



BANK OF ENGLAND

Speech

Lambda

Speech given by

Mark Carney, Governor of the Bank of England

London School of Economics

Monday 16 January 2017

I am grateful to Ben Nelson and Iain de Weymarn for their assistance in preparing these remarks, and to Alex Haberis, Clare Macallan, Michael Mcleay, Kate Reinhold, and Matt Waldron for background research and analysis.

Good evening. It is a pleasure to be at the LSE and an honour to share the podium with Amartya Sen.

Professor Sen is rightly recognised for his many contributions, not least to welfare economics and social choice theory. He has posed, and in many cases answered, some of the most fundamental questions facing economists. For example, given the diversity of people's preferences, is it possible to arrive at coherent aggregative judgements about how society is arranged? If there are as many preferences as there are people, is reasonable social choice possible at all? And how do we value public goods?

His work underscores the value of empirics – informational broadening – to help make the interpersonal comparisons necessary to understand, and act upon, the force of public concerns about poverty, inequality, even tyranny. And his insights have been applied to the most pressing economic and ethical questions, such as the prevention of famines.

His insights are also relevant to social choices about macroeconomic stabilisation, including inflation control, and what society is prepared to do to achieve it.

What level of inflation does society wish to achieve? How aggressively should inflation stabilisation be pursued when doing so imposes costs in terms of lost output and higher unemployment?

Low, stable and predictable inflation is a public good. It is not merely that rising prices mean households have to shop around or businesses have to update their prices periodically. High inflation hurts those, particularly the worst off in society, who don't hold equities or property as well as those whose incomes are fixed in nominal terms. It distorts price signals, inhibits investment, and can ultimately damage the productive potential of the economy.

Equally, deflation can imperil growth and employment. In a highly indebted economy, deflation raises real interest rates, increases debt burdens, lowers wages, and reduces growth. In the extreme, these can morph into debt deflation, causing very high and persistent unemployment and financial collapse.

The happy medium is low, stable, predictable inflation over the medium term. A little inflation 'greases the wheels' of the economy, for example by helping real wages adjust more smoothly.¹ Moreover, a positive inflation rate gives monetary policy space to deliver better outcomes for jobs and growth when shocks hit, without the distortionary costs of high and volatile inflation.²

Recognising the social value of inflation control is one thing, delivering it is quite another.

In the past, many societies could not. This is because the instrument that affects inflation most powerfully – monetary policy – also affects output and employment, at least in the short run. That influence tempted

¹ See Tobin (1972), "Inflation and unemployment", *American Economic Review*, 62.

² Inflation can have large redistributive effects. Volatile inflation arbitrarily re-distributes wealth across borrowers and savers. And most recently, there has been a debate about how the stance of monetary policy has affected the distribution of wealth and income. In a series of recent analysis and speeches, members of the MPC have shown that, while distribution is not an objective of monetary policy, the evidence strongly suggests that policy since the crisis has been associated with better outcomes.

authorities to promise low inflation in the future, but then to renege in order to boost activity. Electoral cycles reinforced this predisposition.³ Firms and households began to anticipate these incentives, however, and eventually pre-empt them. The economy ended up in a worse equilibrium with higher inflation and unemployment.⁴ Such time-inconsistent policies contributed to excessive inflation and higher structural unemployment in the UK during the 1970s and 1980s (when inflation averaged 9.5% and unemployment over 7.5%).

In light of similar experiences, many societies came to recognise that macro outcomes can be improved by having society first choose the preferred rate of inflation and then delegate operational responsibility to the monetary authority to take the necessary monetary actions to achieve that objective.⁵ By “tying the hands” of authorities, time inconsistency is resolved and better outcomes for both inflation and unemployment become possible.

Inflation targeting, as practised in the UK, represents the most comprehensive adoption of these insights. By delegating monetary policy to an operationally independent central bank, inflation control becomes a more technical, ‘engineering’ problem, as described in Professor Sen’s work on ethics and economics.⁶ The committee responsible is given a clear remit – with a lexicographic preference for inflation control – and is charged to do what is necessary to achieve the inflation target over the policy horizon. For example, the Monetary Policy Committee (MPC) of the Bank of England must deliver price stability over the medium term, as defined by 2% CPI inflation. The inflation target is symmetric (meaning we care as much about returning inflation to target from below as from above), and it applies at all times.

To be clear, although the method by which the Bank of England achieves the inflation target may be a technical exercise, the UK’s monetary policy framework is grounded in society’s choice of the desired end. These ethical determinations are encoded in the monetary policy remit, from which the MPC takes its orders and against which it is accountable. Society chooses the ends, and, within pre-set boundaries, the MPC determines the means to achieve them.

Yet even in this framework monetary policymaking will at times involve striking short-term trade-offs between stabilising inflation and supporting growth and employment. In other words, the monetary policy problem cannot be fully contracted *ex ante*. But to re-emphasise, decisions about the precise trade-off that the MPC pursues are subject to the limits provided by its statutory objectives and the Chancellor’s remit. And they are disciplined by the transparency and accountability mechanisms including *Inflation Reports*, the minutes and transcripts of its meetings, the Parliamentary testimonies of its members, the open letter process triggered if

³ Nordhaus, W D (1975), “The political business cycle”, *Review of Economic Studies*.

⁴ See, inter alia, Kydland, F and Prescott, E, (1977), “Rules rather than discretion: the inconsistency of optimal plans”, *Journal of Political Economy*, 85(3), 473-492; Barro, R, and Gordon, D B, (1983). “A Positive Theory of Monetary Policy in a Natural-Rate Model”. *Journal of Political Economy*. 91 (4): 589–610.

⁵ See Rogoff, K (1985), “The optimal degree of commitment to an intermediate monetary target”, *Quarterly Journal of Economics*, November, pages 1,169–90.

⁶ The ‘engineering’ view is primarily concerned “with logistic issues rather than with ultimate ends”, including their ethical characteristics. “The ends are taken ... as given, and the object of the exercise is to find the appropriate means to serve them”. See Sen, A K (1987), *On ethics and economics*, Blackwell.

the inflation target is missed by more than one percentage point, and discussions and debates in public fora such as this one.

The balance of my remarks this evening will concentrate on the experience with such trade-offs, which have become more common of late for two reasons. First, since the global financial crisis, the UK economy has been subject to a series of major supply shocks, which, for reasons that I will explain, create a tension between short-term inflation and output stabilisation. And second, there has been a renewed recognition that, in extreme circumstances, the monetary authority may have to take into account financial stability considerations in setting monetary policy with potential consequences for the short-term path of inflation control.

These are big decisions, decisions subject to considerable uncertainty, and decisions which entail potentially large welfare costs, arising from high unemployment and unevenly distributed costs of inflation. They deserve scrutiny to ensure that they are consistent with society's preferences as expressed in statute and remit.

1. Monetary policy remit

For most of the Bank's history, an assessment of whether it was acting in a way consistent with society's view of the public good could not be made.

As former Governor Eddie George remarked, during the half century that followed its nationalisation in 1946 "the Bank operated under legislation which, remarkably, did not attempt to define our objectives or functions." They were, instead, "assumed to carry over from [the Bank's] earlier long history."⁷ In that regard, the Bank's 'constitution' resembled that of the United Kingdom more broadly, comprising a rich history of law, principle and convention.

All changed with the passing of the Bank of England Act in 1998, which made specific "provision about the constitution, regulation, financial arrangements and functions of the Bank." The Act clarified the Bank's responsibilities and granted independence to the Bank for the operation of monetary policy. In delegating authority to an independent body, the MPC, in this way, the Act ensured the Bank would operate under 'constrained' rather than 'unfettered' discretion.⁸ It would be accountable to Parliament for operating the instruments of monetary policy to achieve the objectives of monetary policy, as determined by the Government.

The operational independence of the Bank of England is an example of power flowing from the people via Parliament within carefully circumscribed limits. Independence in turn demands accountability for the Bank to command the legitimacy necessary to fulfil its mission. By publishing its analysis, giving testimony, and

⁷ George, E. (2000), "Central bank independence", speech at the SEANZA Governors' Symposium, 26 August 2000.

⁸ See, for example, King, M. (2000), "Monetary policy: theory in practice", 7th January 2000. King attributes the 'constrained' versus 'unfettered' distinction originally to Ben Bernanke and Frederick Mishkin.

delivering speeches, the Bank explains how it is exercising its powers to achieve its clearly defined policy remits.

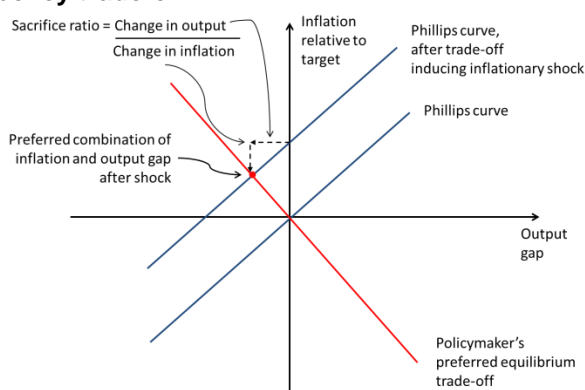
Today, the MPC's monetary policy remit requires it to achieve price stability, defined by the Government to mean 12-month CPI inflation of 2%. The remit recognises that inflation may deviate temporarily from the target on account of shocks, however. Since 2013, the remit has explicitly recognised that in these circumstances, bringing inflation back to target too rapidly could cause volatility in output and employment that is undesirable. The remit requires the MPC to consider, balance and explain such short-run monetary policy trade-offs.

Specifically, in exceptional circumstances, "shocks to the economy may be particularly large or the effects of shocks may persist over an extended period or both". When this is the case, the challenge facing the MPC can be more significant, and the remits directs that "[i]n forming and communicating its judgements, the Committee should promote understanding of the trade-offs inherent in setting monetary policy", including, importantly "the horizon over which the Committee judges it is appropriate to return inflation to the target".

2. Achieving price stability

A simple framework for thinking about how to manage the monetary policy trade-off is shown in **Figure 1**. The red line in the figure illustrates a policymaker's preferred trade-off: mapping the size of the shortfall of output below potential – the output gap – that the policymaker is prepared to tolerate for a given overshoot of inflation from target, and vice versa. The flatter the line the less the weight the policymaker places on output stabilisation and the more they are willing to tolerate large output gaps in order to eliminate small overshoots in inflation.

Figure 1: A stylised depiction of the monetary policy trade-off



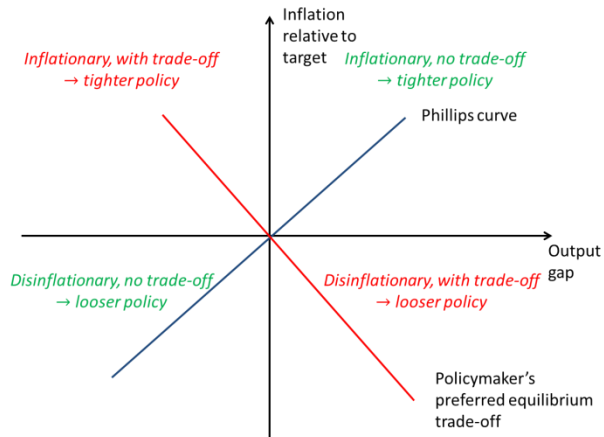
The blue line in the figure is the Phillips curve, which summarises the structure of the economy and, in particular, how changes in demand, via the output gap, affect inflation.⁹ The economy's equilibrium is at the intersection of these two lines.

In the case of an inflationary shock that induces a trade-off, the Phillips curve shifts up, as shown in the figure. Monetary policy needs to tighten to lower pressure on resources and reduce inflationary pressures, such that the economy ends up on the

⁹ Originally formulated by Phillips (1958), albeit as a relationship between wage growth and unemployment. See: Phillips, A. W. (1958). "The Relationship between Unemployment and the Rate of Change of Money Wages in the United Kingdom 1861-1957". *Economica*, 25 (100): 283–299.

red line, consistent with the policymaker's preferences (as shown by the dotted arrows in the figure). The slope of the Phillips curve is crucial in determining in the size of the fall in output needed to reduce inflation to an acceptable level – the sacrifice ratio.

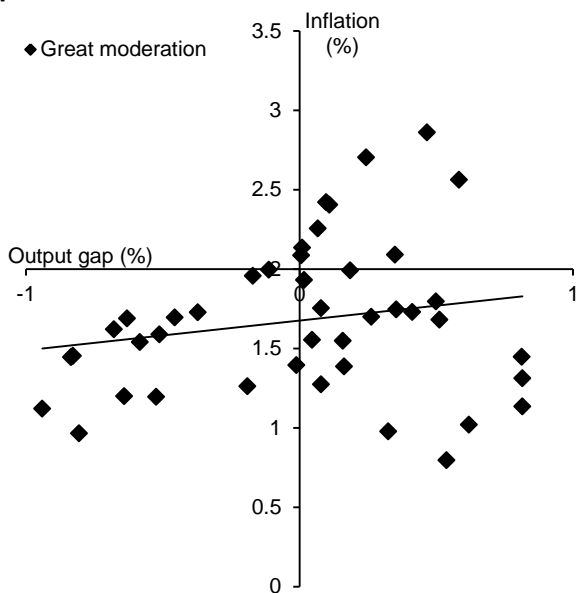
Figure 2: Stylised policy responses



Frequently, the economy experiences shocks that drive inflation and output in the same direction. These shocks to aggregate demand can include variations in government consumption, households' desire to consume, or business' desire to invest. Increases in demand put pressure on the use of resources, causing prices to rise. Because monetary policy can also influence demand it can lean against such shocks. If successful, it can stabilise inflation. In this case, no output-inflation trade-off arises. This is the so-called "divine coincidence".¹⁰

In **Figure 2**, negative demand shocks would create a tendency for inflation and the output gap to move into the south-west quadrant. Similarly, for positive demand shocks, inflation and output gap would move to the north-east. Because in these cases, there is no trade-off between stabilising inflation and the output gap, it is relatively straightforward for monetary policy to keep inflation and the output gap close to the origin.

Chart 1: For much of the great moderation period, data were consistent with predominance of demand shocks...



Source: Bank calculations.


In practice, data in these quadrants reflect lags in the transmission of monetary policy and other uncertainties: it takes time for monetary policy to affect output, and for output in turn to affect inflation. But if monetary policy has been effective, these deviations would not be expected to be either large or persistent. This state of affairs characterised much of the "Great Moderation" in the United Kingdom, from 1993 to 2007 (**Chart 1**), which was effectively a period of demand management.¹¹

Things are different when shocks drive inflation up or down independently of demand. Exogenous changes in firms' pricing power are one example – so-called cost-push shocks. Shocks to the exchange rate, the economy's supply capacity, or commodity prices also

¹⁰ See Blanchard, O., and J. Galí (2007). "Real Wage Rigidities and the New Keynesian Model". *Journal of Money, Credit, and Banking*, 39(1), pp. 35-65.


¹¹ See Broadbent, B (2013), "Conditional guidance as a response to supply uncertainty", speech at the London Business School.

Chart 2: ... whereas after the financial crisis, the economy experienced a challenging monetary policy trade-off



Source: Bank calculations.

Chart 3: US has continued to experience predominantly demand shocks, not trade-off shocks, since the global financial crisis...




Source: Bureau of Economic Analysis, CBO and Bank calculations. The measure of inflation is the four-quarter change in the personal consumption expenditures (PCE) deflator. The output gap is calculated using the CBO estimate of potential output from its August 2016 Update to the Budget and Economic Outlook.

have this flavour.¹² Because monetary policy's influence on inflation is predominantly an indirect one, via demand, in such circumstances inflation can only be controlled by delivering an opposing movement in aggregate spending. If something pushes up on inflation directly, monetary policy can only bring inflation back down by causing a reduction in spending via higher interest rates. The speed with which this adjustment is delivered is determined by the monetary policy maker guided by their remit.

Such circumstances have characterised the period in the UK since the global financial crisis (**Chart 2**), which entailed a large adjustment to the supply side of the economy, meaning a lower exchange rate, lower growth, and higher inflation.

Chart 4: ... as has the euro area



Source: Eurostat, IMF and Bank calculations. The measure of inflation is the four-quarter change in the harmonised index of consumer prices (HICP). The output gap is the IMF's estimate from the October 2016 WEO at an annual frequency, interpolated to create a quarterly series.

¹² Exchange rate shocks also have demand effects including via net trade. For example, a higher exchange rate lowers imported price pressures, reducing inflation directly. It also makes domestic goods less attractive on world markets, lowering the contribution of net trade to demand, which has second-round effects on inflation via the reduced demand of exporters for labour. At the same time, real household income rises with a higher exchange rate, providing some countervailing support to demand and domestically-generated inflation.

In contrast, the US and the euro area have seldom faced a trade-off between output and inflation stabilisation, even since the global financial crisis (**Charts 3 and 4**).

A simple way to represent more formally how the policymaker optimises the trade-off is in “linear-quadratic” form – a set of linear constraints describing the behaviour of the economy, and quadratic preferences that penalise deviations of inflation from its target and output from its potential. The relative weight the policymaker places on output stabilisation, relative to inflation stabilisation, is often denoted λ – or ‘lambda’. The policymaker’s “loss function” in any particular time period is:

$$Loss_t \equiv (\pi_t - \pi^*)^2 + \lambda(y_t - y_t^*)^2$$

where π_t is inflation, π^* is the inflation target, y_t is output and y_t^* is ‘trend’ output, so that $y_t - y_t^*$ is the output gap, and t subscripts denote time periods.¹³ The policymaker’s objective is to minimise the discounted sum of these losses over time. In this formulation, a lambda of zero would imply no weight on the stabilisation of real activity – so-called “inflation nutter” preferences. When lambda is positive, in contrast, the policymaker is willing to strike at least some trade-off between output and inflation stabilisation, as directed by the MPC remit.

This objective function is optimised subject to the constraints implied by the aggregate behaviour of the economy, including the relationship between interest rates and activity; and the relationship between activity and inflation – the Phillips curve shown in Figure 1. In essence, these suggest lower interest rates raise activity;¹⁴ and higher activity generates higher inflation.

Under this simple framework, in certain circumstances, the optimal policy balances inflation overshoots with shortfalls of activity relative to potential, and vice versa for inflation undershoots. (A derivation appears in the Annex.) The relative size of these two deviations is governed by the strength with which higher output translates into higher inflation; and the preferences of the policymaker, or:

$$\pi_t - \pi^* = -\frac{\lambda}{\kappa}(y_t - y_t^*)$$

where κ is the effect of the output gap on inflation. From this, it is clear that the higher is lambda, the greater the weight placed on output stabilisation and the more a given shock is allowed to flow through to inflation. As lambda shrinks to zero, the policymaker becomes an inflation nutter, with all of the adjustment to a shock forced through the output gap in order to keep inflation very close to the target.

¹³ A “micro-founded” loss function resembling this can be derived in New Keynesian models, in which lambda is a function of the underlying deep parameters in the economy (such as those governing the degree of price stickiness). These are typically based on the representative agent assumption. A newer literature drops representative agent assumption. In this case, the “micro-founded loss function” contains additional terms capturing some aspects of heterogeneity between households. For example, in a model with borrowers and savers, a term in the loss function containing the “consumption gap” arises – measuring the gap between the consumptions of the two types of household. The reason terms like this appear is that dropping the representative agent assumption tends also to introduce additional frictions, like borrowing constraints. In the example given here, the “consumption gap” represents the welfare loss due to imperfect risk-sharing (borrowers can’t smooth consumption as easily). Strictly speaking, this is something the “social planner” should care about, even if the central bank, with a narrowly-defined inflation mandate, does not.

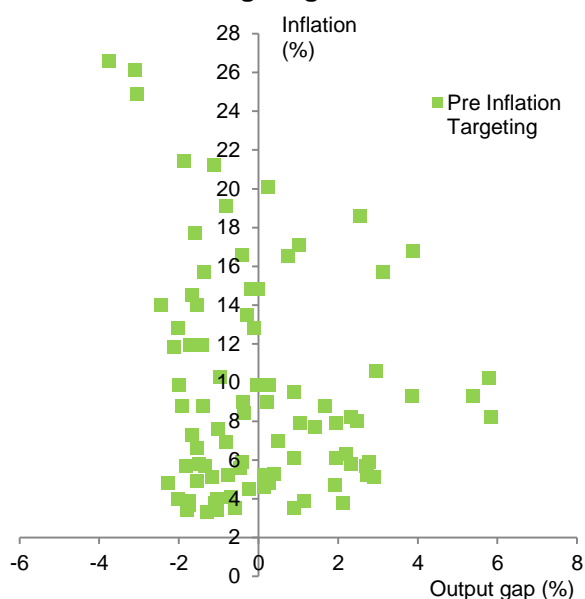
¹⁴ Lower interest rates, that is, relative to the equilibrium or “natural” interest rate.

The slope of the trade-off is also affected by that of the Phillips curve. When the Phillips curve is flatter, output gaps of a given size result in larger inflation over- or under-shoots.¹⁵ Intuitively, a flat Phillips curve makes the output costs associated with moving inflation around higher. As a result, a flatter Phillips curve means the policymaker will optimally choose to allow inflation to be further away from the target than if the Phillips curve were steeper.

3. Lambda, lambda, lambda

What does past behaviour reveal about the preferences of policymakers – the value of lambda?

Chart 5: Inflation very high and very volatile prior to inflation targeting



Source: Bank calculations. Inflation is the annual growth rate of the retail price index (RPI) until 1975, and the annual growth rate of the retail price index excluding mortgage interest payments (RPIX) after 1975. The output gap is the percentage deviation of quarterly GDP from its trend, where the trend is measured by a Hodrick-Prescott filter.

In the 1970s and 1980s, monetary policy was, for the most part, trained neither squarely nor credibly on inflation control. As a result, the inflation outcomes over this period ranged from poor to risible (**Chart 5**).

From 1971 to 1992, retail price inflation averaged 9.5%, with a standard deviation of 5.6%. Over the same period, unemployment averaged 7.6%, with a standard deviation of 2.8%. Soon after the adoption of inflation targeting, unemployment fell by one percentage point and its standard deviation was cut in half. The volatility in economic growth rate in the 1970s and 1980s was also one-and-a-half times higher than after the adoption of inflation targeting. In sum, high, volatile inflation likely contributed to a higher natural rate of unemployment and poor supply side performance.

Once the UK adopted inflation targeting and granted operational independence to the Bank, consistent with the theory of constrained discretion, the UK moved towards a better equilibrium characterised by low, stable and predictable inflation, and relatively stable employment. Moreover, by successfully anchoring inflation expectations around the target, inflation targeting helped lower UK borrowing costs for households and businesses, and has allowed monetary policy to promote short-run stabilisation of the economy. In other words, some trade-off between inflation and output has become possible and lambda has become relevant.

As I mentioned earlier, the period of inflation targeting can be divided into pre- and post-crisis periods. During the Great Moderation, the economy's supply capacity grew at a steady pace, and, for the most part,


¹⁵ In New Keynesian models, the slope of the Phillips Curve is related to the deep parameters governing the frequency with which firms are assumed to be able to re-optimize their prices, following the formalisation of Calvo (1983). See Calvo, G (1983), "Staggered prices in a utility-maximising framework", *Journal of Monetary Economics*, 12.

deviations of inflation from target reflected temporary disturbances to demand that were expected to drive output and inflation in the same direction (**Chart 1**).¹⁶ Generally, monetary policy faced few trade-offs; divine coincidence largely ruled.

The crisis brought shocks both to demand and supply, higher inflation and large shortfalls in output and employment. The era of divine coincidence was over and the MPC had to balance a series of challenging trade-offs (**Chart 2**).

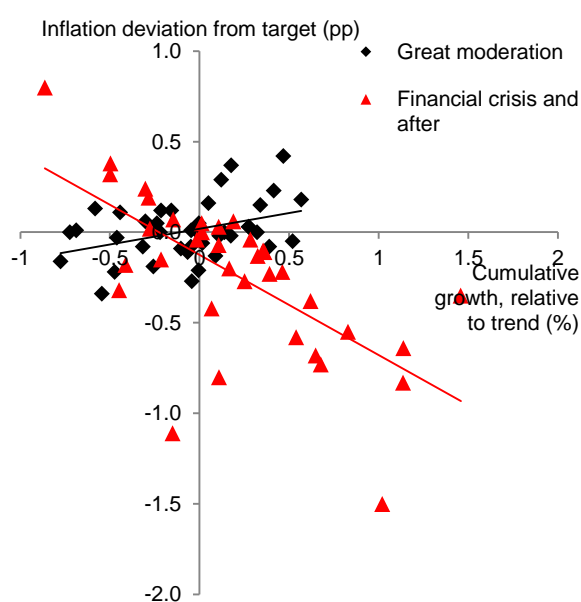
What can be learned about lambda from this era? Taking these trade-offs at face value suggests a value of lambda of between 0.1 and 0.2 for the inflation targeting period as a whole, on the basis of data outturns (**Chart 6**) and on the basis of MPC forecasts (**Chart 7**).¹⁷ For the post-crisis period alone, when the monetary policy trade-off was particularly stark, lambda is estimated to be around one-quarter, also on the basis of both outturns and forecasts.¹⁸ In other words, the MPC tolerated a period of above-target inflation in order to support the real economy. When combined with an average Phillips Curve slope, this implies deviations of the output gap from zero that are around half the size of those of inflation from the target.

Chart 6: Data outturns are consistent with a value of lambda of between 0.1 and 0.2 for the inflation-targeting era as a whole, and around ¼ for the post-crisis period...



Source: Bank calculations.

Chart 7: ... with similar values implied by MPC forecasts



Notes: Chart shows inflation deviation from target at the 2-year forecast horizon against cumulative growth, relative to trend, as a proxy for the output gap.

¹⁶ Typically, these demand disturbances were well signalled by survey measures of activity, resulting in a high correlation of these and Bank Rate (Broadbent, *ibid.*).

¹⁷ Inferring policy preferences from the MPC's published forecasts in this way is not necessarily straightforward, however, because the MPC's forecasts are conditioned on the market-implied path for interest rates. Implicitly, then, the assumption would be that the MPC's forecasts and reaction function are, on average, well understood by the market.

¹⁸ This assumes a Phillips Curve slope, κ , of $\frac{1}{2}$. The lines of best fit in Charts 1 and 2 have slopes of around minus 0.3 for the period as a whole and minus 0.55 over the post-crisis period. With a value of $\kappa = 0.5$, this suggests λ of between 0.1 and 0.2 for the whole period and λ of around $\frac{1}{4}$ for the crisis period alone.

To put this into some context, consider the stated and revealed preferences of the FOMC. Chair Yellen has referred to a “balanced approach” to trade-off management, implying equal weights on inflation deviations and the unemployment gap.¹⁹ Econometric studies for the US provide a range of estimates for lambda, centred around $\frac{1}{2}$ albeit with significant estimation uncertainty, with some studies finding evidence that the FOMC’s implied lambda has also varied over time.²⁰

4. Why there isn’t a simple rule to achieve society’s inflation objective?

This exercise illustrates how the MPC has sought to do what society has asked. Within limits, it has tolerated periods of above-target inflation in order to support the real economy, while at all times respecting the primacy of the inflation target.

That said, it is important not to over-interpret these results. Specifically, the MPC’s average revealed lambda in the past does not reveal a simple monetary rule that can govern decisions regarding inflation control in all scenarios in the future.

I stress this point because some have suggested monetary policy ought to follow, or be evaluated against, a simple rule, such as a Taylor (1993) rule.²¹ But it is important to recall that, while such rules were estimates of *actual* stances of past policy – positive descriptions of central bank behaviour – they have been re-interpreted as guides for what the central bank should do in all circumstances – normative prescriptions. Taking the past as a strict guide to the future is to assume that the nature shocks does not change and that the structure of the economy remains constant.

Such stability is hard to square in an era of financial crises; deep, variable technological change; and potentially large shifts in openness. In a changing world, monetary policy decisions would more properly emerge from optimal control approaches to monetary policy, which respect the different nature of shocks and circumstances which the monetary authority must weigh in setting its strategy.

The time-variation of lambda is illustrated in part by the fact that the regression estimating the MPC’s lambda leaves around two-thirds of the variation of inflation over the output gap unexplained, even in expectation and over the post-crisis period. So why isn’t the estimated lambda fixed, and what considerations inform its value in practice?

¹⁹ Yellen, J (2012), “Revolution and Evolution in Central Bank Communications”, November. Converting this to an equivalent value for the weight on the output gap depends on the Okun’s law coefficient one assumes, estimates for which vary.

²⁰ Ilbas (2012) estimates the Smets-Wouters model under optimal commitment policy for the Volcker/Greenspan era and find a value of lambda on output growth of around $\frac{1}{2}$. Dennis (2006) uses a semi-structural model and finds a weight of around 3 on the output gap, whereas Givens (2012) estimates a value of around 0.1 over a similar sample. Givens and Salami (2015) estimate a small New Keynesian model on post-Volker data and find a lambda on the output gap of around $\frac{1}{4}$. Sack (2000) finds a weight of around 0.8 on the unemployment gap over the period 1984-1998 using a VAR to describe the economy’s dynamics. Lakdawala (2016) allows for time variation in central bank preferences and finds that the weight on inflation relative to the output gap to have fluctuated over time, falling to low levels in the 1990s and early 2000s (implying a high relative weight on the output gap), before rising in the period leading up to the financial crisis.

²¹ Taylor, J (1993), “Discretion versus policy rules in practice”, Carnegie Rochester Series on Public Policy, 39.

First, any lambda that can be estimated from simple scatter plots is actually an amalgam of several elements. Most fundamental is the policy remit itself, which frames the ‘true’, underlying preferences that the monetary authority must pursue, as I discussed earlier. As important, however, is that the actual trade-off struck will be influenced by individual MPC members’ views of both the nature of the shocks hitting the economy and the transmission mechanism of monetary policy – in other words, why the economy is behaving as it is, and how interest rates affect output and inflation. Put another way, if members have, for example, different views of kappa – the slope of the Phillips curve – their lambdas will appear to be different. Indeed, it is quite possible for the Phillips curve slope to vary over time, and some have argued that it may have flattened in recent years.²²

A close reading of the MPC’s minutes since its inception reveals the vast majority of differences of opinion regarding the path for monetary policy have originated in this second set of factors (the nature of the shocks, the structure of the economy and the monetary policy transmission mechanism). Disagreements about the desired path for policy have arisen because of different views of how the economy is acting. This is as should be expected because society, through Government, has chosen the degree of inflation control and provided guidance via the remit on how it should be achieved.

Second, the nature, scale and persistence of the shocks affecting the economy matters for the trade-off that is pursued at any given time.

Consider the case of commodity price shocks. For the most part, these pass through to the level of consumer prices relatively quickly – changes in petrol prices being one example. Once these shocks have been reflected in the price level, the inflation rate, which is the change in the price level, would typically return close to its baseline path. In this case, monetary policy makers may be more inclined to ‘look through’ the direct, temporary effects of commodity price shocks on inflation, which means they do not intend to engineer output gaps to offset the inflationary effects in response. Because, over the short term, inflation would be allowed to vary while the real economy remained relatively stable, an observer would thus infer a higher value of lambda.

This may be less the case for inflationary shocks of a more persistent nature, such as those arising from movements in the exchange rate.²³ When shocks have persistent effects on inflation over the policy horizon, the monetary policymaker may have to balance them with surfeits or shortfalls of activity around potential. Because this would call for greater variation in employment and growth, an outside observer would, all else equal, infer a lower value for lambda.

This is when different views regarding the transmission mechanism of monetary policy become relevant. Those who regard the impact of monetary policy on the economy as relatively rapid may be more inclined to

²² See e.g. Blanchard, O, Cerutti, E and Summers, L (2015), “Inflation and activity – two explorations and their monetary policy implications”, NBER working paper No. 21726.

²³ At root this is because, in aggregate, retailers whose costs rise adjust their prices only gradually, perhaps reflecting strategic, competitive concerns, or because many importers, over the short term at least, hedge their foreign exchange exposure. See Box on page 28 of the November 2015 *Inflation Report* for a review.

act in response to shocks whose effects are relatively short-lived. All else equal this would tend to produce a shallower observed relationship between inflationary shocks and output gaps – a lower lambda. If, on the other hand, monetary policy is viewed as operating with a substantial lag, it may be that the more enduring influences on inflation are the ones that weigh most heavily in the monetary policy decision. In this case, this would tend to produce a steeper observed relationship between inflationary shocks and output gaps – a higher lambda.

The third reason why lambda may vary over time is that the transmission mechanism of monetary policy has evolved with the changing nature of the financial sector, the openness of the economy, and the indebtedness of households. Moreover, the tools in the monetary arsenal have expanded since the crisis, including large-scale asset purchases whose transmission is less well understood than conventional measures. When such uncertainties are particularly high, it is natural to take them into account when setting monetary policy. That can motivate setting monetary policy in a cautious way,²⁴ favouring gradualism. In effect, this adds a desire for interest rate stability to the monetary loss function, which has the effect of reducing the relative weight on output stabilisation.

Fourth, given the primacy of price stability in the Bank's monetary policy remit, the credibility of the monetary framework is paramount. This MPC carefully reviews household, business, and financial markets' expectations regarding future inflation. Any perceived risks to medium-term inflation expectations could justify more concerted efforts to return inflation to the target than usual, and therefore lower lambda.

Fifth, the role of unknowns in the economic outlook may be so prominent that they don't constitute risks, but "radical uncertainties".²⁵ These are uncertainties that are at once powerful and profound, yet impossible to quantify over the monetary policy horizon, even with subjective probabilities. Such uncertainties do not fit neatly into the "linear-quadratic" approach to policy which embeds certainty equivalence and sets policy with regard to the mean outcome for the economy. Their presence can raise the question whether policy should be conducted in a way that insures against them by trying to avoid particularly bad outcomes, rather than trying to optimise the mean.

5. Beyond lambda

Some of the most important tail risks relate to financial stability. It is now clear that, although it secured fifteen years of inflation control and steady economic growth, price stability, narrowly defined, morphed from healthy focus to a dangerous distraction. The financial crisis that exploded the Great Moderation was a powerful reminder that price stability is not sufficient to maintain macroeconomic stability, and that the

²⁴ As per Brainard, W (1967), "Uncertainty and the effectiveness of policy", *American Economic Review*, 57(2). There is a literature exploring how this result varies with other forms of uncertainty about the transmission of monetary policy. E.g. Söderström, U (2002), "Monetary policy with uncertain parameters," *Scandinavian Journal of Economics*, vol. 104(1), pages 125-145. Interest rate smoothing can also be motivated to the extent that it generates greater purchase over long-term interest rates, e.g. Woodford, M (2003), "Optimal interest rate smoothing", *Review of Economic Studies*, 70(4).

²⁵ See King, M (2016), *The End of Alchemy: Banking, the Global Economy and the Future of Money*, Little Brown.

evolution of credit and indebtedness can build up major macro problems for the future.²⁶ The fact that such financial cycles tend to be slow-moving – operating at longer frequency than the monetary policy horizon or even the business cycle – considerably complicates trade-off management.²⁷ Although by no means the sole culprit, the absence of a body with a remit and tools to promote the stability of the financial system as a whole was a major gap in the UK’s policy framework.²⁸ In a sense, prior to the global financial crisis, the macroeconomic loss function was incomplete.

The principal remedy has been the establishment of the independent Financial Policy Committee (FPC) at the Bank, with a remit “to protect and enhance the stability of the financial system of the United Kingdom.” The Chancellor’s remit to the FPC defines this as relating “primarily to the identification of, monitoring of, and taking of action to remove or reduce systemic risks”, as long as its actions would be unlikely “to have a significant adverse effect on the capacity of the financial sector to contribute to the growth of the UK economy in the medium or long term.”²⁹ To achieve these objectives, the FPC has a range of tools, including the countercyclical capital buffer banks are required to maintain, together with the ability to set limits on mortgage underwriting standards, such as maximum permissible debt-to-income or loan-to-value ratios.

The FPC is the first line of defence against financial stability risks. But the remits recognise that circumstances are possible “in which attempts [by the MPC] to keep inflation at the inflation target could exacerbate the development of imbalances that the FPC may judge to represent a potential risk to financial stability”. In such a case, the MPC “may wish to allow inflation to deviate from the target temporarily, consistent with its need to have regard to the policy actions of the FPC.”

In the event of such a judgement by the FPC, the MPC will have to weigh the appropriate monetary policy response within the context of its inflation target and its secondary objective to support the Government’s economic policy objective to achieve strong, sustainable, balanced growth. In effect, the MPC’s loss function could be modified to include proximate indicators of financial stability (that are considered by the FPC):

$$Loss_t \equiv (\pi_t - \pi^*)^2 + \lambda(y_t - y_t^*)^2 + \mathbf{1}_{FPC}\beta(\mathbf{s}_t - \mathbf{s}_t^*)^2$$

where \mathbf{s}_t is a vector of financial stability indicators, \mathbf{s}_t^* are the levels judged consistent with stability, and whose differences are given a weights of β . The indicator variable $\mathbf{1}_{FPC}$ takes a value of one if the FPC judges there to be potential risks to financial stability, and is zero otherwise.

Because its remit requires it to support strong *yet sustainable and balanced* growth, in exceptional circumstances, and in consultation with the FPC, the MPC could have to optimise – subject as always to society’s risk preferences – over a longer time horizon in order to avoid inter-temporal disequilibrium, with the

²⁶ E.g. Mian, A and Sufi, A (2014), *House of debt*, Chicago.

²⁷ See for example, Aikman, D, Haldane, A and Nelson, B, (2015) “Curbing the credit cycle”, *Economic Journal*, 125(585), June.

²⁸ Though not the only one. Others included the absence of a credible resolution regime for failing institutions, inadequate capital and liquidity standards, excessive complexity and opacity in the financial system, and cultural failings alongside other perverse incentives for excessive risk taking.

²⁹ The full FPC remit can be found here:

<http://www.bankofengland.co.uk/financialstability/Documents/fpc/letters/chancellorletter160316.pdf>

result that inflation could deviate below target temporarily in order to achieve a more sustainable return to the target and more stable growth beyond the conventional policy horizon.


6. The Current Trade-off

How do these factors relate to the recent conduct of monetary policy?

In May, the MPC said that the effects of a vote to leave the EU on inflation would depend on the balance of its effects on demand, supply and the exchange rate. Having observed the behaviour of the economy in the period since the crisis, the Committee was aware of the possibility that these effects could, in combination, mean a period of weaker growth but higher inflation and thus create a trade-off for monetary policy to balance.

The Committee made its first full assessment of these effects in August. Over the three-year forecast period, uncertainty was expected to weigh on demand, investment spending by firms in particular, and a period of resource reallocation as the UK's trading relationships were re-configured was expected to weigh on supply. At the same time, the sharp fall in the sterling exchange rate following the vote to leave would imply higher imported inflation, meaning upward pressure on inflation. The MPC responded to this outlook by reducing

Chart 8: No stimulus in August would have been consistent with a weight of zero on stabilising the output gap at the 2-year forecast horizon



Source: Bank calculations. A value of 0.1 gives an approximate lower bound on estimates of lambda over the post inflation targeting period (see text). A value of unity would correspond to a "balanced approach" which weighs inflation stabilisation and output gap stabilisation equally.

Bank Rate by 25 basis points, announcing £70 billion of asset purchases, and by launching a Term Funding Scheme to ensure banks passed through lower interest rates to end borrowers. A majority of the MPC also indicated that, if the economy evolved in line with its forecast, they were prepared to take further action.

Suppose the MPC had not responded in this timely, coherent, and comprehensive way. **Chart 8** shows that, given the constellation of shocks affecting the economy, opting not to provide monetary stimulus would have meant, in all likelihood, inflation around the target at the two-year point of the forecast, but an output gap of some 1½%, implying around 1/4 million lost jobs. In other words, fully offsetting the persistent effects of sterling's depreciation on inflation would have required exerting further downward pressure on domestic costs. And that would have meant even more lost output and a total disregard for higher

unemployment. Given our remit, that would have been undesirable. Moreover, because spare capacity would have persisted in such a scenario, this strategy would have been unlikely to generate a *sustainable* return of inflation to the target once the effects of higher import prices had become fully reflected in the price level.

Chart 8 shows how, with the stimulus measures in place, the MPC balanced the trade-off we faced by placing some weight on output stabilisation consistent with the Committee's remit.

Over the autumn, demand growth remained more resilient than had been expected, particularly consumer spending and to a lesser extent the housing market. That contrasted with the less sanguine assessment of financial markets as illustrated by the further, sharp decline in the sterling exchange rate over the autumn.

Ultimately, the tension between consumer strength on the one hand and the more pessimistic expectations of markets on the other will be resolved.

In the MPC's November projections, this resolution is expected to occur as imported inflation begins to weigh in coming months on people's real incomes, slowing consumption growth. This moderation in household spending reinforces the cumulative effects of a pick-up in uncertainty on investment. As a consequence, growth is expected to remain below past averages for the next few years.

One corroborating indicator of this potential deceleration is that the UK expansion is increasingly consumption-led. Evidence from the past quarter century across a range of countries suggests episodes of consumption-led growth tends to be both slower and less durable. This is because consumption growth eventually outpaces earnings growth, increasing debt and making demand more sensitive to changes in employment and income.

At present, households appear to be entirely looking through Brexit-related uncertainties. The saving rate has fallen towards its pre-crisis lows, and consumer borrowing has accelerated notably. In the year to November, total household borrowing rose 4%, and consumer credit rose over 10%, the fastest rate since 2005. How household spending evolves, and the inter-temporal trade-off that households strike, will be important considerations over the next year. The MPC will continue to monitor these dynamics.

In November, the MPC reiterated that we were choosing a period of somewhat higher consumer price inflation in exchange for a more modest increase in unemployment. But it also noted that there are limits to the extent to which above-target inflation can be tolerated.

Those limits depend on the cause of the inflation overshoot, the extent of second-round effects on inflation expectations and domestic costs, the scale of the shortfall in economic activity below potential, and any risks around the development of imbalances that could threaten a sustainable return of inflation to the target. Recently, there have been signs of continued solid consumer momentum domestically and a stronger growth outlook globally. The MPC will monitor developments in the light of its inflation tolerance, and will explain its assessment and policy stance accordingly.

It remains the case that the outlook for inflation will depend on the evolution of the prospects for demand, supply and the exchange rate. Monetary policy can respond, in either direction, to changes to the economic outlook as they unfold to ensure a sustainable return of inflation to the 2% target.

7. Conclusion

The people of the UK have chosen price stability as the primary objective for monetary policy, having reason to value inflation control.³⁰ Their preferences have been encoded in the MPC's inflation remit.

Although those preferences are set in a statute, there remains an element of discretion in how the MPC delivers its inflation objective. That is because the people of the UK also have reason to value stable growth, jobs and incomes. And in exceptional circumstances, trade-offs between real stability and inflation can arise that monetary policy is required to balance.

This is now the case given the decision of the people of the United Kingdom to leave the EU. In the coming years, the UK will redefine its openness to the movement of goods, services, people and capital. The flexibility and dynamism of this economy will help it adjust as its relationship with the EU becomes clearer and new opportunities with the rest of the world open up. Over the next few years, the magnitude of the effects of this adjustment on the economy's supply potential, domestic demand, and the value of sterling will be somewhat uncertain; and this process will have a significant bearing on inflation.

Whatever transpires, the MPC will manage monetary policy to achieve the inflation target in a sustainable manner consistent with the preferences and instructions of the people of the United Kingdom.

³⁰ Sen's capability approach emphasises "capabilities of [people] to lead the kind of lives they value – and have reason to value". As discussed in the text, people have reason to value inflation control, together with macroeconomic stability, hence these objectives of macroeconomic policy.

Annex: The simple optimal targeting rule

The condition describing the trade-off between inflation and the output gap in equation (1) of the text is known as an optimal targeting rule:

$$(1) \quad \pi_t - \pi^* = -\frac{\lambda}{\kappa}(y_t - y_t^*)$$

It comes from the first-order condition of a policymaker acting under “discretion” (i.e. setting monetary policy afresh each quarter, without the ability to commit to future policy actions, and therefore manipulate future expectations) in a very stylised model of the monetary policy problem. In this problem, the policymaker minimises a loss function by choice of the nominal interest rate, formally solving:

$$\min_{\{i_t\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t [(\pi_t - \pi^*)^2 + \lambda(y_t - y_t^*)^2]$$

where π_t is inflation, π^* is the inflation target, y_t is the level of output, y_t^* is the natural level of output, so $y_t - y_t^*$ is the output gap, and i_t is the nominal interest rate – the policy instrument. β is the discount factor. The policymaker solves this problem subject to the “constraints” imposed by the structure of the economy. In the simplest New Keynesian model, these are, first, the “IS” curve, which traces a negative relation between output and the (ex-ante) real interest rate ($i_t - E_t\pi_{t+1}$):

$$y_t = E_t y_{t+1} - \frac{1}{\sigma}(i_t - E_t\pi_{t+1})$$

This equation comes from the optimisation problem of households who choose consumption to maximise their lifetime utility. E_t denotes the mathematical expectations operator. The term in expected output, $E_t y_{t+1}$, appears on account of households’ desire to smooth consumption. The real interest rate enters this condition because ultimately households care about their real consumption levels and it is the real interest rate that is relevant for transferring these consumption levels, via saving, across time. The parameter σ , the inter-temporal elasticity of substitution, controls the size of this effect.³¹

In the simplest monetary models, the natural level of output (or “trend output”, if you like), is exogenous (e.g. is determined by TFP alone).³² An expression equivalent to the IS curve above can then be used to define the natural interest rate, r_t^* , which is the interest rate that would prevail in the absence of any nominal rigidities:

$$y_t^* = E_t y_{t+1}^* - \frac{1}{\sigma} r_t^*$$

Using this, the IS curve can be re-written in output gap form, using the output gap definition $x_t \equiv y_t - y_t^*$ and the definition of the natural interest rate to give:

$$(2) \quad x_t = E_t x_{t+1} - \frac{1}{\sigma} [(i_t - E_t\pi_{t+1}) - r_t^*]$$

³¹ In other words, the simple model here abstracts from all channels of monetary transmission other than the inter-temporal substitution channel. Of course, building in addition frictions and model features, like borrowing constraints and open economy features, can restore the roles of the other channels.

³² In particular, in the simple model, one can show that $y_t^* = \psi a_t$, where a_t is the log-level of TFP, assumed to follow an exogenous process.

This says that the output gap is negatively related to the sequence of “real rate gaps” $(i_t - E_t \pi_{t+1}) - r_t^*$ delivered by monetary policy. For example, when monetary policy delivers real interest rates equal to the natural rate, so $i_t - E_t \pi_{t+1} = r_t^*$, the output gap is zero at all times. If there were only “demand” shocks in the economy, so only shocks to r_t^* , then inflation could always be stabilised by setting the ex ante real interest rate equal to r_t^* , and keeping the output gap closed. This special case is called “divine coincidence” in the literature.

However, trade-offs can arise via the second key “constraint” the policymaker faces: the relation between output and inflation, i.e. the Phillips curve. In the simplest model, this takes the form:

$$(3) \quad \pi_t - \pi^* = \beta E_t(\pi_{t+1} - \pi^*) + \kappa x_t + \varepsilon_t^\pi$$

where β is the household’s rate of time preference and the slope κ can be related to the frequency with which firms re-set their prices.³³ In (3), ε_t^π is a “cost-push” shock that drives inflation away from the target for a given output gap, and is the source of monetary policy trade-offs. The policy problem is therefore to solve (1) subject to (2) and (3).

Note that there is a one-to-one mapping between the interest rate i_t and the output gap x_t . This means it is ‘as if’ the monetary policymaker chooses the output gap directly. (One can then infer the nominal interest rate needed to deliver this output gap from the IS curve). Using this transformation, the policy problem can be re-stated as:

$$\begin{aligned} & \min_{x_t} (\pi_t - \pi^*)^2 + \lambda x_t^2 \\ & s. t. \quad \pi_t - \pi^* = \beta E_t(\pi_{t+1} - \pi^*) + \kappa x_t + \varepsilon_t^\pi \end{aligned}$$

Because the policymaker cannot commit to future policy actions, she takes expectations as given (i.e. does not internalise the possibility that, by committing to future policy actions, expectations can be manipulated having an effect on outcomes today). The first-order condition in this case is:

$$2(\pi_t - \pi^*) \frac{\partial \pi_t}{\partial x_t} + 2\lambda x_t = 0$$

Using $\frac{\partial \pi_t}{\partial x_t} = \kappa$, this simplifies to:

$$\pi_t - \pi^* = -\frac{\lambda}{\kappa} x_t$$

giving equation (1) after using $x_t \equiv y_t - y_t^*$. This optimality condition has three intuitive features. First, positive inflation gaps should be matched by negative output gaps, and *vice versa*. Given a quadratic loss function, it makes sense to “spread the cost” associated with shocks across both target variables, rather than loading all the cost on to one alone. Optimal policy therefore returns inflation and output back to their target

³³ In particular, if firms re-set their prices with probability $1 - \theta$ each quarter, then one can show that $\kappa \equiv \frac{(1-\beta\theta)(1-\theta)}{\theta}$, such that as θ rises (i.e. as prices become more sticky), the slope of the Phillips curve falls. The average duration of a price is $(1 - \theta)^{-1}$. For example, if the average duration of a price is expected to be ½ a year, then $\theta = 0.5$; and if the β , the discount rate, is consistent with a steady state annual real interest rate of 3%, say, (i.e. $\beta = (1.03)^{-\frac{1}{4}}$) then $\kappa \approx 0.5$.

levels from opposite directions. Second, the larger is the preference for output stabilisation, λ , the larger should be the absolute size of inflation deviations from the target, so that less of the cost of adjustment is borne by output. Finally, the flatter is the Phillips curve, κ , the larger should be the inflation deviations from target. This is because flatter Phillips curves mean you have to move output by more to achieve a given inflation objective.