



BANK OF ENGLAND

Speech

The distributional implications of low structural interest rates and some remarks about monetary policy trade-offs

Speech given by

Ben Broadbent, Deputy Governor Monetary Policy, Bank of England

At the Society of Business Economists Annual Conference

18 November 2016

I would like to thank my fellow MPC members and other Bank colleagues for their 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.

Good morning, and thank you for having me back.

This is a talk with two distinct halves.

When I was last in front of this audience, a couple of years ago, I talked about the causes of structurally low real interest rates. I also said a little about their potential distributional consequences.

I'll begin, shortly, with something quite unrelated – the interaction between monetary policy and the exchange rate. The focus in this first section will be on the trade-off between inflation and output stabilisation and how that affected the MPC's view of policy in its November Inflation Report.

But those earlier issues haven't gone away, it seems, and today I'll discuss again what low interest rates do (or do not do) to the distributions of income and wealth.

These are important topics and therefore intensively researched ones as well. I'll be making only a few, very high-level points. It's worth summarising them up front:

One is that a decline in structural interest rates needn't do much to the distribution of income, which is determined mainly by that of wages (rather than asset income). At least in the UK, summary measures of income inequality – in particular the so-called Gini coefficient – have been broadly unchanged since real interest rates began to decline a quarter of a century ago (Chart 1).

The evolution of the distribution of wealth is harder to document because the data aren't as good. Its interaction with the structural level of interest rates is also complex.


There are big disparities in the distribution of assets, in particular financial securities such as stocks and shares. In principle, therefore, declines in structural interest rates can increase the value of these disparities if they also lower the discount rate on (i.e. raise the price of) risky assets.

But “can” and “if” are the operative words. The fact is that equities have performed relatively poorly over much of this period. They rose strongly during the 1990s, as profits did the same and as longer-term interest rates started to fall. But they've since under-performed and their real price is currently no higher than it was twenty years ago. Nor is it any higher than before the financial crisis (Chart 2). Equities are a more important asset for richer than poorer households (they account for a higher share of their total wealth). So this under-performance¹ will have tended to compress the distribution of wealth, however that's measured.

¹ Equities have under-performed bonds; equity prices have also fallen relative to wages.


Prices of houses, which account for a much larger share of household wealth in aggregate, did a lot better over this period, particularly in the decade before the financial crisis. In real terms, and partly because of the fall in real interest rates, they more than doubled in the decade leading up to the financial crisis (Chart 2 again).

Chart 1: Income inequality broadly flat since real rates began to decline in early 1990s



Source: World Bank and Bank of England

Chart 2: Real equity and house prices no higher now than a decade ago




Source: Datastream and Bank of England

The distributional impact depends on which measure of inequality you use. In absolute terms, and also relative to household incomes, the jump in house prices widened the distribution of wealth. Some people own more housing than others; the rise in house prices increased the value of these differences (the red line in Chart 3).

The most significant effect, in this respect, was across generations: those that happened to enter the housing market before the mid-1990s gained relative to younger households (Chart 4).


Somewhat confusingly, however, you saw the opposite trend in proportionate measures of inequality, including the standard Gini coefficient (the blue line in Chart 3). That's because housing accounts for a greater share of wealth for those in the middle of the distribution than for those at the top end. So when real house prices went up the median household experienced faster proportionate rises in its wealth than those in the upper part of the distribution, even if the gains in pounds were smaller.

Chart 3: Increases in house prices benefit the median household



Source: British Household Panel Survey (Institute for Social and Economic Research) and Wealth and Asset Survey (ONS); dotted line is for years for which the WAS doesn't have matching income data.

Chart 4: Age profile of wealth steeped in decade to 2005, broadly unchanged since



Source: British Household Panel Survey (Institute for Social and Economic Research) and Wealth and Asset Survey (ONS); * adjusted to correspond to WAS using overlapping periods.

It takes a little thought to get one's head around these differing trends. But the important point, I think, is that the significant movements happened quite some time ago. Because after that big jump, either side of the millennium, real house prices in the UK stopped rising. They fell after the financial crisis and have recovered since. But, like those of equities, they're no higher now than they were a decade ago. And the result is that, whichever measure you use, the distribution of wealth has been broadly flat over that period.

None of this is to deny that, at the margin, discretionary monetary policy can have distributional effects, at least for a period of time². There are also plenty of other aspects of these data one might be concerned about. Conceptually, for example, it's not clear that a pure rise in house prices strictly qualifies as extra wealth, since the only way to realise it is to sell one's house. Equally, there are arguments for including in wealth the capitalised value of unfunded pensions (if so they'd also count as a liability for those that fund them, largely future taxpayers). And whatever the exact definitions involved these summary measures of inequality only get one so far. There are other cuts of the distribution, and large movements within it over time, that may matter.

But the simple Gini coefficients do tell a story – or rather, perhaps, the absence of one. Structural interest rates don't do much to the distribution of income, most of which reflects the spread of wages. They can affect the distribution of wealth, but only via prices of risky assets. That's what happened in the housing market, prior to the crisis, and the jump in prices had significant inter-generational effects. But during the period since – which included QE – real asset prices have been flat, as have the distributions of wealth and income.


² My colleague Jon Cunliffe made the same point earlier this week (Cunliffe (2016)).

Monetary policy and the exchange rate

Anyhow, I'd better keep at least some of that back for later. Let me start, as promised, with some remarks about sterling.

Exchange rates are generally volatile, particularly in relatively small, open economies like this one. With the UK's (pre-1992) history of high and variable inflation, that was even more the case for sterling's exchange rate in the period before inflation targeting.

Chart 5: Sterling depreciation on par with IMF and 2008 financial crisis



Source: Hills, Thomas, Dimsdale (2010) and Bank of England

If you focus specifically on changes in commodity prices, this “passthrough” to the retail level looks to be fairly rapid. Petrol prices respond pretty quickly to changes in the price of oil, in both directions, and their direct impact on the annual rate of inflation therefore peters out after only a year or so. Given the time it takes for monetary policy to work – the peak impact on inflation is generally reckoned to take 18-24 months to come through – it’s not clear it makes sense to respond to changes in oil prices unless they have knock-on, “second-round” effects on underlying inflation (wage inflation in particular).

One often hears the same argument about changes in the exchange rate: monetary policy should “look through” big moves in sterling because their effects on inflation come and go before policy has a chance to do much about them.

There’s some truth to that – the impact on inflation does eventually dissipate. That’s partly why inflation was already set to rise later this year, and into 2017, even before the recent depreciation. Before that happened, sterling had experienced a significant (albeit more drawn out) appreciation. Together with the precipitous drop in commodity prices, this pushed CPI inflation well below the 2% target over the past couple of years.

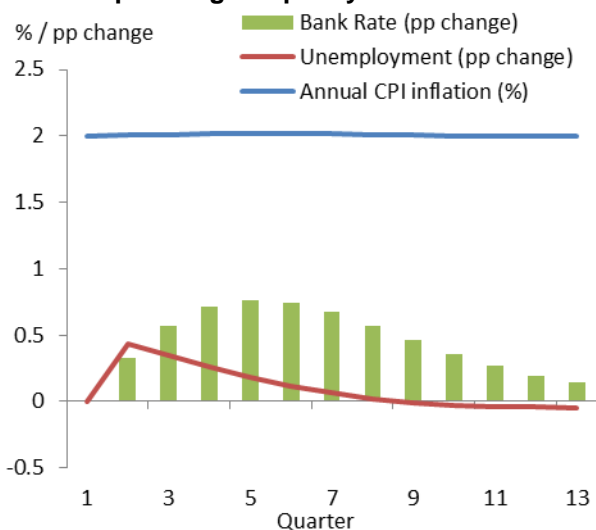
But even by those standards the moves over the past year have been significant. The 11% fall against the US dollar, in the two days immediately after the EU referendum, was the sharpest since the devaluation of 1967. The 20% drop in the trade-weighted index, in the year to the November Inflation Report, is on a par with what happened at the time of the IMF crisis and the 2008 financial crisis (Chart 5).

Over time, this is likely to put significant upward pressure on import prices and, because around 30% of consumption is imported, consumer prices as well.

But those things were only ever going to affect inflation for a period of time. In its May Inflation Report, which was conditioned on a “remain” vote (and therefore a significantly stronger exchange rate), the MPC projected a significant pick-up in inflation through to the end of 2017, towards the 2% target, even without any fall in the currency. Not every economic event is due to “Brexit”.

But the decline in the currency is clearly going to add to that pick-up. And the point I was getting to is this: unlike those in the price of oil specifically, changes in the general level of sterling import prices seem to take quite a long time to feed through to the retail level – long enough, in principle, for monetary policy to do something about them. The evidence suggests that exchange rate changes have a direct bearing on inflation at horizons two to three years ahead. So it’s not clear policy is impotent in the face of such events; nor is it clear, even in the absence of “second-round” effects on inflation expectations, that policy makers should be indifferent to them.

Chart 6: Offsetting impact of pure exchange rate shock requires tighter policy



Source: Bank of England calculations

So what would it actually take, after a significant depreciation, to prevent inflation from rising?

Chart 6 gives you an idea. Using one of the Bank’s economic models³, it plots the simulated paths of interest rates and unemployment required to offset in full the impact of a 10% exchange rate depreciation on inflation. You can see that the interest rate would have to rise by around ½% point. This dominates any expansionary impact of the lower exchange rate on demand and unemployment too goes up by close to ½%. As the impact of the depreciation begins to top out, a year later, interest rates and unemployment start to fall back.

So it’s technically feasible, at least according to this model, to keep inflation pretty close to target after a significant depreciation in the exchange rate. It’s just rather costly in terms of unemployment.

This is not what the MPC is meant to do. The inflation target prevails “at all times”. But our remit also recognises that there are instances in which “attempts to keep inflation at the inflation target...may cause undesirable volatility in output due to the short-term trade-offs involved, and the MPC may therefore wish to

³ I’m grateful to the Bank’s Monetary Analysis Directorate for its help in examining optimal policy simulations where the policymaker chooses their preferred interest rate each period (ie, under discretion) to minimise a loss function defined in terms of deviations of annual CPI inflation from target. The optimal policy solution under discretion is calculated using the algorithm of Dennis (2007). The model used for the simulation is a version of the Bank’s COMPASS model, set out in Burgess et al (2013). It has been augmented with a version of Okun’s Law, which assumes a linear relationship between the output gap and the deviation of unemployment from its equilibrium rate.

allow inflation to deviate from the target temporarily”. It then goes further: “in exceptional circumstances the MPC is likely to be faced with more significant trade-offs between the speed with which it aims to bring inflation back to the target and the consideration that should be placed on the variability of output”.

This flexibility isn’t unique to the UK. All inflation targeting regimes allow (or rather demand) that policy should pay some heed to real as well as nominal stabilisation. As the economist Lars Svensson puts it⁴, “inflation targeting is never strict but always flexible...inflation-targeting central banks aim at stabilizing inflation around the inflation target but also put some weight on stabilizing the real economy”.

Sometimes that distinction doesn’t amount to much. A negative shock to aggregate demand, for example, will depress output and subsequently inflation as well. It pushes them both in the same direction and the appropriate response is to ease monetary policy whatever weight one puts on the two objectives. At other times, however – after a shock to the economy’s supply potential, for example – the two can come into conflict and policy has to “trade off” one against the other. Charts 7a and 7b give a (highly stylised) representation of this difference.

Now it’s not clear that, on its own, a drop in the exchange rate would produce such a conflict. Depending on the response of wages, higher import prices are likely to dampen real wage growth, and with it the growth of consumer spending. Set against that, however, a weaker exchange rate normally boosts net export demand, by enough to outweigh the negative effect on consumption. At the margin, it should therefore be net positive for demand. If so – and if the depreciation comes entirely out of the blue, with nothing else going on in the economy – the correct response might still be to tighten policy somewhat, even if not to the same extent as in Chart 6.

But this is where another complication comes into play: what has caused the depreciation in the first place? Skittish they may be, but a big change in financial market prices is rarely unprovoked. There’s usually some sort of reason. And whatever the underlying cause, that will have its own effect on inflation to take into account. As the last Governor Mervyn King used to point out, the overall impact of changes in asset prices depends critically on “the source of the shock”.

⁴ Svensson (2008)

Chart 7a: Aggregate demand shock pushes objectives in same direction

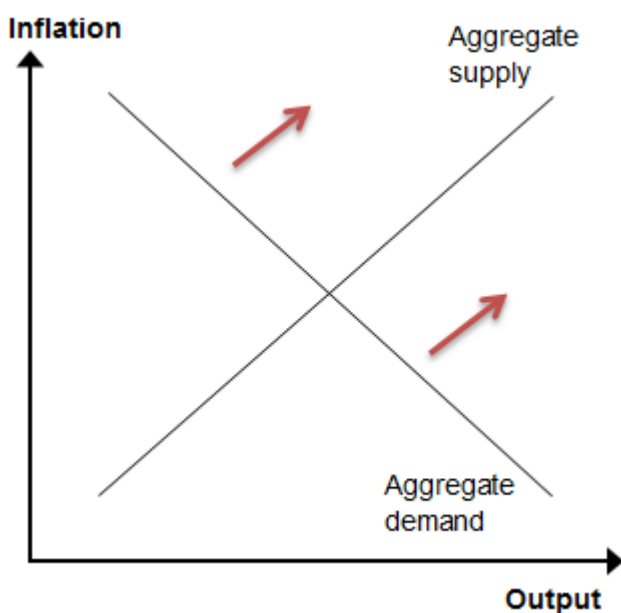



Chart 7b: Shock to supply pushes them in opposite directions, creating a trade-off



In principle, any number of things can cause exchange rates to fall. An expected drop in productivity in the tradables sector; unfavourable news about the terms of trade (the risk of future export tariffs, for example); higher risk premia for domestic assets; or a simple decline in domestic demand. These are all things that would weaken the fair value of the currency. (Nor is the list restricted to domestic factors: by symmetry the same shocks among one's trading partners would also shift the exchange rate in the opposite direction.)


It's hard to say which has been uppermost in the minds of global investors over the past year. The timing of the depreciation, in its various stages, indicates pretty clearly a link with the outcome of the referendum. The view in the market is that, if only for a period of time, exiting from the EU is likely to inhibit the UK's openness to trade. That could affect the currency through any or all of these channels.

And it does matter a bit which is at work. Suppose, for example, that the exchange rate falls purely because the market expects a hit to the level of UK's productivity, and that this does actually occur (it goes without saying that the currency market, which attempts to price the future, could turn out to be wrong). In that case unemployment wouldn't necessarily rise and inflation would end up higher than in the simulation in Chart 6, pushed up in the main by domestic factors, rather than the currency. Even a flexible inflation targeter would probably have to tighten policy.

If, on the other hand, it's purely the prospect of weaker demand that's driven a depreciation, with no impact on supply, the outcome would be significantly less inflationary. Unemployment would tend to rise, domestically generated inflation to fall and, notwithstanding the impact of higher import prices, the right response would be to ease policy aggressively (Charts 8).

As I say, the forecasts made by the currency market could well turn out to be wrong. In some sense this happens every minute of the day – otherwise the exchange rate wouldn't ever budge.

Chart 8: When depreciation reflects only hit to domestic demand, looser policy required



Source: Bank of England calculations

The MPC has to make its own judgements. Since the referendum result we have made various adjustments to our forecasts. The main thing, recognising the heightened uncertainty about this process, has been to lower our projections for business investment.

This weighs somewhat on supply. Productivity is still expected to accelerate – in the central forecast annual growth is around a percentage point faster than it has been since the crisis. But the extent of that projected acceleration has been toned down.

The more sizeable effects, however, are on demand. Coupled with the impact of lower real wage growth on consumption, this would mean slightly higher unemployment. And it's this projected rise in the degree of spare capacity in the economy that has to be balanced against the rise in inflation. As the latest *Inflation Report* puts it "the limits to which above-target inflation can be tolerated...depend on the scale of the shortfall in economic activity below potential". If we hadn't judged that demand growth was likely to soften – that the slide in the currency partly reflects that prospect – the yield curve on which the forecast was conditioned would probably have looked too shallow.

As it was, we judged the projected trade-off between the two to be broadly in the right place. We can tolerate high inflation because the alternative is a larger rise in unemployment, and weaker wage growth; put another way, we expect to have to accept those things (to some extent) in order to prevent inflation from rising further above the target. That's the unavoidable nature of the trade-off that monetary policy occasionally faces.

Low structural interest rates and the distributions of income and wealth

So much for monetary policy. Let me now turn to low structural interest rates.

I think it's right to distinguish the two. In this same venue, a couple of years ago, I explained why the low level of real interest rates around the world had more to do with real forces – things that raise the desire to save, and/or the hurdle for new investment – than with the whim of monetary policy makers. For the most

part central banks have been forced to accommodate these trends; they have not determined them. Because with open capital markets real interest rates are determined by global forces, the same applies that much more strongly to any individual central bank, particularly in a small open economy. It's not clear that anything in the UK – whether monetary policy or anything else – could have prevented the long decline in real interest rates over the past twenty five years⁵.

These points have been recognised by many economists over many years, and I won't re-hash them today. Besides, your next and very distinguished speaker will address the very same issue. What I do want to do, however, is to make some very broad points about the effect of lower structural real interest rates on inequality.

I emphasise again these are pretty high-level remarks. There is an enormous economic literature on inequality, including its interaction with returns on capital (Piketty (2013) and the many responses to his work are only one strand of that). But it may be worthwhile to go over some basic points.

Little sign of any impact on distribution of income

The first is that lower real interest rates needn't have much bearing on the distribution of income. It's certainly striking, in this respect, that conventional measures of income inequality in the UK have been broadly stable over the past quarter century, precisely the period over which real interest rates have been declining.

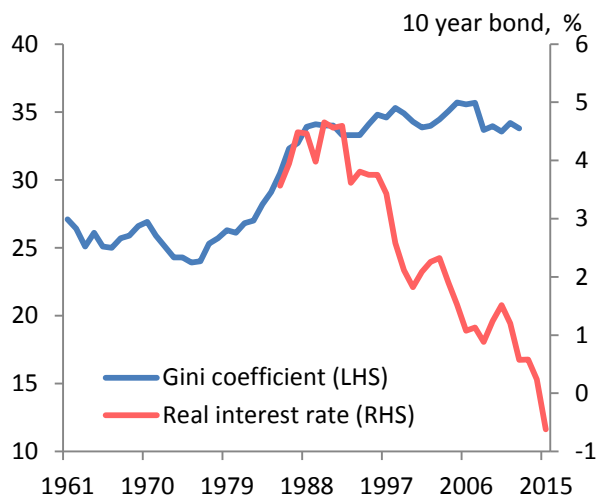
I say broadly: the standard Gini coefficient (the blue line in Chart 1) has certainly moved around over that period, partly for cyclical reasons⁶. It's well documented that employment is more cyclical for the less well paid. So income inequality tends to rise in recessions and fall back in expansions. This helps to explain the drop in the UK's Gini coefficient for income in the past 3-4 years. It's also true a single summary measure can only take you so far. Many researchers have also paid particular attention to the rise in incomes at the very top of the distribution. The income share of the top 1%, and the ratio of the top 10% relative to the bottom 10%, are plotted in Chart 9. There's an upward trend in both from the mid-1970s to end of the 1980s; the share of the top 1% carried on rising through most of the 1990s as well. But both have stabilised since (the 90:10 ratio has actually fallen back a bit). And the most commonly used summary measure of income inequality, the Gini coefficient, has been pretty much flat since the early 1990s.

⁵ Long-term real interest rates aren't identical across countries, even those that share an open capital market, as investors may demand differing risk premia, even on "safe" assets like government bonds. Prior to inflation targeting, for example, real sterling interest rates looked to be markedly higher than in most other developed countries. This may have reflected the volatility of inflation and the higher risk premium this induced. But large movements in long-term real interest have clearly have been common to all developed economies.

⁶ The Gini coefficient measures the distance from uniformity of a statistical distribution.


If this indicates that income inequality is relatively insensitive to changes in structural interest rates – and it does constitute something of an experiment⁷ – I’m not sure we should be terribly surprised.

Chart 1 (again): Income inequality broadly flat since real rates began to decline in 1990s



Source: World Bank and Bank of England

Chart 9: Share of top 1% rose through 90s, since stable; like Gini, 90:10 ratio stable since early 90s



Source: IFS calculations using the Family Resources Survey, 2014–15 (Belfield et. al. 2016)⁸

For one thing, most income accrues to labour and inequality of income reflects mainly the spread of wages. Chart 10 plots the distribution by broad type of income. Government benefits go disproportionately to the less well off, helping to flatten the distribution. Asset income, mostly via private pensions, goes mostly to the better off. But the bulk of household income, and therefore the bulk of its variation, comes from wages. And there’s no reason to think the distribution of labour income specifically should be affected by the long-term level of real interest rates.


Second, while capital income is clearly skewed to the better off, it hasn’t risen any faster than wages over the period as a whole. The blue line in Chart 11 plots corporate profits generated in the UK relative to aggregate wage income. Some of these belong to foreign-owned companies and their shareholders; UK households receive income on foreign assets (again, mainly via pension funds and insurance companies). So the red line in Chart 10, profits plus net overseas income, is a better measure of the capital income of UK residents. It went up strongly in the 1980s and 1990s. But it’s since fallen back and, relative to labour income, is no higher now than it was thirty-five years ago. It’s quite a bit lower than in the early 1990s, when real interest rates started to fall⁹.

⁷ Some economists argue that any causation could run the other way around – that rising inequality has contributed to global savings demand and thereby the decline in structural interest rates. If the UK provides any sort of test of this relationship it can only be from interest rates to inequality as the UK is too small to make much difference to global saving trends.

⁸ I am grateful to the authors for sharing with me the data behind Charts 9 and 10.


⁹ This sits rather uneasily with the hypothesis of Piketty (2013), who suggested that inequality had risen – and that it would rise further in an environment of slower growth – because the share of capital income was also going up. Acemoglu and Robinson (2014) present a more thorough critique.

Chart 10 Most income variation from wages



Source: IFS calculations using the Family Resources Survey, 2014–15 (Belfield et. al. 2016). All income sources are measured net of any taxes paid on them

Chart 11 No upward trend in share of national income going to capital



Source: ONS

These two facts – that most income comes from wages and that, while skewed to the better off, capital income hasn't risen any faster than wages – can help explain why the decline in real interest rates appears not to have done much.

This, of course, is just the UK experience. In many other developed countries – notably the US, but also Sweden and Germany – income inequality has continued to rise over the past twenty-five years. That's not a universal trend, however (in France and Portugal, for example, income inequality has fallen). Across the world as a whole strong economic growth in poorer countries has meant that inequality at a global level has declined¹⁰. And even in countries where it has risen, researchers have focused not on any role for interest rates but on the nature of technical progress and why that appears to have complemented higher-skilled jobs more than those in the lower-to-middle part of the distribution¹¹.

That's understandable: as in the UK, most income (and therefore most income inequality) comes from wages.

Why the distribution of wealth hasn't widened in recent years

Interest rates may be a side-show when it comes to the distribution of income, but perhaps people's concerns relate more to the effects on wealth.

¹⁰ Bourguignon and Morrisson (2002)

¹¹ See, for example, Autor (2006), Violante (2008) and Card and DiNardo (2002)

In principle, one can understand why. Capital income in the UK mightn't be any higher than it was a generation ago, relative to wages. It might not grow any faster over the future. But if the prices of financial assets represent the present value of that income – something (very roughly) like this:

$$\text{asset price} = (\text{expected capital income})/(\text{discount rate})$$

– then one might well have expected lower interest rates to boost all asset prices, simply by lowering the rate at which that future capital income is discounted. And because the ownership of assets is highly uneven (to a greater degree than income) this in turn would have widened the dispersion of wealth.

But there's no inevitability about the effect of lower interest rates on prices of risky assets. The discount rate for an asset whose income is uncertain depends not just on risk-free interest rates but on the risk premium, which can also move around. And unlike government bonds, which offer a fixed and certain coupon, the income generated by such assets (and therefore their expected income as well) can also vary significantly over time. A slightly richer version of the relationship would be this:

$$\text{asset price} = (\text{capital income})/(\text{risk-free real interest rate} + \text{risk premium} - \text{expected growth rate of income})$$


I've re-produced Chart 2 below. It plots what's actually happened to the real prices of risky UK assets (equities and houses) over the past twenty-five years. As you can see, equities performed strongly during the early part of that period. Profits were strong, long-term real interest rates were falling and, it appears, the risk premium and expected future growth in profits were relatively stable.

But equity prices have under-performed since then. Profits have disappointed and, in some combination, the risk premium has gone up and expected future profit growth has fallen.

When I was last here I suggested that concerns about downside economic risks might help to explain both why risk premia had risen and why the risk-free rate of interest had fallen. But whatever the explanation, the plain fact is that, despite the downward trend in real bond yields, real equity prices are now where they were in 1997.


Chart 2 (again): Equity and house prices no higher now than a decade ago

Source: Datastream and Bank of England



This is far from unprecedented. Chart 12 plots estimates of the 20-year growth in real equity prices going back 150 years or so. There've been several protracted periods over which real UK equity prices have fallen, notably after the world wars and also in the 1960s and 1970s. And perhaps, after the sharp run-up over the following two decades, during the 1980s and 1990s, a period of stabilisation (or worse) was warranted. I'm certainly making no comment about whether shares are currently over or under-valued. I'm simply pointing out that the relationship with bond yields isn't precise and that, in the past 15-20 years, the one appears not to have benefited the other.

Chart 12: Periods of protracted falls in real UK equity prices



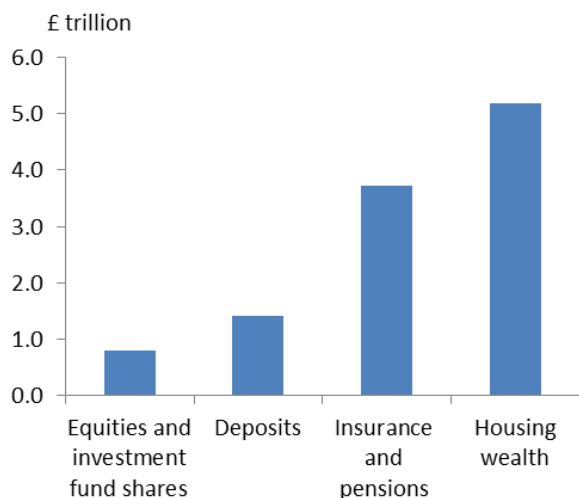
Source: Hills, Thomas, Dimsdale (2010), Bank of England and ONS

House prices were a different story, at least in the years leading up to the financial crisis. Having stagnated for the first half of the 1990s they all but tripled in the ten years thereafter. And although tightening supply had something to do with that, there's no doubt that falling long-term interest rates also played their part.

But that's not true of the years since the crisis. Real house prices fell back quite sharply in 2009 and 2010 and, while they've since recovered, they've not regained those earlier peaks. Over the last decade as a whole real house prices have been flat.

What has all this done to the distribution of wealth? It's not easy to tell because the data aren't that good. We have some idea of the distribution of freely held assets from the quinquennial British Household Panel Survey, first conducted in 1991. But this doesn't include the value of claims on future pension benefits, whether funded or unfunded. The ONS's Wealth and Asset Survey is more thorough, and does at least cover funded pension schemes, but this only began in 2006.

Chart 13: Housing the largest component of UK household wealth



Source: ONS

Another is that a good part of the dispersion of housing wealth relates to age. Currently, for example, the ratio of gross housing wealth to income for those under the age of 40 is 1.7; for people older than 40 it's 4 (the respective ratios for net housing wealth, deducting mortgages, are 0.2 and 3). The dispersion of housing wealth within age cohorts is notably smaller than across the population as a whole (0.58 versus 0.65)¹².

Third, as well as the spread in aggregate wealth there are noticeable differences in its composition. As a proportion of overall wealth, housing is most important in the middle of the distribution. Fewer at the lower end own their own houses. The ownership of financial assets is skewed towards wealthier households. So while they own more housing in absolute terms than those in the middle of the distribution, it accounts for a lower proportion of their wealth.

Armed with these facts what might one expect to have resulted from the trends in asset prices over the past twenty years?

First, that the most important movements were probably determined by house prices. Equities matter for richer households and for those with funded pensions. The under-performance of equity prices is therefore likely to have compressed the distribution of wealth somewhat, however that's measured. But the dominant asset class is housing, which is more evenly held.

Second, the big jump in house prices either side of the millennium is likely to have widened the distribution of wealth in absolute terms, particularly across age groups. But because it accounts for a greater share of their

¹² See Crawford et al. (2016)

At any rate, it's useful first to go through a couple of preliminary points.

One is to note the importance of housing in aggregate. The housing stock is currently worth close to £5tn, freely held financial assets, excluding bank deposits, only £0.8tn (Chart 13). Households have a much greater exposure to financial securities via their claims on insurance companies and funded pension schemes. But even allowing for that their houses are worth more.


wealth, relative to those at the upper end, you might well expect rises in house prices to compress the distribution when you measure that in proportionate terms, as the standard Gini coefficient does. If prices of houses grow faster than those of financial assets the median household should experience faster growth in wealth than those at the upper end.

Third, it's not clear the distribution of wealth will have changed much in the last decade, given the flatness of real asset prices, housing as well as shares.

By and large, and allowing for the slightly patchy data, from two different sources, this is exactly what you do see.


Chart 3 plots two summary measures of the dispersion of household wealth, one relative to incomes the other a purely proportionate measure (red and blue respectively). From the mid-nineties to the mid-noughties the absolute measure went up. The distribution of wealth-to-income ratios widened. Those that happened to buy houses before that big surge period gained relative to younger people. As Chart 4 demonstrates, the age profile of wealth steepened considerably over that period.

Chart 3 (again): Increases in house prices benefit the median household



Source: British Household Panel Survey (Institute for Social and Economic Research) and Wealth and Asset Survey (ONS); dotted line is for years for which the WAS doesn't have matching income data.

Chart 4 (again): Age profile of wealth steepened in decade to 2005, broadly unchanged since



Source: British Household Panel Survey (Institute for Social and Economic Research) and Wealth and Asset Survey (ONS); * adjusted to correspond to WAS using overlapping periods.

But the standard (proportionate) Gini coefficient declined. And – broadly speaking – not much has happened to either distribution, or to the age profile of wealth, since then. The absolute inequality measure has edged down (though it's possible this is an artifice of using two slightly different datasets). The proportionate

measure has been flat. Apart from a rise in wealth for over-65s the age profile is also pretty much the same as it was a decade ago.

As I say, the household-level wealth data aren't that good in the UK (certainly less good than those for the distribution of income). There are also several conceptual questions about what really counts as wealth. For example, there's reason to doubt whether higher house prices can be counted as extra "wealth" in the conventional sense. The only way you can realise such a gain is by moving into a less valuable house, i.e. by cutting your consumption of housing services. So if the strict definition of "a rise in wealth" is something that can fund extra consumption across the board, including that of housing, it's not clear a pure rise in the price of housing qualifies¹³. Besides, some of these intergenerational inequalities will in time get unwound via bequests.

Equally, these surveys fail to cover something that arguably should be counted as wealth – at least for those entitled to them – namely unfunded pensions. Their capitalised value is significant. In 2014/15 the present value of pensions paid to current and former government employees was estimated to be £1.5tn¹⁴; in all likelihood the present value of the universal state pension is more than that. This isn't true, "outside" wealth – there's an offsetting (off-balance sheet) liability for future taxpayers. Taking account of both would almost certainly flatten the true distribution of wealth.

But those are, in the current context, relatively minor points. The main things I want to get across here are (i) that there's no reason to expect long-term interest rates to affect the distribution of income and (ii) the interaction between asset prices and the distribution of wealth – and it's not an entirely straightforward interaction – depends not on real interest rates per se but the price of riskier assets – shares and above all housing.

The evidence of the past decade bears this out. A further decline in bond yields there may have been. Monetary policy, conventional and unconventional, may or may not have had some role to play in that and, at the margin, QE may therefore have supported risky asset prices as well. After all, that was one of its primary purposes, or at least an important transmission channel. In doing so, it would also have helped to prevent the economy from weakening and unemployment from going up – which, by the by, would have been regressive for incomes.

But to the extent it did support assets it was only offsetting other things. Real asset prices have basically been flat over the past ten years. And, perhaps unsurprisingly, so have the distributions of wealth and income. Indeed the main summary measure of income inequality in the UK has been broadly unchanged for the past twenty-five years, precisely the time at which real interest rates began to fall.

¹³ The box on page 18-19 of the latest Inflation Report sets out the evidence against the idea that the empirical correlation between house prices and consumption reflects a "wealth effect". For my part I find that evidence quite persuasive. See also, for example, Buiter (2008).

¹⁴ H.M.Treasury (2016)

Some concluding remarks

This has been a very long speech. For those that have got this far, I congratulate you. Dedication, indeed.

But after trailing through all these distributional statistics, you might well be asking what any of this has to do with monetary policy.

That's a good question.

The first part of the talk certainly did concern monetary policy. In fact, judging the occasional trade-offs that occur between real and nominal stabilisation is absolutely at the core of what monetary policy, or at least inflation targeting, is about. My aim here was to add some colour to what the MPC said in the last Inflation Report, namely that the tolerance of inflation overshoots depends in part on "the scale of the shortfall in economic activity below potential". The MPC would be less inclined to accommodate above-target inflation if it didn't also expect demand to fall slightly short of that supply potential; it would be able to do more prevent any rise in unemployment if it weren't for the inflationary pressure brought about by the fall in the currency.

When it comes to these distributional matters, however, I see much less of a link to monetary policy. That's because I doubt that any independent decision of monetary authorities, the MPC included, has that much bearing on the behaviour of real asset prices over long periods of time¹⁵, or any distributional consequences that follow.

And the only reason for raising these issues, therefore, is the apparent concern that the opposite is true – that looser monetary policy, unconventional policy in particular, is having material and lasting effects on the distributions of income and wealth.

I don't think the evidence gives much support to that view.

Thank you.

¹⁵ There might have been a second-order effect, during the 1990s, from the introduction of inflation targeting itself. To the extent that led investors to expect a more stable macro-economy, you'd've expected it to contribute to a flatter yield curve, a stronger currency and lower risk premia in housing and equity markets (see also footnote 4). Perhaps that was one factor in the growth of UK asset prices during that decade. But in an open capital market, trends in longer-term real interest rates are more likely to be caused by global, and structural factors.

References

- Acemoglu D and J. Robinson**, 2014, "The rise and decline of general laws of capitalism", *NBER Working Paper*, No. 20766.
- Autor, D., L. Katz and M. Kearney**, 2006, "The polarisation of the U.S. labour market", *NBER Working Paper* no. 11986.
- Bastagli, F and J. Hills**, 2012, "Wealth accumulation in Great Britain 1995-2005: The role of house prices and the life cycle", *Centre for Analysis of Social Exclusion, London School of Economics*.
- Belfield C., J. Cribb, A. Hood, R. Joyce**, 2016, "Living Standards, Poverty and Inequality in the UK: 2016", *The Institute for Fiscal Studies*.
- Bourguignon, F. and C. Morrisson**, 2002, "Inequality Among World Citizens: 1820-1992", *American Economic Review* Volume 92 Issue 4.
- Buiter, W. H.**, 2008, "Housing Wealth Isn't Wealth", *NBER Working Paper* No. 14204.
- Burgess S, E. Fernandez-Corugedo, C. Groth, R. Harrison, F. Monti, K. Theodoridis, and M. Waldron**, 2013, "The Bank of England's forecasting platform: COMPASS, MAPS, EASE and the suite of models," *Bank of England Working Paper* No. 471.
- Card D and J. E. DiNardo**, 2002, "Skill-Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles", *Journal of Labour Economics*, vol 20, no.4.
- Crawford R, D. Innes and C. O'Dea**, 2016, "Household Wealth in Great Britain: Distribution, Composition and Changes 2006–12", *The Institute for Fiscal Studies*, Vol. 37, Issue 1, Pages 35–54.
- Cunliffe, J.**, 2016, "Why are interest rates low?", speech given at Manchester University, <http://www.bankofengland.co.uk/publications/Documents/speeches/2016/speech935.pdf>.
- Dennis R**, 2007, "Optimal policy in rational expectations models: new solution algorithms," *Macroeconomic Dynamics*, *Cambridge University Press*, 11(37), 31-55.
- Hills S., R. Thomas and N. Dimsdale**, 2010, "The UK recession in context - what do three centuries of data tell us?", *Bank of England Quarterly Bulletin*, 2010 Q4.

HM Treasury, Whole of Government Accounts for year ended 31 March 2015 (2016),
<https://www.gov.uk/government/publications/whole-of-government-accounts-2014-to-2015>

Piketty T., 2013, “Capital in the Twenty-First Century”, *Harvard University Press*.

Svenson L., 2008, “Inflation Targeting”, *The New Palgrave Dictionary of Economics, Second Edition*.

Violante G., 2008, “Skill-Biased Technological Change”, *The New Palgrave Dictionary of Economics, 2nd Edition*.