past

Not everyone is afflicted with such doubts. Since touching on the subject in a speech last year, I've read a number of pieces arguing that the solution to the productivity puzzle is “simple”: it's “obviously” the result either of labour hoarding by firms and/or greater flexibility in pay. And because these are just cyclical phenomena, it must follow that productivity will rise, along with aggregate demand, as the economy recovers.

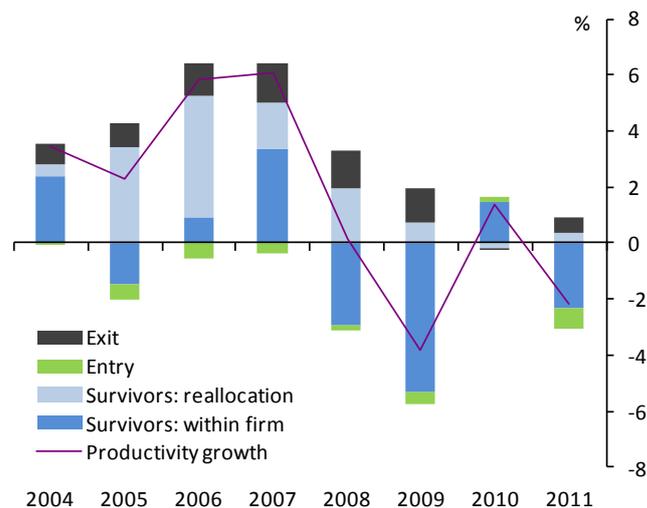
These factors help to explain why measured productivity is, in general, cyclical. They have almost certainly been at work in recent years. Some employees will have been retained in expectation of a pick-up in orders in the future, rather than to meet any today. Relative to the 1970s and 1980s, at least, real pay does indeed look more flexible than in the previous downturns.

Chart 3: Productivity unusually weak even allowing for scale of downturn



Source: ONS, CBI, Markit Economics, Bank calculations

Chart 4: Productivity growth involves significant reallocation across firms^{5 6}



Source: ONS and Bank calculations.

But it's a real stretch, on any closer examination, to explain the divergence in Chart 3 using these things alone⁷. If, more generally, "aggregate demand" alone is responsible for weak growth it's hard to explain why domestic cost and price growth has also been much stronger, relative to output, than in the past. And, as I suggested last year, there are also positive indications that the economy has been slow to adapt to significant shifts in relative economic conditions, across different sectors and firms. This will have reduced effective supply. What I want to do now is to present a little more evidence for this channel. I will then qualify it by admitting that, even to the extent it matters over the recent past, it may not tell us much about how productivity behaves in future.

⁵ Chart 4 covers the entire population (~2m per year) of UK private non-financial companies, excluding those in the agriculture, mining and utilities sectors. Decomposition methodology is based on Baily, Bartelsman and Haltiwanger (2001). See Barnett, Barriel, Chiu and Franklin (2013), *Bank of England Working Paper*, forthcoming.

⁶ This work contains statistical data from the ONS, which is Crown Copyright. The use of the ONS statistical data in this work does not imply the endorsement of the ONS in relation to the interpretation or analysis of the statistical data. This work uses research data sets that may not exactly reproduce National Statistics aggregates.

⁷ On wage flexibility, for example, the impact on productivity is likely to be around one half the impact on pay itself – see note 1 in the Appendix. While real pay does look to have fallen by more than in the pre-1990s recessions, all else equal, the gap is only around 2%, smaller relative to the data in the recession in the early 1990s, and may in any case be due (at least in part) to a drop in underlying, effective supply. So my guess is that extra wage flexibility can account for no more than 1% point of the productivity shortfall, probably quite a bit less. The arguments against a big increase in labour hoarding – the strength of gross hiring, the pessimism of employers about future orders and the absence of any rise in hiring costs – I went through last year. They still apply.

Like Roman Gaul the evidence is in three parts: a general point about the importance of cross-firm reallocations in the economy, a bit on what relative price changes tell us and some results on the behaviour of investment. The first and third come from work led by colleagues here at the Bank⁸.

Look first at Chart 4. This uses individual firm-level responses to the ONS surveys of output and employment to break down changes in aggregate productivity into various component parts – in dark blue, those that occur within individual firms; the effect of moving employment across existing firms, from the less to the more productive (light blue); the contribution of new entrants (green) and that of firm exits (in black).

There are a few interesting things here. First, firms leaving the population are on average less productive than others (that's why the black bars are positive). This is why the surprisingly low level of firm exits in and since the last recession may not be an unmitigated good. Second, in the (brief) period before the financial crisis, only a minority of productivity growth took place within existing firms. Most was due to reallocation of one sort or another, principally across the existing population of companies. Third, that reallocation has since slowed down.

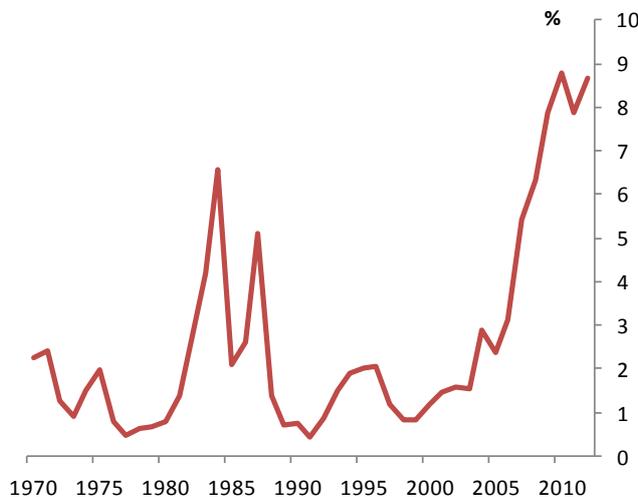
This is hardly conclusive evidence. For one thing, the dark blue bars too have fallen significantly since the crisis, so it's not as if slower reallocation, measured this way, is the only thing going on. Also, this is just a mechanical, accounting breakdown: we can't be sure that we can legitimately imagine changes in one set of bars without affecting the others. And the sample only covers a few years. Ideally, one would want to know how the same breakdown looked in earlier recessions.

The other two pieces of evidence do cover earlier recessions. As I explained a year ago, shifts in relative demand won't affect aggregate productivity as long as productive resources move seamlessly in response. But if that does occur, relative prices won't have to change either: it's the fact that capital can move from one sector or firm to another that keeps marginal costs (and therefore prices) stable in the face of relative demand shifts. That's why Chart 5 is interesting. It plots the variance of real prices, across 29 sectors of the economy, compared with their pre-crisis trends.⁹ And it shows a significant increase in the dispersion since the crisis. This suggests marginal costs too have become more dispersed. To the extent this is due to (relative) demand shocks, it would necessarily imply a misallocation of resources and a loss of aggregate productivity.

⁸ Barnett, Barriel, Chiu and Franklin (2013), *Bank of England Working Paper*, forthcoming.

⁹ There's an important technical point here. If relative prices aren't mean-reverting, at least relative to a fixed trend, you'd expect to see increasing divergence over time and the pattern in Chart 5 would be an artefact of the way we'd constructed the data. But according to statistical tests, this is not the case: nearly all these relative prices are, indeed, "trend stationary" prior to 2008.

Chart 5: Significant increase in price dispersion^(a)



Source: ONS, EU KLEMS and Bank calculations

^(a) Size-weighted variance of real prices, relative to pre-crisis trends, across 29 sectors. See appendix for details.

Table 1: Unlike in US, UK investment less sensitive to profitability than in earlier recessions

Q: Did investment respond positively to rates of return?		Pre-crisis	2007 recession	1980's & 1990's recession
UK				
Firm-level	ONS Annual Business Survey	Yes	No	-(^a)
	Thomson Reuters WorldScope	Yes	No	-(^a)
Industry-level	ONS sectoral data	Yes	No	Yes
US				
Industry-level	EU KLEMS sectoral data	Yes	Yes	Yes

Source: Bank of England

^(a) “-” indicates periods during which data were not available.

Indeed, one can have a go at translating the one into the other (see note 2 in the Appendix): to a first order approximation, the loss in productivity would be around a half the variance in prices. Taken at face value, therefore, Chart 5 suggests that, compared with the pre-crisis period, slow reallocation might have knocked up 3%-4% off aggregate productivity¹⁰.

The suggestion here is that the crisis may have impaired the economy's ability to channel finance away from areas where demand and profits have fallen and towards those where they've both risen. The third piece of evidence, presented in Table 1, gets at this more directly. Using three separate datasets for the UK, and one for the United States, it presents results of firm and industry-level regressions of investment on profitability. It shows that the two were positively related in the pre-crisis period UK data, including in prior recessions, but not in the more recent period. The same has apparently not happened in the United States where, as it happens, productivity growth has been much less weak than in this country.

¹⁰ This is the effect on the assumption that there are only demand shocks in the world, whether at the aggregate or sub-aggregate level. In practice, the dispersion of prices could reflect other things, including cross-sectoral shifts in supply. However, it would be odd to rule out the possibility of supply shocks at an aggregate level, as an explanation for weaker productivity, only to reintroduce them at a sectoral level. That's why the appropriate “null” hypothesis is a world without such supply shocks.

Why we are uncertain about productivity over the future

I find this evidence reasonably persuasive. I also think it's significant that our experience is far from unique: it is generally the case that, during and immediately after big financial crises, and even allowing for the depth of the downturns, productivity growth is lower than in other recessions. It seems only reasonable to infer that this might have something to do with the performance of the financial system specifically. We cannot hope to understand this effect if we stick rigidly to a view of the world in which variations in output are caused only, always and everywhere, by changes in "aggregate demand", treatable with the traditional tools of monetary and fiscal policy.

But that doesn't mean this is the only thing going on – it's surely likely that demand and supply-side factors have both contributed to weak productivity growth – or that one can be at all confident about the precise scale of these various effects. And even if one could be, I'm not sure whether it would help much in forecasting how productivity behaves over the future. Suppose there were a sudden (and independent) improvement in demand. To the extent it's been weak only because of labour hoarding, or more flexible pay, productivity would automatically improve¹¹. But operating profits in the financial sector would also grow, as would the cash-flow and confidence of self-financed small businesses. So, although it might take a little longer, one should probably expect these re-allocative frictions, such as they are, to dissipate as well.¹² Either way, effective supply would, to a degree, react "endogenously" to demand¹³. That's why I'm happy with the MPC's central forecast, in the latest Inflation Report, that foresees an acceleration in productivity – and a correspondingly gentle decline in the rate of unemployment – as the economy recovers. (On that note, it's worth pointing out that, even after growth of around 2% in the number of private-sector jobs over the past year – a rate that the latest surveys suggest we might reasonably expect to be maintained for a while yet (Chart 6) – the rate of unemployment has fallen by only 0.3% points. This is because employment in the public sector is still falling and, more importantly, the active workforce is growing – see Chart 7¹⁴).

My main reaction to all this, however, is to feel very cautious about the likelihood of any particular path for potential output. There is a close connection between our ability to understand the past and the confidence we should allow ourselves about our predictions of the future. If we don't really have a good explanation for the way something has behaved in recent experience we should surely be less confident about our ability to forecast it over the future.

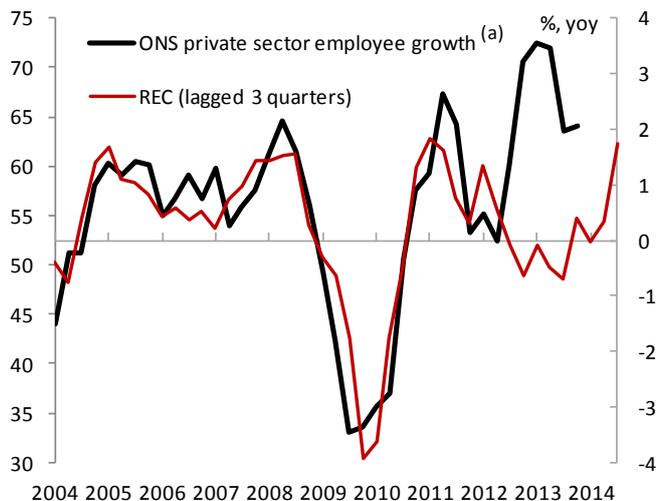
¹¹ Though they may imply similar things for productivity these two explanations may have different implications for pay growth during a recovery. Specifically, if pay is more flexible downwards it might also be more flexible upwards. If so, then productivity growth would accelerate more than usual through the recovery but unit costs – wages divided by productivity – would behave as they have during other recoveries.

¹² In this regard, it's notable that, even after more normal recessions, bankruptcies and firm failures tend to peak only after the economy has begun to recover.

¹³ There are several ways in which an increase in aggregate demand, originally unrelated to underlying TFP, might nonetheless produce an "endogenous" response in supply. By increasing confidence, for example, it might boost investment. My colleague David Miles, for example, suggests in a speech he will give later this week that this mechanism, if strong enough, could give rise to "multiple equilibria" in the determination of output and that strong demand might be the trigger that prompts a return to a better (and higher-productivity) state, without the need for a big increase in employment (Miles (2013), forthcoming).

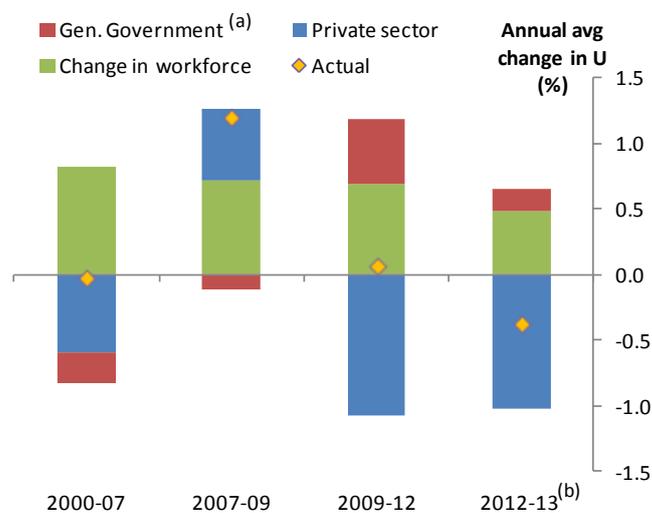
¹⁴ Chart 6 plots growth in private-sector employees, excluding the self-employed. The blue bars in Chart 7 plot changes in total private-sector employment, including the self-employed, relative to the active workforce.

Chart 6: Surveys suggest private-sector employment will continue to grow for a while yet



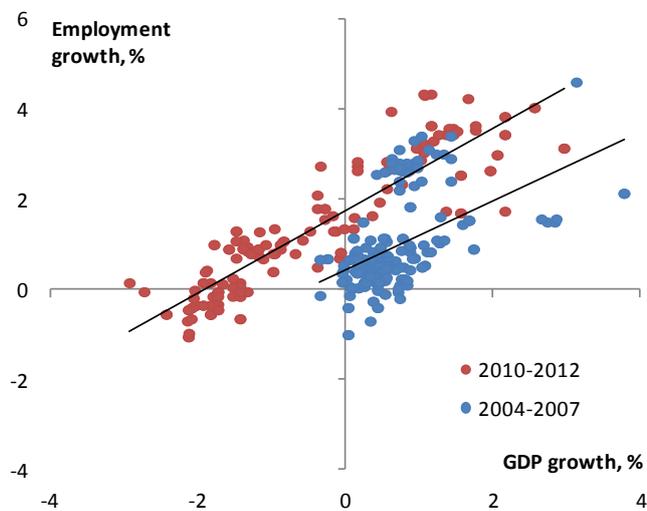
Source: REC employment survey and Bank calculations
 (a) Employee growth for Q3 2013 approximated by that between April and July.

Chart 7: Contributions to changes in unemployment rate



Source: ONS and Bank calculations
 (a) Includes government supported training schemes and unpaid family workers . (b) Year to May-July 2013.

Chart 8: Professional forecasters' forecasts of GDP and employment growth pre- and post-crisis^(a)



Source: HMT
 (a) Forecasts taken from June and July of each year for GDP and employment growth in the following calendar year. Blues dots are forecasts made for calendar years 2004-07; red dots for 2010-12.

And the connection runs both ways: something that is obvious after the event should probably have been forecastable before the fact. So perhaps I can end this section with Chart 8, which plots the difference between out-turns and prior forecasts of output and employment growth. The data are from the Treasury's survey of professional forecasters and cover predictions made by all the main UK forecasters since 2003. There is a clear divide in the performance around the time of the financial crisis. In normal times, there is no obvious bias in forecasting employment once you take account of any error in forecasting GDP (the regression line goes close to the origin). But there is a clear bias after the crisis: employment has been consistently stronger than expected, given GDP. Put another way, something unusual has happened to

productivity. And it would in my view be brave, not to say heroic, wholly to ignore the fact.

Supply uncertainty and conditional guidance

Now you may or may not buy any of this. And perhaps the best way to approach the question of how policy should behave when there is uncertainty about potential output is to start with the case where there is none.

We will do so imagining an economy in which inflation is determined by the underlying pressure on resources, in particular employment, and where this in turn depends, albeit with something of a delay, on aggregate demand relative to underlying productivity (the “output gap”). The faster demand grows, relative to productivity, the greater the demand for employment from firms and the greater the upward pressure on wages and inflation. If the “output gap” declines, whether because demand is weak or supply is strong, employment and inflation are more likely to fall.

When there is no uncertainty about supply growth – when the only thing giving rise to unpredicted changes in output are shocks to “aggregate demand” – setting monetary policy is reasonably straightforward. If output growth strengthens you immediately tighten policy: you know, without having to wait for confirmation from the labour market, that stronger growth will lead to more pressure on resources and inflation. So you’re better off responding straightaway. Conversely, and because it necessarily means employment and then inflation will subsequently fall, weaker output growth leads you immediately to loosen policy¹⁵.

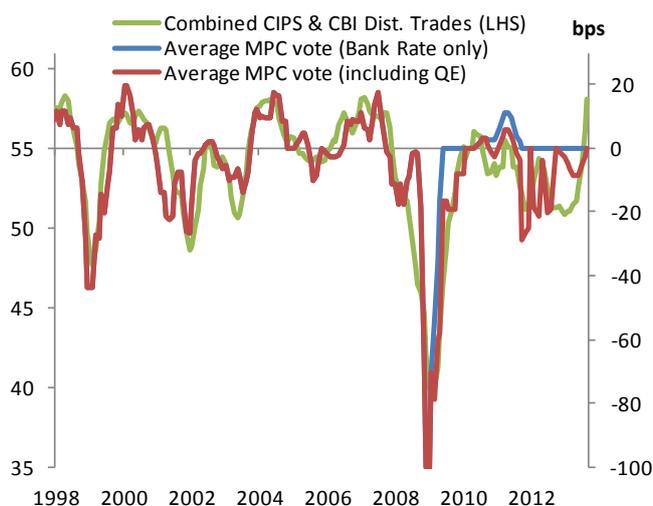
Perhaps this abstract world sounds too good to be true. But I think it’s not a bad description of those halcyon, pre-crisis days known as the “Great Stability”. After years of steady growth and inflation, it was probably reasonable to conclude that the economy’s productive potential was chugging along, at a steady rate of 2½% or so, in the background. Certainly that’s what most people assumed. And consistent with that assumption, UK monetary policy responded sensitively, immediately and (more or less) uniquely to actual growth in output. You can see this in Chart 9. The red line is the same composite CIPS index plotted in Chart 1. It’s a timely and (as a 3-month moving average) a relatively smooth indicator of private-sector output growth. Rather than using actual changes in interest rates, which is a more jumpy series, the blue line plots the average vote on official interest rates on the MPC. The two are closely correlated. The reason, I believe, is that every acceleration in output was thought to represent a rise in output gap, signalling higher inflation risks, every drop in growth the opposite.

Unfortunately, however, you can’t make that deduction when supply, not just demand, is unpredictable. Suppose the policy maker sees strong growth. This may turn out to require extra resources, including higher employment, in which case it is likely to add to inflationary pressure. But it may not: it’s also possible the upturn is being accompanied by faster productivity growth, in which case it would be wrong to tighten policy. What, under these circumstances, should one do? The answer is to respond less sensitively to output and

¹⁵ Note that you behave this way whether you were trying simply to stabilise output itself or, ultimately (and via the output gap), inflation. So it doesn’t even matter much what weight you put on these two objectives. Everyone, hawk or dove, has an incentive to lean against pure demand shocks.

more to developments in the labour market, even if you have to wait for a while to see them. Changes in unemployment are now a more reliable measure of what's happening to the degree of slack in the economy than economic growth alone. And what the policy maker loses in response time he or she makes up in a greater degree of clarity about the degree of spare capacity. The optimal policy – the relative weight one puts on unemployment relative to output growth – trades off one against the other.

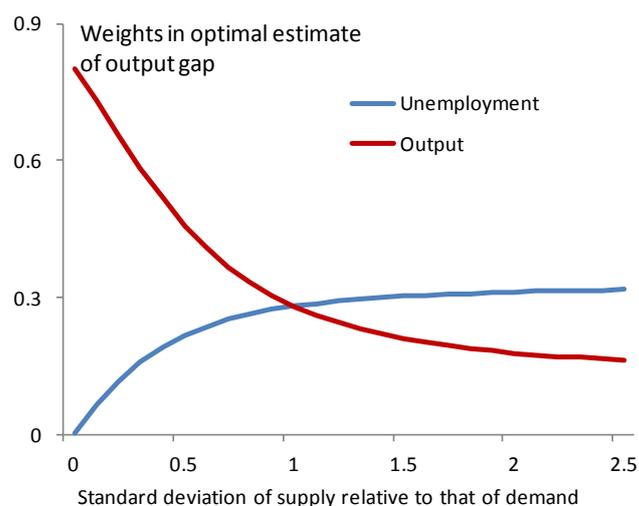
Chart 9: Average MPC votes against composite PMI indicator^(a)



Source: ONS, CBI, Markit Economics, Bank calculations

^(a) Average MPC vote calculated as the average voting result on Bank Rate and QE, with every extra £25bn of QE treated as if cutting Bank Rate by 25bps. See Joyce, Tong and Woods (2011).

Chart 10: Greater uncertainty about supply means unemployment a more valuable indicator for policy, output growth less so



Source: Own calculation (see Appendix)

I made this point briefly – that uncertainty about productivity means we should pay more attention to developments in the labour market – in a speech last year. The result of a more formal representation, in a simple economic model¹⁶, is in Chart 10. Against the variability of supply (relative to that of demand), on the horizontal axis, it plots the weights optimal policy should put on output growth and unemployment¹⁷. As one might expect, there is no point in paying any attention to unemployment when there are only demand shocks. Because it responds only slowly, and with some noise, to swings in demand, you might as well respond as soon as you see those in the output data. But as you become less confident about future supply, the unemployment data become more informative.

This, in my view, is a key rationale for conditioning policy on the rate of unemployment, as we have in our policy announcement in August. Reasonable people can disagree about how to interpret changes in output growth. Given the experience of the past few years, and the apparent breakdown in the relationship

¹⁶ See note 3 in the Appendix and, for a closely related result, Swanson (2004).

¹⁷ These are actually the weights on these observed series in the optimal estimate of the unobserved “output gap”. But a generalised version of certainty equivalence means they are also the relative weights in the policy rule (Svensson and Woodford (2003)).

between output growth, on the one hand, and employment and inflation on the other, that's only to be expected. But what we can say, with a greater degree of confidence, is that it would be appropriate to consider withdrawing some of the monetary easing in place if we see a marked decline in the single most reliable (albeit lagged) measure of spare capacity, namely the rate of unemployment.

Summary and conclusion

It is common for people to complain that the predecessors (and their successors too) had it easier. One example of this, in the world of central banking, is the often-heard line that the job of monetary policy is harder because the economic outlook is "particularly uncertain". As an outside observer, in my last job, I used to take this with a pinch of salt: after all, when has one ever heard the opposite ("happily, we can report that the outlook is unusually certain")?

True to type, however, and now that I'm on the other side, that's exactly what I'm trying to argue here: that the world is, indeed, "particularly uncertain" – or, at least, that we have to be more cautious about a key assumption, the rate of growth of the economy's productive potential, that we used to take somewhat for granted. The right response to this uncertainty is to attach less weight to output growth itself and more to direct indicators of economic slack, even if these emerge only with a delay. That judgement is central to the policy of "conditional guidance".

For the most part, the policy has been well understood. But it has, in some quarters, been mistaken as a commitment to keep policy unchanged for a particular period of time, come what may. Some press reports, for example, have argued that, because unemployment could fall faster than in central projection in the latest *Inflation Report*, the MPC might be "forced" to consider an earlier rise in interest rates than it has "promised".

I have to say that I find this argument decidedly odd. There is no promise unconditionally to keep interest rates fixed for a particular length of time. What we have pledged to do – and the clue is in the word "conditional" – is to examine the case for a withdrawal of monetary stimulus only after a significant fall in unemployment and as long as the inflation and financial stability "knock-outs" have not been breached.

That is only reasonable. Interest rates, like any other policy instrument, are not an end in themselves: they are a means to an end. If unemployment falls faster than we're expecting either because productivity does less well than in our central projection, or because demand grows more strongly, it would be right to ask whether we should think about withdrawing some of the monetary stimulus currently in place. Nor would we be displeased with such an outcome (what's not to like about lower unemployment?). If unemployment declines more slowly it would be right to leave the monetary stance unchanged for that much longer. This is something the MPC, along with the rest of the world, will reassess over time; in the August *Inflation Report* we viewed these two possibilities as equally likely. But then that's the thing about something that's uncertain: it could go either way.

References

Barnett, A., M. Barriel, A. Chiu and J. Franklin (2013), *Bank of England Working Paper*, forthcoming.

Broadbent, B. (2012a), "Productivity and the allocation of resources", Speech given at Durham Business School, available at www.bankofengland.co.uk/publications/Documents/speeches/2012/speech599.pdf.

Broadbent, B. (2013), "Forecast errors", Speech given at The Mile End Group of Queen Mary, University of London, available at www.bankofengland.co.uk/publications/Documents/speeches/2013/speech653.pdf.

Joyce, M., M. Tong and R. Woods (2011), "The United Kingdom's quantitative easing policy: design, operation and impact", *Bank of England Quarterly Bulletin*, 2011Q3, available at www.bankofengland.co.uk/publications/Documents/quarterlybulletin/qb110301.pdf.

Miles, D. (2013), forthcoming speech.

Swanson, E.T. (2004), "Signal Extraction and Non-Certainty-Equivalence in Optimal Monetary Policy Rules", *Macroeconomic Dynamics* 8(1), January 2004, 27-50

Svensson, L. and M. Woodford (2003), "Indicator variables for optimal policy," *Journal of Monetary Economics*, Elsevier, vol. 50(3), pages 691-720, April.