ESSENTIAL READING FOR BENCHMARK

How should the fall in the market curve since August affect the forecast?

(CAPD and MFAD)

The most important piece of news to the forecast since last round has been the fall in the market curve, which we take as the conditioning path for Bank Rate. In the Benchmark forecast the stimulus from lower interest rates is only partly offset by a weaker world forecast, so that inflation ends the forecast at 2.4%. We will use the second part of the Benchmark meeting to ask whether you are content with the impact of the fall in the market curve on the forecast or whether you would like to make a judgement to change it.

- The fall in the market curve means that monetary conditions in the Benchmark forecast are substantially looser than in August. All else equal the fall in the curve boosts GDP by 0.6% by the third year of the forecast and raises inflation by 0.3pp; in addition there are further small positive contributions from the falls in the exchange rate and the oil price. Partly offsetting this is a drag from a weaker world outlook: but the total drag from these channels is not nearly enough to offset the boost from the yield curve itself. So inflation reaches 2.2% at Year 2 and 2.4% at Year 3.
- The judgement you will need to make in this forecast round is whether the low current path of Bank Rate – which ends the forecast at just 1.3% – provides the appropriate amount of monetary stimulus; or whether, as in the Benchmark Forecast, it would be likely to lead inflation to overshoot the target, and, if so, by how much.
- At a high level we think there are three broad options: stick with the interpretation of the news and the effect of the fall in the market curve that is currently in the Benchmark; take a different view on the interpretation of the data news, or on the outlook for the output gap or inflation as a whole; or take a different view on the amount of stimulus provided by the fall in the market curve. This note focuses on the third of these options.
- We are not proposing that you make a judgement to take a different view on the amount of stimulus provided by the fall in the market curve to what we have assumed in the Benchmark. But there are some arguments that could be made in favour of doing so.
- One argument for this would be if you thought some feature of the news in the market curve meant it should have a different effect this time to what we would typically assume. The evidence on what explains the fall in the curve suggests it reflects a combination of lower central expectations, less weight on upside risks and lower term premia. In principle if the news was disproportionately concentrated in term premia then that might lead you to put in less of an effect through current asset prices, since what evidence there is suggests that changes in term premia have less powerful effects than changes in rate expectations. It isn't clear, though, that the contribution of term premia is particularly large relative to a normal fall in the curve.
- Another argument would be if you thought we should be using different assumptions about the effect of changes in the market curve in general. It is very hard to find meaningful or clear 'apples to apples' comparisons of our own standard multipliers against estimates in the literature. To the extent that we can find such comparisons, our multipliers do not look out of line with other estimates.

Introduction

The most important piece of news to the forecast since last round has been the fall in the market curve, which we take as the conditioning path for Bank Rate. In the Benchmark forecast the stimulus from lower interest rates is only partly offset by a weaker world forecast, so that inflation ends the forecast at 2.4%. Publishing such a forecast would send a signal that the market curve was significantly lower than would be appropriate given the outlook for demand and inflation.

Given that, it is important that you are happy with how we have interpreted the fall in the market curve and where the forecast has ended up as a result. At a high level we think there are three broad options: stick with the interpretation of the news and the effect of the fall in the market curve that is currently in the Benchmark; take a different view on the interpretation of the data news, or on the outlook for the output gap or inflation as a whole; or take a different view on the amount of stimulus provided by the fall in the market curve. We will use the second part of the Benchmark meeting to focus on the third of these, and ask whether you are content with the impact of the news in the market curve on the forecast or whether you would like to make a judgement to change it.

The rest of this note therefore looks at the reasons why you might or might not want to aim off our assumptions about the effect of the fall in the market curve in principle, and considers how much evidence there is to support such a judgement. Section 1 of the note summarises what has happened to the market curve this round and what effect that has had on the Benchmark forecast; Section 2 considers the message this forecast would send, and sets out three options you could take. Section 3a looks at whether we should make a one-off judgement to aim off the effect of the fall in the curve this round, while Section 3b looks at whether our standard assumptions about the effects of changes in the curve should themselves be different.

Section 1: What has happened to the market curve?

Since the August *IR* the market curve has fallen substantially, with three -year forward OIS rates around 50bps lower (**Chart 1**): they now reach only 1.3% at the end of the forecast. The curve now crosses 75bps in 2017Q2, compared 2016Q2 in the August *IR*.

The main falls in UK OIS rates occurred in the period following the devaluation of the renminbi and (after recovering somewhat on the UK wages data) again immediately following and in the days after the September FOMC meeting (**Chart 2**). These international factors appear to have been generally more important than domestic influences over the period.

Chart 1 The fall in the market curve

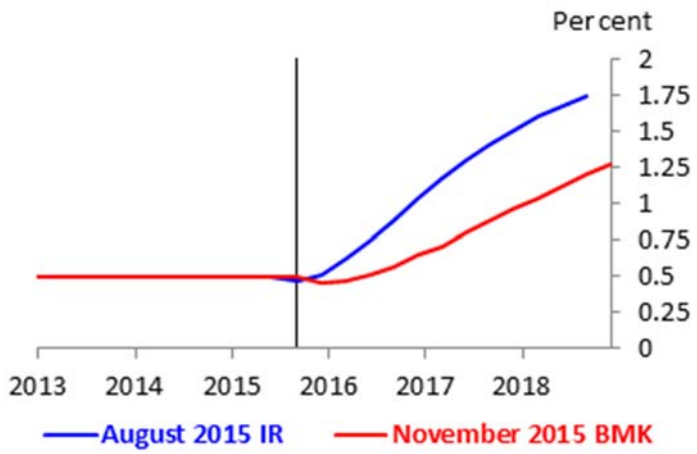

Over the same period, "risky" assets have sold off and "safe" assets have rallied: equity and corporate bond prices have fallen around the world, the VIX has risen and commodity prices are lower. Market contacts attribute these changes to increasing worries about global growth and spillovers from emerging markets to advanced economies. [1]

On the face of it, therefore, the curve appears to have fallen primarily in response to worries about a weaker world outlook. The question for the Benchmark forecast is whether that fall delivers the appropriate amount of stimulus to offset any additional weakness in the world.

Chart 2: Cumulative change in UK 3 year spot OIS rate since the August *IR*

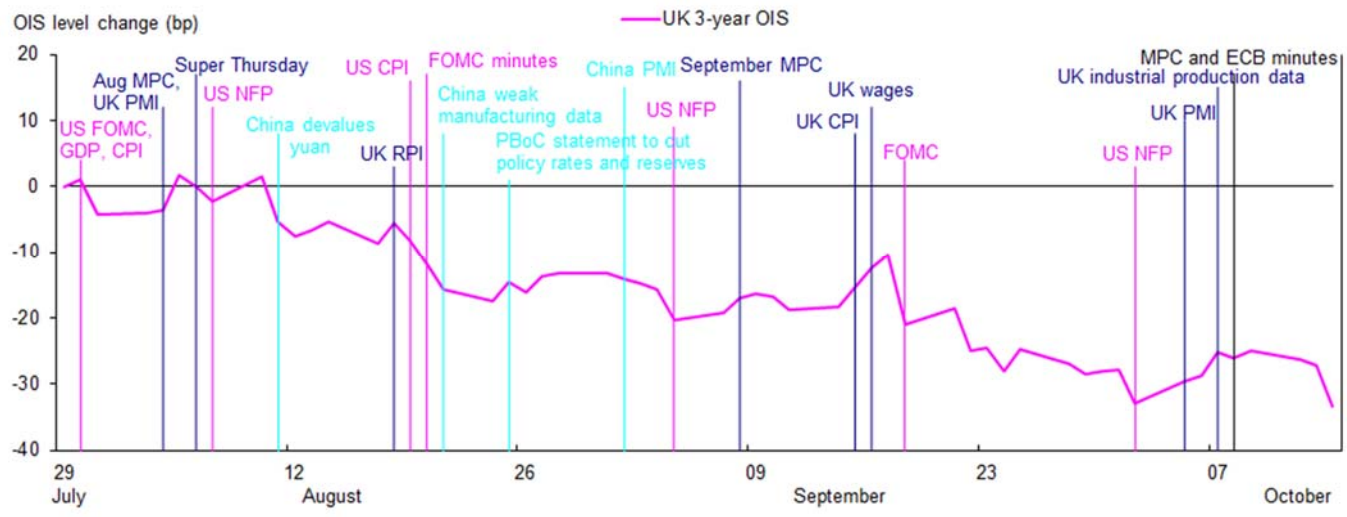

How has the fall in the curve affected the Benchmark forecast?

The fall in the market curve means that monetary conditions in the Benchmark forecast are substantially looser than in August. All else equal the fall in the curve boosts GDP by 0.6% by the third year of the forecast and raises inflation by 0.3pp (**Chart 3**); in addition there are further small positive contributions from the falls in the exchange rate and the oil price.

Partly offsetting this is a drag from a weaker world outlook. The downgrades we have made to our world forecast take 0.3% off GDP and 0.1pp off inflation via direct trade channels; on top of that we have built in a near term drag from the rise in uncertainty over the quarter. There is also some extra drag from higher bank funding costs; but this is broadly offset by a judgement on banking sector competition, leaving the net impact from credit conditions close to zero.

Chart 3: News to GDP and inflation since the August *IR*

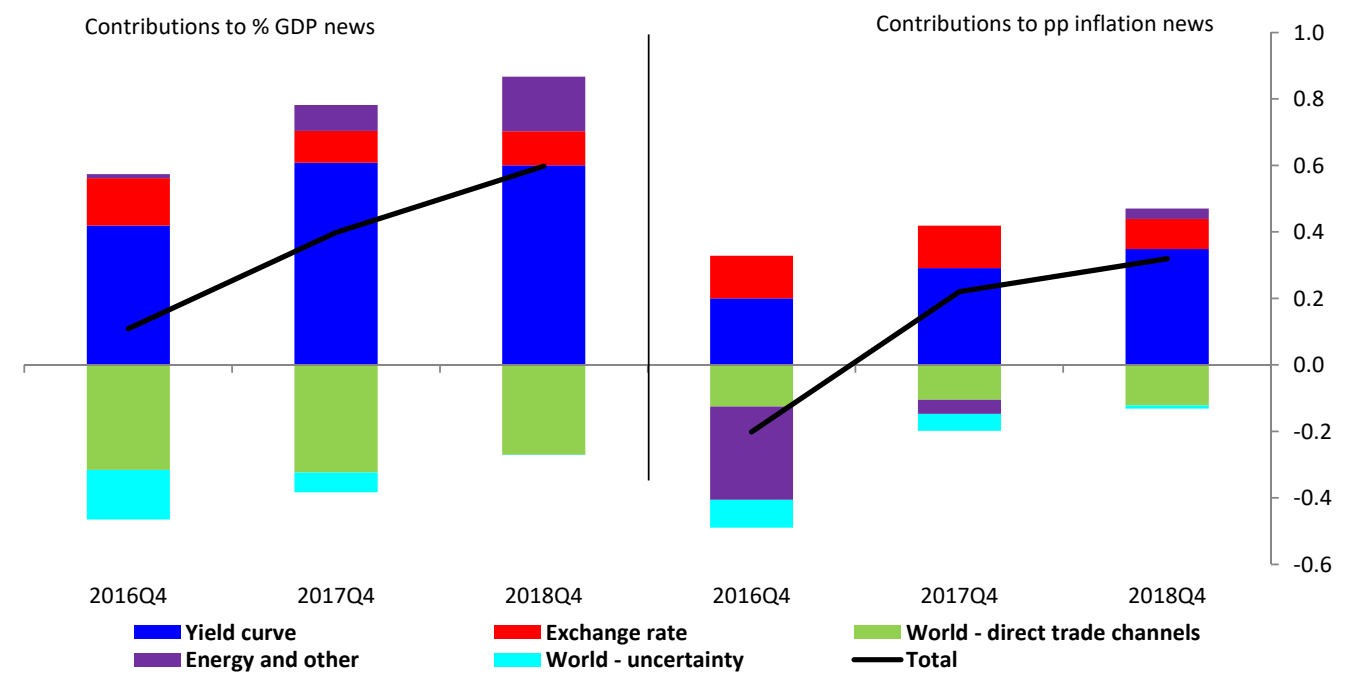

The net effect of these two things – the boost from looser monetary conditions and the increased drag from the world – is GDP ends the forecast 0.6% higher and inflation 0.2 and 0.3pp higher at Years 2 and 3 respectively than in August. So conditional on the current market curve inflation reaches 2.2% at Year 2 and 2.4% at Year 3. Unemployment, meanwhile, falls to 4.4% and the output gap finishes at +0.5%. In other words, the Benchmark forecast implies that the market curve has fallen by much more than would be needed to offset the increased drag from the world, and that following the market curve would lead the economy to begin to overheat.[2]

Section 2: Are you happy with our interpretation of the news in the Benchmark forecast?

Publishing a forecast in which inflation ended at 2.4% would be a departure both from recent communications – we have had Year 3 inflation at 2 or 2.1% in every forecast since February 2014, indicating each time that the MPC viewed the prevailing market curve as broadly appropriate given the outlook. And it would be unprecedented – we have never previously published a market rate forecast in which Year 3 inflation was above 2.3%. So publishing this forecast as it stands would be a significant step.

Of course none of these are reasons not to publish such a forecast. If you were content with our interpretation of the news since last round and agreed with the proposition that following the current market curve would probably leave inflation well above target, then publishing this forecast and sending a signal that the curve was inappropriately low could well be the right thing to do.

But as always other interpretations of the news since August are possible, and you may disagree with how we have treated the data news or how we have interpreted the fall in the market curve; if so, that might lead you to want to publish a different forecast and send a different policy message. The judgement you will need to make in this forecast round is whether the low current path of Bank Rate – which ends the forecast at just 1.3% – provides the appropriate amount of monetary stimulus; or whether as in the Benchmark Forecast it would be likely to lead inflation to overshoot the target, and, if so, by how much.

At a high level we think there are three broad options:

1. Stick with the interpretation of the news and the effect of the fall in the market curve that is currently in the Benchmark. This would be the option to choose if you agreed with the interpretation of the data news in the Benchmark forecast, and wanted to send a signal that the market curve was significantly lower than would be appropriate given the outlook for demand and inflation.
2. Take a different view on the interpretation of the data news, or on the outlook for the output gap or inflation as a whole. This would be the option to choose if, for instance, you thought we should have taken more of a downside signal for world demand, or if you thought our forecast for domestic demand was too strong even given the fall in the curve.
3. Take a different view on the amount of stimulus provided by the fall in the market curve. We think there are two distinct potential reasons why you might want to do this:
    a. If you thought some feature of the news in the market curve meant it should have a different effect this time to what we would typically assume. For instance if you thought that term premia were playing a bigger role than usual, and that changes in term premia had smaller wealth effects than changes in expectations, that might lead you to want to put in a smaller boost to demand from higher asset prices.
    b. If you thought we should be using different assumptions about the effect of changes in the market curve as standard. For instance you might think our current multipliers are too large at the policy-relevant horizon.

[2] For more details see the November 2015 Benchmark Forecast Note

As usual this forecast round you will have the opportunity at the Benchmark meeting to discuss how we have interpreted the news in the data. And we intend to cover the outlook for the world and for domestic demand in detail in the two Key Issues meetings, before coming back at Draft to discuss whether you would like to change the forecast in view of any of these. So the purpose of this note is to allow you to consider the third option: we will use the second part of the Benchmark meeting to ask whether you are content with the impact of the fall in the market curve on the forecast or whether you would like to make a judgement to change it.

To be clear, we are not proposing that you make such a judgement. But there are some arguments that could be made in favour of making such a judgement. The rest of this note therefore looks at the reasons why you might or might not want to aim off our assumed impact of the fall in the market curve in principle, and considers what evidence there is in practice to support such a judgement. The next section looks at whether we should make a one-off judgement to aim off the impact of the fall in the market curve this round (option 3a); the final section looks at whether our standard assumptions about the impact of changes in the curve should themselves be different (option 3b).

Section 3a: Should we assume a different effect from the fall in the market curve this time?

One potential reason to aim off the effect of the fall in the market curve in the Benchmark forecast would be if you thought something about this particular fall meant its effects would be different to those we assume for a "typical" fall in the curve.

We think there are two possible reasons why you might think that was the case. This section sets out those possible reasons, and then considers what evidence we have that either of them is relevant to this particular fall in the curve.

Reason 1: If you thought the change in the curve didn't reflect true changes in expectations

If you thought the fall in the curve didn't reflect a true fall in market expectations, that might lead you to want to aim off the curve itself. In general, true market expectations of Bank Rate are not necessarily the same as the market curve. The market curve is the mean, not the mode, of the priced distribution of interest rates; so *if* you wanted to condition on modal expectations then that would imply using a different path. In theory, the priced mean might also not be equal to actual mean expectations – in other words the observed price might include non-zero term premia. So again you might want to correct for this by adjusting the market curve.

In practice, the fall in interest rates appears to reflect a combination of falls in all three – mean expectations, the balance of risks, and term premia:

- Our model decomposition attributes around half of the decline in 3-year government bond yields to term premia and half to mean expectations (**Chart 4**). That said, we typically assume that term premia are likely to be small at short horizons, and tend to have less faith in the performance of these models for shorter term yields, particularly so close to the lower bound. Furthermore, 3-year government bond yields have only fallen by around 20bp, so the results may not necessarily apply to OIS rates (for which we do not have a similar decomposition).
- The fall in the market path reflects declines in both real rates and inflation swap rates, suggesting there has been a decline in market-implied inflation expectations (**Chart 5**). That appears consistent with the view that market participants believe there has been a weakening in the outlook for inflation which led to a change in mean policy expectations. But surveys of market participants have moved by much less than short-term interest rates since August (**Chart 6**). Together these facts again imply that changes in mean expectations and term premia both account for at least some of the move.
- The distribution implied by options linked to three-month LIBOR rates in 6 months' time has shifted to the left, and become more strongly centred on the mode (**Chart 7**), suggesting that modal expectations have at least to some extent moved in line with mean expectations. But there is also less weight in the upper tail, suggesting that part of the fall in the curve reflects a downward shift in perceptions of the balance of risks. There is a limit to the

Chart 4: Term structure model decomposition of changes in 3-year spot government bond yield

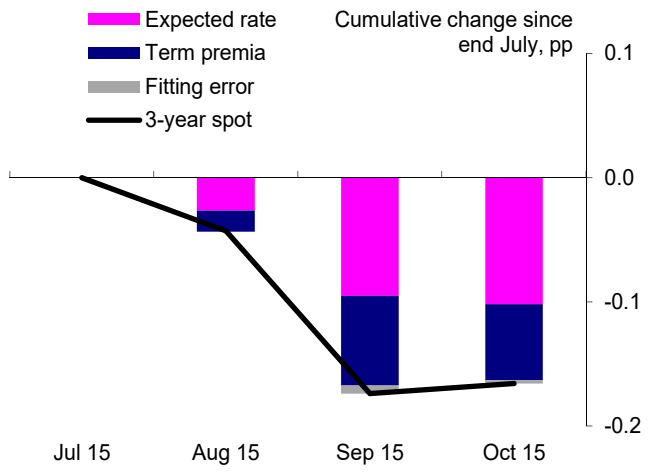

Source: Bloomberg, Reuters

Chart 5 Real / inflation decomposition of changes in short term forward rates since August *IR*

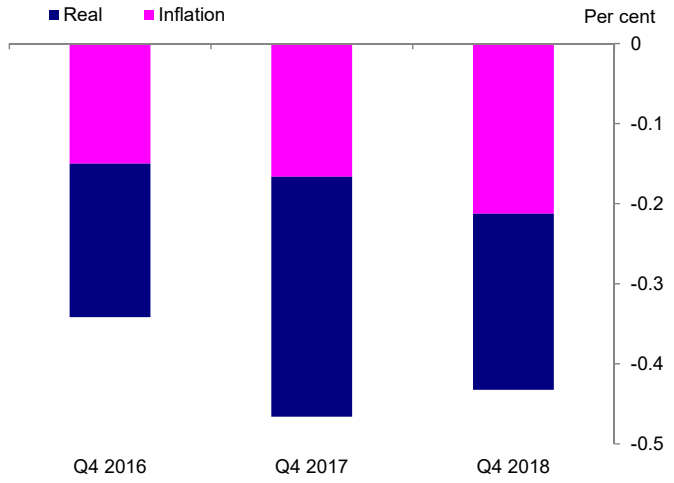

Source: Bloomberg, MFAD

Chart 6: Date at which OIS swap rates reach 0.75%, and expected time of the first rate rise from Reuters survey

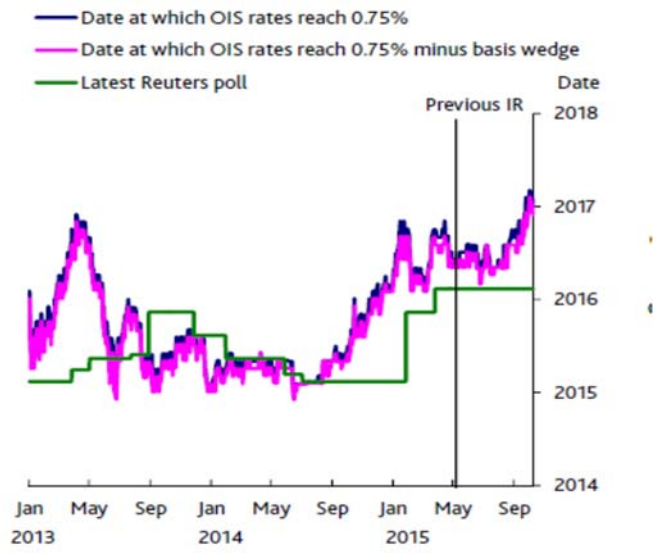

Source: Bloomberg, Reuters

Chart 7: Option-implied weight of 3m LIBOR sterling rates at 6m horizon

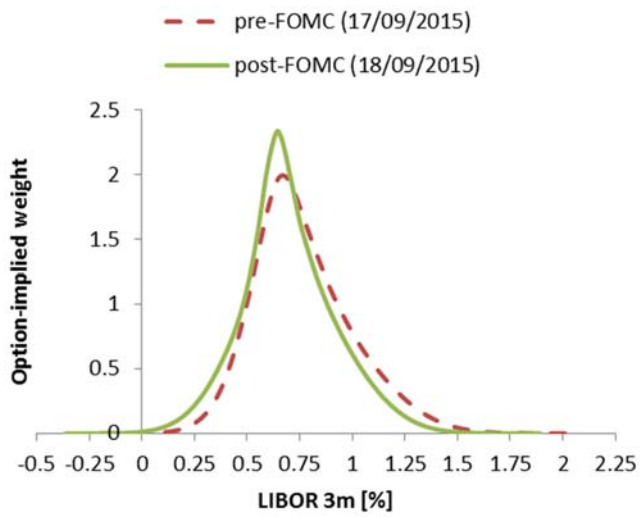

Source: Bloomberg, MFAD

extent to which the mean versus mode question can be viewed separately from the expectations versus term premia question – these options are based on Libor, and will themselves be distorted by risk premia. Indeed in Barro-type rare disaster asset pricing frameworks it is changes in the probability of disaster that drive changes in term premia.

As a practical matter, however, we think the case against assuming a different path for Bank Rate from the market curve is quite strong. The convention of assuming Bank Rate follows the market curve is very well established, so that departing from it might well be seen as a significant change in communications. Moreover, while the arguments for assuming a different path are based on trying to find a better measure of expected Bank Rate, the *Inflation Report* doesn't claim to be conditioned on expectations: it simply says that Bank Rate is assumed to follow the market curve.[3] So aiming off the fall in the curve because we thought it didn't accurately represent expectations would also

[3] Indeed if term premia were negative, that could in theory explain situations like the recent one in which our inflation forecasts have been persistently above target with little market reaction. By publishing above-target forecasts we would try to signal that markets should expect

be a break with the past, and would perhaps beg the question of why we hadn't made such adjustments before. On top of that, using the market curve has the benefit of transparency; any adjustment would be model- or judgement-dependent and harder to explain to the outside world.

Of course none of these arguments prevent you from departing from convention – you may well feel that a change in convention is warranted, and what goes in the *Inflation Report* is, of course, completely up to you. We think, however, that the cost of communicating such a change would be quite high, and are doubtful in this case that the benefits would justify the cost.

Reason 2: If you thought the effects of the fall in the curve should be different because it reflected different shocks

Even if you did not want to aim off the curve itself, you might think that the fact that the fall in the curve appears to reflect changes term premia and the balance of risks as well as market mean expectations might mean it would provide more or less stimulus than usual. To believe that you would need to think both that term premia or the balance of risks were contributing an unusual amount to the news this round, and that this implied a different transmission to the economy.

Chart 8: 3 year gilt rates decomposed into expectations and term premia

There is some evidence to support the first of these, though it is not strong. While the correlation between changes in our estimates of term premia at these horizons and changes in expectations is negative, it is quite weak, and it is not that unusual for both to move in the same direction (**Chart 8**). Changes in the skew of the Libor distribution, meanwhile, are not significantly correlated with the mean.

Why might this matter for the effect of changes in the market curve? As with any piece of data news, the effect on the forecast will depend on the source of the shock. A change in market rates that reflects changes in risk aversion might transmit differently to the economy to a change in market rates that reflects expectations: for instance if people perceive changes in term premia as transient then they may look through the change in rates. Kiley (2015) finds for the US that a persistent move in interest rates caused by term premia has about half the effect of an equivalent-sized move caused by expectations (although his model strictly applies by construction only to term premia on long rates, and so may not carry over to the short rate case we're interested in here).[4]

Making a judgement like this would again be a departure from previous forecast rounds: we have not made such a judgement before, even for somewhat larger changes in term premia (**Chart 8**). So again communicating why we had made such a judgement this time might prove challenging. To the extent that term premia are below average now, this line of logic might also lead you to question whether we have assumed too much stimulus from the low absolute level of the market curve by not making similar judgements in the past.

To give some idea how important this might be for the forecast in practice, **Charts 10 and 11** show a scenario in which we assume that the fall in the market curve since August has only half the effect via current asset prices. In that scenario GDP growth peaks at 2.7%% rather than 2.9% and inflation reaches 2.1% and 2.2% in Years 2 and 3, respectively, rather than 2.2 and 2.4%.

Bank Rate to be above the market curve. But, knowing that term premia were negative, they would already have expected that; so our signalling would tell them nothing new, and the market curve would remain in the same place.

for a discussion of this paper

Chart 9: GDP growth with a smaller effect via current asset prices

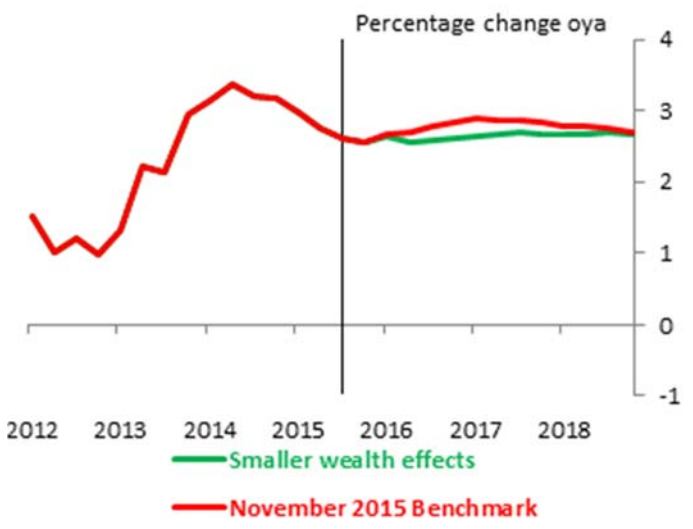

Chart 10: Inflation with a smaller effect via current asset prices

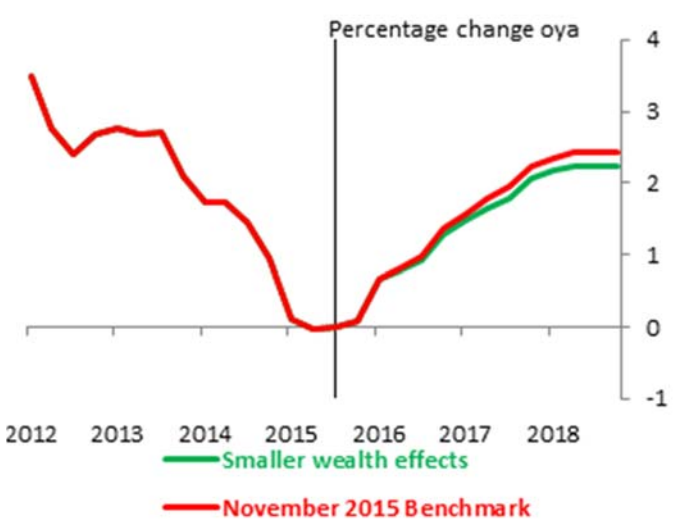

Section 3b: How reasonable are our standard assumptions?

The previous section considered whether we should aim off our usual assumptions about the effects of the fall in the market curve as a one-off this this round given what's driven the fall. This section goes a step further and considers whether our standard multipliers themselves seem reasonable.

There are two reasons why you might think our assumptions should be different. One would be if you thought something about the current conjuncture meant the transmission of a given change in Bank Rate would have a larger or smaller effect. The other would be if you simply thought we had not got the magnitude or timing of some of their effects quite right. As mentioned in the previous section, you discussed the first option extensively last year, following the Special Project on policy normalisation; at the time you concluded that while there were risks in both directions, you did not feel they warranted aiming off our usual assumptions.[5] This section therefore focuses on the second possible reason, by comparing our assumptions with the other available evidence on the effects of Bank Rate changes to see to what extent it supports out current assumptions.

What are our standard multipliers?

When producing the forecast, we process news in the market curve using a suite model rather than the mechanisms built into our main forecast model, COMPASS. We do this because we struggle to obtain realistic responses using COMPASS to the effects of changes in interest rates, and particularly in expected future interest rates (**Chart 11**). Since monetary policy works through expectations, changes in future rates should affect behaviour now; and in principle COMPASS should be able to capture this effect using "anticipated shocks". But, in practice, it tends to produce implausibly large and front-loaded responses in activity and, especially, pricing – this is a general feature of all DSGE models: the "forward guidance puzzle". Historically we worked round this by using unanticipated shocks. But that became increasingly untenable as the market curve became flat for a lot of the forecast period, with most variation in longer term rates, which was why we introduced this alternative approach.

The suite model we use to process interest rate news is derived from a version of MASD's sectoral model.[6] Specifically**,** the suite model estimates the effect of interest rate news in a bottom-up way as the sum of estimates of the effects of the channels described in the previous section:

[5] For more details

[6] See Staff Working Paper No. 556 by Cloyne et al.

- Wealth and cost of capital effects. These, together with an exchange rate channel, correspond to the asset price effects channel from the previous section. We estimate the effects of the change in the yield curve on other asset prices, then use the sectoral model to obtain the effect of those on activity and prices.
- An exchange rate channel. The effect on the exchange rate is calculated using a half UIP half random walk relationship. We ignore the exchange rate effect in this note because in practice when we process news from round to round it is subsumed in the news to the exchange rate conditioning path – the only place it appears independently is in the constant rate forecast, where we allow the exchange rate to respond to the change in the interest rate conditioning path.

Chart 11: Impact of the yield curve news since Aug *IR* on level of GDP using anticipated shocks in COMPASS

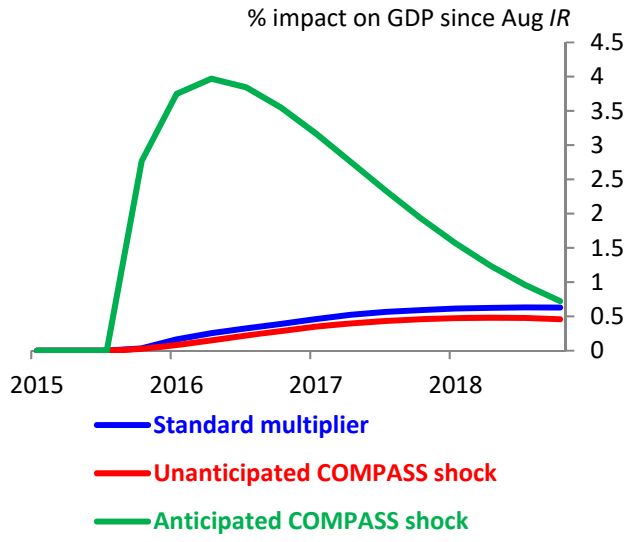

- A cashflow effect channel. As described in the previous section, this captures the effects of changing interest payments on households' and firms' disposable income. The size of this effect is calibrated to match the empirical evidence on how much of a change in Bank Rate actually comes from these effects.

When we talk about our "standard multipliers", what we are really referring to is the output from this suite model.[7] One thing worth bearing in mind is that although this model is designed to mimic some forward-looking behaviour, it does not rely on an unusual degree of forward-lookingness among households and firms. The sectoral model upon which it is built is estimated over 1977-2013, and so will simply reflect the average response in the past of activity and inflation to changes in asset prices.

The remainder of this section considers the standard multipliers in the context of the other evidence we have on the effect of changes in Bank Rate. It first looks at the magnitude of the effects on GDP and inflation; then at the timing of the effects; and finally at the slope of the reduced-form Phillips curve that they imply.

The magnitude of the effect

MASD's recent summary of the empirical literature on the effects of monetary policy shocks provides a useful source of information against which to assess our current multipliers.[8] That said, there are two reasons why our ability to draw strong inferences from such a comparison may be limited. First, it is of course difficult to identify a genuine monetary policy 'shock' and the standard techniques used to do so in the literature (e.g. sign restrictions, zero restrictions or event studies) will all suffer from limitations that might make us doubt that a genuinely 'exogenous' change to policy rates has been detected. Second, beyond the first period impulse, the behaviour of monetary policy – and therefore the multiplier – will depend on the assumed monetary policy rule. And so if multipliers vary across studies, that may reflect a differences in the assumed monetary policy rule rather than a given change in policy rates having a different effect. For both these reasons, it is hard to make a meaningful or clear 'apples to apples' comparison of our own standard multipliers against estimates in the literature.

The closest – albeit highly imperfect – comparison is to compare the impulse response function of our standard multipliers with two other models: 1) an (unanticipated) monetary policy shock in COMPASS, and 2) a monetary

policy shock from a time-varying parameter factor-augmented VAR of the UK (updated from Ellis, Mumtaz and Zabczyk (2011)). This is not a typical piece of news for the forecast – the change in interest rates is concentrated in the near term, whereas in the forecast we typically see the whole curve move – but it is the cleanest comparison that we have been able to make.

The impacts on the level of GDP and annual inflation are shown in **Charts 12 and 13** below, respectively. In all cases Bank Rate is 1pp higher than baseline in the first period and then allowed to move in line with the model dynamics (or Taylor rule) thereafter. The standard multipliers have a peak impact on the level of GDP of around 0.4% after the first year, and around 0.2% after three years. That is a larger and more front-loaded initial impact than the FAVAR, but a little smaller in the medium-term: the impact on GDP in the FAVAR has largely come through within 6 quarters (-0.2%) and builds very gradually thereafter. The magnitude over the first year is around double that of the standard multipliers (-0.7%), but the effect in COMPASS unwinds more quickly with output returning to baseline after three years. (As this is just a one-period shock, an unanticipated and anticipated shock coincide in COMPASS because policy reverts back to the Taylor rule from the second period in both cases.)

The peak impact of the monetary policy shock on inflation is -0.2pp for the standard multipliers. That is less than the alternative estimates: COMPASS would predict a peak impact of -0.4pp, while the FAVAR would anticipate inflation being around 1pp lower. The differences in inflationary impacts will of course be strongly linked to the impact on GDP: the monetary policy shock in COMPASS has a larger impact on output than the standard multipliers and so it is unsurprising that the impact on inflation is also larger. As we discuss later, however, the *relative* impacts on inflation and output also differ between the models.

Looking at a broader set of studies in the literature, there range of estimates for the effect of monetary policy shocks is very wide. Our standard multipliers do not appear to be particularly out of line with other estimates for the effect of monetary policy shocks on output or prices (**Charts 14 and 15**).

The experiment using IRFs effectively meant that the shock to Bank Rate was treated as relatively transient, as policy returned back to baseline. In practice, however, the news in the yield curve from round-to-round that we process tends to be quite persistent, with long rates typically changing by more than rates at shorter maturities. To demonstrate the effects of this kind of persistent news, we compare our multipliers against an unanticipated shock in COMPASS to see what they each imply for the impact on GDP and inflation of the fall in the yield curve since the

Chart 12 Impact of monetary policy shock on level of GDP[(a)]

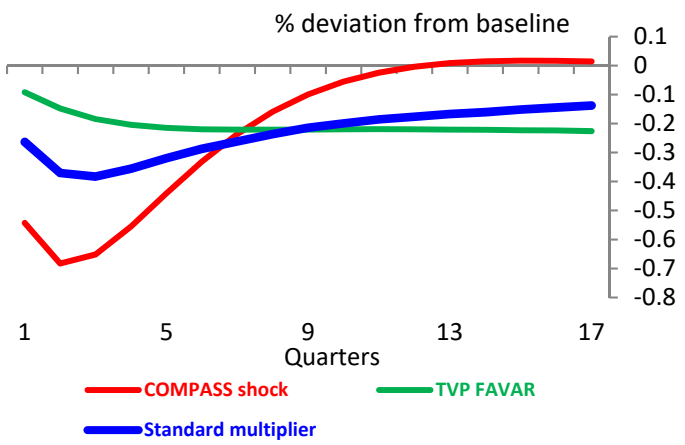

(a) "Standard multiplier" IRFs are based on interest rate path calculated using COMPASS Taylor rule, and include an exchange rate effect for comparability.

Chart 13: Impact of monetary policy shock on annual inflation

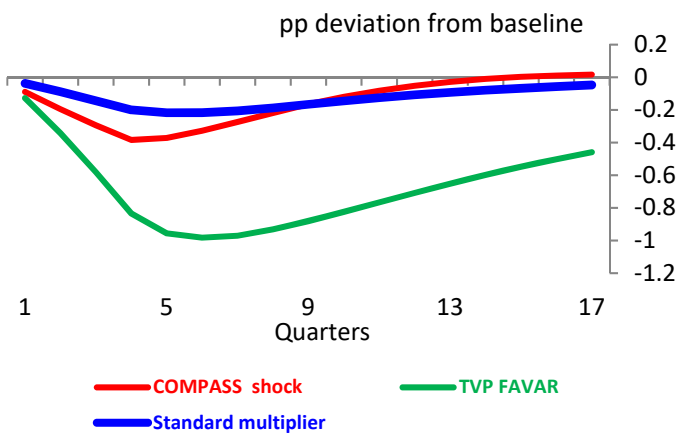

Chart 14: Estimates for impact of monetary policy shock on level of GDP

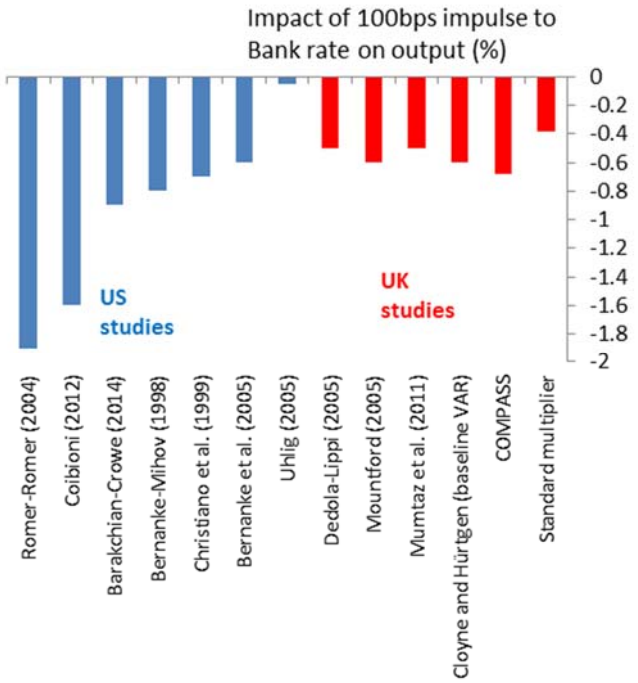

Chart 15: Estimates for impact of monetary policy shock on price level

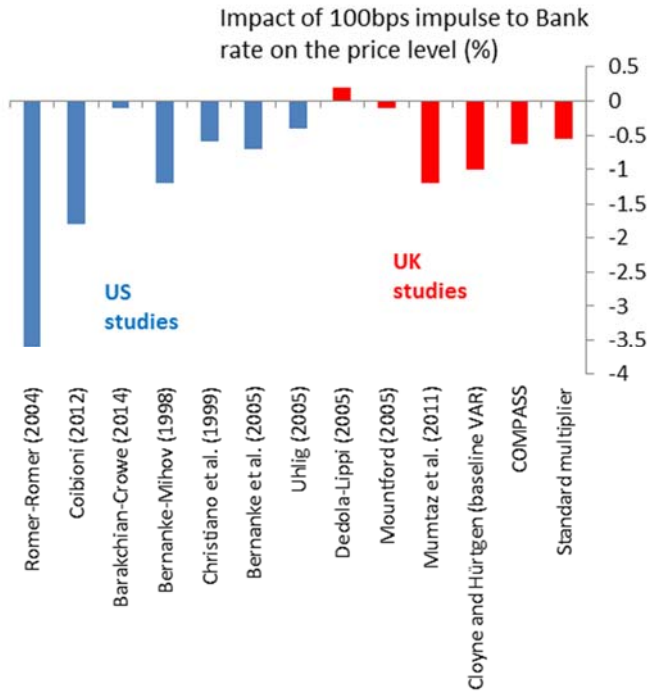

August *IR*. As shown in **Charts 16 and 17**, the standard multipliers embody effects on output and inflation which are a little larger and more frontloaded than an unanticipated shock in COMPASS. The frontloadedness is not a surprise – these multipliers are designed to mimic some forward-looking behaviour, with demand rising in anticipation of future moves.

The timing of the effect

The non-UK literature tends to find a peak effect on GDP and inflation after 4-8 quarters. For US studies, the peak effect tends to occur after 6-7 quarters on GDP and inflation. Despite the peak effect on output typically occurring in the second year, demand does tend to respond relatively quickly and a large proportion of the impact has often come through prior to the peak. For other countries – such as the euro area, Canada and Japan - the peak effect on GDP tends to be a little earlier than in the US and, on average, is early in the second year.

The UK literature on this topic suffers from a relative lack of studies but finds that, on average, the peak effect on output occurs after 5 quarters. The two studies that explicitly assess the impact on UK inflation find a relatively speedy effect, with response peak after either 4 or 6 quarters. The standard multipliers embody a peak effect on output after 3-4 quarters, with a peak impact on inflation after 5-6 quarters. The FAVAR discussed earlier finds that the majority of the effect on output has occurred after 6 quarters, with the peak impact on annual inflation after 4-6 quarters. The unanticipated monetary policy shock in COMPASS has fairly front-loaded effects, with the peak impact on output after just 2-3 quarters and annual inflation after 4-5 quarters. So the timing of peak effects for standard multipliers appears to be broadly in line with the limited UK literature on inflation and in line with or perhaps slightly quick on output.

Chart 16: Impact of yield curve news since Aug *IR* on level of GDP

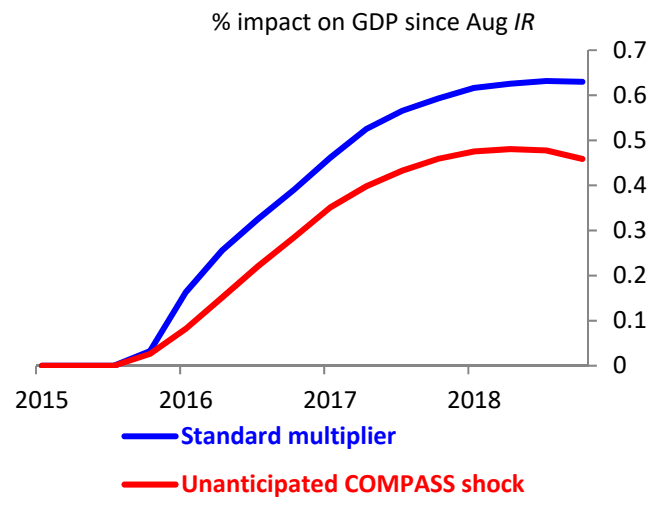

Chart 17: Impact of yield curve news since Aug *IR* on annual inflation

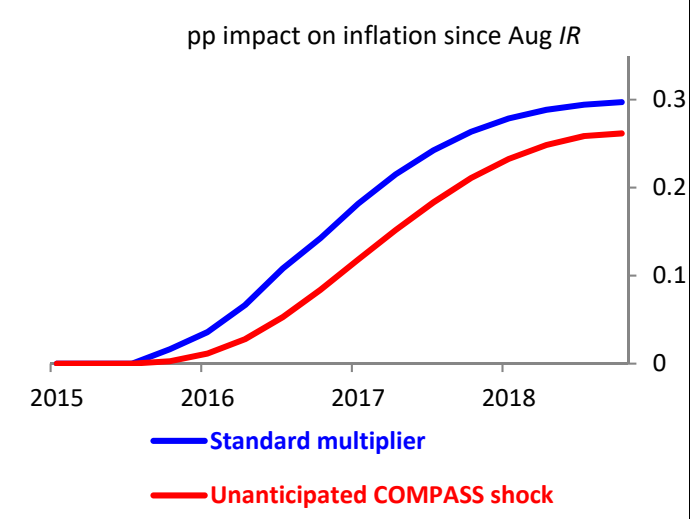

The slope of the Phillips curve

One specific thing that might lead you to question the size of the GDP and inflation effects of the multipliers is the ratio between the two – that is, about the slope of the reduced-form Phillips curve that they embody. One thing we have become aware of as we have used the multipliers is that the slope of their reduced-form Phillips curve is somewhat steeper than the equivalent slope from COMPASS or than that which empirical evidence would point to.

It is hard to be precise about this statement because the structure of models like COMPASS means that there is not really such a thing as a single Phillips curve: different shocks have different relative effects on output and inflation depending on their source and their expected persistence. And the difficulty of estimating even reduced-form Phillips curves means that pinning down a precise number empirically is also very hard. But for what it's worth, a rule of thumb that is broadly consistent with typical COMPASS demand shocks is that all else equal a 1pp increase in the output gap leads to a 0.4pp increase in inflation; in the standard multipliers the ratio is more like 1:0.6. The empirical work contained in SEAD's stocktake of Phillips curve estimates tends to imply a number as small as or smaller than that in COMPASS – though the VECM on which the multipliers are based is also empirically estimated, and MASD's FAVAR model of the transmission mechanism also appears to have a steeper reduced-form Phillips curve. [9]

We also think this can help rationalise some of the large pink "unaccounted for" bars over the forecast in our inflation decomposition (**Chart 18**); since this decomposition assumes a 0.4 mapping from the output gap to inflation, any additional boost from the steeper Phillips curve in the multipliers will mechanically appear in the pink bars rather than the yellow ones. We think the fall in the market curve this round has added about 0.1pp to the pink bars, while the cumulative news since last November has contributed about 0.2pp. Of course that does not mean the inflation profile would be lower if the multipliers had the same Phillips curves as COMPASS – it could equally be that we have the right

Chart 18: The pink bars in our inflation accounting decomposition

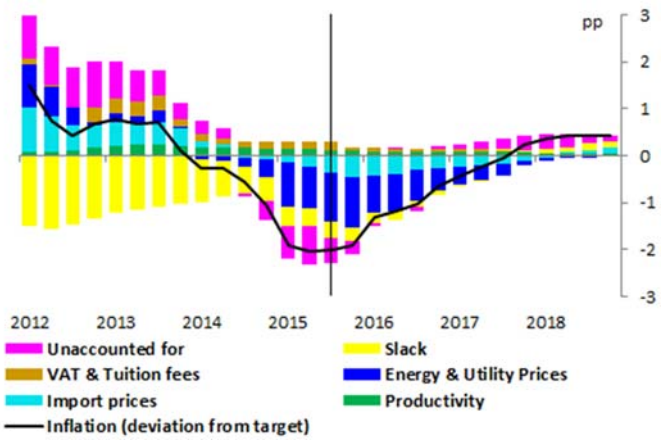

inflation effect but the GDP effect is too low. As shown in the last section, the GDP and inflation effects both look broadly in line with other estimates, and so it is not obvious a priori which if any of these might be out line.

(Of course there are lots of other reasons why pink bars might appear in this chart – this is a simple rule-of-thumb based decomposition of a forecast produced using a complicated and sophisticated model. But these bars are unusually large, suggesting some specific other factor is at work; and we are fairly confident that the multipliers account for at least part of them.)

This difference does not necessarily imply that our multipliers are erroneous: there are some arguments why interest rate changes might in principle have a larger effect on inflation for a given change in demand than other shocks. If you put weight on these, then you might want to keep the slope of the Phillips curve in the multipliers unchanged. For instance, if you think monetary policy acts partly by coordinating expectations, that could lead to an additional effect on inflation beyond the direct impacts via the labour market. DSGE models like COMPASS would also support this, to the extent that anticipated monetary policy shocks have particularly large inflation effects relative to output. We would not put too much weight on this second argument, though, since it is exactly this mechanism that generates the “forward guidance puzzle” and leads us to aim off COMPASS’ responses in the first place!