

MPC MEMBERS

Copies to: MA Leadership Team
ID Leadership Team

From: [REDACTED]
[REDACTED]

ANNOTATED AGENDA FOR MPC SPECIAL TOPIC MEETING ON r^*

This note is essential reading ahead of an MPC special topic meeting devoted to r^* on Monday 20th July. We start by providing a framework for thinking about the determinants of r^* and why it is a relevant concept for monetary policy. We present some estimates of r^* for the United Kingdom, drawing on recent work by the International Directorate and by replicating estimates from the literature. We conclude by discussing some implications for policy.

Framework & terminology

In this section, we present a framework for thinking about r^* and its determinants. Throughout this note, we refer to r^* as the equilibrium real interest rate. We define it as the real rate of interest that is consistent with the full employment of labour and capital resources.

The framework from which this definition arises is quite general and starts from the basic macroeconomic paradigm used in all inflation targeting central banks. The central idea is that prices and wages take time to adjust to changes in the economic environment. This means that monetary policy can affect real spending over the horizon for which prices and wages are adjusting. In the longer run, because prices and wages have fully adjusted, monetary policy can affect the price level and inflation, but not real spending. Following this logic, one way to characterise the framework is to think about various phases of adjustment of the economy after a shock.

- In the **short run** price and nominal wages are slow to adjust to shocks, so that monetary policy can affect the real economy. Because prices do not adjust instantaneously, changes in the short-term *nominal* policy rate affect *real* rates of return on assets and hence real spending.
- In the **medium run** enough time has passed for nominal prices and wages to adjust fully. If set appropriately, monetary policy can achieve an output gap of zero and inflation at target at this horizon. However, the *real* economy may still be undergoing adjustment to long-lived shocks. In particular, stocks of physical and financial assets and their relative prices may still be adjusting. Continued adjustment in these stocks may result in continued adjustment of potential output (e.g. the capital stock may take a long time to adjust). In this phase of adjustment, an appropriate monetary policy can ensure that actual output moves in line with potential output, thus maintaining a zero output gap. This is sometimes called “flow equilibrium”.
- In the **long run** all stocks and flows have adjusted to the disturbance. The economy is in steady state (sometimes called “stock-flow equilibrium”).

The framework implies that the entire ‘term structure’ of the equilibrium real rate matters for monetary policy. In the short run, monetary policy responds to shocks hitting the economy by adjusting the actual real interest rate (via changes in the nominal policy rate). The deviation between the actual and equilibrium real interest rates is often thought of as an indicator of the stance of monetary policy (discussed below). As the adjustment process continues, the output gap and inflation are stabilised. To achieve this, monetary policy is set to ensure that the real rate gap is zero and the path of the actual real rate tracks the adjustment path of the equilibrium real rate to its long-run level.

The framework also helps us to uncover the determinants of the equilibrium real rate at different horizons. Because r^* corresponds to the real rate of interest consistent with full employment of both capital and labour, any factors that move either the demand for investment or the supply of savings could affect r^* . Depending on the persistence of the factors, they may affect r^* in the short, medium or long-term. Note that very persistent factors that cause long-term shifts in r^* are also likely to affect short-run r^* , because the adjustment of the economy to a very persistent shock will likely require changes in saving and investment flows in the short term.

Table 1 provides some examples of factors determining r^* at different horizons. These determinants encapsulate the leading rationalisations for the observed decline in interest rates around the world over the past twenty years or so, as well as hypotheses about the prospects for r^* in the future, such as the secular stagnation hypothesis.

Table 1: Determinants of r^*

Short run	Medium run	Long run
Uncertainty & fear	Deleveraging and debt cycles (inc. fiscal)	Trend productivity growth
Credit supply	Population dynamics	Regulatory environment
World demand	Distribution of wealth	Relative price of investment
Fiscal policy	Financial development	Preferences
	Risk aversion	

Although not immediately apparent from the table, global factors are crucial for the determination of r^* in an open economy like the United Kingdom. The framework set out above can be thought of as applying to the world economy as a whole, so that global r^* in the longer run is determined by the stock-flow equilibrium in the world economy. Such an equilibrium would require, for example, that capital flows between countries had taken place to ensure that no individual country wished to increase or decrease its stocks of assets. Most economic models assume that stock-flow equilibrium is consistent with (near) equalisation of real rates of return across countries and across asset classes. This means that estimates of global long-run r^* are likely to be good estimates of UK r^* for that horizon. Over shorter horizons, however, real rates of return may differ as real exchange rate adjustments take place and there are likely to be barriers to arbitrage across countries, suggesting that both domestic and global factors are important determinants of r^* at shorter horizons.

Just as we routinely use estimates of the equilibrium rate of unemployment to help us to assess the amount of slack in the economy, we can use estimates of r^* to help assess the stance of monetary policy. That is because the gap between r and r^* (the so-called real interest rate gap) is related to the amount of slack in the economy. As with the other equilibrium concepts that we regularly discuss (e.g. the equilibrium rate of unemployment), estimation is necessary because r^* is not directly observable.

Estimates

The framework described above can be used to build models to estimate r^* . Some of the models we consider are very simple and require very few additional assumptions beyond those we made above. Other models include more assumptions to estimate r^* alongside other macroeconomic variables (like the output gap) and to quantify some of its determinants. In this section, we look at estimates from a range of models that vary in the nature of additional assumptions used.

We start with a simple example. Consider forward interest rates on long term nominal government bonds. There is a host of reasons why these rates may not be accurate reflections of the expectations of market participants. But setting those issues aside, these measures can be used to estimate medium to long-run r^* if we assume that market participants are using the framework set out above. Specifically, assume that market

participants expect monetary policy to: (a) close the output gap and return inflation to target in the short run; and (b) keep the output gap closed and inflation on target thereafter. In this case, market expectations of nominal short-term rates at longer horizons, adjusted for inflation expectations (CPI inflation at 2%), provide an estimate of r^* . For example, the (instantaneous) nominal forward rate 10 years ahead is currently just over 3%, suggesting a market estimate for r^* in the medium run of 1%. Similar logic has been used to infer FOMC views of medium run r^* based on the medium term SEP forecasts for the Fed funds rate.

Our r^* estimates for the United Kingdom are presented in Table 2 below. Each of the approaches (many of which have been used at other central banks) is described briefly in an annex at the end of this note. The estimates suggest that r^* is currently very low, but higher at medium and longer horizons. Even though r^* is estimated to be higher at longer horizons, the estimates remain below historical averages of observed real interest rates.

Table 2: UK estimates of r^*

Model/approach	Point %	Range %	Horizon	Type	Open econ?
Taylor rule	-2	-2.0,-0.5	Short	Single equation	No
(2004)	-2.9	-3.3,-1.7	Short	Semi-structural	No
Berger et al (2014)	-1.7	-2.9, 0.8	Short	Semi-structural	Yes
COMPASS	-0.2	-1.3, 0.7	Short	DSGE	Yes
Nominal fwd yield	1	-	Medium	Fin markets	-
No-arbitrage model	0.5	-	Medium	Statistical/fin	-
Investment/saving approach	1	-	Medium-long	Semi-structural	Yes
(2015)	-	-	-	-	-
Hamilton et al (2015)	-0.5	-	Long	Statistical	-

At one level, the current low estimates for r^* are very easy to explain. The policy rate is very low, the output gap is small and inflation is below target. By definition, therefore, r^* must also be low. This logic is embedded directly in the Taylor-rule based estimates and also appears indirectly in the other short-run estimates. But it is also possible to explain at least some of the low level of rates via analysis of the determinants of r^* . Over a medium to long horizon, the investment and saving framework used by [redacted] from the International Directorate points to several secular trends as having been important in driving the fall in global real interest rates since the 1980s. In particular, the ID analysis identifies various secular trends in desired global saving and investment as well as a decline in labour productivity growth since the crisis as being capable of explaining 400bps of the circa 450bps decline in long-run global real rates over that period (Table 3).

Table 3: Secular determinants of r^*

Determinant	Change in r^* using ID's framework
Shifts in desired saving and investment	300bps of which:
Higher desired saving	160bps of which:
- Demographics	90bps
- Higher inequality	45bps
- Savings glut (Asian governments)	25bps
Lower desired investment	140bps of which:
- Decline in relative price of capital goods	50bps
- Preference shift away from public investment	20bps
- Higher credit spreads	70bps
Productivity	100bps
Total	400bps

Other estimates are better suited to explaining short-run movements in r^* . For example, COMPASS explains the estimated low level of r^* largely by weak world and domestic demand. The level of r^* in COMPASS has risen by around 5pp since its trough in 2010/11 and around 3pp since 2013, suggesting that these headwinds are weaker than they were during the crisis and subsequently. (Note that the COMPASS estimate is independent of the forecast, which includes additional information from the suite of models and staff/MPC judgement.)

Implications

The short-run estimates from Table 2 suggest that r^* is currently somewhere between -3% and 0.5% (with a mid-point just below -1%). That compares to an ex-ante short-term real interest rate (incorporating the effect of QE) of around -2% in the Benchmark forecast.¹ This suggests that the current *real* policy rate is probably slightly below r^* .

As noted above, the gap between the actual and equilibrium real interest rate is often used as an indicator of the stance of monetary policy. This perspective would interpret our short-run estimates as evidence that policy is accommodative at present which, qualitatively at least, seems appropriate given that inflation is below target and we believe that the output gap is negative. However, three caveats should be noted:

- This interpretation is appropriate only if short term interest rates are a sufficient statistic for monetary conditions. For open economies like the United Kingdom, monetary conditions will be influenced by the exchange rate. More generally, monetary conditions are influenced by a range of financial market prices that are unlikely to be perfectly correlated with the short-term real rate.
- As explained in the Annex, most of the models of short-term r^* use macroeconomic relationships to estimate r^* as an “unobserved component”. The relationships between real interest rates and spending in those models imply that each model’s estimate of r^* will be related to its estimate of the output gap. So if these models produce different estimates of r^* , they will also produce different estimates of the output gap.
- The extent to which the equilibrium rate is an appropriate target for the actual real interest rate depends on the nature of the shocks hitting the economy. For example, it may be appropriate to set the real rate below r^* if the economy experiences a negative inflationary shock that induces a trade-off between stabilising inflation and the output gap.

The estimates of r^* at a longer horizon (focusing on those with some behavioural content) suggest that r^* is likely to increase from its current low levels. The message from these estimates is consistent with the MPC’s forward guidance issued in February 2014: Bank Rate is likely to return to a level below pre-crisis averages in the medium term.

Judging where r^* is likely to settle at a longer horizon requires forming judgements on the secular trends identified by ID from Table 3. ID conclude that many of the factors they identified are likely to persist, but that there is a great deal of uncertainty about the magnitudes and their persistence.

It is important to note that our estimates have the advantage of being independent cross checks on the forecast. Different information sets and models have been used. The flip side of this, however, is that each of our estimates is based on a smaller information set than is used in the forecast.

¹ Computed as the shadow policy rate that incorporates the effect of QE (equal to -0.6% in 2015Q3) minus the forecast for annual CPI inflation in 2016Q3 (1.3%).

Annex: Explaining the approaches to estimating r^*

This annex briefly explains each where each of the estimates from Table 2 comes from.

Taylor rule

r^* is the time-varying intercept in a Taylor rule. The logic is that if the output gap is closed and inflation at target, then the real interest rate is equal to r^* . The point estimate comes from a benchmark rule with Taylor's original calibration (coefficient of 1.5 on inflation and 0.5 on output gap). This estimate also happens to coincide with an estimate from a reduced-form state space model in which the parameters of the rule are estimated simultaneously with the unobserved r^* . The range reflects different variants of the rule including: core inflation instead of headline; interest rate smoothing; differences in the treatment of the unobserved r^* . All variants use the shadow policy rate that incorporates the effect of QE as the measure of the policy rate. All variants also use the forecast measure of the output gap and treat it as fully observable.

(2004)

r^* is filtered in an "unobserved components" model. The logic of this model follows our definition of r^* . It comprises a reduced-form Phillips curve, in which expected inflation is a function of lags of itself and lags of the output gap, and a reduced-form IS curve, in which the output gap is a function of lags of itself and a real interest rate gap. If the policy rate has fallen, but the output gap and inflation are unchanged, then r^* must also have fallen. The parameters of the model are estimated simultaneously (using Bayesian maximum likelihood) with the unobservable variables, r^* and y^* (which are assumed to evolve as random walks). The model is estimated on data for the shadow policy rate, GDP and inflation. The range of estimates in Table 2 is computed from the 5th and 95th percentiles of the posterior distribution for r^* .

Berger et al (2014)

As in [redacted], r^* is filtered in an unobserved components model. The model is an open economy variant of the Laubach-Williams model that has been estimated for the United States (which is similar to the Larsen-McKeown model). The model is also based around reduced-form IS and Phillips curves, but it extends Larsen-McKeown in two ways: a) it incorporates a common unobserved component that can drive r^* and changes in y^* (which can be interpreted as movements in trend growth); b) it incorporates a real exchange rate gap term in both the IS and Phillips curves. In flow equilibrium, all gap terms must equal zero, but the r^* variables (including the real exchange rate) can move around. The model is estimated in the same way as Larsen-McKeown and the range is computed in the same way.

COMPASS

COMPASS is a New-Keynesian DSGE model. The concept of r^* in NK models is a more precise variant of our definition: r^* is the real interest rate that would prevail in a version of the model in which prices and wages are fully flexible. The logic of this definition follows from the observation that if prices and wages are fully flexible, then they would adjust instantaneously to shocks and so output would be at potential. The associated real interest rate is sometimes called the flexible-price real interest rate. Like most models in its class, COMPASS is linearised around a constant (de-trended) steady state. The steady state real interest rate in COMPASS is assumed to be 2.3%. This means that r^* in COMPASS is determined by transitory (though possibly persistent) shocks. These shocks are identified as the most likely combination of shocks to have driven all of the variables in the model that are mapped into data given the model's structure. For example, a simultaneous fall in inflation, output, hours and the policy rate is likely to be explained by negative shocks to demand, which would push down on r^* . A fall in output with hours and inflation broadly unchanged would tend to be explained by productivity shocks, which would have less impact on r^* .

Nominal fwd yield

A nominal forward rate (at horizons beyond which policy can affect the real economy) can be used to calculate market expectations of r^* at a particular horizon under the strong assumptions that: a) there are no term or other risk premia; b) policy is credible (so that people expect the real policy rate to track r^* in the medium to long run). In the text, we use an instantaneous nominal forward yield 10 years ahead.

No arbitrage model

The main weakness in trying to infer market participants' expectations of r^* from government bond yields is that these yields are comprised of term premia as well as expected interest rates. Term premia cannot be observed directly, so it is necessary to use a model to try and estimate them. The models used in MFAD derive from the affine term structure model of Duffee (2002). They are no-arbitrage factor models, where bond values are exposed to risks and where discount factors of risk-averse investors are modelled as a function of a small number of risk pricing factors. Imposing a 'no-arbitrage' assumption provides restrictions on the model parameters. The models can also be used to produce forecasts of the short-term policy rate over the maturity of the bond, e.g. 10 years. The point estimate in Table 2 is for an instantaneous forward rate 10 years ahead (calculated as the nominal yield minus the estimated term premium minus the inflation target).

Investment/saving approach

In this framework, r^* is the rate that equilibrates desired saving and investment. ID used this framework to identify secular trends that have driven the saving and investment schedules over the past. For example, an uneven population distribution (as caused by the ageing of the large baby-boom generation) would tend to lead to higher desired saving when the large generation is middle-aged, which would tend to reduce r^* . These sorts of trends can evolve slowly over many years and are not well-captured in standard DSGE models (and could be captured, but are left unexplained in semi-structural models like Berger et al or Larsen McKeown, making these models less useful in making forward-looking predictions).

Hamilton et al. (2015)

This approach provides a statistical description of how r^* may have evolved over the past at low frequencies (i.e. stripping out cyclical factors) by fitting a time-varying parameter AR(1) model with stochastic volatility (ARCH(2)) to the historical ex-ante real rates series. The identifying assumption (as far as a read across to r^* is concerned) is that the actual real interest rate has fluctuated around this low-frequency measure of r^* over the past. It is silent on the causes of these low frequency movements and cannot be used for forecasting.