Forecast evaluation report – October 2017

Robert Chote, Chairman, Office for Budget Responsibility

Good afternoon everyone.

My name is Robert Chote, chairman of the OBR, and I would like to welcome you to this briefing on our 2017 *Forecast evaluation report*.

[SLIDE] The FER is a publication we produce every year, looking back on the performance of some of our recent economic and fiscal forecasts relative to the latest outturn data. We produce it in part for reasons of transparency and accountability, to help people understand how we put our forecasts together and to appreciate the uncertainty that always lies around them. But we also use it to draw lessons that we hope will improve our future forecasts.

This year's FER focuses on the economic and fiscal forecasts we made in March 2015 – our last under the Coalition Government – and March 2016 – our last before the EU referendum. But it puts them into the context of our broader forecasting record. In this year's report, we also describe a review of the individual fiscal forecasting models we use and give some early results.

[SLIDE] In my remarks today, I will focus on those parts of the report that have lessons for our forthcoming Budget forecast in November. I will start by looking at our public sector net borrowing forecasts since 2010, focusing in particular on the latest complete fiscal year, 2016-17, and on the unusually large revisions to the outturn data that we have seen since our last forecast in March. I will then turn briefly to our GDP forecasts, before focusing on how that has been reflected in productivity and the labour market. Finally, I will look at the economic assumptions we made regarding the Brexit vote and how they compare with developments to date. The slides and words will be available after I finish.

[SLIDE] Let me start by reminding you of the big picture, as regards public sector net borrowing. As you can see here, the government was running a budget deficit between 2 and 3 per cent of GDP in the immediate run-up to the financial crisis in 2008-09. The deficit then ballooned to 10 per cent of GDP in 2009-10, but has since returned to pre-crisis levels – thanks both to the recovery of the economy and to spending cuts and tax increases. Back in

March we forecast that the deficit would widen a little this year and then shrink to 0.7 per cent of GDP by the end of the five-year forecast in 2021-22.

It is important to remember that this forecast - like all our forecasts - is not our best guess of what will actually happen, but of what would happen under current Government policy and if interest rates move in line with the market expectations implied by interest rate swap prices. Other forecasters may assume changes in government policy or different paths for monetary policy.

[SLIDE] Needless to say, no forecast for the public finances turns out correct in every respect. And this can be for any of four reasons:

- First, the Office for National Statistics may change the definition of the fiscal aggregate that we are forecasting after we have forecast it;
- Second, the Government may (and usually does) change tax and spending policy after we have made our forecast;
- Third, the economy may evolve in ways we did not expect, with automatic effects on revenues, spending and the deficit; and
- Fourth, revenues and spending may differ from our expectations for reasons that have nothing to do with our economy forecast, for example the performance of welfare reforms or changes in the amount and type of alcohol people drink - more gin recently, as it happens.

Much of this report is an exercise in looking at the latest data for the public finances and decomposing the differences with our earlier forecasts into these four factors. We refer to the gaps between them as forecast differences, rather than errors, as it would have been impossible to avoid many of them - particularly the first two - given the information available at the time.

[SLIDE] To give you a sense of how big these differences typically are, here is our latest forecast for public sector net borrowing, from 2009-10 this time, [SLIDE] and our previous forecasts. In order to compare like with like, we have adjusted the earlier forecasts for the largest definitional changes that the ONS has made in this period, namely aligning the data with the 2010 European System of Accounts and reclassifying housing associations to the public sector.

As you can see, the remaining differences between forecast and outturn can be significant. They arise not just because new data differ from the forecast, but also because earlier data are revised with the passage of time.

[SLIDE] Looking at the differences between our forecasts for the budget deficit and the latest outturns over a two-year time horizon, you can see here that some have been too pessimistic and some have been too optimistic - more often the latter than the former, especially up to 2012. As you can also see, revenue and spending outturns both typically differ from the forecasts we make, with larger differences on average for revenues because they are more directly affected by developments in the economy.

[SLIDE] As you can see here, the differences to date between our forecasts and the latest data have generally been smaller than the average differences between Treasury forecasts and outturns over the previous 20 years - as shown by the green blocks. [SLIDE] That is particularly so for our spending forecasts [SLIDE] and to a lesser extent for revenue. But this does not prove that we are better forecasters than the Treasury were, as we are looking at different periods with different economic and policy developments.

[SLIDE] As regular readers will know, we routinely illustrate the uncertainty that lies around our forecasts by drawing probability distributions or fan charts around our central predictions, showing the confidence that you can place in them based on the size and distribution of past forecasting differences.

This chart shows the fan chart around our March 2012 central forecast for the budget deficit, [SLIDE] and the latest outturns. As you can see, the latest data suggest that we underestimated the deficit by around 1.6 per cent of GDP towards the end of the forecast. But the much greater width of the fan at that horizon shows that this is not a particularly big absolute difference - there would have been a more than 60 per cent chance of an even bigger one, based on average forecast performance in the past.

[SLIDE] So now let me now focus in more detail on the deficit last year - 2016-17 - and what implications this might have for future years.

[SLIDE] This chart shows that back in March 2016, just before the beginning of the fiscal year, we predicted a deficit in 2016-17 of £55.5 billion. [SLIDE] By November 2016 we had data for the first six months of the year, shown in yellow, and they were worse than expected. So we revised our forecast for the full year up to £68.2 billion. [SLIDE] By the time of the most recent Budget in March, the ONS had revised the deficit for the first six months down by almost £10 billion and the data for the subsequent four months was slightly better than expected, so we brought the full year forecast down to £51.7 billion.

[SLIDE] When the ONS published its first estimate for the full year a month later, it was just £300 million higher at £52.0 billion, and adjusted for minor differences in classification treatment the difference was only £5 million. But the forecaster's old friend 'offsetting differences' was behind that, as a bigger-than-expected deficit in February and March outweighed a downward revision to the deficit in the previous ten months.

But that is not the end of the story - [SLIDE] over the following five months (after the fiscal year was over) the ONS has revised down its initial estimate by £7 billion to £45 billion. This is a relatively big revision by comparison with previous years and the numbers may well move again - up or down.

One conclusion to draw is that you be should very wary of treating the monthly flow of public finances data through the year - and even the initial outturn for the year as a whole - as a reliable guide to where the numbers will eventually settle down. Accurate estimates of spending, revenue and borrowing by local authorities and public corporations only become available with a long lag. The new accounting treatment for corporation tax receipts means that cash payments as far away as December 2018 could affect the estimate of accrued borrowing in 2016-17 - although the amounts involved would be small by then.

[SLIDE] As we approach the November Budget forecast, the key question for us is whether the recent downward revisions to the 2016-17 deficit are likely to feed through to future years and improve the outlook relative to our March forecast. So where have the differences from our March forecast come from?

As this table shows, £1.3 billion of the downward revision comes from statistical changes, mostly an imputation related to public sector pension schemes. These will feed through to future years, lowering the deficit forecast. Income tax and national insurance receipts were £3.5 billion higher than expected in March, more than half of which reflects stronger-than-expected bonus payments. VAT and corporation tax receipts were both almost £1 billion higher than expected, reflecting strong cash payments at the start of 2017-18 that were accrued back to last year. In each case, you might expect some of this unexpected good news to push through.

On the spending side, central government spending was almost £2 billion lower than expected, largely because government departments underspent their plans by more than we had expected. But there is no mechanical link between this and our judgements for future years – we will need to look afresh at the pressures on departmental budgets closer to the forecast.

The main area where the public finances look weaker than they did in March is for local authorities - they borrowed almost £2 billion more than we expected last year. Some councils have drawn down their cash reserves by more than we thought, especially those with responsibility for social care. But, more significantly, this also reflects higher so-called 'prudential' borrowing than we expected. As we noted in our *Fiscal risks report* in July, this may in part reflect councils undertaking potentially risky commercial development projects to boost their revenue. We will need to look carefully at these two elements to see how much if any of the extra borrowing will be repeated in future years.

So looking at 2016-17 as a whole, what we can say at this stage is that some of the unexpected good news since March is likely to feed through to lower deficits in the current and future years, but by no means all of it.

[SLIDE] As regards the performance of the public finances so far this year, borrowing during the first five months of 2017-18 is little changed from the same period last year. Back in March we forecast a total deficit of £58.3 billion, up £6.5 billion on our then estimate for last year but up £13.2 billion on the latest outturn, following those downward revisions.

We expect the comparison to look less favourable later in the year, as the shifting of dividend income to beat the April 2016 dividend tax increase unwinds and depresses self-assessed income tax receipts in January and February. Hence we expect a 13 per cent fall in SA receipts over the full year, compared to the 6 per cent rise on last year that we have seen to date. Elsewhere, as this table shows, PAYE income tax and NICs are growing more strongly than we forecast for the year as a whole, but this may be because they were depressed in the first half of last year. Growth in VAT receipts and central government spending are both broadly in line with our full-year forecast. Higher RPI inflation means debt interest payments are up 17 per cent on a year ago, close to our full-year forecast of almost 15 per cent.

As well as reflecting the latest outturn data for spending and receipts – and the near-term impact of any movements in the economy – our November forecast will also reflect some changes to individual fiscal forecasting models. Changes to our North Sea oil and gas modelling may push receipts down a bit, but, working in the opposite direction, changes to our modelling of index-linked gilts could push debt interest spending down a bit too.

Other things being equal, all this suggests that we are more likely to reduce our forecast for the deficit this year than increase it come November. But other things may not be equal and, as we have just discussed, within-year data can be a poor guide to the eventual outcome. More importantly, changes to our medium-term view of the economy are also likely to be less favourable for the public finances than the recent outturns, as I will explain in a minute.

Any forecast for the public finances over the medium term has to be based on a view about how the economy is likely to evolve, as public spending and especially tax receipts are closely linked to different components of national income and spending, to asset prices and transactions, to interest rates and inflation. When forecasting the outlook for tax revenues in particular, we are generally interested in cash measures of income and spending. But most public discussion of economic forecasts focuses on real measures - the volume of goods and services produced in the economy. So let us have a look at how real GDP growth has performed relative to our forecasts and how that has been reflected in productivity and the labour market.

[SLIDE] As with the public finances, the differences between our growth forecasts and the latest outturns have generally been smaller than they were on average under the Treasury. Once again this does not prove that we are doing a better job - the outturn data could still move and we have not yet had to forecast over a recession, when the differences tend to be much larger.

[SLIDE] This chart shows the differences between our Budget forecasts for real GDP growth looking one year ahead and the latest outturn data. As you can see, there is no clear bias in either direction - some were optimistic, some were pessimistic and some were right on the money (for now at least).

But a more consistent pattern emerges when we decompose GDP growth into the contributions of employment growth, changes in the average hours people work and changes in productivity – the amount workers produce every hour. Here you can see that whether our growth forecasts were optimistic or pessimistic, [SLIDE] productivity growth has been consistently weaker than expected and [SLIDE] employment growth and hours worked consistently stronger. One of the most important decisions we shall have to make in

November is what to assume about these contributors to GDP growth over the medium term.

[SLIDE] The disappointing performance of productivity growth since the financial crisis is a familiar picture and a much studied one. Having grown by an average of 2.1 per cent a year over the three and a half decades prior to the crisis, productivity has now been virtually flat since 2011. Indeed, hourly productivity is now more than 20 per cent below its pre-crisis trend.

This chart shows our forecasts for productivity growth since June 2010 and the latest outturn data. In the face of repeated disappointments we have always assumed that productivity growth would pick up to something approaching its long-run historical rate over the medium-term, but that the level would not rebound to its previous path as it always had following previous recessions. But even this relatively pessimistic view - that the financial crisis had done permanent and large, but not *increasingly* large, damage to the economy's productive potential - has not been pessimistic enough.

[SLIDE] Our assumption that productivity growth would return to a more normal rate within a few years reflected a judgement that whatever factors were depressing it in the wake of the financial crisis would fade as it receded further into the past. But as the period of weak performance gets longer, the explanations that people pointed to immediately after the crisis look less convincing and others seem more plausible. However, the chances are that different factors have played more or less important roles at different times.

Let's have a look at some of the explanations that have been put forward.

- In the early post-crisis period it seemed plausible that firms were holding on to labour in the face of temporarily weak demand, because it would be expensive to let people go and then have to rehire them again. But as the economy began to recover – and as hiring picked up alongside the recovery in demand – that looked less plausible.

- Later we put weight on the argument that weak balance sheets and other problems in the financial system were preventing the efficient reallocation of labour and capital from weaker to potentially stronger firms. This meant that 'zombie' firms depressed productivity for longer, while potentially productive ones could not grow as quickly as they otherwise might. But this seems less plausible today, as the financial system is in better health and credit is more readily available - thanks in part to the policy measures put in place by the Bank of England.

- Some economists have argued that very low interest rates may themselves be depressing productivity growth, even if they support demand, by helping weaker firms service their debts and slowing the reallocation of resources. Interest rates are only expected to rise very gradually, and to remain low by historical standards, so if this is an explanation it looks likely to persist.

- More recently, productivity growth was expected to pick up as a tighter labour market put upward pressure on wages and forced firms to seek productivity gains rather than hiring more people or increasing their hours. But even though unemployment has now fallen to its lowest rate in decades - and below the rate we have assumed consistent with stable inflation - wage growth remains subdued. With many other indicators pointing to a tight labour market it isn't clear whether the response has simply been delayed or if something more fundamental has changed.

- [SLIDE] The fall in business investment during the financial crisis - in both intellectual property and physical assets - has slowed the pace of capital deepening and depressed productivity growth. Business investment today is just 5 per cent above its pre-crisis peak almost a decade ago, compared with increases of 60 and 30 per cent following the 1980s and 1990s recessions. If investment remains weak, this would continue to depress productivity growth relative to the pre-crisis period.

[SLIDE] Looking internationally, it is also striking that the recent weakness of productivity growth is not unique to the UK. This chart shows that our counterparts at the Congressional Budget Office have repeatedly had to revise down their forecasts for productivity growth in the United States.

[SLIDE] This chart shows the OECD's November 2014 forecasts for productivity growth in 2015 and 2016 in each country of the G7. [SLIDE] And these are the latest outturn figures - worse than expected in every country, and in some cases by more than in the UK. Weak investment and abnormally low interest rates are plausible contributors across several of these countries. Some economists even believe that advanced economies have entered an era of permanently subdued productivity growth, either because there is less to be gained from technological progress or because central banks cannot deliver real interest rates low enough to deliver adequate growth in demand.

[SLIDE] So where does this leave us come the November forecast? This chart shows the latest outturn data for annual productivity growth and our forecasts for the next five years, back in March. When we made this forecast, output per hour had grown by 0.4 per cent per quarter on average in 2016 and was 1.5 per cent higher than a year earlier by the final quarter. So, at the time, it did not seem unreasonable to assume growth rates of close to 2 per cent a year over the medium term - reflecting both a gradual pick-up in potential productivity growth and some cyclical rebound towards that trend.

But the data since have disappointed once again. The ONS has revised average quarterly growth through 2016 down to 0.3 per cent and annual growth in calendar year 2016 to just 0.2 per cent. Productivity actually fell by 0.7 per cent over the first two quarters of this year, although that may be a temporary effect of the Brexit vote and the uncertainty it has generated. This all implies that the apparent pick-up through last year was another false dawn, like that we saw in the year to mid-2015 that had evaporated six months later.

All this suggests that it would no longer be central to assume that productivity growth will recover to the 1.8 per cent a year we assumed in March over a five year horizon. We will take a final decision in November, on the basis of all the data available to us then. But for now we are minded to revise down potential productivity growth significantly, but without going so far as to assume there is no recovery at all from the very weak performance of recent years.

Weaker potential productivity growth would imply weaker growth in real GDP if the Bank of England manages demand in the economy to meet and maintain the inflation target. That in turn implies weaker growth in tax receipts and a significantly less favourable outlook for the public finances.

As I noted earlier, productivity growth may have been a consistent disappointment in recent years, but the impact on GDP growth has been at least partly offset by stronger-than-expected growth in employment and by people working more hours on average than we expected. We will need to take these trends into account in our November forecast too.

[SLIDE] As this chart shows, the unemployment rate has fallen more rapidly than we expected since mid-2013. Indeed, it has continued to fall even after we expected it to stabilise and then rise somewhat in our more recent forecasts.

Back in March we assumed that the sustainable rate of unemployment, at which inflation and wage growth would remain stable, was around 5 per cent of the labour force and would rise slightly to 5.2 per cent over the next five years as the National Living Wage rises more quickly than productivity and prices some people out of work. But with the unemployment rate already down to 4.3 per cent, and no sign of serious upward pressure on wages and domestically generated inflation, this looks too high. For what it is worth, the Bank of England assumes a rate of around 4½ per cent. A lower sustainable rate of unemployment might imply more spare capacity in the economy and therefore more scope for the economy to grow. This would offset some of the hit from assuming weaker potential productivity growth.

Another possibility - not mutually exclusive - is that the relationship between unemployment on the one hand and wage and price inflation on the other is weaker than it used to be – or, as economists would say, that we now have a pretty flat Phillips Curve. In that event you would not expect inflation and wage growth to rise very much even if unemployment did fall below its sustainable rate. But if inflation did end up above target it would be costly to bring it down again.

[SLIDE] As regards average hours worked, these have been on a long-term downward trend for centuries, with people choosing to take some of the fruits of rising real living standards as leisure rather than income. But average hours rose following the financial crisis, as people chose to protect their spending power against the impact of weaker earnings growth and reduced wealth.

In our forecasts to date, we have assumed that the long-term trend would reassert itself within the five-year forecast. But if we take a more pessimistic view of productivity growth over the medium term - and therefore of growth in earnings - it would be logical to assume that this too will take rather longer. Like a lower sustainable unemployment rate, this would boost growth over the medium term and help offset a weaker outlook for productivity.

[SLIDE] Before wrapping up the various portents for November, let me briefly remind you of the key economic assumptions that we made when reflecting last year’s Brexit vote in our November 2016 and March 2017 forecasts - and how they look in the light of subsequent data.

We made five main judgements:

- First, we assumed that the fall in the pound would squeeze real incomes and growth in real consumer spending, weakening GDP growth. This judgement seems broadly on track, although real consumption held up slightly better than expected as households ran down their savings.

- [SLIDE] Second, we assumed that businesses would postpone and cancel some investment because of uncertainty about the outcome of the vote and then the outcome of the negotiations. As I noted earlier, investment has been weak for some time, but the latest data show it falling slightly in the year before the vote and rising slightly thereafter. However, this in part reflects investment in aircraft that has long lead times. And, more generally, the aggregate data are so volatile and prone to revision that it is hard to place much weight on the precise quarterly profile.

- [SLIDE] Third, we assumed that the depreciation of sterling would boost net exports in the short term, partly offsetting the hit to GDP growth from consumption and investment. Net trade has indeed boosted GDP growth, but by less than we expected in our forecast last November.

- Fourth, we assumed that leaving the EU would reduce export and import growth in parallel over the transition to a less trade-intensive economy. It is too early to assess this judgement.

- Fifth, and finally, we assumed that net inward migration would fall. To begin with this would reflect a weaker 'pull factor', for example as the fall in the pound reduces the value of UK wages in would-be immigrants' home currencies. Later we assumed that the UK would move to a tighter migration regime. [SLIDE] The latest data do indeed suggest that migration is falling in response to weaker pull factor, but the Government has yet to set out a future migration regime.

[SLIDE] Taken overall, the initial Brexit effect does seem to be broadly in line with expectations. GDP growth has slowed noticeably coming into 2017 – slightly later than we expected in November and slightly sooner than we expected in March. We do not expect to make any significant changes to our Brexit assumptions come November, as there remains no meaningful basis to predict the precise outcome of the negotiations and no statement by the Government of the changes in domestic policy that would accompany it.

[SLIDE] Finally, let me conclude. The Budget is now just six weeks away and we are already working on the forecast that we will publish alongside it. There are numerous moving parts and it is far too early to say exactly what it will show.

That said, there are some lessons from today's report that are likely to feed into it. On the one hand, we expect to revise down our assumption for potential productive growth, which would worsen the outlook for the public finances. On the other hand, we expect to revise down our estimate of the sustainable rate of unemployment, to revise up our projection for average hours worked and to push some of the downward revisions to last year's budget deficit through into this and future years. All these would improve the outlook for the public finances. That said, it is the downward revision to productivity growth that is likely to have the biggest overall impact.

And, with that, we are happy to take any questions.