

Uncertainty in macroeconomic policy making: art or science?

**David Aikman, Philip Barrett, Sujit Kapadia, Mervyn King, James Proudman,
Tim Taylor, Iain de Weymarn and Tony Yates
(Bank of England)**

Paper delivered to The Royal Society Conference on

“Handling Uncertainty in Science”

Monday 22 March 2010, London

Uncertainty in macroeconomic policy making: art or science?

1. Uncertainty and public policy

Most questions of public policy relate to uncertainties. Answers depend upon an ability to understand and evaluate those uncertainties. Yet many commentators and members of the public want to believe in certainties. They want to cut through the thickets of caveats and technical difficulties to the “bottom line”. We see this in economic policy, weather forecasting and the interpretation of secret intelligence.

Five years ago, the Bank of England and the Met Office held an informal seminar to exchange ideas on how to communicate measures of uncertainty to the general public. Although we failed to persuade broadcasters to explore the issues of economic and weather forecasting in more depth, the bilateral relationship which developed between the Bank and Paul Hardaker, who moved from the Met Office to take over the Royal Meteorological Society, flourished. Today’s conference is a tribute to Paul’s persistence. It could hardly be more timely. At the heart of debates about climate change, the Chilcot Inquiry into the war in Iraq, and the speed of the necessary fiscal consolidation of our public finances, lie questions about how to evaluate and respond to considerable uncertainties.

At that seminar in 2005, we discussed three questions. First, in modelling the phenomenon of interest, what is the right balance between simplifying assumptions and technical rigour? Second, how open should we be about what we don’t really know? Third, in communicating uncertainty, how can we prevent the intended audience from focussing excessively on one or two parameters of the underlying distribution at the expense of the whole picture? The discussions compared the Bank of England’s fan charts for the outlook for inflation with the Met Office’s plume charts for forecasts of temperature. This paper updates that discussion, and presents some new ways of communicating uncertainty which are under evaluation by the Monetary Policy Committee.

Reference to economics and meteorology is not just because the interests of some of us span both. They are two disciplines united by a common need to grapple with complex systems and communicate forecasts to a wide audience. It is not enough to explain only to experts. Much of the value of forecasts is in their being understood – in all their subtleties – by the general public, or at least sub-groups such as the agricultural industry in the case of meteorology and those involved in wage bargaining in the case of economics.

The use or misuse of forecasts has caused problems for both the Bank of England and the Met Office in recent years. The Bank’s forecasts of inflation were criticised for overstating the size of confidence intervals during the period from 1995 to 2007, the so-called “Great Stability”, and then understating them once the financial crisis occurred. And the Met Office has recently announced that it will no longer publish forecasts for the weather in forthcoming seasons after it “failed” to predict the cold winter or the wet summer (in some parts of the country) last year. The seasonal forecasts will be replaced by monthly predictions. In a statement, the Met Office said, “by their nature, forecasts become less accurate the further out we look. Although we can identify general patterns of weather, the science does not exist to allow an exact forecast beyond five days, or to absolutely promise a certain type of weather”. But to publish forecasts only when the outcome is virtually certain would be an admission of at least partial defeat.

If descriptions of probability distributions are misinterpreted by the media or the public, how can we discover which messages are robust and which are too fragile to survive the process of communication to a wide audience? First, we need to consider how we should form a view of the likelihood of different outcomes.

2. Unpredictability in economics and science

Unpredictability is a challenge faced by practitioners in many disciplines. As Kenneth Arrow, Nobel Laureate in Economics, said about his four-year experience as a weather officer in the US during World War II, “one thing I learned from meteorology is that being an actual science was no guarantee of exactness”.

Economics may or may not be an actual science, but it certainly suffers when trying to predict the future. There are at least three explanations for this unpredictability that are common to economics and other scientific disciplines.

First, assigning probabilities is particularly difficult for infrequent high-impact events such as financial crises, tsunamis, or climate change, for which there are few precedents. Because the samples of such events are so small, as we learn about these processes our assessment of their likelihood will change, possibly sharply.

Second, even systems that are deterministic rather than stochastic can be very unpredictable. Small differences in starting conditions can imply very different outcomes. Weather systems display chaotic dynamics, so that small initial forecast errors may lead to large revisions to the forecast further out. Economists, in contrast, often rely on simple linear models, but with additive stochastic errors. In modelling any phenomenon there is a choice, and among the experts in the area often a live debate, about whether it is more useful to use deterministic but chaotic or stochastic models.¹ But the motivation in both cases is to bring out the importance of unpredictability.

Another feature that hinders predictability is that *a priori* indistinguishable shocks can have very different effects in some systems. If you drop a grain of sand randomly on a sand pile, drop a lit match at random in a forest, or infect a random person with a disease, it may have very minimal, localised effects or may lead respectively to an avalanche, wildfire or epidemic. The outcome depends on the state of the system when the shock hits (the steepness of the sand pile, dryness of the forest, or density of the population) and the precise point at which the system is hit.²

Third, systems can be subject to sudden transitions between locally stable states. In both economics and science, we can find examples of systems that have multiple stable states.

And switches between these may be hard to predict. Models of the circulation pattern in the North Atlantic Ocean suggest that it can either be vigorous, helping to warm the northern hemisphere as water is transported from lower latitudes, or can shut down almost entirely. The magnetic properties of iron provide another example of a system with different stable states. When exposed temporarily to an external magnetic field, iron retains its magnetism even when the external field is removed: the same external conditions are consistent with different states of the system, depending on its history.³ Many economic systems also exhibit forms of hysteresis or path dependence. For example, changing social attitudes towards personal bankruptcy, inflation or unemployment may affect equilibrium outcomes for those variables.

3. Policy responses to unpredictability

Predicting the precise timing and dynamics of instability in all of these settings is very difficult. It is, however, possible to identify a sand pile or forest as being prone to large-scale instability and determine the factors that contribute to that instability. Analogously, many people, including at the Bank of England, did identify the vulnerabilities of the financial system before the crisis without actually being able to predict when or how the crisis would begin.

Economists have been able to learn from other disciplines about how to cope with these types of instability. For example, insights from ecology and epidemiology have been applied to financial networks (for more discussion, see Haldane 2009), and economists have learned much from engineers about how to control dynamic systems.

The key is for policymakers to focus on making the structure of the underlying system more robust to shocks. For example, in avalanche areas the snow may be ‘seeded’ so that, by inducing small avalanches, the chance of a large avalanche is mitigated. And in fire-risk areas, controlled burns or fire breaks are sometimes used to limit the risk of a large fire. In the context of the financial system, policymakers could impose a ‘Glass-Steagall’ style separation between the payments system network and risky activities.

Actions can also be taken to mitigate the impact of a bad outcome. For example, this might mean building an avalanche net or fence; funding a fire-fighting service; or

launching a lifeboat. In the context of financial regulation, this might entail reducing the costs associated with resolving financial institutions, for example by requiring banks to develop resolution and recovery plans, or requiring forced subsidiarisation of foreign bank branches.

As always, however, there is the risk that such interventions can have unintended consequences: shortly after the Titanic sank, a steamer, the S.S. Eastland, capsized on Lake Michigan, with the loss of around 850 lives, because it had been destabilised by the extra lifeboats required by post-Titanic regulations!

4. Decision making under uncertainty and the role of beliefs

The analogies between economics and science can be taken only so far. A key difference between economics and the physical sciences is the role played by active decision-makers – such as households and businesses – whose presence complicates substantially the dynamics of the system.

In particular, economic outcomes are sensitive to the way people behave under uncertainty, and to their beliefs about the past, present and future. Perceptions of uncertainty affect behaviour. If enough people are uncertain about their future job prospects and save more as a precaution, that may lead firms to cut employment, so increasing the chance that jobs are indeed lost. But uncertainty about whether it will rain today does not change the weather.

Beliefs adapt over time in response to changes in the environment; and this in turn affects how economic systems behave. Because the surrounding environment can affect economic decision-making, there are probably few genuinely ‘deep’ (and, therefore, stable) parameters or relationships in economics. In contrast, in many settings in the physical sciences there are stable ‘rules of the game’ (e.g. the laws of gravity are as good an approximation one day as the next).

The study of decision making under uncertainty has a distinguished ancestry in economics. Frank Knight (1921) wrote about the difference between risk (where frequencies could be used to calculate probabilities) and uncertainty (where there was no objective basis on which to derive probabilities). Of course, a true Bayesian would

assume that people would be able to construct subjective probabilities, and their behaviour would be governed by those beliefs. But Keynes (1936) made much of the pervasive nature of uncertainty in his explanation of business cycles because of the instability of those subjective probabilities. Occasional sharp changes in expectations could, and did, result. That is exactly what happened in the Autumn of 2008, with the sudden and synchronous collapse in business confidence around the world.

So the behaviour of economies depends in part upon the beliefs about future events held by the various players: consumers, businesses and governments. Outcomes are endogenous to beliefs about the likelihoods of the outcomes. And in some contexts there may be self-fulfilling equilibria.

Sudden shifts in beliefs are particularly relevant when many decision-makers interact. The behaviour of a bank's depositors has many equilibria, each of which depends upon beliefs about whether other depositors are likely to run on the bank. Changes in perceptions of uncertainty may be transmitted rapidly, and amplified, with the potential to affect behaviour in a coordinated way.⁴ A restaurant may be much busier than an identical competitor if the first customers of the evening dine there and the next customers suspect the first may know something they don't, and so follow, with others subsequently joining the herd. The reverse may occur on the following evening.

Disentangling the incremental role of shifts in beliefs is often difficult. For example, a fire sale of assets by one bank may have a mechanical amplifying impact on other banks through its effect on market prices, but could simultaneously affect other banks via confidence and beliefs, if people become less optimistic about valuations of other assets when one asset class is in distress.⁵

The role of beliefs and behaviour in affecting outcomes means that there is a premium on understanding decision-making under uncertainty. In economics, broadly speaking there are two approaches to modelling decision making under uncertainty: the standard, 'rational' approach of '*homo economicus*' which assumes that beliefs are based on an efficient processing of all available information, and the 'behavioural' approach, informed by the psychology literature, which highlights psychological traits that violate the assumptions underlying the 'rational' view.

The two approaches yield different conclusions on both the desirability of policy intervention and the form it should take. In practice, the applicability of each approach depends on the context. Rationality may be a reasonable approximation when decisions are carefully considered (perhaps because the returns to a ‘good’ decision are high). But behavioural economists and psychologists contend that people often make decisions using simple or ‘fast and frugal’ heuristics because of limited time, information, and cognitive capacity (Gigerenzer and Todd 1999).

There is strong experimental evidence that people are subject to biases when collecting and processing information under uncertainty, and that such processing occurs only intermittently.⁶ The evidence suggests that people are also frequently incoherent in assessing, assigning and processing probabilities – even in highly artificial settings where information is complete – and find it particularly difficult to assess probabilities which are very low or very high, partly because there is less scope for learning in these cases. There is also evidence that, when making complex decisions or without full information, such as when deciding whether to have eye surgery, people often use rough heuristics, or go on first impressions, appearances, gut instinct or intuition.

We want to highlight four main aspects of the behavioural view:

(i) Perceptions of risk have been shown to be unduly influenced by recent or personal experience. After an earthquake, demand for earthquake insurance at first increases and then declines, while the probability of the next large earthquake may grow over time.⁷ So-called ‘Depression babies’ seem to take less financial risk than other generations.⁸ Such an effect may explain why the fear of financial crises declines over time as memories of the last crisis fade.

(ii) Decisions under uncertainty have been shown to be sensitive to the way questions are presented. People understand probabilities better if they are expressed in terms of natural frequencies (1 in 10,000) rather than percentages (0.01%). This is one symptom of a wider confusion over probabilistic concepts. Research by Gerd Gigerenzer and others shows that doctors, for example, are often misled by information presented in the form of probabilities rather than sample frequencies. The statement that a diagnostic test has a 5% probability of a false positive is misinterpreted more often than the statement that 5 patients out of every hundred taking the test will show a false positive.⁹ And this is

compounded by the fact that many forecasts are not of the kind where reference back to a sample frequency is informative. Forecasts made today of the weather or the inflation rate at some date in the future are unique; they are not the outcomes of a repeated experiment. Decisions about saving for future pension provision, for example, are difficult to make the subject of experimentation, although we can learn from the experience of others.

(iii) People tend to follow the actions of others. That can be an effective strategy. For example, on the game show “Who Wants To Be A Millionaire?”, the “Ask the Audience” lifeline has a 91% success rate, showing that there can be wisdom in crowds.¹⁰ But experiments also show that a substantial proportion of people choose a glaringly incorrect answer to a simple question when told it is the answer others have chosen (even though the same question is answered correctly in isolation).¹¹ In a sporting setting, goalkeepers facing penalties tend to dive, rather than stay in the middle of the goal, more than the direction of penalties suggests is appropriate.¹² Such types of behaviour reflect the fact that people are unwilling to risk failure by pursuing strategies that are different from the norm. The penalty of failure may be less if others fail in the same way – as Keynes pointed out was true of banks.

(iv) People have excessive faith in their own judgements and are subject to confirmation bias and wishful thinking. An experiment, well known in the literature, asked respondents to report 90% confidence intervals around our estimates of a set of unknown quantities (e.g. the diameter of the moon).¹³ But the true answer lay outside the reported confidence intervals half of the time. This reflects a general result that people are overly confident in their own judgements. This effect suggests that there may be a tendency to understate uncertainty and could explain why the ‘severe’ stress testing scenarios used by banks before the crisis were so mild.

5. Policy implications

These observations about decision making under uncertainty in economics have implications for policy along three dimensions.

First, decision making under uncertainty affects the dynamics of the economy. Some of the psychological biases discussed above can offer an explanation for ‘herding’

behaviour and ‘risk illusion’ (a tendency to over/underestimate likelihoods of different type of events). Taken together, these biases could offer an explanation for periodic bouts of exuberance in the economy and the financial system in particular.¹⁴ As such they should bear on our assessment of the likelihood of the range of possible outcomes. Our ignorance of how people behave under uncertainty, together with the presence of feedbacks between behaviour, outturns and beliefs, highlights some of the difficulties and pitfalls of economic policy making.

Second, the recognition that beliefs and behaviour play an important role in the dynamics of the system affects the type of policy actions that are desirable. Ensuring that panic did not spread is one reason why guaranteeing all bank deposits following the run on Northern Rock in September 2007 was necessary.

Third, communication is part of the policy response to uncertainty. In this respect, economic and financial systems differ from physical systems. In the latter, there is no-one listening. The control strategy in a physical system is defined by a scheme that relates the settings of policy instruments to certain features of the environment (e.g. fire left booster rocket if rocket veers to left of target). But in an economic system – because people’s beliefs about future policy actions affect their behaviour today – the control strategy has to comprise not only a plan for setting the instruments (e.g. raise Bank Rate if expected inflation rises) but also a plan to condition beliefs about how policy makers would respond in future. That involves communication.

6. Communication

A communication strategy involves deciding how much information to communicate, and in what form. Revealing information entails costs and benefits. An optimal communication strategy balances those costs and benefits.

Consider first the benefits. Information is valuable and can improve the quality of decisions. Many government agencies have valuable information that private citizens cannot easily obtain for themselves. An example is the information collected by the

Office for National Statistics on the economy and social trends. Central banks, too, may have information about the shocks hitting the economy, or about how they will be propagated, or indeed about the way in which policy is likely to respond. Provision of information about these issues may help people to reduce their forecast errors and take more informed decisions.

As we discussed earlier, complex economic systems may have multiple equilibria. Movements between them depend on expectations and beliefs about others' beliefs. That is especially so if there are strategic complementarities where the benefits of a strategy depend on the choice of strategies by others. Communication might help to coordinate beliefs so that a good rather than bad equilibrium results. Suppose that private investors begin to doubt the health of a bank which the central bank knows to be fundamentally solvent. A signal from the central bank that the bank is solvent may help to change beliefs, thus ensuring the continued liquidity and hence solvency of the healthy bank. But it may not be clear that the bank is solvent, and in practice there is often no smoke without fire.

Communication may help to improve the credibility of public policy. Economic policymakers often confront the problem that a future policy setting that is optimal today is not likely to be the policy that it would be optimal to follow tomorrow. Before central bank independence, a problem for monetary policy was that a government promise to deliver low inflation was not credible. Wage and price setters knew that if contracts were based on the assumption that future inflation would be low, the government would not be able to resist the temptation to generate inflation higher than promised because that would temporarily boost employment and output. So the initial promise was not believed. Once expectations of high inflation were embedded into prices and wages, it was at that point rational for the government to tolerate high inflation, because the cost of reducing it would be a deep recession. The idea of giving central banks independent control over monetary policy was to remove this temptation to inflate.¹⁵

Communication is a potential solution to the credibility problem. If policymakers are transparent about the objectives of a policy – and the strategy to be followed in pursuit of those objectives – it will be easier to hold them accountable for their actions. This increases the incentives for policy makers to meet their objectives.

Despite all these benefits, it does not always follow that more communication is beneficial. More communication can be costly and self-defeating. First, private participants face costs in collecting and processing information that could mean important messages are lost. The Bank of England could publish all of the many documents which it creates internally each month. But that would simply make it harder for the really important information to get through. Less can be more. We could have minute by minute revisions to weather forecasts or terrorism threat estimates, but that might overwhelm genuine flood or bomb warnings to which we want people to respond.¹⁶

Second, communication can sometimes aggravate market failures. Suppose the RSPB knows a secret location where rare birds are nesting. To preserve biodiversity, there may be a great benefit to ensuring the birds can nest undisturbed. If the RSPB reveals its secret, bird watchers may converge on the location. Here, less than full transparency mitigates the problem.

An example in the same spirit is whether or not use by commercial banks of the Bank of England's lending facilities should be concealed. If use of the facility were made public, depositors would ask themselves why it had been used: if they attached even a small probability to the possibility that the bank was in trouble, it would be rational for them to withdraw their deposits – triggering a run on the bank. Anticipating this reaction, there is a risk that banks would choose not to use the facilities. Further, some market participants might use the information that the facilities were being used to spread malicious rumours in the hope of profiting from the market reaction. Such rumour-spreading would be privately beneficial for the rumour-spreaders, but socially harmful. That is why when designing its discount window lending facilities, the Bank decided only to disclose average usage over a quarter, and then only with a lag.¹⁷

Communication of forecasts is not easy. Probabilities are hard to understand. Yet a forecast is not a number but an estimated probability distribution.¹⁸ And we cannot assume that people obtain information in a neutral form. Indeed for most people information is conveyed to them via the media which may filter out or reinterpret information. This leads to an incentive problem. Instead of simply reporting the information about future prospects made available in the forecasts, there is an incentive to summarise the prospects in a single number or to describe one part of the probability distribution in a way that makes for a good story. That suggests there is merit in trying to communicate directly to the public.

Forecasts, or probabilistic statements in general, can easily be misunderstood or misinterpreted. As a result, forecasters may decide to withhold information, even when, correctly understood, there is something useful in the forecast – a case in point being the recent Met Office withdrawal of medium term seasonal forecasts.

When experts retire from the field, it is left wide open to quackery – and there is no shortage of quacks whether in medicine, economics or meteorology. So there is a major challenge in explaining that, although we may not have a precise prediction of the effects of a particular policy action, nor are we wholly ignorant of its consequences. Finding illuminating ways of expressing those judgements is difficult but important. We must not pretend to know more than we do, and we must avoid the delusion of spurious precision.

7. Communication in practice: fan charts

It is time to put some flesh on the bones of these abstract discussions. It is time in fact for a show – a slide show of the graphical representation of the uncertainty about future inflation which the Bank of England uses to communicate its views about the outlook.

The Bank of England publishes a range of documents which describe the uncertainty about our analysis of the economy, and hence the uncertainty inherent in setting policy. Our longest-running, and perhaps the best known, is the quarterly *Inflation Report*, which the Bank has been producing since 1993. In it, the Monetary Policy Committee

communicates its views about an uncertain future using methods that draw on the ideas discussed above.

In the *Inflation Report*, the MPC sets out its latest projections for key economic variables. The most important of these is the projection for inflation. The Committee has a clear target to meet, set by government. Inflation, as measured by the change in the Consumer Price Index, should be 2% at all times. But because interest rates take time to affect inflation, forecasts of inflation are integral to setting monetary policy. And because the MPC's forecast is a probability distribution, rather than a single number, it is communicated in a fan chart. That depicts the MPC's view of the probability of a range of different outcomes. Each forecast covers a three year forecast period, because that is the horizon over which we expect the current stance of monetary policy to have the majority of its effect. Here is an example fan chart

(see <http://www.bankofengland.co.uk/publications/speeches/2010/speech432.pps>).

It shows the inflation projection published in the February 2010 *Inflation Report*. Each fan chart is conditioned on a particular path for the Bank's policy instruments. This one assumes that Bank Rate will follow the path expected by financial markets. The MPC also publishes another fan chart conditional on a different policy assumption, namely that Bank Rate remains at its current level throughout the forecast.

We believe these charts have been helpful in communicating the big picture that the MPC needs to get across. The fan chart works by showing what the Committee believes would happen to inflation if today's economic circumstances were repeated on 100 occasions. In any given quarter in the future, CPI should lie somewhere in the red shaded area on 90 out of 100 occasions: 10 times in the darkest red band, and 10 times in each pair of lighter red bands. The darker shades refer to increasing probability densities at any given point. Of the remaining 10 out of every 100 occasions inflation is expected to lie outside the shaded part of the fan altogether. But the distribution of these tail events is not explicitly specified, as to do so would require a spurious degree of precision on the part of the MPC.

The fan chart is designed to communicate specific concepts important to policy. Next, we will illustrate these using the February fan chart. But first, let's briefly recap the recent history of inflation. Here, you can see that over the past couple of years inflation has been volatile, first rising sharply to around 5% and then falling back recently <Slide 2>. Looking ahead, the MPC felt it was likely that inflation would rise again in the first quarter of this year <Slide 3>. You can see that in this projection the entire 90% of the distribution covered by the shaded fan is higher than the last data point. Remember, though, with 10% of the distribution not specified, the MPC were not ruling out the possibility that inflation might have remained unchanged or fallen.

Since this projection was published, new data have been released for inflation in January, shown by the green marker <Slide 4>. This showed that inflation had indeed risen. So it now looks even more likely that inflation in the first quarter of 2010 will be close to the middle of the MPC's fan chart. Thereafter, the MPC expected that inflation was more likely than not to fall back and be below target for a period <Slide 5>. You can see that the most likely outcome, shown in the darkest band, was for inflation to stay below the target at the end of the forecast period. But the MPC also saw upside risks to that projection, stemming from the possibility that the current high level of inflation might become embedded in inflation expectations, so making inflation itself more persistent. To make these risks easier to discern, the MPC also makes available charts of the cross section of the probability distribution at particular points on the forecast horizon, like the one shown in Slide [6]. The upside risks mean that at the end of the forecast horizon, the MPC judged that inflation was as likely to be above the target as below it. We can see this in the chart, because the narrower bands towards the left hand side mean that the mass below the vertical line marking 2% is about the same as the mass above it.

But why do we show our projections like this? Many other economic forecasters publish their forecast as a single, most likely outcome. In our view, this sort of forecast is not very informative, as the actual probability of it being correct is very close to zero. In contrast, the whole distribution contains much more information. What matters for policy is the entire distribution of outcomes. The fan charts are useful because they contain information relevant to the likely distribution of Bank Rate. To explain its

current policy setting, and to provide information on the likely path of policy, the MPC needs to communicate its judgement of the outlook as a whole.

Most other central banks also try to illustrate uncertainty around their projections, albeit in slightly different ways. The European Central Bank, for example, presents its forecasts as a range defined as twice the mean absolute forecast error over the past <Slide 7>.

The Federal Reserve produces projections that show the range of individual [most likely] forecasts – as viewed by each Federal Reserve Bank President and Governor. This is shown by the spread of the whiskers <Slide 8>. They also produce a measure of “central tendency” which truncates that distribution by excluding the three highest and three lowest forecasts, here represented by the boxes. This representation of the differences of view within the Federal Reserve is very different from a probability distribution of future outcomes. In our experience, communication of the uncertainty about future inflation is important in explaining monetary policy.

The Swedish Riksbank adopts a similar approach to ours and use fan charts as their key tool for communicating views on the outlook. <Slide 9>.

Looking back, we can see how our communication strategy has evolved in the light of experience over the past 17 years. Slide [10] shows that even in the early days, we tried to get across uncertainty around our projections, illustrating this with a “trumpet” of uncertainty in the left hand panel. It was not until 1996 that we introduced the more sophisticated fan chart, on the right, as our main device for communicating about the outlook.

As part of that evolutionary process, and to complement our projections for inflation, the MPC has also produced fan charts for activity in the economy – specifically GDP – since November 1997. Ordinarily, we have shown this as a projection for the four-quarter growth rate, which for most commentators, is the standard way of looking at GDP. Slide [11] shows the recent data for GDP, portraying a sharp decline in output over the past two years. There are two points to highlight with this chart. First, there is a fan

around the historical data <Slide 12>. That reflects the fact that economic data are frequently revised over time – sometimes heavily. CPI, in contrast, is rarely revised. The chart shows that there is quite a lot of uncertainty about the past data. The MPC believes it is more likely than not that the current vintage of data will be revised up slightly. But the big picture of a large decline in output is likely to remain even in the mature data. Uncertainty about the past is an important component of our uncertainty about the future, and is reflected in the forward-looking part of the GDP fan <Slide 13>. The second point is that the forecast part of the chart shows that a return to growth is likely. But if you just had this chart to look at you might feel that the MPC was quite optimistic, and after a painful, deep recession it was expecting the economy to rebound.

In some ways, that is true. But by itself that would be a misleading inference. The reality is that even if *growth* rebounds, the level of activity is still very likely to remain weak for a considerable period compared to the peak at the start of 2008. The economic environment is likely to continue to feel far from normal for some time. After all, it is ultimately the level of activity, rather than its growth rate, that matters for employment and the degree of inflationary pressure. It was precisely to draw out this point that recent editions of the *Inflation Report* have included an extra chart <Slide 14>. This shows the level of output corresponding to the growth chart. In practice, we have found that, even though these charts show the same information, the levels chart can be more effective at getting that big picture message across.

Despite these efforts, there are still nuances that we have found challenging to communicate. The fan chart conveys a rich view of the uncertainties in the current forecast. But it does so at the cost of it being hard to see how the distribution has changed compared with the preceding forecast. So perhaps it is unsurprising that some commentators have found it hard to grasp the subtleties of our forecasts. That has led some to suggest that we publish tables of key summary statistics on the same day as our forecasts – the mean, median and mode of the distribution.

We fear that if we make such statistics available, all of our hard efforts to communicate the outlook as a whole get washed away in an extreme focus on point estimates. We do

not want media commentators to summarise that outlook in a single number, preventing us getting the broader view across.

As well as leading to over-simplification, publishing the information can also promote over-interpretation of our forecast. Whilst the fan chart represents the Committee's probabilistic view of the future, it is important to recognise that it has only incomplete information on that distribution. We therefore should not overplay the particular parametric form of the chart. As a result, it makes little sense to place a lot of weight on particular summary statistics of the distribution. Rather, what the Committee can do is make broad statements – such as “inflation is more likely than not to be above the target”. Being more precise risks overstating the MPC's degree of knowledge about the distribution.

So the MPC's concern is that early publication of the numerical parameters risks not only trivialising the policymakers' overall view of the economy, but misrepresenting that view by being overly precise. The fan chart is designed to strike the right balance between these two hazards, by giving a broad brush articulation of the distribution that the Committee judges appropriate.

But there is always room for improvement. So we turn to some of the ideas we have considered for providing additional information on particular aspects of the distribution, to help deal with some of these difficulties. For example, the next chart shows how the distribution has changed since the previous forecast by placing the two distributions side by side, as in [Slide 15]. This chart shows that the most likely outcome in the dark red band has shifted down compared with the earlier forecast. But that change in the most likely outcome is not reflected to the same degree in the rest of the fan chart – and that is because there has been an offsetting upwards shift in the balance of risks. This is most clearly seen by noting how the dark line is flatter than the light one on the right hand side of the chart, and steeper than it is on the left.

As we discussed earlier, the aim of MPC communication is to convey useful information about the broad distribution, without being driven to spurious accuracy. There are many ways we can do this. One idea that we are investigating is to show charts of the MPC's view of the probability that inflation will be above the 2% target at a particular point in time. This is shown in [Slide 16]. Such a chart does speak to the whole distribution, and it facilitates easy comparison with the previous forecast round < Slide 17>

But it still has the problem that it assigns a precise probability to a given outcome. A solitary Bayesian policymaker may be comfortable with that. But in reality there is no unique answer that encapsulates the collective judgement of the MPC. In the Bank's experience, MPC members find it more important to focus on distinguishing between the different broad orders of magnitude of future events, and are less interested in debating whether something will happen 49 times out of 100 or 52 times.

One answer to this tension is to describe the problem in words. But that carries its own difficulties. An alternative is to design a graphical device which prevents the illusion of spurious precision. Simple graphical changes to the chart such as <Slide 18> emphasise that the MPC's view is not best expressed by a precise numerical for the probability. But a downside of this sort of representation is that it does not readily permit comparisons between successive forecasts.

Instead, a further alternative that the MPC is considering is to present this as a "probability ribbon" <Slide 19>. This shows a range estimate of the probability that inflation will exceed target, defined by the edges of the relevant band of the fan chart. In doing so it preserves the broad impression of the Committee's overall view, yet filters out the artificial accuracy. For instance, the Committee thought that the chance that inflation would be 2% or higher in 2011Q1 was between 15% and 20%.

We can once again use the February forecast to highlight a few of the key messages that the MPC wanted to communicate. In the early part of the forecast period, the probability of inflation being above the target is very high. This is because the data for inflation are

well above target at present, and also because the fan chart here is relatively narrow. But the projection for inflation falls over the forecast period. This implies a fall in the probability of inflation being above target. Towards the final part of the projection, the probability of inflation being above the target rises back to being broadly one half. This chart emphasises that, in spite of the most likely outcome for inflation being below target, the nature of the distribution of risks implies that the probability of exceeding the target is about one in two.

We can also combine the usual top-down fan chart with the information in the cross-sectional densities to produce a 3D fan chart <slide 20>. The particular advantage of this chart is that it shows how the overall density changes over time. For instance, when the distribution is very narrow, such as at the start of the forecast period, the probability density at the centre of the fan is higher, simply because the MPC is more confident about its central forecast here.

Overall, however, both from the point of view of how policy is determined, as well as how it is communicated, it is the broad distribution that matters.

8. Conclusions

Self-confidence is infectious. It can also be dangerous. How often have we drawn false comfort from the apparent confidence of a professional advisor promising certain success only to be disappointed by subsequent performance? Uncertainty pervades almost all public policy questions. Economics and many other disciplines are united by a common need to grapple with complex systems. As the crisis of the autumn of 2008 showed, such systems can sometimes be subject to abrupt changes, the precise timing of which cannot easily be identified in advance. But policy-makers are often expected to anticipate the unpredictable. How can they retain the trust of the public while being open about the true degree of uncertainty? Or, to put it another way, what is the appropriate bedside manner for policy-makers when dealing with complex uncertain problems in a public debate? We have explained how the Bank of England tries to deal with some of those challenges

in the context of monetary policy. We are sure we have much to learn, and our efforts to improve communication will continue for a long time to come.

References

Anderson, R and May, R (1991), "Infectious diseases of humans: dynamics and control." Oxford University Press.

Arrow, K (1995), Interview with "The Region" magazine, available at http://www.minneapolisfed.org/publications_papers/pub_display.cfm?id=3675

Asch, S E (1951), 'Effects of group pressure upon the modification and distortion of judgment'. In Guetzkow, H (ed.) "Groups, leadership and men" Pittsburgh, PA: Carnegie Press.

Banerjee, A (1992), "A Simple model of herd behaviour." *Quarterly Journal of Economics*, 107(3), 797-817.

Bar-Eli, M, Azer, O, Ritov, I, Keidar-Levin, Y, Schein, G (2007) 'Action bias among elite soccer goalkeepers: the case of penalty kicks', *Journal of Economic Psychology*, vol 28 No 5, pp 606-621.

Bikhchandani, S. Hirshleifer, D. and Welch, I. (1992), "A theory of fads, fashion, custom and cultural change as informational cascades." *Journal of Political Economy*, 100(5), 992-1026.

Caballero, R and Krishnamurthy, A (2008) "Collective Risk Management in a Flight to Quality Episode." *The Journal of Finance* Vol. 63 No. 5, October 2008.

Dale, S, Orphanides, A and Osterholm, P (2008) 'Imperfect central bank communication versus distraction', *IMF working papers 08/60*, IMF.

Fischer, S (1994) 'Modern central banking', in Capie, F. Fischer, S. Goodhart, C. and Schnadt, N. [eds] "The future of central banking", Cambridge, Cambridge University Press.

Gigerenzer, G and Edwards, A (2003), 'Simple tools for understanding risks: from innumeracy to insight', *British Medical Journal*, vol 327, pp 741-744.

Gigerenzer, G and Todd, PM (1999) Simple heuristics that make us smart, OUP, New York.

Haldane, A G (2009), "Rethinking the Financial Network", available at <http://www.bankofengland.co.uk/publications/speeches/2009/speech386.pdf>

Kahneman, D, Slovic, P and Tversky, A (1982), “Judgment under Uncertainty: Heuristics and Biases”, Cambridge University Press.

Keynes, JM (1936) “The general theory of employment, interest and money”, London: MacMillan (reprinted 2007).

Knight, F (1921) “Risk, Uncertainty and Profit.” Houghton Mifflin, Boston.

Kydland, F and Prescott, E (1977), ‘Rules rather than discretion: the inconsistency of optimal plans’, *Journal of Political Economy*, vol 85, pp473-490.

Malmendier, U and Nagel, S, (2007), ‘Do macroeconomic experiences affect risk taking’, mimeo.

Met Office (2010), ‘Changes to Met Office forecasts’, Press Release, 5 March

Morris, S. and Shin, H.S. (2002), “The Social Value of Public Information,” *American Economic Review* 92, 1521—1534)

May, R. and Arinaminpathy, N. (2010), “Systemic risk: the dynamics of model banking systems.” *Journal of the Royal Society Interface*. [In Press]

Palm, R (1995) ‘Earthquake insurance: A longitudinal study of California Homeowners’

Plous, S (1993), “The psychology of judgment and decision making”. Philadelphia Temple University Press

Scheinkman and Woodford (1994), “Self-Organized Criticality and Economic Fluctuations.” *American Economic Review*, Vol. 84, No.2.

Sims, C (2003) ‘Implications of rational inattention’, *Journal of Monetary Economics*, vol 50, no 3, pp665-690.

Sugihara, G and May, R (1990) ‘Nonlinear forecasting as a way of distinguishing chaos from measurement error in time series’, *Nature* , 344, 6268, 734-741

Surowiecki, J (2004), ‘The Wisdom of Crowds’.

Tucker, P (2009) “The Debate on Financial System Resilience: Macroprudential Instruments” available at

<http://www.bankofengland.co.uk/publications/speeches/2009/speech407.pdf>

Accompanying slides

<http://www.bankofengland.co.uk/archive/Documents/historicpubs/speeches/2010/speech432.pps>

Endnotes

¹ See, for example, Sugihara and May (1990).

² In relation to the latter, key nodes or players may be instrumental in the spreading of shocks (Anderson and May, 1991). This thinking has been applied to the financial system showing that a priori identical shocks can have differing effects depending on where the shock hits. And, inspired by the literature on sand piles, Scheinkman and Woodford's (1994) model of inventory dynamics shows how small shocks may cause aggregate fluctuations.

³ This particular example is fortuitous for earth science and biology, since it helps us date rocks and the fossils they contain and chart continental drift, based on the history of changes in the magnetic field of the earth.

⁴ For example, Caballero and Krishnamurthy (2008) show how liquidity shocks in the presence of Knightian uncertainty can cause investors to hoard liquid assets. This occurs because all firms assume a worst-case scenario – that they will be unable to access liquidity when required – and this causes them to hoard liquidity. Alternatively when information is dispersed across a large number of people, each with a noisy signal about economic fundamentals, outcomes may be determined by an 'information cascade', the outcome of which is hard to predict (Banerjee, 1992; Bikhchandani et al, 1992).

⁵ See May and Arinaminpathy (2010) for a model of the financial system which makes this distinction.

⁶ See the collection of articles in Kahneman, D, Slovic, P and Tversky, A (1982).

⁷ For example, see Palm (1995). She found that those most affected by earthquakes were most likely to take out insurance afterward. But this effect faded as time elapsed after an earthquake.

⁸ For example, Malmendier and Nagel (2006) report that those who experience high stock returns during their life are more likely to hold stocks and report lower risk aversion; those who have experienced high inflation are less likely to hold bonds.

⁹ See, for example, Gigerenzer, G and Edwards, A (2003).

¹⁰ See Surowiecki (2004).

¹¹ See, for example, Asch (1951).

¹² See Bar-Eli *et al* (2007).

¹³ See Plous (1993).

¹⁴ Tucker (2009) discusses how this might create a friction that could be countered through the use of so-called macroprudential instruments.

¹⁵ For an exposition of the credibility problem in monetary policy, see Kydland and Prescott (1977); for an account of the benefits of central bank independence, see Fischer (1994).

¹⁶ This idea has received theoretical attention recently following the encouragement of Sims (2003). He sketches a model in which agents are 'rationally inattentive', focusing only on information which they diagnose makes an important contribution to minimising forecast errors.

¹⁷ Morris and Shin (2002) study a case in which a central bank can only reveal a noisy signal. That signal, which is unavoidably inaccurate (e.g. a central bank forecast), is given an inappropriately high weight in private decisions because each private participant expects others to use it, and each wants to ensure that his action is not too far out of line with others'. This is an example where a public communication coordinates on an inferior equilibrium. Morris and Shin also show that public information can crowd out private information. By ensuring that all agents overweight public information in their decisions, that means that other agents observing those decisions can infer less from them about the private information others have that motivated them. In the limit, if everyone took decisions solely on the basis of public information, they could reveal nothing about some private information they may also have, but did not act on.

¹⁸ Dale, Orphanides and Osterholm (2008) illustrate how welfare can fall if the private sector overestimates the accuracy of central bank forecasts.