



BANK OF ENGLAND

Speech

Telling tails of oil and global inflation

Speech given by

Martin Weale, External member of the Monetary Policy Committee

Association of European Conjuncture Institutes conference, London

21 May 2015

I am grateful to Nicholas Fawcett, Tom Key and Lotte Adams for their help in preparing the text. Some of the material in this speech draws on research conducted with Nicholas Fawcett and Andrew Harvey. I would like to thank Andy Blake, Frank Eich, Kristin Forbes, Andrew Harvey, Tsveti Nenova, Ivan Petrella, Nyssa Roberts, Greg Thwaites, Seb Walsh and Tomasz Wieladek for advice. I am particularly grateful to Bob Gilhooly and Lucia Quaglietti for their help with international data, and to Andrew Patton for sharing computer code. The views expressed, though, are my own, and are not necessarily those of the Monetary Policy Committee or the Bank of England.

Introduction

May I start by thanking you for inviting me to speak here today. I would like to talk about two important international issues, both of which are highly relevant for monetary policy-making in the United Kingdom. The first is the effect of the recent sharp fall in the oil price on both GDP and inflation. The second is the more general question of whether or how far underlying UK inflation is influenced by inflation abroad.

Addressing these issues, I find we must pay particular attention to two key facts. First, the price of oil is erratic, and secondly inflation both in the UK and in other advanced countries is unusually low. In both cases I am trying to understand the possible effects of “tail” events. This may sound like an unnecessary complication, but, as was demonstrated by the shock of the 2008-09 financial crisis, people frequently underestimate the probability of tail, or extreme events, and their consequences.

The impact of oil price shocks on the economy


The power of oil to shock is demonstrated by Chart 1. This shows headline real GDP growth and CPI inflation in the OECD since 1970. I have also highlighted periods when the oil price moved sharply over a particular quarter, with rises of over 20% marked in red, and falls of more than 20% in green. The effect of the two big price increases of the 1970s on both growth and inflation is very clear. More generally there is a suggestion that periods of rising oil prices were followed by periods of weak growth and rising inflation. The reverse is also true at least to some extent, after periods of falling oil prices. What does economics tell us about the likely impact of oil price movements?

Household spending is likely to have been stronger in the face of cheaper oil, as consumers have been left with more money to spend. This in turn should add to demand overall. For oil importers, there is a favourable movement in the terms of trade and in real national income; so the increase in spending is able to take place without any increase in borrowing from abroad. The picture is slightly different for a country which produces its own oil. Though a fall in the price of oil does not affect national income, it does depress the profits of oil well owners. And since everyone else sees an increase in their real incomes, overall spending can still rise. At the same time businesses face lower costs as a result of the lower oil price and also enjoy higher demand for their output. This combination is likely to encourage investment spending, raising the demand for capital goods. Once again, in countries which are oil produces, this effect may be offset by a reduction in investment by the oil industry.¹

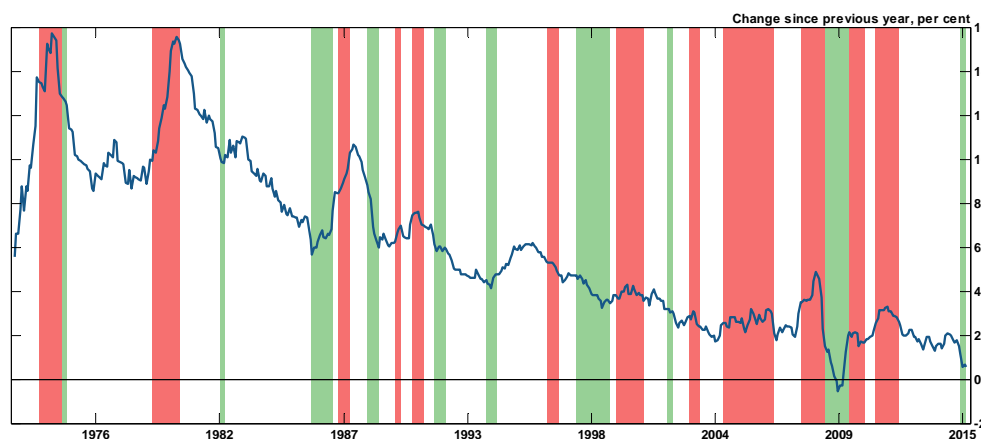
¹ Although in the UK the import content of oil investment is high. This effect, while it may be important for investment, is probably less important for GDP.

Chart 1: Growth and headline inflation in the OECD

(a) OECD GDP growth



(b) OECD headline CPI inflation



Source: OECD and Bank calculations. **Note:** shaded areas denote periods in which the oil price changed sharply during a quarterly period. Red denotes an increase of more than 20% over the quarter, and green a fall

Expectations play an important role in determining the size and timing of the increase in spending. If the fall in oil prices is expected to be temporary, then spending is likely to be boosted less than if people see it is as permanent. Spending may build up only gradually if people take a while to believe that a movement is permanent.

Whether or not the price fall is permanent, it has an immediate and one-off effect on overall inflation. An important part of consumer spending goes, directly or indirectly, on oil products or goods and services produced using oil. First to move, typically, is the price of petrol, followed by other energy costs. As this feeds through the supply chain, it has an impact on the price of other goods and services, but this takes longer. The MPC has, however, taken the view that such effects are largely likely to work through relatively quickly, so that a fall in the price of oil has little direct effect on inflation at a two to three year horizon.

But there is always a risk that the effect of inflation may be more persistent. The fall in actual inflation may influence expectations of future inflation or influence future pay bargaining, simply because pay is set with reference to recent inflation. In that case second-round effects could mean that a fall in oil prices affects inflation for longer than the initial effects would imply. Of course what matters more than simply whether such effects are present is their magnitude.

Table 1: The Effects of a Forty Per Cent Reduction in the Oil Price on GDP and Inflation (percentage points)

	Euro Area	UK	USA	World
Impact on the level of GDP				
Azreki and Blanchard (2014)	0.5		0.5	
Barrell and Pomerantz (2004)	0.5		0.6	
Carabenciov et al (2008)	0.2		0.8	
Cashin et al (2014) Supply Shock	0.3		0.3	
Cashin et al (2014) Demand Shock	-0.2		-0.3	
EC (2008)	0.2		0.0	
Hervé et al (2010)	0.8		1.2	
IMF (2015) Full Pass-through	0.4		1.6	1.2
IMF (2015) Limited Pass-through	0.4		1.6	0.8
Jimenez-Rodriguez and Sanchez (2004)	0.0		1.6	
Kilian and Vigfusson (2014)			0.8	
Kirby and Meaning (2015) Permanent Fall	1.8	1.2	3.2	
MPC (2015)		0.5		
Peersman and Van Robays (2012) Supply Shock	0.3		1.8	
Peersman and Van Robays (2012) Demand Shock	-2.2		-2.0	
Average impact of studies simulating oil supply shocks	0.5		1.3	
Impact on the annual rate of inflation				
Barrell and Pomerantz (2004)	-0.4		-0.7	
Carabenciov et al (2008)	-0.6		-1.1	
Cashin et al (2014) Supply Shock	-0.2		0.0	
Cashin et al (2014) Demand Shock	-0.2		-0.1	
EC (2008)	-0.5			
Hervé et al (2010)	-1.2		-1.6	
IMF (2015) Full Pass-through				-2.8
IMF (2015) Limited Pass-through				-2.0
MPC (2015)		-0.6		
Average impact of studies simulating oil supply shocks	-0.5		-0.7	

Note. The figures show, where relevant, the maximum effect. The table draws heavily on Riksbank (2015). Where sources quote the effects of a price change different from 40%, the figures have been scaled pro rata to give those in the table.

Table 1 lists a range of studies that have estimated the effect of oil prices on GDP and inflation. Since they each consider a slightly different size of oil price change, I have converted them as best I can so that the numbers report the effects of a forty per cent fall in the price of oil on GDP and inflation; I have chosen this point of comparison because the oil price currently stands around 40% lower than it was in June last year.

Three points stand out from these studies.

First, although they point to a wide range in estimates of the effect the general impression is that the impact on GDP is thought to be greater in the United States than in the euro area. For the effects on inflation the differences are much smaller; only slight larger effects are shown for the United States. The estimates do depend on how far the recent drop is expected to persist. Thus the IMF assumes that, after a sharp fall, the real price of oil rises over time so that after five years the initial real fall is halved, while Kirby and Meaning (2015) demonstrate that much smaller effects are expected if the price fall is temporary.

Secondly, the IMF's estimates of the impact on world inflation point to much stronger effects than any of the other numbers. As they observe, domestic oil prices in some developing countries are controlled by the government, which is likely to dampen the effect of the price fall (limited pass through), and in that case the impact on world prices is weaker than with full pass-through. That suggests that the full-pass through estimates offer a better guide to what they expect to happen in developed countries.

Thirdly, the cause of the change in oil prices matters a great deal. Peersman and van Robays (2012) distinguish the effects of an oil price change arising from oil supply changes from one resulting from global demand effects. The latter are not associated with much movement in GDP, although it is difficult to disentangle global demand effects from oil price effects; the movement in the oil price on its own may still have a powerful effect on GDP. The recent weakness in the oil price could be the consequence of oversupply or of weak demand. McCafferty (2015) suggests that the recent fall in prices is largely a consequence of rising supply.

There is also a suggestion that economies are now less sensitive to oil price movements than they were in the 1970s and early 1980s. Blanchard and Gali (2007) reported that, although the oil price increases of the late 1990s and mid-2000s were, in percentage terms, similar in magnitude to the first two oil crises, OECD output declined sharply during the first two crises but not in the second. Surging inflation was similarly much more of a problem during the 1970s and early 1980s than later on. They suggested a number of reasons for this. First of all, the crises of the 1970s were associated with other co-incident shocks; for example there were financial crises in 1973/4. Secondly, monetary policy became more credible later. Thirdly, labour markets have become more flexible and finally, the importance of oil in the economy had declined. While much of their analysis focuses on the United States, many of these observations certainly also apply to much of Europe. Most recently, the development of the fracking industry in the United States might suggest that the impact there of the price fall on GDP will be much dampened by lower investment by the oil industry.

Tails of the unexpected

In view of these sharp swings in oil prices over the past, making sense of the relationship between oil and the macroeconomy relies on being able to distinguish signal from noise. We want to pay attention to changes which send a strong signal about the economy, but place little weight on random noise. We can't do this properly, however, without first understanding how much noise there really is in the data. Though it might seem like a technical point, it is in fact fundamental: if you seriously underplay the chance of relatively extreme events happening, then not only will you be more surprised when they do happen, but you may be tempted to read too much into them.

If this seems too abstract, consider a concrete example relating to quarterly oil price changes over the past twenty-five years. A common assumption – underpinning many of the models referred to above – is that these follow a normal distribution, the shape of which follows a bell curve.² This implies that extremely large swings in the oil price are relatively rare. For example, the chance of a big quarterly fall – say 50% – is only 0.06%. This is small, though still about nine thousand times more likely than winning the UK's national lottery. Assuming that the probability stays the same over time, you would expect it to happen about once every four hundred years.

This prediction doesn't seem reasonable, however, when you consider that twice in the past twenty-five years alone, the oil price has changed by more than 50% in a quarter. Events that the bell curve says should be extremely rare in fact happen with much greater frequency. Indeed, the use of the bell curve lay behind the sense of complacency that built up ahead of the financial crisis of 2008-09. It was not, to use Donald Rumsfeld's language, even the 'unknown unknowns' that took people by surprise: it was that people grossly underestimated the chance of 'known unknowns' happening. Theory and evidence seem badly mis-matched.

One practical, if not perfect, solution comes from an unlikely source: in the words of the classic slogan, we must 'Ask for Guinness'. W.S.Gosset, an employee of Guinness in Dublin and later Chief Brewer in London, developed the *t* distribution in 1908 as a means of studying the variability of samples of barley.³ He needed an alternative to the normal distribution which recognised that the spread of data is greater in small samples than in large samples. In our own context, the chance of observing a large change in the oil price is much higher under the *t* distribution than it is with a bell curve.

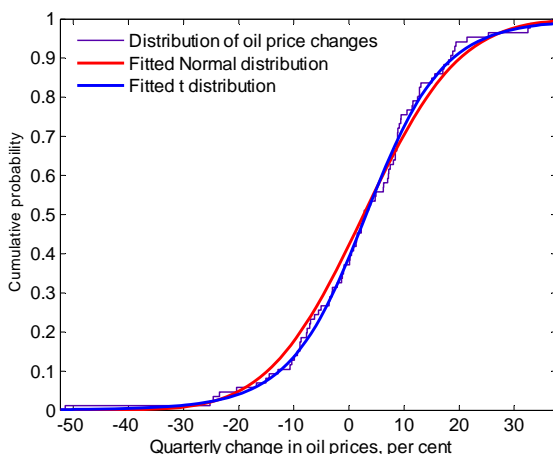
Looking at the same past changes in the oil price through the lens of the *t* distribution, I find that the probability of a 50% drop is now 0.6% – ten times greater than under the bell curve, or ninety thousand times more likely than winning the lottery. As another Guinness slogan put it, 'Good things come to those who

² As my colleague Andy Haldane (2012) has observed, this is also sometimes referred to as a Gaussian distribution, after Gauss (1809).

³ His classic paper, Student (1908) was published under a pseudonym to preserve the anonymity of his employer.

wait', and under the t , you would only have to wait forty years for it to happen, not four hundred. Even this seems to understate the actual likelihood, but it is an order of magnitude closer than the normal.

Chart 2: Distribution of oil price changes



The actual distribution of changes is shown in the jagged line, and the fitted distributions in red (normal) and blue (t).

The chart clearly shows that, while neither is perfect, the t distribution provides a closer fit to the data than the normal. The blue line runs through the data reasonably well, whereas the red line strays quite far. Just as I observed earlier, the reason for this is that oil prices are much more likely to change sharply from one quarter to the next, than the normal distribution allows.

This non-normality also extends to other variables of interest. The distribution of UK GDP growth, which many models typically assume is normal, turns out to have a t distribution as well. This adds to the complexity of reading the macroeconomic data and making decisions about monetary policy, as the task of differentiating outliers from more normal changes is more difficult. And the problem may only become worse if we use models that are based on incorrect assumptions, such as that all variables are normally distributed.

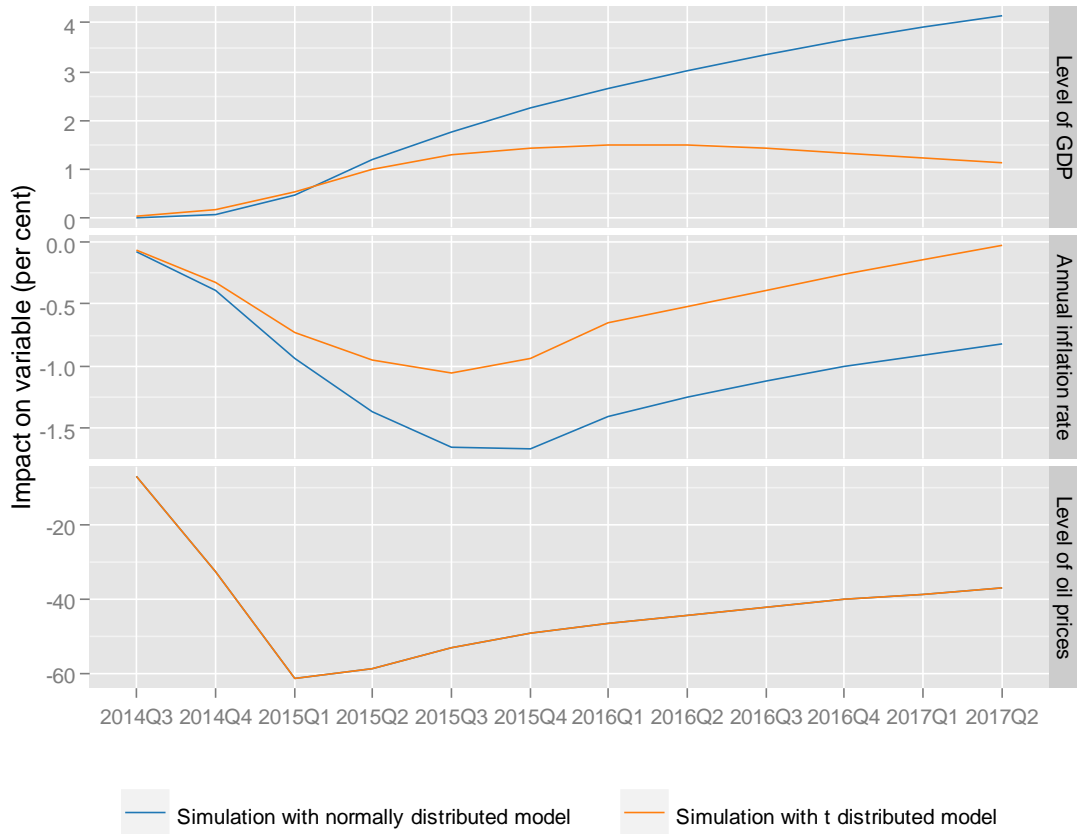
Seeing through the morass: the impact of the recent oil price fall

The near-universal popularity of the bell curve arises because of its great computational convenience. Fortunately, however, recent advances in computing have allowed us to explore economic relationships without being bound by restrictive assumptions about normality. In a set of important contributions, Harvey (2013) and Creal, Koopman and Lucas (2012) have developed a framework that helps us filter out random noise in the data, when the data may be t distributed, and outliers therefore more likely than in the

With that in mind, which of the two distributions, normal or t , more closely resembles the data of immediate interest? I can answer this by comparing the distribution of oil price changes that we have observed since 1993, to the predicted distributions from the normal and t distributions. These are shown in Chart 2. On the horizontal axis is the observed per cent change in the oil price from one quarter to the next, and on the vertical axis, the probability of a change at least as high as that. So, for example, the probability of having growth less than or equal to 0% is about 0.4, indicating that oil prices fell in about 40% of the observations

traditional case.⁴ We can get a better sense of genuine signals in the data, and understand how they feed through into the rest of the economy.

Chart 3: Impact of the oil price fall on UK GDP and inflation



Note: the top two panels show the impact on the level of oil prices and GDP respectively; the third shows the impact on the rate of annual CPI inflation

I can use this more robust approach to examine how far the oil price fall has affected GDP and inflation in the UK over the past nine months. The first step is to establish the relationship between oil prices, GDP growth, inflation, and domestic interest rates.⁵ Then, by superimposing the fall in the oil price, and subsequent but partial recovery, that we have witnessed since June last year, and together with some assumption about its future path, I can estimate its impact on the variables. Obviously the future path of the oil price is unknown. I use the path indicated in the futures market, taken at a point during the most recent MPC forecast round which, whilst it varies from one day to the next, provides as good a projection for the price of oil as any.

⁴ The shape of the t -distribution depends on the number of degrees of freedom. This is estimated as a part of the statistical analysis. For a high number of degrees of freedom the t -distribution merges into the normal distribution. Thus the technique does not exclude the bell curve if that is the most appropriate assumption.

⁵ The Appendix provides details of the model and data used.

I mentioned earlier that there were questions about both the extent to which spending adjusts to changes in incomes and about the extent to which the impact on inflation follows through into second round effects. The modelling approach implicitly deals with these points by assuming that the response to an oil price shock will be the same as it was in the past. The channels are not studied explicitly; rather everything is rolled up into an analysis of the relationship between movements in GDP and inflation on the one hand and the price of oil on the other hand.

Chart 3 shows the estimated impact of the oil price fall on GDP and inflation. The model suggests that GDP by about 1.5% by 2016, although this effect is likely to drop off thereafter. Inflation is around 1% lower by the end of 2015. But it then recovers, and the effect has faded almost completely by early 2017. While this is not, of course, the only way of looking at the issue, it does imply that there is a risk that GDP growth will be stronger, and inflation a little weaker, than the MPC assumed.

Given that I stressed the importance of the choice of assuming that variables were normally or *t* distributed, I find it very interesting that the estimates of the oil price impact under the *t*, in yellow, are very different to those obtained from a normal distribution, in blue. The GDP impact in particular, at over 4% in two years' time, seems to me to be implausibly large, and substantially higher than any of the estimates presented in Table 1. So I am more confident that the estimates from the model with a *t* distribution are closer to the likely effect, notwithstanding the uncertainty inherent surrounding any model estimate. Models can only go so far to capture such complex relationships.

Caveats

An important question I have not addressed here is whether the impact of oil prices is symmetric: does a 10% fall in prices boost GDP as much as a 10% increase might drag on it? One reason why that might not be the case could be that the capital stock adjusts to the expected oil price. When the oil price is expected to be high people use fuel-efficient production methods, while when oil is expected to be cheap, they are less concerned about saving fuel. In either case, a sharp movement in the price of oil in effect renders a part of the capital stock obsolete. This phenomenon was suggested as one factor behind weak productivity growth after the oil crisis of 1973/4. While it reinforces the effect of an increase in the oil price, it offsets to some extent the impact of a reduction.

The evidence of such asymmetry is mixed. On one hand, Hamilton (2003; 2008) finds not only that increases in the oil price seem to be associated with sharper reactions in GDP growth, but that the size of the change also matters. He suggests that price rises act as a brake on growth only if they are unusually large compared to recent changes. In contrast, Kilian and Vigfusson (2011) study the impact of oil price changes on the US economy, and using a different method to Hamilton, find no evidence that the response is asymmetric. This has been echoed in more recent work for a range of countries by Herrera, Lagalo and Wada (2015).

The role of global disinflation

In any case, the fall in the oil price tells only part of the story, as even before it started, inflation was low and falling in several countries. Even after the effects of the price fall fade away, could other global pressures continue to bear down on inflation? How far is inflation in the UK low because underlying world inflation is low.

Chart 4: Core and Headline inflation in the OECD




Source: OECD and Bank calculations.

You might well argue that to look at inflation excluding the items whose prices are falling most rapidly, misses the point. After all, the MPC is very clear that its own target is for inflation as a whole, and not for those components of inflation which deliver an answer close to two per cent.

A reasonable response to this is that food and energy prices are volatile. Movements in these lead to sharp movements in inflation, as Chart 4 shows. To the extent that these movements are the result of level changes in the prices of food and oil, they drop out of the inflation measure after a year. These effects certainly help explain why inflation can be volatile, but, unless food and energy prices are expected to fall steadily, they are not a harbinger of sustained weak inflation. Rather, they are some of the influences that the MPC has, in the past, looked through. Arguably, the measure excluding food and energy is a better indicator of underlying inflationary pressures, at least in normal circumstances.⁶

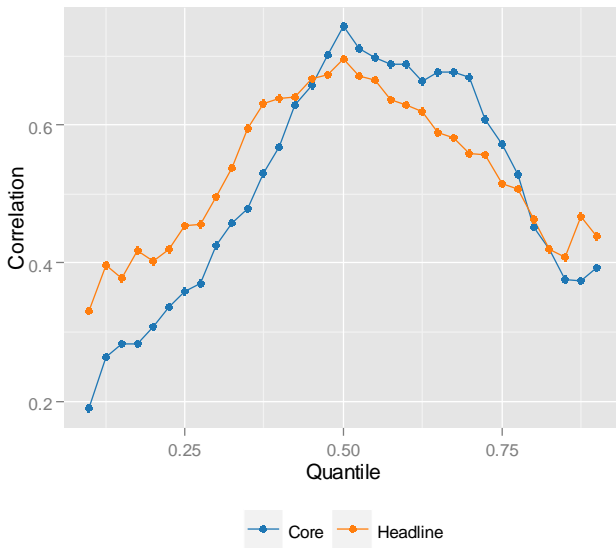
⁶ This is not universally the case. Charlie Bean has pointed out that, in the early 2000s, China's economic development was pushing up on prices of food and energy but also providing cheap imports of manufactures which pushed down on this measure and made it misleading as a guide to underlying inflationary forces.

Chart 5: Core inflation in the UK and rest of the OECD



Source: OECD and Bank calculations.

Chart 6: Quantile dependence for OECD and UK inflation rates



Source: OECD and Bank calculations.

inflation in the rest of the OECD as a whole inevitably feeds through into deflation in the UK.

One might wonder how far inflation in the UK is indeed determined independently of what happens in other advanced economies. In the short run, at least, external shocks to core inflation may have a large influence on the UK. Chart 5 compares core inflation in the UK with the rest of the OECD.⁷ Both seem to have been moving steadily lower over time, and there are periods where a common factor seems to have the same effect on the two.

A simple measure of how the inflation rates have moved together is the linear correlation between the two. Over the past thirty-five years, the correlation is just over ½. So, squaring this, in a purely statistical sense, inflation in the rest of the OECD has “explained” about a quarter of the variation in UK inflation.⁸

But is the correlation the same when inflation rates are generally low, as it is when rates are high? This question matters because some of the drivers of high inflation in the past – such as the oil price shocks of the 1970s – were global phenomena which had a broad international reach. Is the same also true of factors that lead to low inflation?

In effect, I want to see whether there is so-called ‘tail dependence’ – in which, say, an extreme outcome such as very low

⁷ For this chart, and the empirical work that follows, I have excluded the UK from any measure of OECD inflation.

⁸ The square of the correlation coefficient is a measure of the proportion of total variation in one variable explained by another.

The simple correlation between different inflation rates does not provide a complete answer, as it is assumed to be the same for high and low inflation episodes.⁹ So in Chart 6, I illustrate the correlation between different parts of the distribution of inflation rates in both the UK and the rest of the OECD, using monthly data since 1970.¹⁰ I first rank the inflation rates by size, from smallest to largest. Then, for the situations in which UK is very low – for example, in the smallest 10% of observations – I can find out the probability of inflation in the OECD also being similarly low.

What emerges from this picture is that the UK is more strongly correlated with the rest of the OECD when core inflation is higher than average, than when inflation is low. To see why, we can compare the correlation for the smallest 25% of inflation episodes, which is about 0.35, to the largest 25% of episodes, which is about 0.55. Though this may seem like a small difference, in practice it can have a big effect. In the first case, core inflation variability in the OECD excluding the UK explains about an eighth of the variability of UK inflation while in the second case it explains about a third.

Whilst these correlations shed some light on how far UK inflation is related to global inflation, they still don't allow me to explore how unusual it is for both to be as low as they have been recently. To do this, I need to return to my earlier theme about modelling the distributions of inflation rates, in the same way that I was modelling the oil price.

As it happens, the argument that I used there – that it is misleading to assume a bell curve distribution – also applies here. But instead of thinking about 'tail' or extreme events for just UK inflation, now I need to extend my framework to include worldwide inflation too.

This approach echoes that taken by my colleague Andy Haldane, who recently spoke about modelling complexity in nature and economics (Haldane, 2015). It looks at how closely inflation in the UK and the rest of the OECD relate to one another, by estimating a range of alternative models to the simple bell curve or normal case. My aim is to find out how likely it is, under each distribution, to witness the kind of inflation rates that we have seen recently.¹¹

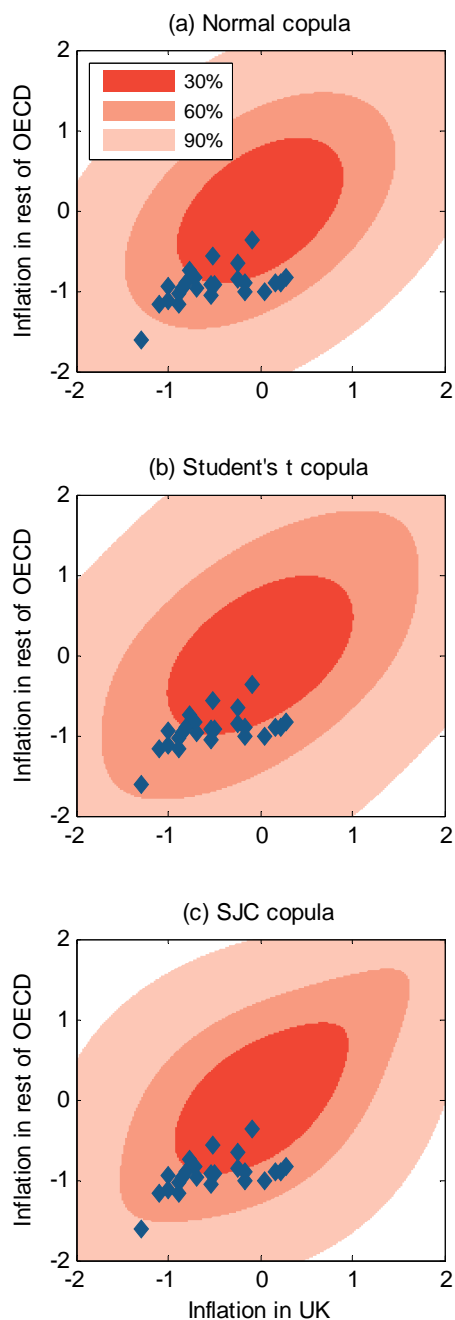
Chart 7 illustrates the results for three different cases. I have shown three different 'probability pools', which indicate how likely it would be to observe a particular combination of inflation rates in the UK and rest of the OECD. Out of 100 pairs drawn at random, around 30 would be in the darkest red area, another 30 in the next darkest, and a further 30 in the lightest shade. The remaining 10 would lie in the white space. Inflation rates for the UK, on the horizontal axis, and the rest of the OECD, on the vertical, are shown in terms of deviations from their average rates, so the point (0,0) denotes inflation at its average for both. The chart shows the relationship for the period of generally low inflation, from January 1980 to March 2015.

⁹ In technical terms, it assumes the relationship between the two inflation rates is linear, whereas in reality it may be non-linear.

¹⁰ I have used seasonally adjusted data from January 1980 to March 2015, to calculate the percentage change in the price index from one month to the next, for the UK, and OECD as a whole. Details of the data sources can be found in the Appendix.

¹¹ Technical details of these models are provided in the Appendix.

Chart 7: Modelling the relationship between inflation in the UK and OECD



Note: Each chart shows the relationship between inflation in the UK and the rest of the OECD. Both axes are measured in terms of deviation from average. The shaded areas show how likely a particular set of inflation rates is; out of 100 pairs drawn at random, around 30 would be in the darkest red area, another 30 in the next darkest, and a further 30 in the lightest shade. The remaining 10 would lie in the white space.

Details of the models estimated are provided in the Appendix.

¹² These are the month-on-month growth rates in the seasonally adjusted price index for the UK and the rest of the OECD.

In the first, (a), I have assumed that the inflation rates have a conventional normal distribution. Since the correlation between UK and OECD inflation rates is positive, the contour lines have an elliptical shape which appears to have been stretched along a South-West to North-East axis. You can see that this model implies that, when inflation is high in the OECD as a whole, it is also likely to be high in the UK. Nevertheless, as I noted above, the correlation of just under 0.5 implies that less than a quarter of the variation in inflation in the UK can be related to movements in inflation in the OECD as a whole.

Against this backdrop, I have plotted the past two years' inflation rates for both UK and OECD, in blue diamonds.¹² Like the probability pools, these are shown in terms of deviations from average, so a negative number means that inflation is below average, and not necessarily that any country is experiencing deflation. What is striking is that most of the recent inflation rates are below average for both the UK and remaining OECD. But on the whole, they all lie relatively close to the top of the 'probability hill'.

Compared to this, panel (b) plots the results from a model with *t* distributed dependence. Although the estimated correlation is similar, at around 0.5, the spread of likely values is much more stretched along the same axis. This, of course, reflects the fact that extreme inflation rates, whether above or below average, are more common under the *t*.

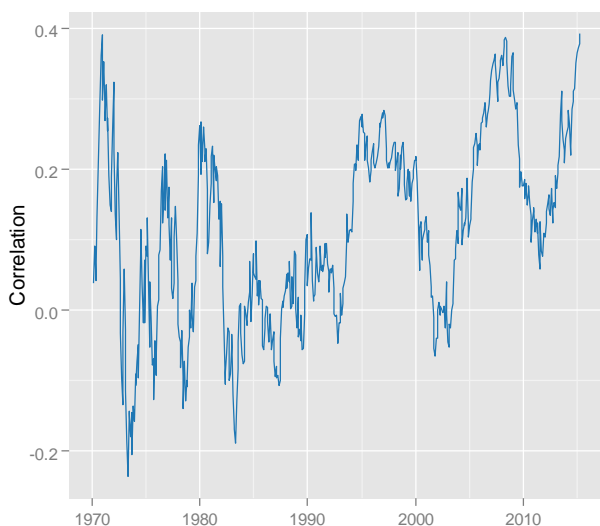
Given that both of these relationships have been estimated using data that stretch back to 1980, one might wonder how far the results are being distorted by the comparatively high inflation of the 1980s and early 90s.

An answer to this comes from panel (c), which plots the results from a model which can allow for the range of

possible inflation rates to be skewed either in the direction of high inflation rates, or low inflation rates. As you can see, in this case there is quite a large skew as the north-eastern tail is stretched much further than the south-western one. This suggests that the past high inflation rates may be affecting the results quite significantly.

There are good economic reasons to believe that countries' inflation rates, and the relationship between them, may have changed over time. Since 1993, the UK has pursued a monetary policy aimed at an inflation target, and since then a number of other OECD members have done the same. So relationships that might have held before this change in regime, may not provide the best guide to the influence of global inflation trends on the UK.

Chart 8: Time-varying correlation between core inflation in the UK and the rest of the OECD



Source: OECD and Bank calculations.

Fortunately, I can use the same modelling techniques I presented earlier in my discussion of oil, to study how the correlation between UK and OECD inflation has varied over time.

Chart 8 traces out the relationship between core inflation in both country blocks since 1970. As you can see, the path has varied considerably over time, and although the correlation has generally been quite low, over the recent past it has risen sharply. That said, a correlation of just under 0.4 does not suggest a strong influence of the rest of the OECD on UK prices.

If these models and methods seem overly technical, let me spell out the most important messages.

First, over the past forty-five years, it has been relatively unusual to see low inflation in both the UK and the rest of the OECD. In contrast, high inflation in one is more likely to go hand-in-hand with high inflation in the other.

Secondly, very high or very low inflation rates are more common than conventional models allow. This is true for both the UK and the rest of the OECD.

Thirdly, since the UK started to target a particular rate of inflation, in 1993, the correlation between UK and remaining OECD inflation rates has been higher in general than it was previously.

Finally, the correlation since the 2008-09 financial crisis has been particularly high, and this is especially true of the past eighteen months. Nevertheless, the tie is not tight. Chart 8, suggests that the variability of core inflation in the rest of the OECD can account for only about a seventh of the variability of UK core inflation.

International Prices as a Long Run Anchor?


My analysis so far has focused on short-term movements in prices, and the relationship between those in the UK and those in the rest of the OECD. It is, however, possible that even if there is a substantial amount of independent movement in prices in the UK in the short term, prices in the rest of the world exert a long-term drag on those in the UK. People who are concerned about low inflation in the UK being imported from the other advanced economies would be right to think about this as well as about the nature of short-term movements.

Chart 9 shows the way in which core UK prices have moved relative to core prices in the rest of the OECD since 2005.¹³ I have focused on this period because it makes the issue clear; if I go back to 1993 when inflation targeting started in the UK, the overall conclusions are not greatly affected. After a period in which prices rose less rapidly than in the rest of the OECD, we experienced several years where they rose faster. From late 2012 onwards, the curve levelled out; our inflation rate moved back in line with that in the OECD. Most recently core inflation in the UK has been slightly slower than in the rest of OECD.

If prices in the OECD exerted a drag on prices in the UK we would expect to see clear upward movements when the line is low and clear downward movements after high values. Someone who wanted to believe the hypothesis that UK price inflation was set by the other advanced countries might well argue that they can see that. After all, the differential was negative, became positive and now is falling again. But this in fact the classic pattern of a random walk: there is no clear equilibrium towards which the relative price is moving, so the path it has taken is largely independent of the initial point.

¹³ The chart shows the difference in the logarithm of the seasonally adjusted price indices, with the gap set to zero in January 2005. The value of -0.02 in January 2009 indicates that prices in the UK had risen approximately two per cent less than in the rest of the OECD between January 2005 and January 2009

Chart 9: Consumer Prices in the UK relative to the rest of the OECD



Source: OECD and Bank calculations.

establish the impact of the price differential on the median value for UK inflation – that is the estimate of the value such that the outturn is equally likely to be above it or below it. I find only a very marginally negative coefficient of -0.002. But, relative to the estimate of uncertainty surrounding it, it is scarcely different from zero. So statistical analysis confirms my earlier observation, that the difference between UK core prices and those in the rest of the OECD is a random walk. UK prices are not, in the long run, drawn towards world prices.

This of course is entirely consistent with the view that the rate of inflation in the UK is set by UK monetary policy, while the rate of inflation in other countries is set by their monetary policies.

Policy Conclusions

To the extent that I can pinpoint the effects of the oil price fall for my policy decision, three things spring to mind from this analysis.

The simulation suggests that oil will push down on inflation by more than the MPC estimated in February, and that the effect will persist well into next year. It suggests a downside risk to our most recent forecast of inflation for next year. By early 2017 I believe the effect will have faded. So the best response is not to worry about this risk should it materialise; after all, policy is set now in order to deliver inflation at target in two years' time, by which time the effect is likely to have gone.

To confirm this I can estimate a regression equation in which the monthly rate of inflation in the UK is explained by the monthly rate of inflation in the rest of the world and also by the differential as shown in Chart 9. If there is a tendency for OECD inflation to drag on UK prices, I would expect a negative coefficient on the differential plotted in Chart 9. When our prices are high relative to those in the rest of the UK, there should be downward pressure.

I drew attention, earlier to the risks of assuming that variables are normally distributed when that is not, in fact, the case. A robust alternative to traditional regression is provided by median regression. This allows me to

Against the prediction of a boost to GDP from the oil price fall, growth in the first quarter was disappointingly weak, at 0.3 per cent. On its own, this does not make me question whether the boost will come at all; both the MPC's analysis and my own profile suggest that nearly all of the impact is yet to come. Over the course of the year, evidence of this may become clearer.

A particular point about the impact of oil on growth is that the effect may be asymmetric. Sharp oil price hikes can, according to some, put a brake on growth. But the boost may be smaller in reverse. Further work is needed to explore this possibility in the modelling framework I described. Asymmetry would suggest a GDP effect weaker than I have shown.

Looking beyond oil, core inflation in the UK is currently quite strongly correlated with inflation elsewhere in the world. We should not be surprised by this – our experience since 2008 should serve as a reminder, if one were necessary, that the UK economy is exposed to developments in the rest of the world.

But the MPC must weigh the need to respond to these international factors, against the desire to provide some stability in the level of interest rates and output. I think the Committee is quite right to let the short-term effects of external shocks feed into inflation, even if this pushes it far from target, whether on the downside as now, or on the upside as four years ago. To do otherwise, and tighten or loosen aggressively, would do little to help inflation in the short term, but would risk a lot with unwanted gyrations in output. This would – with good reason – be highly unpopular with the public.

None of this means, however, that the MPC is incapable of delivering inflation at its target of two per cent. The UK, with its own currency, is free to determine its own monetary policy; and it is this, not short-term external factors, that drives inflation in the medium term. I am confident that this is also true of other countries and currency areas. Inflation is a monetary phenomenon and can be determined by monetary authorities.

References

- Barrell, R. and O. Pomerantz (2004).** "Oil prices and the world economy", NIESR Discussion Paper, no. 242.
- Blanchard, O.J. and J. Gali. (2007).** "The Macroeconomic Effects of Oil Shocks: Why are the 2000s so Different from the 1970s". NBER Working paper No 13368. <http://www.nber.org/papers/w13368>.
- Carabenciov, I., I. Ermolaev, C. Freedman, M. Juillard, O. Kamenik, D. Korshunov, D. Laxton and J. Laxton (2008).** "A Small Quarterly Multi-country projection Model with Financial-real Linkage and Oil Prices", IMF Working Papers 08/280.
- Cashin, P., K. Mohaddes, M. Raissi and M. Raissi. (2014).** "The differential effects of oil demand and supply shocks on the global economy", *Energy Economics*, Vol. 44, No. 1, pp. 113–34.
- Creal, D., Koopman, S. J. and Lucas, A. (2012),** 'Generalized autoregressive score models with applications', *Journal of Applied Econometrics*.
- European Central Bank (2015).** "Effects of the Falling Oil Price on the Global Economy". *Monthly Policy Report*. pp 45-50.
- European Commission (2008).** "Recent economic developments and short-term prospects", Quarterly Report on the Euro Area, No. 7.
- Fan, Y. and Patton, A. (2014).** 'Copulas in Econometrics', *Annual Review of Economics*, 6, pp.179-200.
- Gauss, C.F. (1809).** *Theoria motus corporum coelestium in sectionibus conicis solem ambientium*.
- Haldane, A. (2012),** 'Tails of the unexpected', speech given at the University of Edinburgh Business School
- Haldane, A. (2015),** 'Microscopes and telescopes', speech given to the Lorentz centre workshop on socio-economic complexity, Leiden
- Hamilton, J. (2003),** 'What is an oil shock', *Journal of Econometrics*, 56 (6), 766-779
- Hamilton, J. (2008),** 'Oil and the macroeconomy', in *The New Palgrave Dictionary of Economics*
- Harvey, A. (2013),** *Dynamic Models for Volatility and Heavy Tails*, Cambridge University Press
- Hervé, K., N. Pain, P. Richardson, F. Sédillot, and P. Beffy (2010).** "The OECD's new global model", *Economic Modelling*, vol. 28, pp. 589–601.
- International Monetary Fund (2015).** *World Economic Outlook*. April. Chapter 1. Scenario Box 1. Page 7. <http://www.imf.org/external/pubs/ft/weo/2015/01/pdf/c1.pdf>
- Jiménez-Rodríguez, R. and M. Sánchez (2004).** "Oil price shocks and real GDP growth: Empirical evidence for some OECD countries", European Central Bank Working Papers, No. 362.
- Kilian, L. and R.J. Vigfusson (2011).** "Are the responses of the U.S. economy asymmetric in energy price increases and decreases?", *Quantitative Economics*, Vol. 2, No. 3, pp. 419–53.
- McCafferty, I. (2015)** 'Oil price falls – what consequences for monetary policy?', speech given to Durham University Business School.
- Monetary Policy Committee (2015).** *Inflation Report*. February. pp 32-33.
- Patton, A. (2013).** 'Copula methods for forecasting multivariate time series', in *Handbook of Economic Forecasting vol. 2B*, Elsevier.

Peersman, G. and I. van Robays. (2012). "Cross-country Differences in the Effects of Oil Shocks". *Energy Economics*. Vol 34. pp. p132-1547.

Riksbank (2015). "Effects of the Falling Oil Price on the Global Economy". *Monetary Policy Report*. February pp 45-50.
http://www.riksbank.se/Documents/Rapporter/PPR/2015/150212/rap_ppr_ruta2_150212_eng.pdf

Student [Gosset, W. A.] (1908) 'The probable error of a mean'. *Biometrika*, 6 (1): 1–25.