BANK OF ENGLAND

Speech

Productivity and the allocation of resources

Speech given by

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

Durham Business School

12 September 2012

I would like to thank Alina Barnett, Adrian Chiu and Amardeep Parmar for research assistance and I am also grateful for helpful comments from other colleagues. 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.

All speeches are available online at www.bankofengland.co.uk/publications/Pages/speeches/default.aspx

The UK economy has set some unenviable records in recent years. The recession of 2008-09 was the deepest since the 1930s. The period since then, a phase of the cycle when output normally expands rapidly, has instead seen little, if any, growth. As a result, the economy is still 3% smaller than it was in mid-2007. If you exclude the period immediately after the Second World War this is the worst five-year growth rate since the first half of the 1920s (Chart 1). Remarkably, output is 15% smaller than it would have been had it grown at its previous (post-war) trend rate.

Just as remarkable, however, is the comparative strength of employment and inflation. The empirical link between output and inflation is imprecise – Chart 2 plots a Phillips-curve relationship between the two, and the pre-crisis R-squared is 0.83[1] – but, taken at face value, and given the scale of the downturn, it would have led one to expect a steep fall in inflation in recent years.

Chart 1: Slowest growth since post-war period

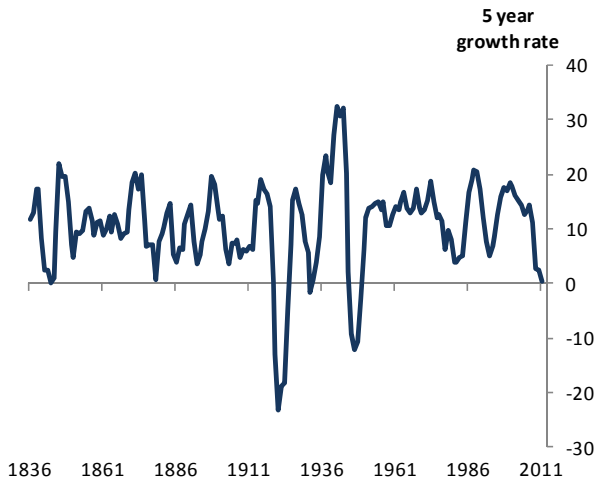

Source: ONS and Bank of England

Chart 2: Inflation surprisingly high relative to detrended output

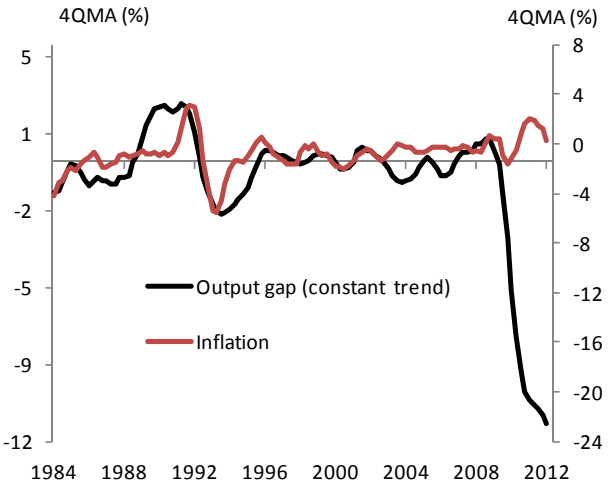

Source: ONS, NIESR and Bank of England calculations

[1] This refers to the R^2 on the fitted "New Keynesian" Phillips curve $\pi = \alpha\pi_{-1} + (1-\alpha)\pi^e + \beta y_{-1}$, where π is inflation (ex energy prices), π^e expected inflation, y (the log of) an "output gap" and the subscript -1 denotes a lagged variable (see, for example, Gali (2008)). Chart 2 plots $\pi - 0.6\pi_{-1} - 0.4\pi^e$ against the difference between output and a straight-line trend, where both the trend and α are estimated on pre-crisis data.

Chart 3: Employment much higher than past relationship with output would imply

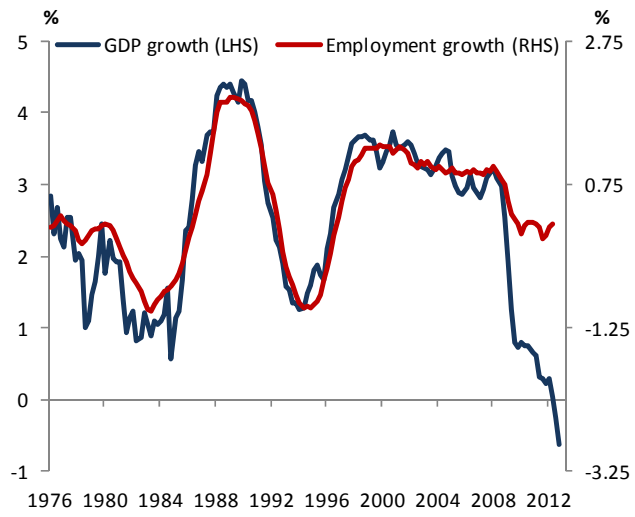

Source: ONS and Bank of England calculation

The pre-crisis correlation between GDP and employment growth is much tighter (Chart 3) and the divergence from it all the more striking. In more normal times, the 3% drop in output since mid-2007 would have been associated with an 8% decline in the number of jobs. Instead, and despite a contraction in the public sector, aggregate employment has risen slightly over the last five years. The likelihood that the past relationship could explain more recent trends is 0.2%, a one in five hundred chance.

One explanation for this is labour hoarding. Because it is costly both to fire and to re-hire employees firms may have hung on to staff in an expectation that demand will pick up. Another is that underlying productivity growth has slowed. In that case, weak output may not be the result purely of weak demand. It may instead reflect an independent hit to supply, one that explains why employment and inflation have been relatively high.

As extremes, these two hypotheses can have very different implications. If the first is true, then without a strong pick-up in economic activity, employment and inflation are likely to decline. Finding themselves in the position of a cartoon character that has run off the edge of a cliff, but is not yet aware of the fact, firms will at some point have to face reality and trim their under-employed workforce. Cost growth and inflation will then fall. If, instead, supply is (and remains) the problem, we could continue to see below-par output growth without any such implications for employment or inflation.

They are not mutually exclusive, however, and my own view is that both have a part to play. Because most people are more familiar with the demand-side explanations – we are unused to thinking about supply factors as a determinant of output except over some vaguely-defined “long run” – I will concentrate my efforts on discussing (i) why it’s hard for demand alone to account for the data and (ii) how an independent hit to supply might plausibly have occurred. In particular, I will suggest that our predicament may be partly caused by significant cross-sectoral disturbances – things that have had differing effects on different parts of the economy – and a failure of the financial system to reallocate resources in response.

If this is true, the economy’s lost potential is the result of a misallocation of capital, rather than any form of “technical regress”. In that case it needn’t have been lost forever. In time, as the financial system heals, and

investment starts to flow, the economy could well expand at an above-trend rate (without generating inflation), catching up some of the ground lost over recent years. But knowing when that might happen is very hard to judge. It's not at all easy, therefore, to say how productivity will evolve over the MPC's forecast horizon. That's why, in its last *Inflation Report*, the Committee widened the width of the "fan chart" around its forecast for economic growth. In the concluding section I will discuss briefly how monetary policy might be set in an environment of heightened uncertainty about productivity growth.

Why aggregate demand alone is unlikely to account fully for slow growth

That demand is independently weak is, I think, indisputable. We know there are signs of spare capacity, both within firms and in the labour market. Nor is there any shortage of candidates when it comes to potential hits to spending: fiscal tightening, slow (UK-weighted) global growth and (via their effect on household income) rises in commodity prices are all things that can reduce demand and output independently of the economy's capacity to supply it.

To that extent, the comparative strength of employment may simply be a reflection of "labour hoarding", or something akin to it[2]. If so, then some of those currently employed are actually "hidden unemployed", economically idle but kept in jobs by firms' expectations that output will at some point pick up and by the costs of hiring and firing. If these expectations are disappointed – unless economic activity does recover – this surplus labour will eventually be discarded. Even before that happens, it's possible that the "hidden unemployed" could exert downward pressure on wages in the same manner as people that are formally unemployed.

For several reasons, however, I find it hard to see how this can be the only explanation for what's happened to the economy[3].

First, there's the scale of the divergence in Chart 2. Note that the pre-crisis relationship in Chart 2 already allows, implicitly, for any stable degree of labour hoarding over that period. So it's not enough for such a phenomenon to exist. For labour hoarding alone to explain Chart 2, the incentives for it have to have become suddenly and dramatically stronger than in the past, sufficient to add an additional 8% to employment.

Yet it's hard to see why that might have happened. After all, it's not as if businesses have become more optimistic about future growth over the past three years (one thing that would encourage them to hang on to

[2] There are other phenomena that have similar effects. For example, it might be that, working in thin markets, firms are having to spend more time trying to drum up business than doing things that are genuinely productive (see, for example, Rios-Roll et al (2011)). Like labour hoarding, however, it seems unlikely this could account for much of the strength of employment. It's hard to see why increasing returns, on such a dramatic scale, can have emerged only recently. And the numbers of people employed in sales functions (as defined by the Standard Occupational Classification) actually fell between 2007 and 2011, both in absolute terms and as a share of total employment.

[3] My colleague Paul Tucker made similar points in a speech last November, in which pointed to the strength of private-sector hiring and wondered whether this was a symptom of the economy's need to rebalance. See Tucker (2011).

surplus staff). Nor, to the extent one can measure them, does there appear to have been a material increase in the costs of changing employment.

Chart 4: Private-sector job creation strongest since series began

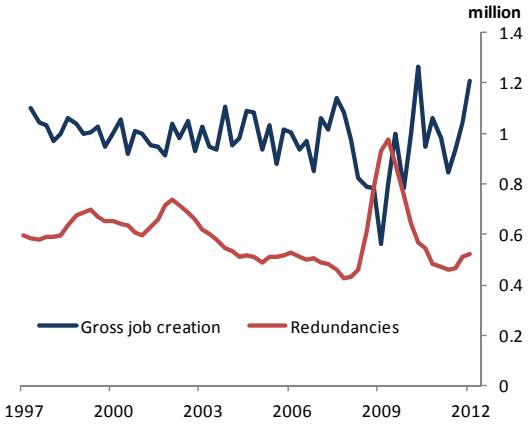

Source: ONS and Bank of England calculation

Second, what's striking about the employment data is not so much the weakness of lay-offs but the strength of new hires (Chart 4). The private sector has created more new jobs in the past two years than in any two-year period (at least) since the data on gross flows began in 1997. But if firms are hoarding labour – if even their existing employees are surplus to current requirements – why hire new people at all?

Third, though there are signs of spare capacity they are more moderate than you might expect if the economy's supply potential had been growing at some uninterrupted rate since the crisis. Chart 5 plots a swathe of business-survey measures of spare capacity against deviations from a fixed-growth trend in GDP. This suggests that spare capacity within firms is considerably smaller than would be the case had capacity continued to grow at its pre-crisis rate.

The same is true of the simplest measure of spare capacity in the labour market, namely the rate of unemployment. This comparison is obviously inappropriate because it's precisely the strength of employment we're trying to explain: to assume the unemployment rate is always an accurate measure of slack in the labour market is to assume away the possibility of the "idle employed". But on this point it's worth drawing attention to Chart 6. Using annual data since inflation targeting began, in 1993, it plots wage growth, less expected inflation (as measured in the gilts market), against the rate of unemployment. If hoarded labour amounts to hidden unemployment you might expect it to give an additional downward push to wage growth – i.e. for the recent data to have come in below the pre-existing line[4]. But there's no sign of that at all. At least in how they affect wage growth, these "hidden unemployed", if they exist, look just like any other employed people. This doesn't mean they don't exist. But it does suggest that they would actually have to get laid off before they exert any downward pressure on pay and then on price inflation.

[4] For reference the graph also includes data for the years 1991 and 1981. The natural rate of unemployment – defined here as the rate necessary to keep real wages growing in line with some fixed (trend) rate of productivity growth – looks to have fallen sharply between these two dates and then again after 1993.

Chart 5: Spare capacity within firms smaller than that implied by uninterrupted growth of potential output

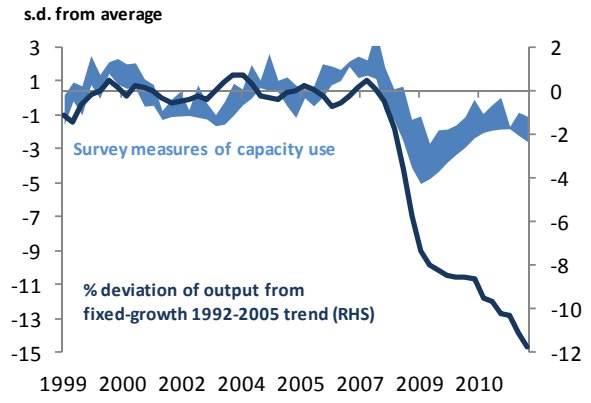

Source: CBI, BBC, CIPS, ONS and Bank of England calculations

Chart 6: No sign of structural change between unemployment and wages

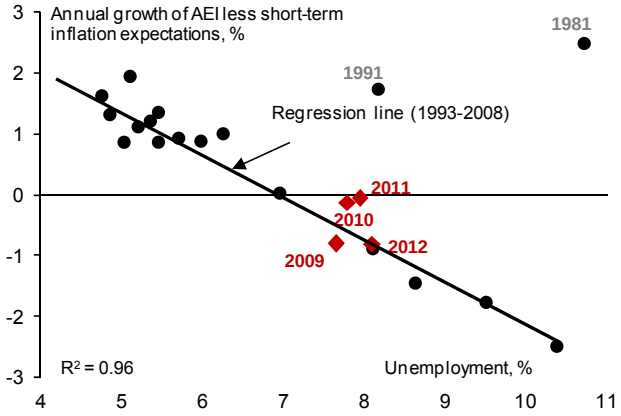

Source: ONS and Bank of England

The supply side: why productivity depends on the cost and mobility of capital

If the weakness of productivity is not due solely to the combination of weak demand and something exceptional in the labour market, what else might be at work? I think we should look instead at the capital market – in other words the way in which productive assets get financed and allocated.

I say this partly because it is not uncommon, during and immediately after financial crises, to see weak productivity growth and (relative to output) higher-than-usual inflation. During the 61 banking crises covered by Bordo and Eichengreen's dataset[5], annual productivity growth is on average 1.4% points lower than at other times. Table 1, which plots the results of a regression on the same dataset, shows that, for given output growth, inflation is higher during financial crises. It is also less well correlated with output growth[6]. Thus the empirical "Phillips curve" looks both higher and flatter than during normal periods (though the latter effect is statistically insignificant). These patterns suggest that the UK's recent experience has something in common with other financial crises and, therefore, something to do with finance.

[5] The dataset covers 35 industrialised countries between 1980 and 2008. We complemented this dataset by adding employment and extended it to 2010 using data from the OECD. The number of banking crisis before 2007 is 48.
[6] The results are from a regression of inflation on lagged inflation, output growth and those same terms times a "dummy variable" that takes the value 1 during banking crisis and 0 at other times. The estimate coefficients on these dummy terms therefore tell you how the intercept and slope differ during crises.

Table 1: Empirical "Phillips curve" looks higher during financial crisis

Inflation	Coefficient	P values
Inflation (t-1)	0.33	0.00
GDP growth (t-1)	0.21	0.06
Banking crisis dummy	0.53	0.38
GDP growth * banking crisis dummy	-0.14	0.27
No of observations		905
No of countries		22

The cost of capital: I say it too because there are plausible mechanisms by which crises might have such effects. One is that finance becomes more expensive. This encourages firms to substitute towards labour (which is now relatively cheaper), reducing capital and output per employee. Chart 7 plots an estimate of the gross cost of capital for UK-quoted firms[7]. If anything, this understates the rise in financing costs because it doesn't allow for credit rationing and excludes those firms (generally small, unquoted companies) that are likely to be most affected by it.

Yet even on this measure the cost of capital is currently close to 200bp higher than over the pre-crisis decade, a proportionate increase of almost a quarter. On standard models, this would be enough to reduce the level of productivity (relative to trend) by around 6%[8]. The predicted impact would almost certainly be bigger if we had an accurate measure of the cost of credit rationing.

On the face of it that goes some way to explaining what's happened, and my own belief is that dearer finance probably has contributed to slower productivity growth. But, in truth, it's hard to build a strong case on the basis of the data we have and the simplest of economic models. For one thing, that 200bp gap has emerged only recently. For much of 2009 and early 2010 (after the aggressive easing in monetary policy but before the start of the Euro area crisis) the cost of capital in Chart 7 was **lower** than before the crisis. Yet the productivity numbers looked odd even back then.

It's possible that credit rationing effects, unaccounted for in Chart 7, were particularly strong at the time. But even if that were true, there's a second and bigger problem regarding the strength of investment. At least in this simple, one-sector model, the rise in the cost of capital affects productivity entirely through a collapse in capital formation. It's this that reduces capital intensity (as the existing stock depreciates) and, as a result, output per employee as well.

[7] Chart 7 plots corporate earnings (after tax but gross of depreciation and interest) divided by the market value of the firm plus an estimate of expected long-run growth rate of earnings (the latter generated from consensus long-run forecasts).
[8] For any constant-returns production technology a proportionate change of 1% in the cost of capital reduces productivity by $\alpha\sigma/(1-\alpha)$, where α is the share of national income going to capital and σ is the elasticity of substitution between capital and labour. Reasonable empirical estimates of these parameters across the UK economy as a whole are around one third and one half respectively, implying that the elasticity of productivity with respect to the gross cost of capital is around one quarter.

But that's not what's happened. Though it fell sharply after 2007, private-sector investment certainly didn't come to a halt (Chart 8). As a result, the aggregate stock of capital, conventionally measured, hasn't fallen nearly fast enough to account for the drop in productivity (the red line plots what would have to have happened to investment for this to be the case[9]).

Chart 7: Cost of capital has risen

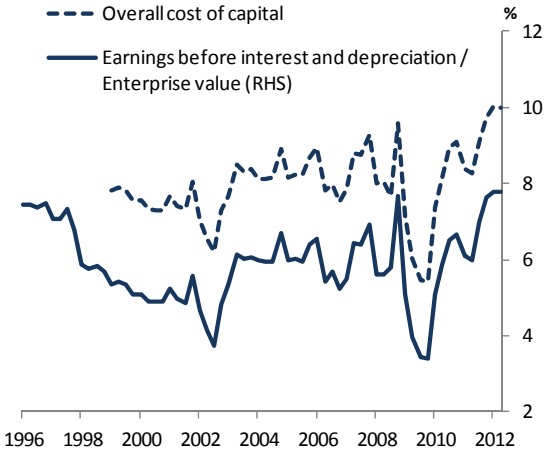

Source: Consensus Economics, Thomson Reuters Datastream and Bank of England calculations[10]

Chart 8: But investment not weak enough to explain productivity slowdown on its own

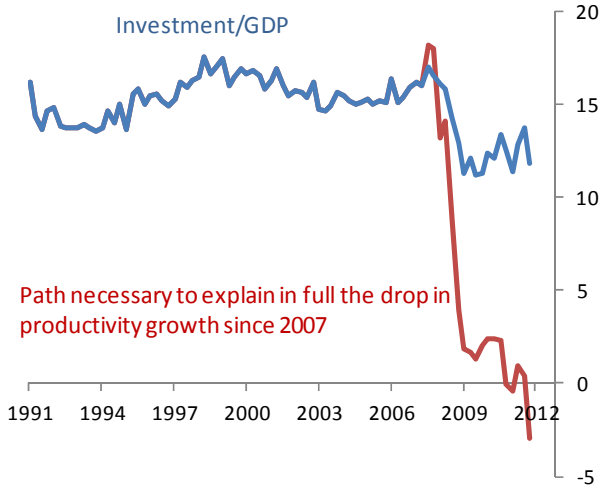

Source: ONS and Bank calculations

The allocation of resources: If weaker demand and more costly capital aren't enough, between them, to explain the slowdown in productivity growth, what else might be at work? One possibility has to do not with the rate of accumulation of capital resources in aggregate but their **allocation** across the economy.

In the real world, the economy is made up of lots of different sectors and firms. As macroeconomists we often ignore this texture and treat the economy as if it were made up of a single, "representative" firm. We talk about "the economy", for example, and say "it is" operating above or below potential. Similarly, we often argue that "firms are" doing this or that, lumping them together into a single, homogeneous whole.

This is fine if the shocks hitting the economy are the same for everyone (uniform across sectors). It's also reasonable, it turns out, if the resources used in production – capital and labour – can be easily moved from one sector to another[11]. In that case, a purely relative shock – a drop in the demand for sector A's output, say, and an equivalent rise in the demand for B – would do nothing to aggregate output or productivity. An

[9] In the language of economics, there seems to have been a sharp drop in the contribution to productivity growth of "Total Factor Productivity", what's left once the estimated impact of capital deepening has been stripped out.

[10] Ratio of earnings before interest and depreciation to enterprise value calculated for all UK listed companies, as defined by Datastream code TOTMKUK; enterprise value sums the market values of firms' equities and outstanding debt. The overall measure is calculated by adding an estimate of expected long-run growth of earnings.

[11] Technically, we can add up many sectoral "production functions" into a single aggregate as long as (i) they all exhibit constant returns to scale and (ii) inputs are fully mobile (so that factor prices are equalised across sectors). See Samuelson and Swamy (1974).

incipient rise in the relative price of A would increase returns in that sector, drawing resources away from B until factor prices (wages and the rate of profit) are re-equalised. At that point A's price falls back and, even if its mix has been altered, the **combined** output of the two sectors is unchanged.

Note, however, that during this interim period, when resources are still moving, aggregate output and productivity is **lower** than it otherwise would be. This is not because of an inadequacy of aggregate demand. Nor does it reflect a loss of technical know-how in any individual sector. In this experiment, at least, what prevents the economy producing as much as it did, given the shift in relative demand, is a misallocation of resources. Following that shift in more needs to go to A and less to B; only if that re-allocation of capital and labour occurs is the economy capable of producing what it did; a failure to re-allocate resources therefore impairs the economy's productive potential; and the rise in A's price, and the disparity in returns across sectors, is a measure of the cost of such a failure.

Evidence of capital mismatch

So much for the theory. I think there is evidence of this effect at work in the post-crisis data in the UK, particularly when it comes to capital mobility.

We begin with Charts 9 and 10, which plot cumulative changes in output volumes and prices respectively – all as standard deviations from a pre-crisis mean – across 81 sectors of the economy. It's clear there's been a marked divergence in growth rates and price inflation, at a sectoral level, since the crisis hit in mid-2007. Chart 11 summarises these graphs by plotting the cross-sectoral spread in output and prices at each point in time.

Chart 9 Marked increase in dispersion of output across sectors since the crisis

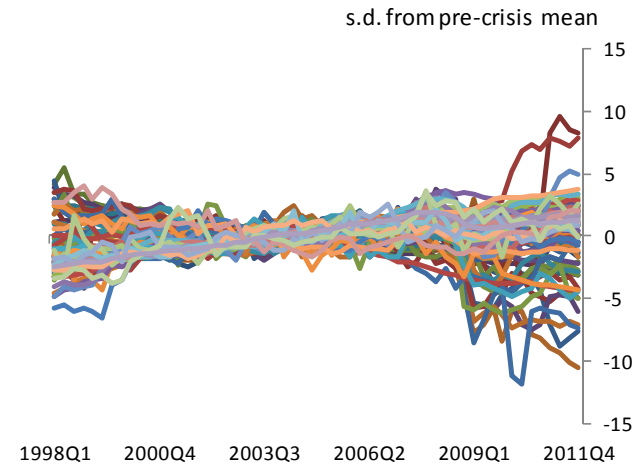

Source: ONS

Chart 10: Relative prices also much more volatile

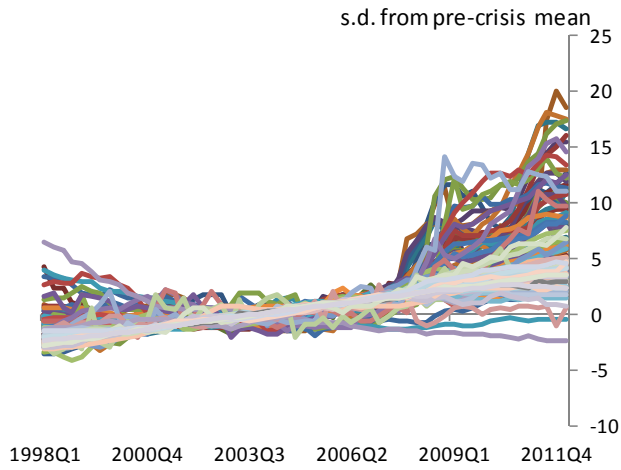

Source: ONS

I won't discuss here why exactly these shifts have occurred, but it's easy enough to think of examples. The rise in commodity prices has had a big impact on the aggregate economy (via a generalised reduction in household income, for example). But it is also likely to have affected some sectors much more than others. Similarly, the credit crunch has hit everyone. But the fact that the UK has relatively large banks, and that it was already a net borrower even prior to the crisis, probably meant that demand in purely domestic-facing sectors was always more vulnerable than others to the ensuing contraction in the supply of credit.

Anyhow, whatever the underlying causes, the fact of this sectoral dispersion is undeniable. And what has been required, to avoid any impact on aggregate productivity, is a similarly sharp reallocation of resources. All else equal the rise in energy prices would normally prompt more investment (and employment) in the more energy-efficient parts of the economy. The relative hit to domestic-facing business requires a similar shift – so-called "rebalancing" – towards tradable output.

As far as the labour market is concerned, the signs are that this is precisely what has happened. The divergence in output growth has been matched by an equally steep increase in the dispersion of employment across sectors (the blue line in Chart 12). The fact that employment has been able to migrate across sectors has been sufficient, it seems, to prevent any widening in the spread of wages (the red line). This suggests that labour has been relatively mobile[12].

Chart 11: Significant sectoral shocks since the crisis

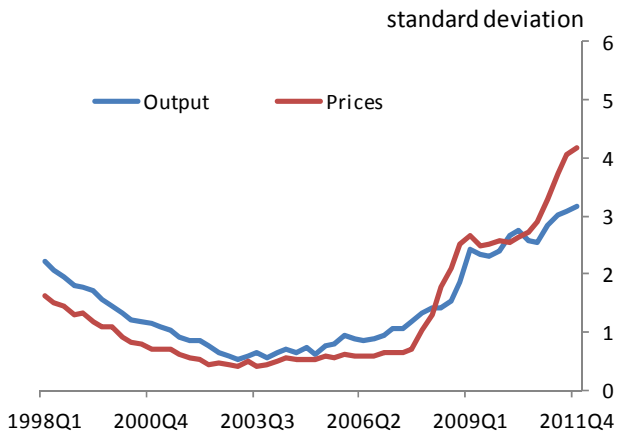

Source: ONS and Bank of England calculations

Chart 12: Low dispersion of wages suggests labour is mobile across sectors

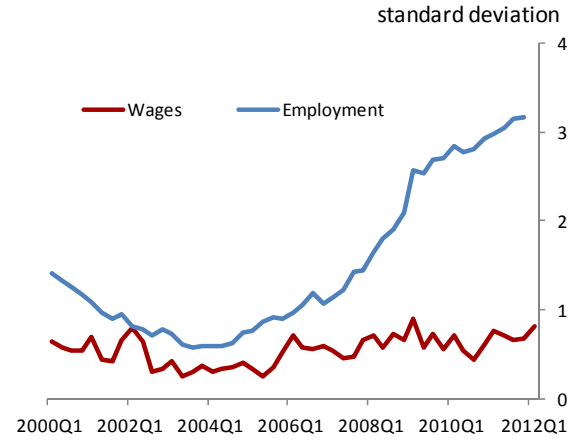

Source: ONS and Bank of England calculations

[12] This in turn suggests that the hiring and firing costs are not that significant, making it harder to explain Chart 2 purely in terms of "labour hoarding": if it's easy for jobs to move across sectors one suspects it can't be that difficult shed excess employment in any given sector.

Chart 13: More cross-sectoral volatility in capital returns

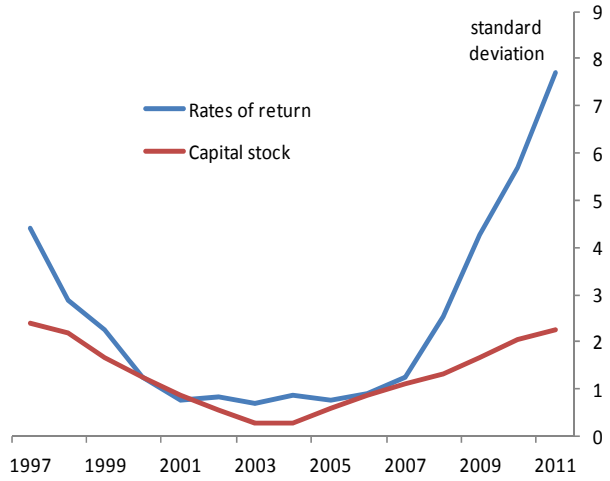

Note: excludes two-digit SIC 2007 sectors D, E, M, N and R.
Source: ONS

Chart 14: Rates of corporate failure and start-ups have been low

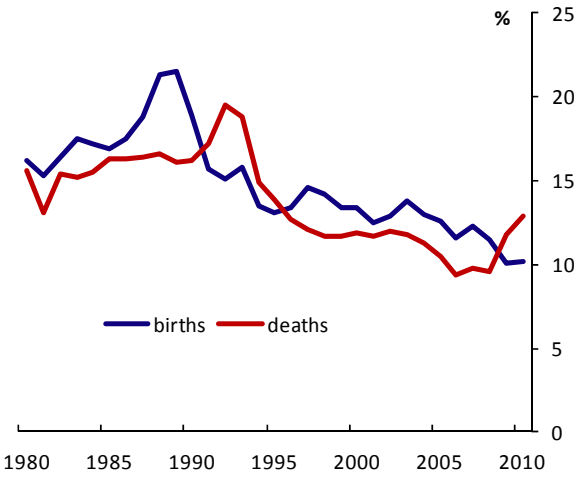

Source: ONS

The same cannot be said for capital. There certainly seems to have been some movement. Despite shrinking in aggregate, investment in some sectors has risen. As a result, the dispersion in the estimated capital stock – this time measured across only 15 sectors[13] – has also gone up (this is the red line in Chart 13). But it's not risen as much as the dispersion of employment (Chart 12). Nor has it been enough, more to the point, to put much of a dent in the spread of rates of return (relative to the pre-crisis averages) that has emerged since the crisis. This suggests that there may still be unexploited opportunities from capital reallocation, and an equivalent cost to the productive potential of the economy. Firms whose rates of return have risen – those in the hypothetical "sector A", where demand is relatively strong – find themselves capital constrained. They would ideally like to invest at a faster rate, but are forced instead to hire more employees to meet demand. By contrast, "Sector B" firms are kept in operation despite making lower returns for investors.

More sectoral dispersion or less capital mobility?

The way I've described it, this "capital mismatch" would occur **either** if the economy became less able to reallocate resources **or** if it faced a greater need to do so. So it's possible that what characterises financial crises is not so much lower capital mobility but an exceptional degree of cross-sectoral volatility. There are also other channels through which this kind of "churn" might affect productivity (so even if capital markets were working perfectly). For example, some economists have argued that, because employees get more skilled the longer they've been in a job, sectoral re-allocations directly (if temporarily) reduce average skill levels[14].

[13] Sectoral data on investment and capital are less refined than those for employment.
[14] See, for example, Abraham and Katz (1986) and Hamilton (1999). Pratap and Quintin (2011) present evidence that this effect has been particularly important, at least in emerging economies, after financial crises.

I certainly don't want to discount these possibilities, any more than one can rule out the possibility of an exceptional degree of labour hoarding. If ever we arrive at a satisfactory explanation of the slowdown in the UK's rate of productivity growth it's likely to include several factors.

But I do think that the spread of profits across the economy, especially when compared with the relative stability of wage spreads, is suggestive of problems in the allocation of capital specifically. Others have suggested this may have been important both in Japan's "lost decade" in the 1990s, and even in the Great Depression in the United States[15]. And there is also anecdotal evidence, at least, of an unusual dispersion of credit conditions in the UK.

We have all heard accounts of how small businesses with profitable business ideas have been unable to finance them. At the same time, the FSA has suggested that there's widespread "forbearance" by banks of loans in breach of pre-existing covenants (commercial property lending in particular[16]). It's hard to know how significant these things are for the macroeconomy[17]. But they would help to explain the strikingly low rate of corporate births and deaths through this cycle (Chart 14). And they would certainly have contributed to the dispersion in profits we observe in the data.

Summary and conclusion

As empirical macroeconomists we are brought up to believe that economic cycles are determined purely by demand-side factors – consumer or business confidence, for example, or conventional macroeconomic policy (whether fiscal or monetary) – while supply-side factors (the size and skills of the workforce or things that determine allocative efficiency) matter only for the long-run level of output.

For the most part this is justified by empirical evidence. It is hard to get the simple models that are good at describing income determination across decades (and countries) to match the data at higher, cyclical frequencies. Those short-run relationships – the pro-cyclicality of employment, unit labour costs and inflation, for example – seem to point to independent movements in demand along fixed, upward-sloping supply curves (for goods and labour).

The period since the financial crisis began in mid-2007 confounds these patterns. Output has shrunk, the first time that's happened over any five year period since immediately after the Second World War, and before that the 1920s. Yet employment, seeming to defy not only the historical data but gravity itself, has risen. Unit

[15] Caballero et al (2001 and 2007) argue persuasively that low productivity growth in Japan owed something to the "forbearance" of non-performing loans by banks. Chari et al (2002) make a similar point about the Great Depression.
[16] According to last December's Financial Stability Report, the FSA estimates that banks are forbearing on 8% of residential mortgages and 33% of commercial property loans.
[17] Though these complaints are legion, there isn't much quantitative evidence, mostly because there aren't good data on UK small businesses to begin with. Note, however, that most (around 80%) business investment is done by larger businesses, those with 250 or more employees, and that many of these can in principle by-pass the banks by issuing securities. This suggests that financial constraints on investment may be caused not just by deleveraging in the banking system but by a more widespread rise in the risk premium. This was something I addressed in a speech in June.

costs (average pay divided by productivity) have grown faster than over the previous five years, and faster than 2%. The same goes for the CPI, even if you strip out energy prices and VAT. On the face of it, this looks more like a contraction in supply along downward-sloping demand curves.

Some of the slowdown probably reflects special factors. Perhaps it's harder to depress nominal wage growth the closer it gets to zero. This could help to explain why the Phillips curve looks flatter not just in the recent UK data but in other financial crises too[18]. Some of the strength of employment may well be due to an exceptional degree of "labour hoarding", as would be expected if firms are unusually uncertain about demand prospects. It's also possible that output growth has been under-measured.

But the greater the scale of the divergences in Charts 2 and 3, and the longer they go on, the harder it is to believe that there hasn't also been a genuine hit to the economy's effective supply capacity. There needn't have been any "technical regress" for this to have occurred. A combination of uneven demand (across sectors) and an impaired financial system, one that is unable to reallocates capital resources sufficiently quickly to respond to such shocks, is enough to reduce aggregate output per employee. Such a process would also give rise to precisely the volatility in relative prices and the widening sectoral dispersion of profitability we've observed in the data. Some firms, it seems, are kept in business (and retain employees) despite making relatively low returns. Others, able to expand but unable to obtain the finance to do so, are forced to substitute labour for capital.

It's hard to imagine this dispersion in returns can persist indefinitely. Assuming the underlying shifts in relative demand are permanent, the economy must in the end adapt to them. Indeed, at some point, once the financial system returns to health, one could imagine exactly the reverse process: a long period of above-trend productivity growth. It would certainly run against the grain of post-war experience if output-per-hour grew permanently at different rates in Europe and the US.

But predicting when that might happen is extremely difficult. It's this uncertainty about future productivity growth that persuaded the MPC to widen further the width of its GDP "fan chart" in the last *Inflation Report*. It also matters how (if ever) it does so. If productivity improves because investment begins again to flow into the higher-yielding sectors of the economy, policy wouldn't need to do much. If, on the other hand, it recovers because funding for the lower-yielding sectors gets pulled, it would do so via lower aggregate employment. This would play out much like the end of labour hoarding and would require easier monetary conditions to avoid an undershoot of the inflation target.

How, amidst all these uncertainties, should monetary policy be set? Precise rules are best avoided in these circumstances – we are painting with a pretty broad brush – but it's worth making a couple of observations.

[18] See Chart 2 and Table 1. I discussed this in a speech last September, "Rebalancing and the real exchange rate". For more on the international evidence see also Meier (2010)

First, we should set policy not just on its ability to affect demand but its capacity to improve the flow of finance in the economy as well. Much of this lies outside our control, and depends, as I argued in a speech earlier this year, on events in the Euro area. But that doesn't mean domestic policy is powerless. If, as the MPC expects, the Funding for Lending Scheme helps to promote the supply of finance across the economy it's likely it will also improve its allocation.

Chart 15: The UK employment- inflation relationship relatively stable

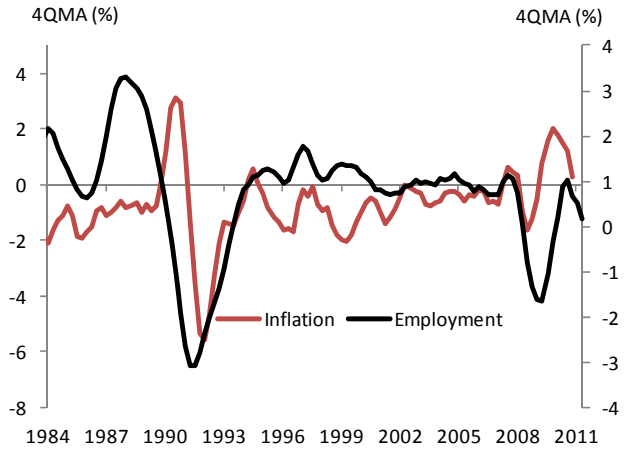

Source: ONS and Bank of England calculations

Second, as monetary policymakers, we should probably pay less attention (than we normally do) to movements in output and relatively more to changes in employment[19]. We may be less confident than usual about the origins of a given change in output (i.e. whether it's "demand" or "supply" led). What we do know is that, whatever the original disturbance, and as long as the supply of labour (and therefore the "NAIRU") is relatively stable, we're likely to want to ease policy if employment falls (because either demand has weakened or underlying productivity has risen) and to tighten it if employment growth improves.

This isn't simply a theoretical point. As Chart 15 demonstrates, the relationship between employment and inflation has proved more stable through the crisis than those between either of those variables and output (Charts 2 and 3).

Finally, if I have to take anything else away from all this, it's the realisation that we have much still to learn!

Thank you.

[19] My fellow MPC member Spencer Dale made a similar point in a speech last weekend. Noting that uncertainty about productivity made the appropriate response to slower output growth much harder to judge, he argued that "the Pavlovian-like response of some commentators to call for more monetary stimulus each time they observe weak output growth is not sensible".

References

Abraham, K.G. and L.F. Katz (1986)**,** "Cyclical Unemployment: Sectoral Shifts or Aggregate Disturbances?", *Journal of Political Economy*, Vol. 94, No. 3, Part 1, pp. 507-522.

Bai, Y., Rios-Rull, J. V. and K. Storesletten (2011), "Demand Shocks as Productivity Shocks", Federal Reserve Bank of Minneapolis and CEPR.

Broadbent, B. (2011), "Rebalancing and the exchange rate", Speech given at Thomson Reuters, London, available at www.bankofengland.co.uk/publications/speeches/2011/speech520.pdf

Caballero, R. J. and Hammour, M. L. (2001) **"**Institutions, Restructuring, and Macroeconomic Performance", in Advances in Macroeconomic Theory, ed. J.Dreze.

Caballero, R. J., Hoshi T and Kashyap A. (2007) **"**Zombie Lending and Depressed Restructuring in Japan", American Economic Review 98.

Chari, V.V., P.J. Kehoe and E.R. McGrattan (2002), "Accounting for the Great Depression", *American Economic Review*, Vol. 92, No. 2, Papers and Proceedings of the 114th Annual Meeting of the American.

Dale, S. (2012), "Limits of monetary policy", Speech given at the 44th annual Money, Macro and Finance Conference at Trinity College, Dublin, available at http://www.bankofengland.co.uk/publications/Documents/speeches/2012/speech597.pdf

Gali, J. (2008), "Monetary Policy, Inflation, and the Business Cycle: An introduction to the New Keynesian Framework", Princeton University Press.

Hamilton, J.D. (1988), "A Neoclassical Model of Unemployment and the Business Cycle", *Journal of Political Economy*, Vol. 96, No. 3, pp. 593-617.

Meier, A. (2010), "Still Minding the Gap—Inflation Dynamics during Episodes of Persistent Large Output Gaps", International Monetary Fund.

Pratap, S and E. Quintin (2011), "Financial crises and labour market turbulence", Journal of Monetary Economis, Vol 58, pp. 601-615

Samuelson, P. A., and S. Swamy (1974), "Invariant Economic Index Numbers and Canonical Duality: Survey and Synthesis", American Economic Review 64 (Sept. 1974), 566-593

Tucker, P (2011), "A few remarks on current monetary policy in a rebalancing economy", speech given at The Joint 1900/City Club Lunch in London, available at http://www.bankofengland.co.uk/publications/Documents/speeches/2011/speech532.pdf