BANK OF ENGLAND

Speech

The transition to a new normal for monetary policy

Speech given by

David Miles, Monetary Policy Committee, Bank of England

Mile End Group, Queen Mary College, London

Thursday 27 February 2014

I would like to thank Jochen Schanz and Chuan Du for research assistance, and others for helpful comments. The views expressed are my own and do not necessarily reflect those of the Bank of England or other members of the Monetary Policy Committee.

What might be the new normal for monetary policy? Let me put that question rather more precisely: what level of Bank Rate can we expect to see once we have got back to a position where output is more in line with capacity and inflation is in line with target? That is a position we have not been in for a long time in the UK – but we may get there in the not too distant future. In the February Inflation Report the Monetary Policy Committee showed a forecast based on a path for Bank Rate implied by (the then) market rates which had a central prediction for inflation around three years ahead that was only a little below target and where slack (the output gap) was also small. Suppose we did get to that place; what level of Bank Rate could – on average – keep you there? It is useful to give that level of Bank Rate – which is not likely to be constant over time – a name; let us call it the neutral or the equilibrium rate (names which I will take to be synonymous).

The answer to this question may to some extent depend on how the balance sheet of the Bank of England evolves and what happens to the stock of assets that the Bank has bought since early 2009 as part of its monetary policy. In the Inflation Report we used as a conditioning assumption an unchanged stock of assets. It is not very likely that the stock will stay at that level indefinitely. But I think that there are reasons to believe that the balance sheet of the Bank – and in particular the stock of gilts that it holds – is not likely to be a major factor in determining what I call the neutral, or equilibrium, level of Bank Rate. So I am going to proceed by first considering that interest rate in isolation from the issue of what happens to the stock of assets the Bank has accumulated in recent years. I will then consider what the implications of the size and the nature of the Bank's stock of assets might be and argue that this probably has a rather small impact on the equilibrium, or neutral, rate.

Where has the neutral rate been?

I used to think that a neutral level of Bank Rate might be around 5%. That was about the average level of Bank Rate in the period between when the MPC was established in 1997 and the eve of the financial mess in 2007 – a period when inflation was always close to the 2% target and when it seemed that output was fluctuating in a fairly narrow band around normal capacity. As it happens 5% is also very close to the average level of interest rates set by the Bank of England since it was created in 1694[1] (the 1694-2014 average is 4.8%). Inflation fluctuated a great deal from year to year over that period. But it appears that the average annual rate of increase in consumer prices over this 320 year period is – rather surprisingly – almost exactly 2%, which is both the current target and the current inflation rate.[2] 5% was also the level of Bank Rate during the longest period when the Bank of England left the rate unchanged – the 103 year period from 1719 to 1822[3].

[1] Though the policy rate set by the Bank of England has not always been called Bank Rate, and at times usury laws meant that there was an effective ceiling on interest rates.
[2] Based on estimates since 1750 from O'Donoghue et al (2004), Table 1. Estimates since 1949 use RPI.
[3] Had there then been an MPC meeting each month to decide the level of Bank Rate they would have had over 1200 consecutive meetings at which they voted to leave the rate unchanged. The Bank was prohibited by usury laws from raising Bank Rate above 5% for some of this period.

Figure 1: The Bank of England's main policy rate

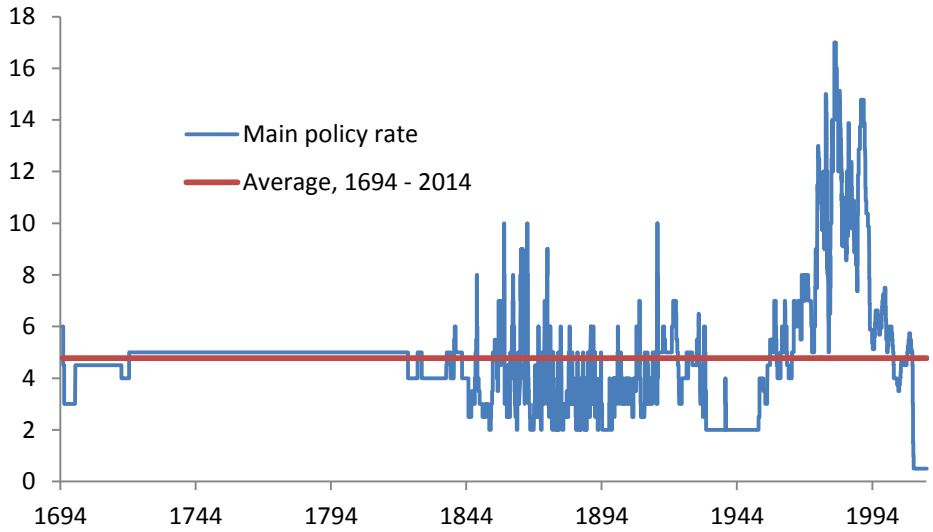

Source: Bank of England, http://www.bankofengland.co.uk/statistics/Documents/rates/baserate.pdf. The chart shows the relevant Bank of England's key lending rate over time (Bank Rate, Minimum Lending Rate, Minimum Band 1 Dealing Rate, Repo Rate, Official Bank Rate).

So the 5% number has some claim on our attention.

But that level of Bank Rate consistent with keeping demand in line with supply and inflation at target – starting from a position where they are at those levels[4] – is not likely to be a constant. People give this rate many different names – neutral rate, natural rate, equilibrium rate, r*, even r**. Names don't matter so much, but what does matter is that we are clear about what we are talking about. So, to be clear: I mean that level of the policy rate which if set by the central bank from a position with inflation at target and output in line with potential then – on average – output would evolve in line with supply capacity and – on average – inflation would stay at the target level. One can call this an equilibrium rate; in doing so there is no implication that this rate is unchanging. Even if 5% was roughly the equilibrium rate for the decade before the recent financial crisis (indeed even if it was the *average* equilibrium rate for the 300 years before the crisis) this does not mean it will be the right level for the period ahead.

Indeed there are reasons to think that for some time to come the level of Bank Rate that will keep demand and supply consistently in balance and keep inflation at the target rate is likely to be below (maybe well below) the 5% figure.

[4] Which is NOT the position we are at in the UK today since it seems likely there exists a significant amount of spare capacity, or slack. Much of that slack exists in the labour market where the level of unemployment is probably around 1% higher than a sustainable rate and where even for those employed the aggregate amount of hours worked is below that which is desired.

Drivers of the equilibrium rate

One way to think about the equilibrium rate is that it is the sum of an equilibrium safe real rate of interest and the inflation target. One reason to focus on safe rates is that the policy rate set by a central bank is usually close to the rate it will pay on some of its liabilities, and those liabilities are safe because they are the ultimate means of payment and a central bank will not default on them. This is certainly true for the Bank of England; Bank Rate is the interest rate paid on reserves held by commercial banks at the Bank of England.

The real safe rate might be measured by the return on inflation proof government bonds which have little chance of default. The UK government (and those of several other countries) have issued such bonds for several decades now. And the trend in the level of real yields on such safe assets has been downwards for at least two decades. Figure 2a (which uses data from King and Low (2014)) shows estimates of the annual spot yield on 10-year inflation-indexed bonds issued by the UK and, where data was available, other G7 countries. If we take the expectations theory of the yield curve as a rough approximation, then the yields shown in Figure 2a are roughly the average equilibrium short-term real safe rate that market participants believe will prevail over the following 10 years.

Figure 2a: An estimate of the real risk-free 'world' interest rate

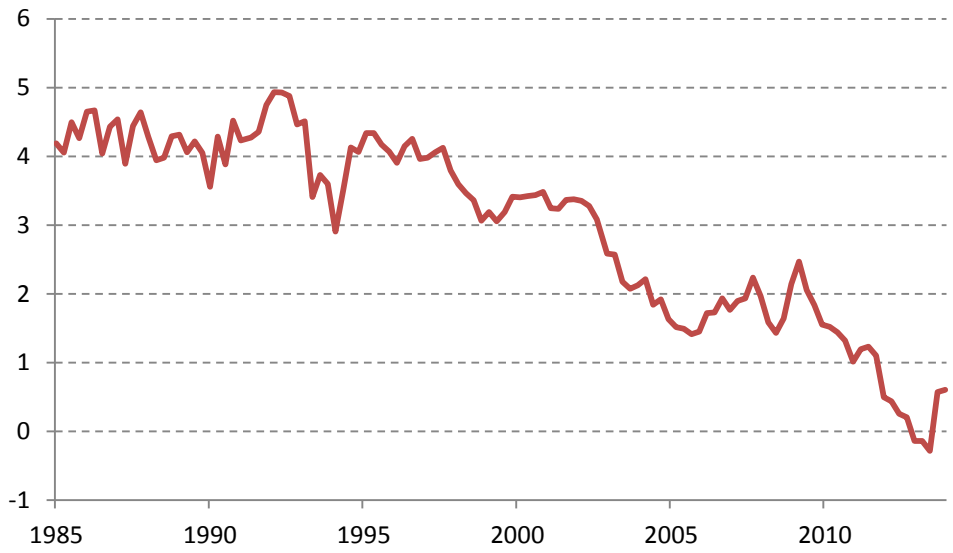

Source: King and Low (2014): a simple average of the estimated spot yields on 10-year inflation-indexed bonds, averaged across all G7 countries for which data are available (except Italy). Estimates before 1999 are based on UK yields only. Estimates are obtained from inflation-indexed bonds using Svensson's (1994) yield curve model.

There may well be risk premia that drive a wedge between these 10 year spot rates and expectations of average short term real rates set by a central bank. But if we use measures of (instantaneous) forward real

rates the picture that emerges is very similar. Figure 2b shows instantaneous forward rates 10 years ahead implied by yields on indexed linked gilts. Figures 2a and 2b reveal a very similar pattern of declining real rates which recently dipped down quite sharply.

Figure 2b: Instantaneous forward real rate 10 years ahead based on index linked gilts

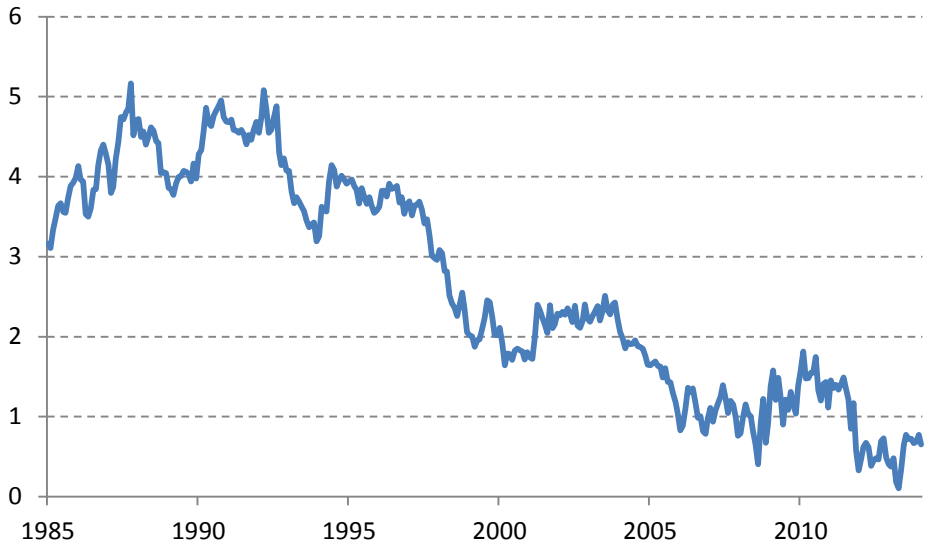

Source: Bank of England.

There are different explanations for the declines in safe real rates, with different implications for whether they will persist. One is that the supply of savings in emerging markets has steadily increased without meeting a corresponding demand for investment goods. This could have driven the return on safe assets lower. But this may change: some developing countries may increase their investment spending, for example to improve their infrastructure and housing. In others, such as China, the average age of the population will likely increase; which could reduce their supply of savings.

Another explanation for the decline in real interest rates on safe assets over the past two decades is that expected future GDP growth (and growth in consumption) has been falling. Lower growth in consumption means that equilibrium returns on safe assets might fall as people need less inducement to save towards a future which is less affluent. GDP growth is likely to depend on trends in the labour force and labour productivity. Both labour force and labour productivity growth indeed appear to have slowed somewhat during the two decades ahead of the crisis in the G7 countries (Figure 3).

Figure 3: Labour force and labour productivity growth, G7

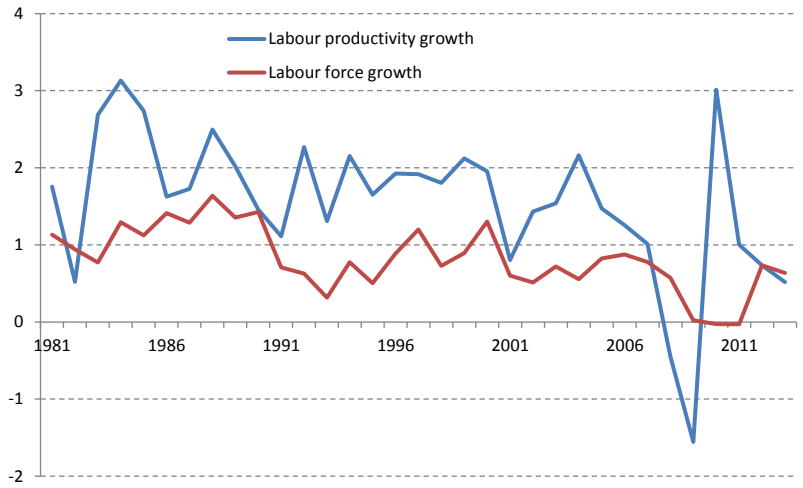

Source: IMF WEO Database, October 2013. Data for 2013 includes estimates.

So perhaps longer term real interest rates are currently so low because investors believe that this decline will continue. I am not convinced this is really the story behind the recent sharp lurch lower in real safe interest rates and I am more optimistic that growth will pick up. But others may be less optimistic and what matters is the wealth-weighted average of opinions; and my weight is not very high.

But there is one factor that I believe has been behind the recent very sharp fall in real yields and which I think will be persistent. This is the impact of changing perceptions of the level of risk in the world. Rises in perceived risk will increase the attractiveness of assets that generate a real return which has little risk – and indexed bonds issued by governments with very small default risks are probably the best example of such assets. Such rises in risk will also tend to increase the difference between the returns on safe assets (which are most closely linked to rates set by central banks) and rates offered on other, less certain assets (which are more relevant to many spending and savings decisions in the economy). A rise in that spread between safe rates and rates on riskier assets is likely to mean that the rate set by a central bank should be lower.

A rise in perceived risk has, I believe, been one of the main reasons why real yields on longer term (inflation-proof) government bonds fell substantially after the financial crisis which reached its worse point in 2008. Figures 2a and 2b show that in the few years leading up to the banking crisis in Autumn 2008 real interest rates on government-issued indexed bonds had been fairly flat; from around 2005 there had not been a continuation of the downwards trend in real yields that had started in the mid 1990s. But since the end of 2008 there was then a decline of almost 2pp in real yields – real yields on 10 year indexed bonds are now close to zero but had been around 2% in the years immediately before the financial crisis.

Instantaneous real forward rates on indexed gilts have fallen by a bit less, but for the last couple of years they have been significantly under 1% – markedly lower than in any period up to 2008.
I think it is worth considering carefully how powerful and persistent a rise in perceived risk might be in driving real rates on safe assets. To do that I want to draw on insights from work by Robert Barro on the impacts of relatively rare – but extreme – events.

Big risks and the safe rate of interest

Perceptions about the scale and likelihood of extreme bad events have probably changed significantly since the financial crisis. Events took place in recent years that many thought close to inconceivable: the almost total implosion of the financial sector; a downturn as serious as the Great Depression; and a subsequent recovery (at least in the UK) weaker than seen after the depression in the early 1930's. People probably now think of these events as rare but not inconceivable. People who might have thought of them as one in a 100 year events (or even 1 in a 1000 year events) might now think of them as one in 50 year events – or even 1 in 25 year events.

This has potentially significant implications for asset prices and the safe interest rate. Robert Barro[5] developed a framework for thinking about how perceptions about rare events influence asset prices. He made some fairly standard assumptions about how risk-averse people are and how they are prepared to accept a lower return on safe investments relative to those whose returns are volatile. He made a crucial – and empirically strongly justified – assumption that as well as what you might call standard variability about a trend in economic activity, there was also a chance of extreme bad events.[6] Those fluctuations in economic activity were assumed to be the fundamental drivers of the returns on risky assets. I have generalised slightly his empirical model of the volatility of economic outcomes to best match the history of GDP for a large group of countries over the past 200 years. I have used that to assess how changes in the chances of extreme events affect the return on safe assets.[7] Let me start by sketching the central ideas here and then describe what it implies (more details are in the annex).

Like Barro, I assume that there are two broad asset classes: risky assets with uncertain real returns and 'safe' (that is, government, inflation-linked) bonds. The first type of asset is one whose cash flows are linked to aggregate corporate income, which is itself linked to the evolution of GDP. It is natural to interpret it as a diversified portfolio of equities. The second asset, whose real cash flows do not depend directly on GDP and which are fixed, I interpret as a government bond. Barro makes a natural assumption, which I follow, that the stream of income on equities follows GDP closely. (In other words the share of corporate income in GDP is

[5] Barro (2006).
[6] For example, the decline in UK real GDP from 2008 to 2009 (-5.2%) has a probability of 0.0004%, or one in 240,000 years, according to a normal distribution fitted to UK annual GDP growth rates from 1949 – 2006. (This distribution would have a mean of 2.9% and a standard deviation of 1.8%.)
[7] This section draws heavily on Miles et al (2005).

roughly flat over time, though it may vary from year to year). For a given assumed pattern of risks over future income/GDP, we can use standard asset pricing techniques to value assets.

I generalise the Barro model in one small, but important, way so that it accounts better for one of the key features of the historical GDP growth rates. I allow for there to be occasional good events – positive surprises to output – as well as negative events which can be, on rare occasions, very extreme. So the generalised model allows for two sorts of rare events:

1. **Extreme and rare bad events** — which create a small probability of a very poor outcome for corporate and other incomes but do not cause default on government bonds.

2. **Unusual events that may be good or bad**, but which are less rare, and rather less extreme — there is as great a chance that, if one of these events occurs, it generates a big positive shock to GDP as a big negative shock.

Investors are risk-averse in this model. The more likely are bad events, the more reluctant they are to invest in risky assets and the greater their demand for safe bonds. So the more likely are bad events, the lower the risk-free interest rate will be in equilibrium. By how much depends crucially on investors' risk aversion. The key inputs to the model that need to be calibrated are:

- Investors' risk aversion. I choose a value of the so called risk aversion coefficient of 4, which is consistent with microeconomic evidence.
- The probability of an unusual event that could be either bad or good, with equal chance, and that is smaller in impact than the extreme bad event.
- The average rate of growth and its variability of real per capita GDP in normal times, when no unusual events occur. Setting average growth in normal times to 2.1%, with a 3.1% standard deviation, fits the data well.
- The probability and severity of an unusually bad event, which causes a strong and prolonged fall in GDP and corporate earnings (but does not cause default on government bonds).

The precise calibration of the probability and severity of an unusually bad event is difficult because these events are, by definition, quite rare. The banking crisis that the UK recently experienced has surely been an unusually bad event. Relative to a path that grew at the long run average rate of around 2.5% per year, the level of output in the UK is now around 15% lower than it would have been absent a crisis. And that crisis may have more lasting effects. If labour productivity growth remains 1pp below its pre-crisis trend for the next ten years, losses induced by this banking crisis could reach 30% of GDP. Something which did knock 20-30% off the level of GDP I would certainly call an extreme bad event.

Table 1 shows how the model's prediction of the change in a risk-free rate when investors re-assess the likelihood of an unusually bad event. The results suggest that doubling the estimate of the likelihood of a bad event from 1% (1 in a hundred years) to 2% (1 in 50 years) reduces the risk-free interest rate by around 1.5 – 2% if we consider extreme events that knock between 20 and 25% off the level of GDP. If such bad events knock 35% off GDP then going from a world in which people see them as once in a century events to one when they are seen as coming along on average once every fifty years would knock a huge 4% off the safe real rate. Clearly that could take the safe real rate negative.

Table 1: Response of the risk-free rate to changed perceptions of unusually bad events

Loss in GDP when event occurs	**Change in equilibrium risk-free rate (pp) when perceived likelihood of unusually bad event changes from…**	
	1% to 2%	1% to 2.5%
-20%	-1.4	-2.0
-25%	-2.0	-3.0
-35%	-4.2	-6.2

Historically, the risk-free real interest rate has been around 3%. Adding an inflation target of 2%, would generate a crude measure of the neutral level of the central bank policy rate of around 5%. Relative to those sorts of levels the model's calibration suggests that if investors permanently raise their perceived likelihood of an extreme negative event, the neutral level of Bank Rate might fall considerably.

There is another – in some ways related – way of thinking about the drivers of what I have called the equilibrium, or neutral, level of the policy rate set by the central bank. This focusses not so much on the safe real rate but on the spread over and above that safe rate faced by private agents who want to borrow. This way of thinking about the neutral (or equilibrium) level of Bank Rate also leads to the conclusion that it is likely to be lower in the years ahead. This is because the spread between risky rates and safe rates[8] may be persistently higher.

What matters for conditions in the wider economy, and for the incentives to spend and to borrow, is the rate of interest faced by households and companies. There are of course a very large number of such rates. One could think of each of those rates as determined by the safe rate and a spread that reflects the costs of intermediating funds between savers and borrowers plus a compensation for risk. For example, banks may choose to set the interest rate that they pay on insured deposits below Bank Rate if households find bank deposits to be the most convenient form of investing money safely (and they have no direct access to deposits at the Bank of England paying Bank Rate). And banks set their lending rates above the safe rate to compensate for the risk of not getting all the money they have been promised back.

[8] One could think of the safe real rate as the yield on an inflation proof bond and in nominal terms one could think of the safe rate as Bank Rate – the rate paid on reserve balances held at the Bank of England.

At the start of the crisis, lending spreads increased sharply. There are several reasons for this. The most important one is probably that banks' perception of the risks associated with lending changed. The resulting increase in spreads over the safe rate is exactly what Barro's model predicts. Another reason might be that banks' funding costs increased, and that banks were trying to maintain their profitability by charging higher lending rates to their borrowers. That was also driven by risk – this time the risk that banks themselves could not repay debt. Banks' funding costs increased not only because investors changed their perception of how risky banks' lending was but also because reformed bank resolution arrangements make it much less likely that investors purchasing bank debt will be bailed out when a bank fails.

Figure 4: Household and corporate borrowing rates

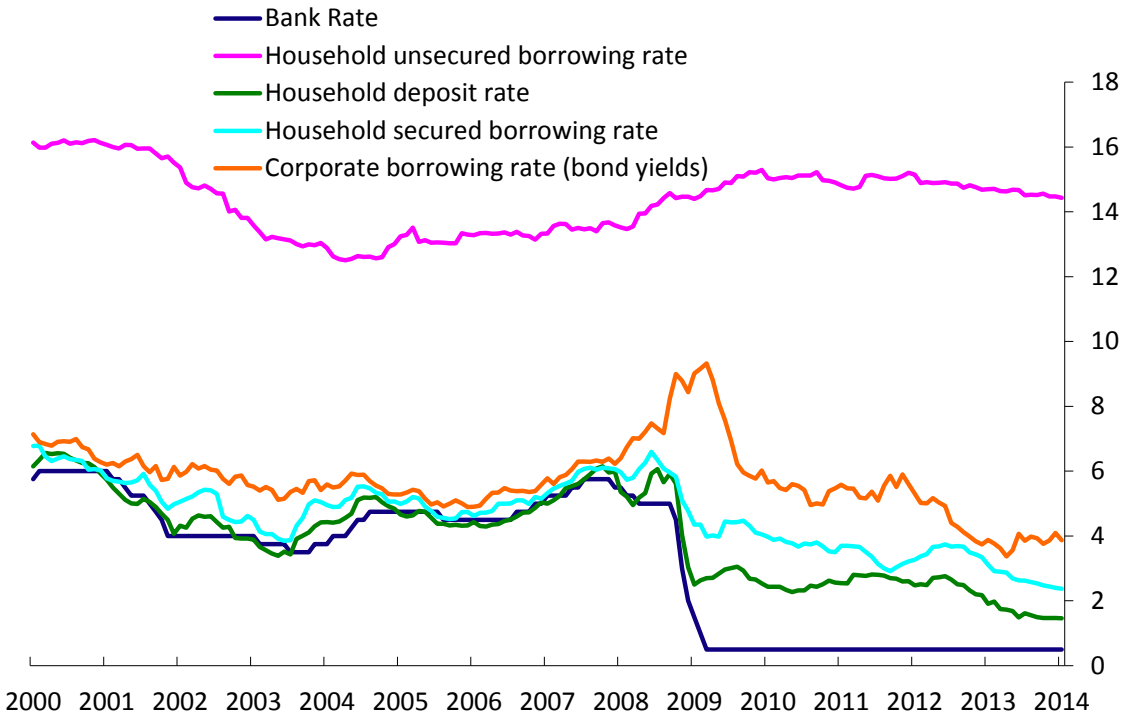

Source: Bank of England and Bank of America Merrill Lynch. The corporate borrowing rate is yield to maturity on uncollateralised bonds issued by companies with a current maturity of 1-10 years. The household unsecured borrowing rate is an average of credit card, personal loan and overdraft rates. The household lending and deposit rate series show data on quoted rates by UK Monetary and Financial Institutions. Household secured rates refer to average 75% LTV, 2 year fixed rate mortgages.

Another reason for higher lending spreads might be that banks competed less aggressively for market share during the crisis, and instead focused on granting loans that looked highly profitable while they were rebuilding their capital position. And banks have stopped cross-selling some highly profitable products, such as payment protection insurance, when granting loans; this will push them towards raising rates on loans.

Spreads between lending rates and Bank Rate may come down somewhat in the coming years once banks have built their capital to more adequate levels and if competition in the banking sector picks up. But spreads on risky lending, whether by banks or by capital markets, are unlikely to fall to where they were

before the crisis. In part this is because where those spreads were before the crisis was quite probably unsustainable. Table 2 illustrates. It shows just how low spreads were during the decade ahead of the crisis compared to the six decades before.

Table 2: Corporate bond and mortgage spreads over the risk-free rate

	Corporate bonds, spread over 5 year government bonds	**Mortgage rates, spread over Bank of England policy rate**
1938-96	1.6	1.2
1997-07	0.9	0.5
2008-13	3.5	2.7
Jan 2014	2.0	1.9

Sources: GFD for corporate bonds yields from 1938-96; Bank of America / Merrill Lynch for 1997-2014. Building Societies Association (1990) for mortgage rates 1938-1988. ONS for 1989-1994. Bank of England for mortgage rates 1995-2014 (average 75% LTV, 2 year fixed rate).

Of course, the structure of the economic environment changes over such a long period. Financial liberalisation may be one factor behind lower spreads during 1997-07 compared to 1938-96. But another likely reason is that in the years before the crisis, lenders and borrowers underestimated the risk that debt would not be repaid. Those risks are now perceived to be significantly higher and are likely to stay higher for many years.

So spreads between risky lending and Bank Rate are likely to be higher than immediately before the crisis. This is another indication that the neutral level of Bank Rate is likely to be lower for the foreseeable future.

There are other reasons why Bank Rate is unlikely to return soon to its pre-crisis average. A number of headwinds to demand are likely to persist for some years. For example, the Government plans to continue to reduce the structural budget deficit for several more years, and growth in euro-area demand may continue to fall short of its pre-crisis average rate. I view these factors, which reflect headwinds to demand in the aftermath of the crisis, as likely to be more short lived than those factors which are driven by higher perceptions of risk. I suspect the memory of the crisis and the effect it has had on the risk perceptions will last longer than the impact on spending and taxes of the need to rebuild balance sheets. But both factors are likely to affect the equilibrium (or natural) rate for several years.

Overall assessment of the neutral rate

The range of factors I have described – a lowering of safe rates because perceived risk has risen; a rise in spreads over safe rates because they may have fallen to commercially unsustainable rates before the crisis and also because perceived risks have risen; a less persistent (though not transitory) fall in some

components of demand as governments try to repair their own damaged balance sheets – are all closely linked to the financial crisis. These effects will last – some more than others. Each of these factors is potentially significant – safe real rates have declined by between 1% and 1.5%; in many cases spreads have risen by about that amount again. There is an element of double counting though in just adding those two factors together because they are both manifestations of the same fundamental factor – namely a rise in perceived risks. To add them together and conclude that the neutral rate may have fallen by of the order of 3% is almost certainly to double count. But for an extended period we might expect the equilibrium policy rate to be as much as 2%, and conceivably a little more, below what we used to think of as normal.

Figure 5: UK instantaneous risk-free forward yields

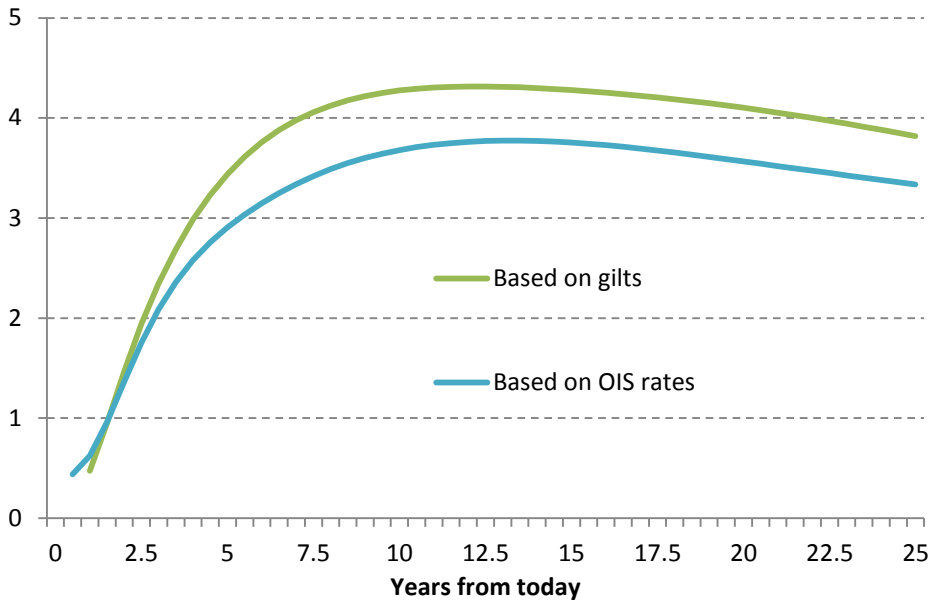

Source: Bank of England. Estimates as of 24 February 2014.

Financial markets participants appear to agree that Bank Rate is unlikely to reach 5% over the foreseeable future. Figure 5 shows that estimates of risk-free short-term interest rates based on OIS contracts are expected to rise from their current level of 0.5% to close to 3.5% over a period of about 8 years, and then remain there.[9] Forward rates based on gilt yields are a little higher.

But might a big part of the reason for why forward rates – either on gilts or OIS contracts – stay well under 5% be that the Bank of England's asset large-scale purchases of gilts continue to depress the term premium substantially? The impact of the Bank holding gilts is something we have ignored so far. Is that a major factor in assessing what a neutral level of Bank rate is?

[9] An OIS contract is an interest rate swap contract; it pays the compounded overnight interest rate prevailing over the term of the contract. The contract is marked-to-market and collateralised on a daily basis. OIS curves are estimated from the fixed rate that sets the contract value to zero. Gilt curves are estimated from the yield of UK government bonds. Each can be used to derive expectations of future risk-free rates. OIS curves may be preferred because they should contain no credit risk premium. However liquidity in OIS contracts beyond the 5 year horizon is relatively limited. Beyond that horizon gilts are more liquid.

The Bank of England Balance Sheet

The Bank of England currently holds £375 billion worth of government bonds on its balance sheet. Empirical models suggest that purchasing these assets may have had an impact on GDP and inflation equivalent to that of lowering Bank Rate by between 200bps and 400bp during the crisis.[10] So should we expect the neutral level of Bank Rate to be affected as long as the Bank of England holds on to these government bonds? And by how much would the neutral level of Bank Rate fall once the Bank of England starts selling these bonds?

Two factors play a role here. One is whether or not the Bank of England will ultimately want to sell all of the government bonds it has acquired during QE. The other is the extent to which varying the stock of assets on the Bank of England's balance sheet still matters for monetary conditions now that financial markets are functioning more normally.

Let us start with the first – and in some ways less fundamental – issue. Figure 6 shows commercial banks' holdings of liquid assets over the past 45 years. Banks' liquid asset holdings shrunk dramatically during the 1970s and 80s, to a minuscule 0.5% to 1% of their total assets in 2006. At that time, banks held around £20bn worth of reserves at the Bank of England. That was a figure thought reasonable in a world where it seemed inconceivable that interbank markets might freeze up – a world where there seemed no need to hold many assets that would remain liquid in the most disrupted markets, and where it seemed very unlikely that a major bank might find that it has lost access to funding markets.

Now, after the crisis, banks' reserves at the Bank of England are about 10 times higher at around £300bn. With the crisis still a painful memory, and with its after effects still very much with us, banks are likely to demand substantially more than the £20bn of reserves they thought appropriate before the crisis. Stricter prudential liquidity rules are strongly reinforcing this increase in demand for liquid assets.

[10] See Joyce et al (2011).

Figure 6: Sterling liquid assets relative to total asset holdings of UK banking sector

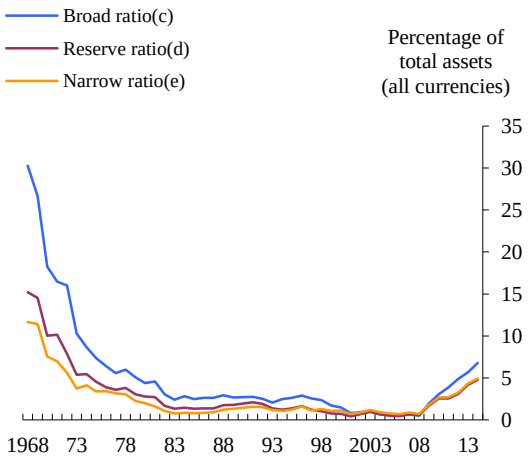

Source: Bank of England. Last observation is November 2013. Data for building societies are included from 2010 onwards. Prior to this, data are for UK banks only. Data are end-year except for 2013 where end-November data are used.
Broad ratio: Cash + Bank of England balances + money at call + eligible bills + UK gilts.
Reserve ratio proxied by Bank of England balances + money at call + eligible bills. Narrow ratio: Cash + Bank of England balances + eligible bills.

So there may be a structural reason – independent of monetary conditions – why the Bank of England might persistently have a much larger balance sheet than before the crisis and may therefore want to keep more government bonds on its balance sheet than before QE started. Banks' reserve holdings are liabilities for the Bank of England. The more reserves banks demand, the more assets the Bank of England needs to hold on its balance sheet, assuming it is willing to supply those reserves. At the moment, these assets largely take the form of government bonds purchased outright from the private sector as part of a monetary policy operation (QE). Some of them could be retained to back the reserves that banks demand in future[11]. That is certainly not inevitable; the Bank will face a choice over what assets to hold against a potentially far higher stock of reserves chosen by commercial banks, and it is not obvious that most of the assets will be gilts. The Bank may opt to back those reserves largely using repos, as it did before the crisis. In this case, there would be no structural reason for the Bank of England to hold on to the government bonds purchased during QE. But commercial banks might themselves be buyers of some of those assets if they were to be held as collateral to be used in repos with the Bank or – more directly – if they were to form a large part of the (much greater) stock of highly liquid assets they want to hold.

A substantially higher demand by the commercial banks for the most liquid assets – reserves and gilts – may mean that both the Bank of England and commercial banks will want to own a greater stock of gilts than they

[11] Strictly speaking the Bank's extra assets that have come about because of asset purchases (QE) consist very largely of a loan (at Bank Rate) to the Asset Purchase Facility, which holds the gilts and is indemnified by HMG. Should they be held by the Bank against a permanently higher level of reserves chosen by commercial banks, those gilts would come to sit directly on its balance sheet. The implications for how many gilts are released on to the market to be held outside the Bank when QE is unwound is unchanged and is as described in the text.

did before the crisis. If it is the Bank that ends up holding a much larger stock of gilts then it might not be a major net seller of the gilts purchased in the QE operations; if it is the commercial banks then the greater demand from them will be an offset to any sales by the Bank[12]. Either way the impact on monetary conditions – and therefore on the neutral level of Bank rate – may well be minimal.

There is a much more fundamental (less institutional) reason to think that if the Bank were to sell gilts at a time when financial markets were operating relatively smoothly the impact on monetary conditions might be small – and much smaller than the impact when the Bank bought those assets in more stressed conditions.

There is an economic literature, going back to Wallace (1981), which shows that if financial markets are working efficiently and the private sector understands that the central bank's balance sheet is really part of the overall balance between assets and liabilities of the public sector, then asset purchases and sales by the central bank have no real effects. In particular, the purchase or sale of assets by the central bank would not really affect monetary conditions and not affect what we have called the neutral, or equilibrium, policy rate of interest. The idea here is really the same as behind the famous Modigliani Miller theorem of corporate finance.

How well financial markets work matters because asset purchases change the composition of the portfolio of financial assets of households, and households' response to this change depends on the state of financial markets. If financial markets work well, households choose between money and bonds on the basis of their payoff properties. These payoff properties do not seem sufficiently different to allow central bank asset purchases to cause bond yields to fall significantly. I provide some evidence for that in a recent paper with a colleague at the Bank.[13] We show that even when households are credit constrained, the impact of asset purchases on GDP is either zero or economically insignificant when financial markets are operating normally[14].

In contrast, if markets are dysfunctional asset purchases are likely to have a somewhat larger impact (see, for example, Woodford and Curdia (2010) and Chen, Curdia, and Ferrero (2012)). Here is one way in which that might work. Suppose that in normal times, companies raise most of their external funds from banks but also issue bonds to tap financial markets directly. This was the situation in the UK before the crisis. Now suppose that banks, concerned about their own ability to refinance themselves, stop lending to companies. If non-bank financial institutions can buy a much greater quantity of newly issued corporate debt

[12] One could see this as creating an asymmetry between the way in which the Bank bought gilts – which were largely from the non-bank private sector – and the way in which they were sold.
[13] Miles, D. and Schanz, J. (2014), 'The relevance or otherwise of the central bank's balance sheet', *Journal of International Economics*, forthcoming.
[14] Others have derived similar results. For example, in Eggertson and Woodford (2003), an infinitely lived, representative household maximizes utility in a world with complete markets and faces no limit on borrowing against future income. It is clear that with these assumptions central bank purchases – which are essentially swaps of assets with the representative agent – can do nothing because that single representative agent owns the balance sheet of the central bank and such swaps do not change its net payoffs.

this could offset the reduced supply of bank credit. But if non-bank financial institutions only have limited capacity to hold long-dated debt, then increased issuance of corporate debt is likely to drive up yields on both long-dated government and corporate bonds, and increase corporates' funding cost. In this situation, the central bank may have a role to play in offsetting this increase in yields by purchasing government bonds, thereby creating space for other long-dated debt, such as newly issued corporate bonds, on non-bank financial institutions' balance sheet. This was the rationale behind the design of QE in the UK during the crisis.

In normal times, banks are active in both the markets for government bonds and corporate bonds as well as in lending. This means that many firms can choose between borrowing from banks and issuing corporate bonds. Now suppose that the Bank of England sold its holdings of government bonds. If the increased supply of government bonds to the private sector meant that non-bank financial institutions had less capacity to absorb newly issued corporate bonds, such that yields on corporate bonds increased, then firms could switch back to borrowing from banks. That would dampen any increase in gilt yields and corporate borrowing costs that the Bank's gilt sales might otherwise have. But this would not work if banks are crippled.

The theoretical literature suggests that purchases and sales of bonds by the central bank should have little or no impact on GDP when financial markets work well. It also makes intuitive sense that a decision from banks about the appropriate level of liquid assets – and in particular the split between them holding gilts, Treasury Bills, other liquid assets and reserves at central bank – should not have a first order impact upon the level of the policy rate the central bank should set so that – on average – it hits its targets. Woodford and Curdia (2010) reach the same conclusion when financial markets are working well.

These ideas suggest that the impact on monetary conditions of gilt sales by the Bank of England could be very different from the effect of its purchases. This is because sales would be conducted in a very different environment. The bulk of the Bank of England's asset purchases were made in 2009 and in 2011 – both occasions when financial markets were seriously disrupted because many financial intermediaries, in particular banks, were reluctant to take any risks because they found themselves short of capital and liquid assets. In contrast, the unwinding of such asset purchases is likely to occur when financial markets are operating more normally. Markets becoming dysfunctional was itself a tightening of monetary conditions, which QE offset.

I think the impact of the Bank of England expanding its balance sheet enormously in the wake of the financial crisis was very significant – and it helped prevent a far sharper decline in activity in the UK that would have had some highly persistent and negative impacts on future levels of GDP. Unwinding the asset purchases is likely to have a much smaller impact. And so I believe the overall operation – even if it is reversed fully (in the sense that The Bank sells all the gilts it bought to the private sector) – will have left the level of incomes significantly higher than they would have been.

But this appears to leave us with a puzzle: Most observers would agree that financial markets were working far better in the middle of 2013 than they were in 2009 or even in 2011 and 2012. Nevertheless, the US Federal Reserve Bank's – very balanced – statements about the possibility of reducing the pace of asset purchases appear to have caused a stir in financial markets and to have had an impact on asset prices that was powerful – perhaps as powerful as the estimated impact of asset purchase announcements made in earlier periods when financial markets were close to dysfunctional. How can one reconcile that with the view that asset purchases should have smaller impacts when financial markets are working better?

The answer, in my view, is that financial markets revised their expectations about the future path of short rates following the Fed's statements. Changes in perceptions about the pace of the Fed's asset purchases are unlikely to have caused the move in forward rates by themselves. Instead, both the changes in forward rates and in the expected pace of asset purchases are likely to have had a common cause: the belief that the Fed had become more optimistic about the economy.

Together, the theoretical literature and the Federal Reserve's recent experience suggest that the Bank of England may be able to unwind its asset purchases without much impact on interest rates when financial markets work well if two conditions are met: First, market participants should have a good idea about the likely future path of interest rates before the Bank announces its plans for reducing the stock of its assets. And second, financial markets should work well. If these two conditions are met, an announcement about the scale of asset sales would not change perceptions about future interest rates. Asset sales might somewhat increase long rates, but the increase is likely to be small when financial markets work well.

The bottom line is that I believe that the neutral level of Bank Rate is likely to be quite insensitive to decisions on when to sell government bonds the Bank of England acquired in its monetary ("QE") operations. That is because financial market conditions have gone a good deal of the way towards normalisation.

Conclusion

The new normal for monetary policy will probably involve setting Bank Rate on average at a lower level than before the crisis. One factor behind the recent sharp fall in real yields – changing perceptions of the level of risk in the world – is likely to be persistent. Households, firms, and investors now attach a higher probability to financial crises and sharp, prolonged downturns in economic activity: events that that many may have thought close to inconceivable. This makes assets which generate a real return with little risk more attractive, driving down the real risk-free interest rate. It also makes the wedge between safe rates and the rates of return required on riskier assets greater. Those forces will tend to reduce the neutral level of the policy rate set by the central bank. But the balance sheet of the Bank of England – and in particular the stock of gilts that it holds – is not likely to be a major factor in determining the long-run neutral level of Bank Rate.

References

Barro, R. (2006), 'Rare disasters and asset markets in the twentieth century', *Quarterly Journal of Economics* 121 (3), pages 823-866.

Bordo, M., Eichengreen, B., Klingebiel, D., Martinez-Peria, M. S., and Rose, A. (2001), 'Is the Crisis Problem Growing More Severe?', *Economic Policy* 16 (32), pages 53-82.

Chen, H., Curdia, V., and Ferrero, A. (2011), 'The Macroeconomic Effects of Large-scale Asset Purchase Programmes', *Economic Journal* 122 (564), pages F289–F315.

Eggertson, G. and Woodford, M. (2003): 'The Zero Bound on Interest Rate and Optimal Monetary Policy', *Brookings Papers on Economic Activity* 34 (1), pages 139-211.

Joyce. M., Tong, M. and Woods, R. (2011), 'The United Kingdom's quantitative easing policy: design, operation and impact', *Bank of England Quarterly Bulletin* Q3, pages 200-212.

King, M. and Low, D. (2014), 'Measuring the 'world' real interest rate', NBER Working Paper No. 19887.
Miles, D., Pillonca, V., and Baker, M. (2005), 'What Should Equities and Bonds Be Worth in a Risky World?', *Morgan Stanley Global Economics*, September 12.

Miles, D. and Schanz, J. (2014), 'The relevance or otherwise of the central bank's balance sheet', *Journal of International Economics*, forthcoming.

O'Donoghue, J., Goulding, L., and Allen, G. (2004), 'Consumer Price Inflation since 1750', Office for National Statistics *Economic Trends* 604, pages 38-46.

Wallace, N (1981), 'A Modigliani-Miller theorem for open-market operations', *American Economic Review* 71(3), pages 267-274.

Woodford, M. and Curdia, V. (2010), 'The central-bank balance sheet as an instrument of monetary policy, *Federal Reserve Bank of New York Staff Reports* 463.

Annex: The Extended Barro Model

The full model is described in detail in Barro (2006). Here we give the key elements.

Asset values depend on the discounted value of the expected payoffs. For 'safe' bonds the payoff is invariant to GDP and corporate income. For risky asset the payoff is the flow of corporate income which moves in line with GDP. The volatility of those risky asset payoffs then mirrors that of total income for the economy and depends on the standard deviation of the unexpected component of growth in normal times as well as the probability and scale of the two types of rare events. The discount rate on future cash flows depends on the rate of time preference and also contains an element that depends on expected growth and its volatility. That second element depends on the degree of risk aversion and the volatility of income flows of risky assets. Since the flow of corporate income is correlated with the flow of all income, in general investors find that asset risk is a risk to all their income, which is why they care a lot about it. The coefficient of risk aversion shows by how much expected returns on risky assets need to exceed a 'safe' rate for each unit of risk.

Investors in the economy are assumed to maximise a utility function that depends on consumption now (C_t) and in the future.

The utility function of the representative agent is:

$$u_t = E_t \sum_{i=0}^{\infty} [u(C_{t+i}) \cdot e^{-\rho i}] \qquad \text{(A1)}$$

Where ρ is the rate of time preference and E_t is the expectations operator. The link between the utility of consumption and the level of consumption is:

$$u(C) = (C^{1-\theta} - 1)/(1-\theta) \qquad \text{(A2)}$$

θ is the coefficient of relative risk aversion. Substituting (A2) into (A1) and differentiating with respect to C leads to the usual first order condition, the Euler equation:

$$u'(C_t) = e^{-\rho} E_t [u'(C_{t+1}) \cdot R_{t1}] \qquad \text{(A3)}$$

This links the marginal utility of consumption in different periods. R_{t1} is the realised return on an asset between t and t+1. We assume that consumption is equal to income and that risky assets generate a flow of cash returns that are a proportion of income.

Generalising the process that Barro used, I assume that total incomes (A), or GDP, follow a random walk with a drift and two distinct random components.

$$\ln(A_{t+1}) = \ln(\mathrm{A_t}) + \gamma + u_{t+1} + v_{t+1} \qquad \text{(A4)}$$

The parameter γ captures productivity growth. The first random component, u, shows the shock in normal times, i.e. it reflects the "normal" level of economic volatility. This shock follows a white noise, normally distributed process (i.i.d.):

$$u \sim N(0, \sigma^2) \qquad \text{(A5)}$$

The other random component ($\mathrm{v_t}$) is zero in normal times, but with given probabilities it takes on significant values. There is a small chance (probability $\mathrm{p_1}$) that v takes on a very large negative value, equal to –b. The parameter b represents the scale of the asymmetric shock; there is no chance of an equally large positive shock. There is a second type of shock, which is symmetric, and whose scale is denoted c. This shock has a higher probability of occurring (equal to $\mathrm{p_2}$), and it is smaller (i.e.the absolute value of b is larger than the absolute value of c). Thus;

$v_{t+1} = 0$ with probability $(1 - p_1 - p_2)$

$v_{t+1} = -b$ with probability p_1

$v_{t+1} = +c$ with probability $p_2/2$

$v_{t+1} = -c$ with probability $p_2/2$

The two random components u and v are assumed to be independent of each other.

We can calculate the moments of the distribution of the change in GDP from the six parameters $\gamma, \sigma, p_1, p_2, b,$ and c.

The mean (i.e. the first moment) is:

$$\mu = \gamma - p_1 b \qquad \text{(A6)}$$

The variance (the second moment) we denote s^2 where

$$s^2 = \sigma_u^2 + (1 - p_1 - p_2)(p_1 b)^2 + p_1[b(p_1 - 1)]^2 + \left(\frac{p_2}{2}\right)(p_1 b - c)^2 + \left(\frac{p_2}{2}\right)(c + p_1 b)^2 \qquad \text{(A7)}$$

Let:

$$\sigma_v^2 = (1-p_1-p_2)(p_1 b)^2 + p_1[(b(p_1-1)]^2 + \left(\frac{p_2}{2}\right)(p_1 b - c)^2 + \left(\frac{p_2}{2}\right)(c+p_1 b)^2 \quad \text{(A8)}$$

$$Skewness = \frac{1}{s^3}\left[(1-p_1-p_2)(p_1 b)^3 + p_1[(b(p_1-1)]^3 + \left(\frac{p_2}{2}\right)[(p_1 b - c)^3 + (c+p_1 b)^3]\right] \quad \text{(A9)}$$

The fourth moment is:

$$Kurtosis = \frac{1}{(\sigma_u^2+\sigma_v^2)^2}\left[(1-p_1-p_2)(p_1 b)^4 + p_1[(b(p_1-1)]^4 + \left(\frac{p_2}{2}\right)[(p_1 b - c)^4 + (c+p_1 b)^4]\right] + 3\left[\frac{\sigma_u^2}{(\sigma_u^2+\sigma_v^2)}\right]^2 + 6\left[\frac{\sigma_u^2\sigma_v^2}{(\sigma_u^2+\sigma_v^2)^2}\right] \quad \text{(A10)}$$

We choose the six parameters to roughly match the four moments based on the observations of historical annual real growth; but we also want to match as best we can the chances of extreme events, based on the frequency of big changes in GDP from a large group of countries over the past 200 years.

Once we have chosen six parameters we are able to solve for expected rates of return on an asset with a constant cash flow (when there is no default):

Replace $\mathrm{R_{t1}}$ in the Euler equation by the risk free real gross return $\mathrm{R_{t1}^f}$; and let $C_t = A_t$; $C_{t+1} = A_{t+1}$ gives:

$$R_{t1}^f = e^{\rho} A_t^{-\theta} / E_t[A_{t+1}^{-\theta}] \quad \text{(A11)}$$

Given the income dynamic in (A4) and the independence between u and v:

$$E_t\left[A_{t+1}^{-\theta}\right] = A_t^{-\theta} \cdot e^{-\gamma\theta} \cdot E_t\left[e^{u_{t+1}(-\theta)}\right] \cdot E_t\left[e^{v_{t+1}(-\theta)}\right]$$
$$= A_t^{-\theta} \cdot e^{-\gamma\theta} \cdot \left[e^{\frac{1}{2}\theta^2\sigma^2}\right] \cdot \left[(e^{-p_1-p_2}) + (1-e^{-p_1}) \cdot e^{b\theta} + \frac{1}{2}(e^{-p_1} - e^{-p_1-p_2}) \cdot \left(e^{c\theta} + e^{-c\theta}\right)\right] \quad \text{(A12)}$$

Where

p_1 is approximated by $(1-e^{-p_1})$

p_2 is approximated by $(e^{-p_1} - e^{-p_1-p_2})$

Therefore given (A11):

$$\ln\left(R_{t1}^{f}\right) = \left[\rho + \gamma\theta - \frac{1}{2}\theta^2\sigma^2\right] - \ln\left[\left(e^{-p_1-p_2}\right) + \left(1 - e^{-p_1}\right)\cdot e^{b\theta} + \frac{1}{2}\left(e^{-p_1} - e^{-p_1-p_2}\right)\cdot\left(e^{c\theta} + e^{-c\theta}\right)\right]$$

(A13)

This is the equation we use to construct the table in the text showing the impact upon the safe rate of changes in perceptions of extreme events.