

Charles Bean: Measuring recession and recovery – an economic perspective

Speech by Mr Charles Bean, Deputy Governor for Monetary Policy at the Bank of England, at the Royal Statistical Society (RSS) Statistics User Forum Conference, London, 27 October 2010.

* * *

Good morning! It is a great pleasure to open this conference dedicated to exploring the way statistics are used in the study and control of the business cycle and, in particular, the issues thrown up by the Great Contraction of 2008–9. Before I turn to that topic, however, it might be useful if I start by saying a few words about the sort of information that the Bank of England's Monetary Policy Committee (MPC) relies on when it sets the level of Bank Rate (or the amount of asset purchases) each month.

First, and most obviously, there is the wealth of official statistics on the economy generated by the Office for National Statistics (ONS) and other national statistical agencies, covering such variables as output, employment, pay and prices. That is the absolute bedrock on which our analysis and policy rests, and without it we would struggle to do our job.

But the picture they provide is rarely perfect. Series often do not quite correspond to the appropriate economic concept, either because they are conceptually difficult to measure or else because only imperfect proxies are available. As an example, take output. In the old days, when the economy was dominated by manufacturing, it might have been relatively easy to measure the volume of output and value added of different industries. But our economy today is dominated by services. Conceptually even defining the value added in, say, the financial sector, let alone measuring it, is not straightforward. Moreover, the meaning of a series often changes subtly over time as the economy evolves. Manufacturing today is a very different animal from manufacturing 30 years ago: for instance, the value added of Rolls-Royce aero engines today comes as much from the after-sales and maintenance as from the engine assembly itself.

Finally, and most saliently for real-time policy analysis, the official data invariably appear with a delay of some months and can be substantially revised as more information accrues and the ONS improve their methodology. For that reason, the Monetary Policy Committee supplements the official data with two other sources of information that are more timely, even if they lack precision compared with the official data. The first are the regular business surveys from the likes of the Confederation of British Industry and the British Chambers of Commerce. As many of these have been going for quite a while, we have in many cases been able to establish their usefulness as predictors of the not-yet-available official data. Second, we also rely on the information provided by our network of regional Agents who have about 8000 business contacts across the country. They can often help the MPC understand what lies behind puzzling movements in the official data or in the business surveys. And, though the evidence they provide is more qualitative rather than quantitative, they can give us early warning of major cyclical developments. The outstanding example here is in the aftermath of the collapse of Lehman Brothers and the subsequent near-seizure of the global financial system in early October 2008, when our Agents reported many of their contacts saying that "orders had fallen off a cliff" – the first real indication of the severity of the contraction that was to follow over the next nine months.

The sharpness of that contraction, as well as its unexpected nature, can be seen in Chart 1, which plots the path of four-quarter real GDP growth on top of the fan chart associated with the MPC's projection for growth, conditioned on prevailing market interest rates, made in August 2008. Given that each shade of green corresponds to 10% of the possible outcomes and the whole fan covers 90% of the possible outcomes, it can be seen that the MPC had

collectively assessed the probability of contraction over the four quarters to 2009 Q2 as about a half, but the probability of growth being less than $-1\frac{1}{2}\%$ as just 1 in 20. Now the MPC does not attempt to calibrate the probabilities associated with tail events precisely. But it is a fair bet that, if we had, then we would have put the chances of a contraction as large as the ONS's current estimate of 6% as being virtually negligible.

It is worth stressing that we were not alone in thinking a very severe contraction in activity was unlikely. Almost all forecasters were in the same boat. For instance, when we prepare our projections each quarter, we also canvass the views of outside forecasting bodies. Their assessment of the probabilities of various ranges of possible outcomes was very similar to ours.¹ Chart 2 illustrates the same point, but puts it into a broader context. The blue line shows the actual path of four-quarter GDP growth over the period since the creation of the MPC, while the pink line shows the Committee's central (mean) projection from a year earlier. The green bars show the range of end-year forecasts for the following year produced by about 25 other organisations, together with their corresponding average (the green diamonds). The striking feature is the close proximity of the MPC's projections and the average of the outside forecasts throughout the sample. For most of the period – one in which, it should be said, growth was unusually stable – those forecasts are close to the subsequent outturns. The two exceptions are the Great Contraction of 2008–9 and the “phantom” slowdown after the 1997–8 Asia/LTCM crisis, when early vintages of the official data for output as well as business surveys pointed to a slowing in growth that was not subsequently borne out in mature vintages of the official data. And for both these episodes, no forecaster came close to predicting the outcome.²

The generic point here – which holds true in fuller analyses over longer time periods – is that the big divergences, when they occur, are not between the central forecasts made by different bodies, but rather between those forecasts and the corresponding outturns. And that is as it should be: deviations of outturns from central (mean) forecasts should be unpredictable if those forecasters are using the information that is available to them efficiently.

Why then, as the Queen famously asked at the London School of Economics in November 2008, did no-one see the Great Contraction coming? Some economic downturns are broadly predictable in nature. In particular, that is the case when they are deliberately policy-induced in order to squeeze inflation down, though even then it may still be difficult to get the magnitude and timing right, and there will always be other unexpected events that perturb the path of the economy. But downturns associated with financial or banking crises are rather different animals. Rapid growth in credit and asset prices can act as a warning sign of building vulnerabilities, but frequently appear to be validated by developments in underlying economic fundamentals. In the case of the present crisis, those developments included: low long-term real interest rates associated with high savings rates in emerging market economies; an unusually stable macroeconomic environment, with an associated apparent reduction in risk; and financial innovation which purported to distribute risk more widely rather than leaving it concentrated in the banking system.

With hindsight, we now know that much of that financial innovation in fact left risk concentrated in the banking sector; it was in effect simply shifting maturity transformation off balance sheet, where it was not subject to the same regulatory requirements. To that must be added: the underestimation of risk, coupled with inadequate risk management; distorted incentives facing the originators of US sub-prime mortgages and the ratings agencies

¹ See page 48 of our August 2008 *Inflation Report*.

² One forecasting organisation in the sample was consistently more pessimistic than others throughout the latter part of the sample. Though that organisation therefore did better at forecasting the Great Contraction when it came, it only did so by systematically underestimating growth in the years preceding the crisis.

responsible for assessing complex structured finance assets. Finally, the complexity of those products and the high degree of interconnectedness between financial institutions meant that the financial system, instead of becoming more stable, had in fact become more vulnerable to failure.

The lesson I want to draw from this is not that the problems in financial markets which began in August 2007 and culminated in the near-meltdown of the global financial system a little over a year later were an inherently unpredictable Act of God. Rather it is that one would need to be endowed with perfect foresight to have been able to predict how the financial crisis would unfold, spilling over from one institution to another, and from one market to another. And who knows what would have happened if, for instance, Lehman Brothers had successfully found a buyer that weekend in September 2008?

Moreover, the impact on the real sector, which is what concerns us here, was also extremely difficult to judge. While it was predictable that the availability of credit would tighten as risk aversion rose and financial institutions rushed to de-leverage their balance sheets, it was far harder to know by how much. And much of the contraction in household and business spending appears to have been down to a sharp rise in precautionary saving and the postponement of investment projects and the running down of inventories as uncertainty rose. Again, while such a reduction in spending was plausible, it was nevertheless extremely difficult to judge its likely magnitude.

The moral from this is that one should not expect to be able to predict the timing and scale of these sorts of events with any precision. But financial institutions and policymakers could – and should – have been more alert to the vulnerabilities that were building during the years leading up to the crisis and therefore to the possibility of a major shock to the financial sector and to the economy more generally. It is somewhat analogous to seismologists trying to predict earthquakes along a fault line. It is impossible to predict the day and magnitude of a shock with any precision, but it may be possible to say something about the likelihood of an earthquake occurring within a given period from seismic measurements and indicators of latent stress.

Now it might seem natural to measure the severity of a recession by the magnitude of the fall in output. Indeed, a recession is sometimes characterised in the media as being two or more consecutive quarters of falling output, with the end of the recession being marked by a return to expansion. But that is not a very useful definition, as the mere resumption of growth will still leave some plants idle and some workers unemployed. From an economic perspective, it should be the margin of unused resources in the economy – the difference between actual and potential output – that we care about. And it is also the margin of unused resources that determines how much the economy can grow during the recovery before it puts significant upward pressure on inflation³, the control of which is the MPC's primary objective.

The problem is that the level of potential output is inherently difficult to measure. For an individual firm, the notion of capacity is slippery. A car manufacturer might be able to say how many units can be produced if all production lines are operating 24/7, but he knows that to do so would prevent essential maintenance and impair the firm's ability to produce output later. And in the service sector, the notion of capacity is even more complex. What is the capacity of a law firm, for instance? One might be able to say how many hours the existing partners could work, but that is not the same as how much output they could produce. Moreover, it is impossible to match unemployed workers to unfilled vacancies instantaneously, so it is not appropriate to ask what level of output could be achieved if all

³ Of course the margin of spare capacity is only one determinant of inflation. Other factors include commodity prices, the exchange rate and inflation expectations. And the rate at which the margin of spare capacity is closed may also matter – so-called “speed limit” effects.

those who wanted a job were actually in one. The frictions and the institutional characteristics of labour and product markets need to be taken into account too.

Economists studying business cycles have used a variety of approaches to construct proxies for potential output. One is to assume that potential output evolves smoothly over time and that actual output fluctuates around this; potential output can then be proxied by a smooth trend line fitted through the more jagged path of actual output. A more structural approach takes measures of capital and the available labour force (after allowing for any labour market frictions) and then assumes something about the nature of the technology with which they are combined to produce output. A third approach eschews direct measurement of potential output altogether and instead uses survey-based measures of the margin of unused resources to draw inferences about the evolution of potential output.

While judging the margin of spare capacity is always a problem for policy makers, it is particularly difficult at the current juncture because a banking crisis accompanied by a deep recession is likely to lead to some impairment of the economy's supply capacity. Moreover, different approaches presently suggest very different degrees of supply impairment. Suppose we use the first of the approaches above, assuming both that the economy was operating around potential in the years leading up to the crisis and that potential output had continued to rise at broadly the same rate throughout the recession, reflecting continued growth in human knowledge. That approach would point to output being a little under 10% below potential at the current juncture, an enormous margin of unused resources.

But we get a very different picture if we look at direct measures of the margin of unused resources. Chart 3 shows a swathe of alternative indicators of capacity utilisation drawn from business surveys. They suggest that the margin of spare capacity within firms is now relatively modest. As far as the labour market goes, the rise in unemployment has been only a little over two percentage points, notably less than in earlier UK recessions even though the fall in output has been greater; see Chart 4. Together, these observations suggest a far more modest margin of unused resources and, by implication, a substantial depression of potential output resulting from the Great Contraction.

There is, however, an important qualification. The counterpart to the modest rise in unemployment is a much smaller fall in employment than would have been expected on the basis of past experience, in part reflecting a high degree of wage moderation in the private sector. This can be seen in Chart 5, which contrasts the evolution of employment and output with the experience during the recession of the early 90s, when output fell less but employment more. The corollary is that productivity growth has been extremely weak during this downturn. That is in stark contrast to the United States, where productivity growth has remained strong and there has been a large shake-out of labour. Now this could be indicative of a fall in potential output. But it could also indicate a considerable degree of labour hoarding, which would be consistent with a rather greater margin of underutilised resources than suggested by Charts 3 and 4.

A number of recent cross-country studies⁴ have documented the substantial and long-lasting effect on output of past recessions caused by banking crises. That is illustrated in Chart 6. The dark blue line shows the mean fall in output relative to a continuation of the pre-crisis trend (calculated omitting the three years preceding the crisis) across 88 previous systemic banking crises in advanced and emerging market economies. The average decline in output, relative to a continuation of the pre-crisis trend, is around 10% and remarkably persistent, suggesting that ultimately aggregate supply follows aggregate demand down, even if it does not fall immediately. The UK experience (red line) has so far been fairly similar to that average past experience, although the most recent quarterly outturns have been somewhat

⁴ See, for instance, the IMF's October 2009 *World Economic Outlook* and Carmen Reinhart and Ken Rogoff, *This Time Is Different*.

stronger than one might have expected. Chart 6 also reveals, however, a very marked heterogeneity across episodes, with the blue swathe covering the central 50% of the distribution of outcomes and the grey swathe 90% of the distribution; a substantial permanent loss in output is by no means pre-ordained.

There are a variety of ways that a banking crisis might impair an economy's supply capacity. Reduced availability of credit and lower activity are both likely to depress investment, which has indeed fallen sharply. They are also likely to result in more business bankruptcies and fewer new firms being formed, though here the evidence is rather more positive, with liquidations rising much less during the recession than one would expect (Chart 7). Finally, as we saw in the eighties, high levels of unemployment can become entrenched if the long-term unemployed are allowed to become disconnected from the labour market. But, as Chart 4 shows, the rise in longterm unemployment has so far been modest, while the exit rate from long-term unemployment into jobs has been broadly flat.

So, at the current juncture, some indicators seem to suggest a rather large margin of spare capacity, while others point to a much smaller margin. Now it may be possible partly to reconcile these apparently conflicting observations once one allows for the fact that in some businesses, especially in manufacturing, it may be possible to shut down some capacity temporarily, reactivating it at relatively low cost once conditions improve. To the extent that the survey responses relate to the immediately operable capacity, rather than potential capacity, it would be possible to have simultaneously a modest margin of spare capacity in the short run, but rather more in the long run. Even so, the generic point remains that we presently have only an imprecise idea of the margin of unused resources in the economy. But this is a key determinant of how much the economy can grow before igniting inflationary pressure, as well of the size of the structural fiscal deficit that needs to be closed.

Would better or different statistics have helped us in anticipating the recession, in monitoring its evolution and in setting policy? Aside from the important, but unfortunately rather challenging, matter of measuring the margin of unused resources in the economy, I do not believe that the recession has highlighted any obvious lacunae in our conventional macroeconomic indicators, though it has certainly presented us with plenty of puzzles. (It has highlighted shortcomings in economic models, but that is a different story.) Rather the key information gaps have been inside the financial sector. Some of these we have been able to fill. For instance, in 2007, following the example of the Federal Reserve and the European Central Bank, we had happily just decided to launch a regular survey of the banks in order to provide more information on credit conditions. That has proved invaluable during the recession as a guide to how the availability of credit was evolving. If we hadn't already decided to launch such a survey, we would certainly have needed to introduce one.

We were already quite well served with information of key monetary and credit aggregates, which have naturally taken a more prominent role in our analysis during the crisis. We did, however, find that following the collapse of Lehman's, we needed a richer and timelier picture of banking sector developments. We therefore instituted an additional regular survey of just the major lenders; some of the resulting data appears in our publication *Trends in Lending*, together with a commentary on credit developments informed by the associated discussions with the lenders. We would, however, have liked access to more information on the evolution of credit by industry and firm size so as to get a better handle on where credit constraints were biting hardest. Gaining access to micro-prudential supervisory data relating to individual institutions collected by the FSA has proved invaluable in assessing the resilience of the banking system and thus its ability to supply the credit to support the recovery.

One area where more information would be useful is on the activities of the shadow banking sector. Almost inevitably, one has less knowledge of these institutions than of regulated entities. But we have seen that serious problems can arise in such hidden corners of the financial sector. High-quality flow of funds data, similar to that available in the United States, would also be valuable. That would facilitate a better understanding of the factors influencing

changes in agents' balance sheets, and thus be valuable in the pursuit of both monetary and financial stability. Cross-sectional information on balance sheets within the household and business sectors could also be useful in evaluating the seriousness of financial strains. Finally, I think it is worth noting that if so-called macro-prudential policies are to be used successfully to improve the resilience of the financial system and to moderate the credit cycle, financial data collection will need to be more flexible in future. Almost certainly the seeds of the next financial crisis will sprout in a different corner of the financial system from this one. Collecting data dedicated to helping fight the last war will not serve, if those emerging risks are to be identified early and dealt with appropriately.

Thank you.