



BANK OF ENGLAND

Speech

Folk Wisdom

Speech given by

Andrew G Haldane

Chief Economist

Bank of England

100th Anniversary of the Bank of Estonia

Tallinn, Estonia

19 September 2018

The views expressed here are not necessarily those of the Bank of England or the Monetary Policy Committee. I would like to thank Shiv Chowla, John Lewis, Jack Meaning and Sophie Stone for their help in preparing the text and to Nicholas Gruen and Matthew Taylor for discussions on these issues. I would like to thank David Bholat, Ben Broadbent, Janine Collier, Laura Daniels, Jonathan Fullwood, Andrew Hebden, Paul Lowe, Clare Macallan and Becky Maule and for their comments and contributions.

I am delighted to be here to celebrate the 100th anniversary of the founding of the Bank of Estonia. It is a particular privilege to be giving this lecture in the Bank's "Independence Hall" – the very spot where, on 24 February 1918, Estonia's Provisional Government was formed. The founding of the Bank of Estonia followed on the Republic's first birthday in 1919.

Reaching your first century is a true milestone for any person or institution. In the UK, when you reach your 100th birthday you receive a signed card of congratulations from the Queen. I am afraid I have no royal birthday card for you today. But I have the next best thing – another speech from another central banker.

Times are tough in central banking. Central banks have borne much of the burden of supporting the global economy as it has recovered from the global financial crisis. In some advanced economies, such as the UK and US where the output gap has largely been closed, that extraordinary monetary support is now being gradually withdrawn. In others, such as the euro area, a gradual withdrawal has been clearly signposted.

The gradual withdrawal of exceptional monetary stimulus is unlikely by itself, however, to remove entirely the weight from central banks' shoulders. Many have acquired new powers and new responsibilities since the crisis, particularly in the regulatory sphere. At the same time, the crisis has dented trust in the financial sector, including in central banks. In combination that has raised expectations, but also some doubts, about the role of central banks in serving society.¹

Central banks have, over the course of the recent past, undergone a revolution in their degree of transparency and accountability. This has delivered huge gains. By improving public understanding, it has helped stabilise the economy and the financial system. And by improving public accountability, it has allayed concerns about central banks being over-powerful and under-accountable to the societies they serve.

But time has moved thing on and, as in the past, central banks will need to move with the times. In response to greater levels of responsibility, but more fragile trust, I will argue a second revolution in central bank practices may be needed. This would need to be every bit as radical as the first, but would require central banks to engage with, and draw on, the general public – their "folk wisdom" – as never before.

Central banks were put on earth to serve the public. The second revolution is about putting the public more firmly into public institutions. At a time of raised expectations and lowered trust, a reinforcement of central banks' social contract with society could make them even more effective in the period ahead, whether you are the Bank of England (into its fourth century) or the Bank of Estonia (commencing its second).

Let me start by saying a word about the central bank revolution of the recent past, before turning to the present and the future and the challenges and changes that may lie ahead in engaging the general public. The Bank of England has already begun taking important steps towards increasing its engagement. Some

¹ For example, Tucker (2018).

of the steps taken by central banks may have relevance to other public institutions, many of which face similar challenges of raised expectations and diminished trust among the general public.

A Central Bank Revolution

You do not typically associate central banks with revolution, unless the “r” is silent. Yet I do not think it is an exaggeration to say that, over the course of recent decades, a revolution has swept through pretty much every central bank on the planet. That revolution has come in the area of external communications, openness and accountability.

It came neither quickly nor easily. It was only a generation ago that, when asked by Mervyn King for a one-word piece of advice, Paul Volcker replied “mystique”. To that point, opacity rather than transparency coursed through central bankers’ veins. That bloodline ran from Montagu Norman a century ago (“never explain, never apologise”) through to Alan Greenspan a generation ago (“I’ve learned to mumble with great incoherence”).

Historically, the case for central bank secrecy was deeply rooted. It rested in part on custom and practice, in part on theology. Central banks historically were rarely subject to external scrutiny, at least by the general public. And central bankers were generally perceived as the high priests of finance. It was said that central bankers, like faith, were better understood by their deeds than by their words.²

That was just as well because central bankers rarely uttered any words at all. A century ago, public utterances by the Governor of the Bank of England amounted to one after-dinner speech each year to a well-watered set of City of London bankers. Heroically, attempts were made by some academics to rationalise this monetary mystique as a necessary means of central banks exerting influence over the economy.³

Times have changed. The central bank secrecy doctrine has been over-turned, in part for conceptual reasons, in part for practical ones. Within academia, economic theorists began constructing models where policy credibility was absolutely central if people’s expectations were to remain anchored and the economy kept stable.⁴ Establishing credibility called for clear policy targets and policy actions.

Secrecy about policy actions, whatever its potential shorter-term benefits, came at a potentially significant longer-term cost. Monetary policy “surprises” might provide a short-term jolt to the economy. These surprises were especially likely when policy was in the hands of politicians, perhaps seeking re-election. But these actions nurtured longer-term suspicions that policymakers might neglect inflation, causing inflation expectations to rise and policy credibility to fall.

² Kynaston (2017) discusses this era of the Bank of England’s history.

³ Brunner (1981).

⁴ Kydland and Prescott (1977), Barro and Gordon (1983).

If so, monetary surprises could prove self-defeating even in stimulating the economy. They would instead merely stoke expectations of higher inflation. In the academic jargon, an “inflation bias” would emerge.⁵ This theoretical story found empirical support in the behaviour of a number of advanced economies during the 1970s, who experienced a so-called “Great Inflation”.⁶

In part in response to the Great Inflation, a number of central banks began to be given new policy responsibilities to reduce the potential for electorally-driven monetary surprises. Perhaps the most important of those was operational independence for the setting of monetary policy. This was an institutional attempt to avoid the temptation to stoke the economy in the short-term, and to anchor inflation expectations in the longer-term, thereby building policy credibility.

With these new policy powers came, reasonably enough, new responsibilities for central banks. For the first time in their histories, central banks needed to be far more open and transparent about their policy targets and actions. Why? Because it was through greater transparency that policy suspicions among the public could be allayed, thereby helping anchor inflation expectations in the economy.

As unelected technocrats acting on society’s behalf, operational independence also called for new means of ensuring central banks were sufficiently accountable to society. Holding central banks’ feet to the flames called for greater public transparency and greater public scrutiny than ever previously. Trust in central banks could not be endowed by statute. It needed instead to be earned through experience and societal scrutiny.

These forces brought about a revolution in central bank practices regarding transparency and accountability. Over the course of the past century, the number of Bank of England appearances before Parliament has risen around 20-fold; the number of publications issued has risen around 600-fold; and the number of words uttered has risen around 1,000-fold. You don’t need Tracy Chapman to tell you I’m talkin’ ‘bout a revolution.

Elsewhere around the world, the pattern has been much the same. We have seen a rising tide of speeches, publications and parliamentary appearances by central banks. The Bank of Estonia has been at or close to the frontier of this revolution. Today it provides a rich daily diet of reports, forecasts, working papers, speeches and tweets.

The 20th century may yet be seen by historians as a Golden Age for central banks. Central bank numbers rose from under 20 at the start of the century to over 175 by its close. Those with operational independence for monetary policy rose from fewer than 40 in 1971 to over 150 today. And the numbers of central banks with responsibilities for regulatory policy has risen sharply too, especially since the crisis.⁷

⁵ Barro and Gordon (1983).

⁶ For example, Nelson (2005).

⁷ Haldane (2017a).

The rise in central bank numbers in part reflects the greater numbers of nation states in the 20th century. But this, too, is revealing. Establishing a central bank has increasingly been seen as an essential building block of nation state-building, a crucial element in the fabric of democratic statecraft. Indeed, that was precisely the Estonian experience, as events in this very room a century ago testify.

The Age of Innocence

This Golden Age of central banking has persisted for a century. As best we can tell, it has served economies and societies well. Prior to the crisis, the global economy entered a golden period of its own with steady growth and stable inflation, known as the “Great Moderation”.⁸ Mervyn King, Bank of England Governor at the time, gave this period a different name – the NICE (Non-inflationary consistently-expansionary) decade.⁹

It is not difficult to see why. Growth in advanced economies averaged 2.5% and inflation 2.2%. The volatility of inflation was around 80% lower, and the volatility of output 5% lower, than in the preceding two decades (Chart 1). With the economy performing NICE-ly, public trust in central banks was running high. Approval ratings for the Federal Reserve, European Central Bank and Bank of England were all at healthy levels.

The global financial crisis has ushered in a new era. Great Moderation gave way first to a Great Recession and then a Grudging Recovery. Advanced economy growth over the past decade has averaged only 1.5%. The volatility of inflation has risen by 40%, and the volatility of output growth by 60%, relative to the pre-crisis decade (Chart 1). The NICE decade has given way to a VILE (Volatile-Inflation Less-Expansionary) one.¹⁰

It should perhaps come as no surprise that one of the casualties of the crisis has been levels of public trust in the financial system, banks and central banks. Measures of satisfaction or confidence in central banks fell in the UK, US and euro area at the time of the global financial crisis.¹¹ They have yet fully to recover and may reflect diminished trust. The situation in Japan is similar, except their crisis came sooner and the bust in trust has lasted longer.

Explaining this increasingly fragile trust in central banks is, in part, a story of the crisis. Pre-crisis, central banks were seen as stewards of the financial system. A failure of that system inevitably then caused some questioning of central banks’ stewardship role. The financial crisis was the catalyst for some loss of faith in the stewards in much the same way the expenses scandal was for British political stewards a few years ago.

But there are also deeper-seated changes underway in the nature of the public’s trust, not specific to central banks, what Rachel Botsman calls the “Trust Shift”.¹² Once upon a time, public trust was endowed on public institutions rather than needing to be earned. Trust was anonymous and centralised. The public did not ask

⁸ Stock and Watson (2003).

⁹ King (2003).

¹⁰ “VILE” was coined by my MPC colleague Michael Saunders (see, for example, Giles and Atkins (2008)).

¹¹ Haldane (2017a).

¹² Botsman (2017).

difficult questions and public institutions, by and large, did not feel the need to provide answers. For public institutions, this was the Age of Innocence.

As centralised and largely anonymous institutions, central banks surfed on this wave of blind trust during their early years. During the Age of Innocence, there was no public expectation of openness or accountability or need to build bridges with the public or parliaments. When quizzed for an explanation by a Parliamentary Committee in 1930, Montagu Norman replied: “I don’t have reasons, I have instincts”.

But the public’s attitudes towards trust have altered and so too has the trust-building process. Faith in public institutions is no longer blind. Trust is no longer endowed anonymously. Instead it is built through direct connections, often personalised and localised. Trust is no longer centralised in institutions. Instead it is built on a distributed basis, peer-by-peer, across the network. In short, the Age of Innocence is over.

As the public’s trust has shape-shifted, central banks and other public institutions have faced new challenges to their trust. Indeed, this helps explain the first revolution in central bank transparency practices. But the crisis has added cyclical push to this structural fall in trust. Central banks are no longer surfing a wave. They face instead a retreating tide. For them, too, the Age of Innocence may be over.

While some loss of trust in central banks might have been expected, it is disappointing that the pre-crisis revolution in transparency practices did not provide a greater degree of insulation against it. Why was central banks’ stock of reputational capital, built steadily through various transparency initiatives, inadequate to protect them from the crisis storm? I think there were two factors at play.

First, central banks have tools, such as interest rates and banking regulations, which support the whole of the economy and financial system. They affect people’s lives greatly, supporting jobs and incomes, a stable cost of living, an adequate supply of credit, confidence in money. But the impact of central banks’ tools is often indirect and lagged; it is difficult for people to observe these benefits in their everyday lives.

When people get a job or a pay rise or a loan, they tend to attribute this to their own good luck or good management. Rarely do they link it to actions taken by central banks, months or perhaps years earlier. The benefits of central bank actions are, for most people, articles of faith rather than fact. These days, with trust needing to be earned, this puts central banks at a natural disadvantage.

We see this pattern in other professions. In surveys, the general public are typically mistrustful of politicians and bankers, who appear close to bottom of their trust league table. Yet their trust scores for *local* MPs and *local* bank branch managers have remained high because of their closer connection with people’s lives. Central banks, operating economy-wide, perhaps face some of the same challenges as bankers and politicians.

The second reason trust proved fragile is because its foundation is public understanding. Once public trust is lost, public understanding becomes crucial to restore it. The hope was that the revolution in central bank transparency initiatives had provided that improved public understanding of the economy and central banks. In practice, the evidence suggests progress towards improving the public's understanding of the economy, finance and central banks has been slow and patchy, despite these initiatives.

To start with the positives, there is concrete evidence of the improved central bank communications having improved the functioning of financial markets, the forecasting performance of economists and even the reporting of the media on the economy.¹³ The transparency revolution improved understanding among those whose day job it was to understand the economy and the role of central banks within it – the experts.

The revolution appears to have had far less impact on understanding among the public – the non-experts. In 1999, the Bank of England began surveying the public to gauge their understanding of the economy and the Bank.¹⁴ Back then, only around 7% identified the Bank's Monetary Policy Committee (MPC) as responsible for setting interest rates. 19 years and one transparency revolution later, the same survey earlier this year suggested around 6% of the general public had heard of the MPC.¹⁵

The failure to boost understanding of the economy and central banks among the general public helps explain why trust in central banks proved more fragile than expected after the crisis. It is for these reasons I have spoken about central banks facing a “twin deficits” problem – a deficit of public *understanding* and a deficit of public *trust*.¹⁶ These twin deficits are clearly complementary and mutually-reinforcing.

These twin deficits have not, so far at least, had especially serious consequences for central banks. There has been no de-anchoring of inflation expectations. There has been no rolling-back of central bank responsibilities nor systematic retreat from central bank operational independence on monetary and regulatory matters. If anything, central banks' sets of responsibilities have recently been added to.

At the same time, more questions are probably being asked about the role and powers of central bank than at any time in a generation. That is in part a reflection of the increased responsibilities central banks have been handed; in part a reflection of the higher expectations being placed on central banks' actions; and in part a reflection of the dent in trust in institutions the crisis has delivered. This is a triple challenge.

A year ago, the Bank of England hosted a conference to mark 20 years since it was granted operational independence for monetary policy. At that event, I was struck by the number of speakers calling time on the

¹³ For example, Gürkaynak, Sack and Swanson (2005), Gürkaynak, Levin and Swanson (2010), Berger, Ehrmann, and Fratzscher (2011).

¹⁴ Bank of England Inflation Attitudes Survey, currently run jointly with TNS.

¹⁵ The same survey also gives the “Bank of England” as option to the question of which group sets interest rates. More respondents identify the Bank of England as being responsible for setting interest rates (around 30%) than the Monetary Policy Committee, but this too has been broadly flat over time.

¹⁶ Haldane (2017b).

Golden Age of central banks, speculating on whether “peak central bank” had been reached. Some asked openly whether central bank independence was appropriate, at least on financial stability matters.

I am not as pessimistic or as fatalistic as that. But I do think some of those concerns about understanding, trust and legitimacy are real. In response to them, new efforts may be needed to close the understanding and trust deficits and bridge the transparency and accountability gaps. Doing so would, I believe, leave central banks even better-informed and better-placed to deliver their core mandates in future. Here is why.

Correcting the Deficits

My answer comes in two, inter-related, parts. Let me take these in turn.

First, *building understanding among the general public in the economy, financial system and central banks.*

Most people would think an improved understanding of the economy among the general public is desirable. It makes for improved decision-making by people on important everyday decisions on spending, saving, borrowing and working. The better those individual decisions (the *micro*), the more stable the economy and financial system are likely to be (the *macro*). Improved public understanding of the economy is in that sense win-win, benefitting policymakers and the public alike.

But it is possible the benefits of improved public understanding in stabilising the economy and financial system have been underestimated. “Popular narratives” among the public might be more important in driving macro-economic behaviour than in the past. And these narratives may be more prone to destabilisation than in the past. If so, too little may have been invested historically, by both the public and policymakers, in informing these narratives and hence stabilising the economy.

Public understanding of the economy and finance starts from a low base. You do not have to take my word for it. That is the strongly-held view of the public themselves. In surveys, almost half of the public say they believe they lack an adequate understanding of economic and financial issues when making crucial everyday decisions.¹⁷ This is a weak endowment.

We now know quite a bit about the costs to individuals of this weak endowment, which appear to be large and long-lasting. For example, poor levels of financial education have been shown to have large and lasting effects on individuals’ health – financial, physical and mental.¹⁸ Financial illiteracy is a recipe for debt, default and depression, whose effects appear to feedback on each another in a vicious spiral.

¹⁷ FCA (2017).

¹⁸ Richardson *et al* (2017), Earwicker (2016).

These individual costs are amplified when they are aggregated up to the macro level. How people's expectations evolve – their degree of optimism or pessimism, exuberance or depression – is crucial for determining their individual decisions. It has long been recognised that these expectations can be shaped importantly by others' expectations. For example, "popular narratives" can emerge which shape collective expectations among the public – optimism or pessimism, exuberance or depression – and which can then drive aggregate economic fluctuations.¹⁹

Recent research has demonstrated just how powerful those emergent popular narratives can be, economically and financially. As one example, the work of Nick Bloom and co-authors has constructed various measures of businesses' degree of uncertainty around key issues, based on the words used in media reporting.²⁰ These narrative-based measures of uncertainty have been found to be important for explaining the investment behaviour of companies over time.

David Tuckett and co-authors have constructed measures of the popular narratives used by participants in financial markets.²¹ They find that these narratives can play an important role in explaining asset price dynamics, over and above the impact of macro-economic fundamentals. In a similar spirit, the work of Michael Bailey and co-authors has found, using Facebook data, an important role for shared social narratives in explaining behaviour in the housing market.²²

At a macroeconomic level, the work of George Akerlof and Robert Shiller has looked at the popular narratives which emerge during periods of boom and bust.²³ Using words extracted from newspapers, they find the prevailing popular narratives about the economy have played a significant role in accounting for the heights of the peaks and depths of the troughs during macro-economic booms and busts. Public expectations, embedded in the stories they tell, are a key macro-economic driver.

The common theme in all of this research is the importance of the public's "narratives" in shaping their behaviour. That is true individually but especially collectively – hence "popular narratives" – with shifts in collective expectations helping account for the width and depth of macro-economic fluctuations. As well as contributing to this research, the Bank has begun using these "narrative" approaches itself to help understand behaviour. Let me give you a topical example.

The Bank's Agents around the UK have a wide network of company contacts who they visit regularly to discuss the economy. These visits are then written up. These write-ups enable a semantic search to be carried out to identify the key themes or narratives emerging in companies' conversations. You will be unsurprised to hear that a key theme among companies over recent quarters has been uncertainty surrounding Brexit.

¹⁹ Shiller (2017).

²⁰ Baker, Bloom and Davis (2016), Bloom (2009) and Bloom, Bond and Van Reenen (2007).

²¹ For example, Tuckett (2011) and Tuckett and Nikolic (2017).

²² Bailey *et al* (2016).

²³ Akerlof and Shiller (2009) and Shiller (2017).

Figure 1 plots a word cloud extracted from those company write-ups. The size of each word connotes the frequency of its use in relation to mentions of “uncertainty”. For a word that did not exist as recently as five years ago, it is striking how Brexit is now dominating companies’ conversations. Currently, it crops up 10 times more often than the word “customer” and almost 7 times more often than the word “staff” in proximity to uncertainty.²⁴

Chart 2 plots a time-series of the incidence of the words “uncertainty” and “Brexit/referendum” in Agents’ company write-ups. The two, unsurprisingly, are positively correlated after the referendum. Indeed, mentions of Brexit have outnumbered mentions of uncertainty by a factor of around 3 since the referendum. While mentions of Brexit have fallen from their referendum peak, it is notable that they have picked up sharply over the past few months, with measures of uncertainty following suit.

This new research strengthens the case for improving public understanding of the economy. An improved understanding could result in better-informed popular narratives, and more stable expectations, among the general. And that, ultimately, would make for a more stable economy and financial system.

In practice, there may be structural factors which could be acting in the opposite direction. The great trust shift may mean people are less willing than in the past to trust expert opinion (such as through mainstream media) and more willing to trust non-expert opinion (such as some parts of social media). Around half of those in the EU use social media as a news source and, among those aged 18-24 in the UK, around a quarter use it as their main news source.²⁵

Yet we know that information from social media sources is typically filtered and personalised.²⁶ That fits people’s (personalised, localised) preferences. But it also results in news narratives which are more likely to be self-reinforcing and self-referential than in the past. The echoes in this chamber are louder, reach further, last longer. They are also less likely to be balanced and objective. More powerful, but less balanced, popular narratives are a potentially destabilising influence on expectations and the economy.

This research points in one direction. The case for improved public understanding of the economy and financial system has always been strong. But in a world where popular narratives are even more important in driving economic behaviour, and where these narratives may be even more susceptible to destabilisation, the case for improving public understanding is significantly reinforced. This is one blade of the scissors of public engagement.

²⁴ This analysis uses company visit write-ups produced by the Bank’s Agents themselves, rather than contacts so is of course only a proxy for the latter’s views.

²⁵ Newman *et al* (2016).

²⁶ For example, Papacharissi (2002).

The second blade is *building central banks' understanding of the economy and financial system through the general public*.

This is newer, but I believe fertile, ground for central banks. Central banks pride themselves, rightly, on their technocratic skills and economic expertise. There are limits to that knowledge, as there are to knowledge in every domain and discipline. Those limits and uncertainties need importantly to be recognised when understanding the economy and when setting policy.

Nonetheless, it would be a false step to dispense with that expertise. No one seriously questions the need for heart surgeons or car mechanics, the need for specialist knowledge and experienced decision-making when repairing hearts or car parts. Human bodies and car engines are complex environments. But so too is the economic and financial system. It too needs specialist knowledge and experienced decision-makers.

It does not follow from this, however, that expertise cannot be augmented or improved on by drawing on as wide a set of opinions as possible, expert and non-expert. Non-experts can provide a different lens on problems and solutions. This diversity of thought has been found to be especially useful in complex decision environments. Colloquially, that is why we call it “folk wisdom”. Back in 2004, James Surowiecki gave it another name – the “wisdom of crowds”.²⁷ There are now many examples of collective wisdom at work.

The most celebrated is 19th century English statistician Francis Galton’s amazement at the accuracy of the average guesses about the weight of an ox among visitors to a county fair. Less celebrated (though, as it turns out, more accurate) was English magician Derren Brown’s attempt to predict the UK National Lottery results in 2009 using the wisdom of crowds. The wisdom of crowds is not magic. Subsequent research has found that it relies instead on two factors: *diversity* and *deliberation*.

To take these in turn, the value of diversity in complex decision-making is fairly well-understood. Diversity of thought simultaneously increases the chances of creative solutions being found, while reducing the chances of terrible solutions being adopted. It fattens the upper tail, and thins the lower tail, of the insight distribution. Diversity delivers higher return for lower risk.²⁸

These benefits arise, in part, from having a wider pool of ideas in which to fish and a deeper pool in which to dilute risks. But they also arise from the collective benefits of having these new ideas and risks challenged and filtered by a group. The benefits of diversity arise as much from interactions within the group, as from the new ideas each person brings. But both are important for good ideas to emerge and for bad ideas to submerge. And it is this which fattens returns and slims risks, whatever the decision.

²⁷ Surowiecki (2004).

²⁸ Haldane (2016).

These benefits have been demonstrated in a number of real-world settings. In his book *The Difference*, Scott Page goes a step further with a “Diversity Trumps Ability” theorem.²⁹ This states that the highest-performing team is very unlikely to arise from bringing together the individually-best experts. Instead it comes from combining insights from diverse sources, expert and non-expert. It comes from a cocktail of experience, some specialist, some generalist.

The returns to diversity are not limitless. When choosing the optimal asset portfolio, most of the benefits of diversification come from the first few assets. The same is true when the assets are people. The diversity benefits in moving from 2 to 4 people are large compared with the diversity benefits of going from 102 to 104 people. There are probably diminishing returns to diversity.

The second factor important for the wisdom of crowds is deliberation. Complex decisions require an investment of time, a period of deliberation. This is needed to understand a complex problem, consider it from different angles, to test it with different solutions. Solving complex problems does not in general call for instinctive, “fast thinking” of the type described by Daniel Kahneman in his famous book, *Thinking Fast and Slow*.³⁰ It instead calls for “slow thinking”, for System 2 rather than System 1 thinking.

Time is also important when harvesting the benefits of interactions within a diverse group. It takes time for groups to explore and test their ideas, winnowing out the winners and filtering out the losers. It takes time for people to change their minds. As with diversity, there are decreasing returns to deliberation time. Research suggests that spending as little as 1 hour on a complex problem can improve performance by 10%.³¹ These benefits are likely to be smaller when moving from 100 to 101 hours spent on the same problem.

While diversity and deliberation are the cornerstones of the wisdom of crowds, it is clear that crowds are not always and everywhere wise. I am rarely much the wiser about the parentage of a football referee for having listened to the collective chants of the crowd. This observation is not a new one; it has a long historical pedigree. In 1841, Charles Mackay published a now-famous book whose title stands in stark contrast to Surowiecki’s – *Popular Delusions and the Madness of Crowds*.³²

Many subsequent authors have provided examples of collective irrationality in crowd decision-making. Its causes have also been extensively studied by psychologists, sociologists and economists.³³ This originates in the contagious spread of opinion through networks. The larger and more connected the network, and the less well-informed its participants, the greater the chances of irrational epidemics spreading.

The behaviour underpinning this group psychology is as old as humankind. The desire to conform to group type has its roots in our hunter-gatherer past. It has been found, time and again, in studies of human

²⁹ Page (2008) and Hong and Page (2004).

³⁰ Kahneman (2012).

³¹ Tetlock and Gardner (2015).

³² Mackay (1841).

³³ For example, Janis (1972), Turner and Pratkanis (1998)

behaviour from the South Sea Bubble to the dotcom bubble, from the Bay of Pigs to the space shuttle disasters.³⁴ This collective irrationality sometimes goes by the name “groupthink”.

More recently, these psychological roots may have been fed and watered by technological advances, such as social media. This has increased our digital connectivity. For reasons set out earlier, it has probably also increased the chances of self-reinforcing and self-referential waves of collective irrationality taking hold.³⁵ If so, these waves of collective irrationality could be becoming larger than ever. While a vestige of our hunter-gatherer past, the madness of crowds can these days spread to a global village.

That leaves us with a conundrum. If decision-makers in complex environments were to draw on a wider set of non-expert opinion, what are they likely to encounter? Is it a case of harnessing the wisdom of crowds or avoiding their madness? Will greater public engagement help inform and improve expert decision-making by widening the pool of lived experience? Or will it hinder that decision-making by polluting the information pool with irrational or irrelevant opinions?

Perhaps the most comprehensive study of the value of expert and non-expert opinion, as it applies to complex decision-making in the social sciences, has been provided by political scientist Philip Tetlock. Over a more than 20-year period, he has researched forecast accuracy and, most recently, run *The Good Judgement Project*. This explored the performance of a set of experts and non-experts in forecasting the outcome of a particularly complex set of decision problems.

So what does Tetlock’s evidence suggest? That the usefulness of non-expert opinion in complex decision settings depends on how exactly this information is drawn in and drawn on. From Tetlock’s book with Dan Gardner in 2015, *Super-forecasting*, I would highlight two key findings. These chime with the evidence outlined earlier on the costs and benefits of expert and non-expert opinion in collective decision-making.

First, there *is* real value in non-expert opinion when it comes to complex forecasting tasks. In a famous example, Tetlock convened a group of 1,000 amateur forecasters and entered them in a prestigious forecasting tournament run by the US intelligence agencies (IARPA).³⁶ This team beat the control group by 60% in year 1 and 78% in year 2. It not only beat expert academics from MIT to Colombia, but professionals from the US intelligence services with access to classified information.

This is what Tetlock calls *Super-forecasting*. The incremental value of non-expert judgements has been shown in a wide array of settings and for a wide range of complex problems. What about judgments on the complex dynamics of the economic and financial system? While we do not have a structured experiment on which to draw, we do have illustrative evidence on the relative forecasting performance of different groups.

³⁴ For example, Roubini and Mihm (2011) and Janis (1972).

³⁵ Shafik (2017) discusses the impact of the rise of social media.

³⁶ Tetlock and Gardner (2015).

Chart 3 considers the forecast errors of five groups – three “expert” (MPC, a panel of professional forecasters and financial markets) and two “non-expert” (companies and households). For each group we consider their forecasts for UK inflation 1, 2 and 3 years ahead. The sample period is relatively short so our conclusions need to be cautious. And forecast accuracy is not the same as policy acumen. Nonetheless, Chart 3 suggests that, in a complex forecasting task, there is little to choose between experts and non-experts.³⁷

Tetlock’s second key point, which qualifies his first, is that not everyone can be a super-forecaster. A randomly-chosen panel of non-experts cannot by itself be expected to outperform the experts. Chart 3 demonstrates that. Super-forecasting is not simply a law of large numbers. Tetlock shows that non-expert super-forecasters have particular characteristics, including inquisitiveness and independence of mind. Only those with these characteristics consistently outperform the experts.

We do not have a clean test of this hypothesis for our inflation forecasts. But just imagine the Bank had been able to identify in advance one “super-forecaster” among the pool of outside forecasters. How much would that have improved its forecasts over the sample period? The gap in two year-ahead forecast errors between the MPC and the best external forecaster in the panel of outside forecasters is around $\frac{1}{4}$ percentage point. 20/20 hindsight is a wonderful thing, but this gives some measure of the scope for forecast improvement.³⁸

If we take Tetlock’s points together, they suggest that there is real value in seeking the views of non-experts to help inform difficult decisions. But there are limits to these benefits. It is unrealistic to expect everyone to become a super-forecaster. For central banks the message is that, by widening their information net, they could stand to benefit from the wisdom of crowds. But there are limits to how and to how far you would wish to widen the net if the crowd is to add signal rather than noise.

To illustrate this point, consider the following thought-experiment. Imagine you have an endowment of time – 10,000 hours – to make as informed a decision as possible on a complex question. I have not chosen 10,000 hours randomly. In his book *Outliers*, Malcolm Gladwell argues, based on a detailed set of case studies, that it takes around 10,000 hours for one person to command true expertise in a subject area.

Now let’s choose a complex, and topical, question: will Artificial Intelligence (AI) cause mass unemployment? It is well-known that the views of experts (and non-experts) differ widely on this complex question and that the outcome is very uncertain. The problem we need to solve is not this but a simpler one: how many people should we choose to help inform our view on the impact of AI on unemployment?

³⁷ The conditioning assumptions (such as the assumed path of interest rates) and purpose of forecasts will vary between central banks and other forecasters.

³⁸ Mankodi and Pike (2018) discuss whether central bankers can become superforecasters.

One possible answer to this question is one person. With Gladwell's 10,000 hours under their belt, this person would be as well-informed on the impact of AI on unemployment as anyone on the planet – a true expert, perhaps even a phenomenon. This is the single person or, if you like, Oracle (in the Ancient Greek sense of the word) solution to this decision problem.

But is it the optimal one? There are reasons for doubt. Evidence clearly suggests that, in a complex decision environment, drawing on non-expert opinion can provide added insight. It offers diversity of perspective, whose benefits are amplified if those diverse perspectives are discussed and debated to arrive at collective conclusions. This is Surowiecki's "wisdom of crowds" at work.

So let's go to the other end of the spectrum and draw instead on the insights of a big crowd – say, 1 million people, roughly the adult population of Estonia. This would then, in effect, be an Estonian referendum on AI and unemployment. Clearly, we would now be drawing on a vast array of expert and non-expert opinion. But that would come at the expense of less time for deliberation. Given our time quota, each person would now have only around 30 seconds to give an answer.

Facing a tight time constraint and a very complex question, most people would struggle. I certainly would. The best a non-expert could probably do is to provide an instinctive, "fast thinking" response. Or, perhaps more likely still, they might jump aboard a collective, but potentially poorly-informed, bandwagon by copying what others say. This would increase greatly the chances of Mackay's "madness of crowds".

Finally, let's try an interior solution – say 1,000 people each spending 10 hours. This is large enough a number to gather most of the benefits of diversity, but small enough to harvest the benefits of deliberation. It would blend expert and non-expert opinion in a way which would allow collective deliberation. This is the "wisdom of crowds" solution.

Figure 2 plots the relationship between insight and the numbers of people providing information in a group, in stylised terms. You might call it the "Insight Curve". The curve (shown in dotted lines) results from combining the effects of a downward-sloping "deliberation" curve and an upward-sloping "diversity" curve. Consistent with the evidence, both exhibit diminishing returns. The inverted U-shaped "Insight Curve" identifies the optimal number of people to draw on to inform a complex decision.

The two corner solutions – the "Oracle" (single person) and the "Madness of Crowds" (population) – involve very different trade-offs between diversity and deliberation. A single source of information gives you plenty of time for deliberation – a hermit-like 10,000 hours. But it leaves you short of diversity of thought. A population-wide pool of information gives us as much diversity as you could ever wish. But it leaves you long fast-thinkers, short slow-thinkers and vulnerable to irrationality.

Both are inferior to an intermediate solution – the “Wisdom of Crowds” solution. This involves combining the insights from both experts and non-experts, to reap the benefits of diversity of thought and action. But this larger pool is not so large as to dilute the time available for slow-thinking and deliberation. It is a cocktail of expert and non-expert opinion, comprising diversity and deliberation in equal measure, that cause good ideas to flourish and bad ones to perish.

A relevant, practical question is where the optimal point along the insight curve lies. What is the optimal number of people from whom to seek views if we are to maximise the benefits of folk wisdom? My stylised example cannot provide an answer to that question. But some studies have looked at this question using experimental methods and, reach interesting conclusions.

Christian Wagner and Tom Vinaimont ask two questions.³⁹ First, how large does a group of non-experts need to be to exhibit expert-like performance? Second, how large does this group need to be to outperform consistently a group of experts? On the assumption experts’ views are ten times more precise than a non-expert, their answer to the first question is around 30 and the second around 1,000. This bookends the range of non-expert views we might need in seeking the “wisdom of crowds”.

Representative Decision-Making

If we take this analysis at face value, the question is how practically central banks can engage a wider, non-expert, audience to help inform their views on the economy and financial system? What new infrastructure or architecture is needed for collecting and assimilating these views? And how do we ensure they balance diversity and deliberation so as to be representative of the general public’s views?

These are deep questions, but they are not new ones. They have been debated by philosophers and political scientists for thousands of years. Alternative models of representative decision-making have been tried and tested throughout the course of history. This experience contains some useful lessons for central banks, and other public institutions, seeking to strengthen the infrastructure linking them to wider society.

Let’s start at the very beginning. The birthplace of democracy is, for many, ancient Greece. Around 2,500 years ago, the city-state of Athens pioneered a particular form of representative decision-making when it came to judgements on some of the most complex, societally-significant, issues of the day. The centrepiece of this decision-making infrastructure was a public body – a Council of Five Hundred, or *boule*.

The *boule* comprised individuals drawn from the general public by random selection (or sortition). They served for a year. The purpose of the *boule* was to debate key issues affecting public life and devise policy proposals. These were then taken to a public assembly for discussion, where a wider set of the general

³⁹ Wagner and Vinaimont (2010).

public participated, often running into the thousands. For a small number of positions requiring specialist skills, public officials and experts were elected or appointed.

The Athenian *boule* was not democratic embroidery, like a modern-day focus group. It had real power. It decided which proposals were taken to the public assembly and which were lawful. Brett Hennig calls it the “nerve centre of power” in ancient Athens.⁴⁰ The *boule* had wide public acceptance and attendance. Participation in public life, through the *boule* or public assembly, was an accepted part of Athenians’ civic responsibilities. Those not taking part were termed *idiotes*, from which the modern word idiot derives.

The infrastructure was far from perfect; the *boule* excluded women, slaves and non-citizens. But it was notable in seeking to strike a balance between the two features subsequent research has shown to be crucial for effective, representative decision-making in complex environments: diversity and deliberation. Indeed, for that reason some have called the Athenian model of decision-making a *deliberative democracy*.⁴¹

Diversity of decision-making was achieved by calling on a rotating, randomly-selected pool of non-expert citizens, augmented by experts in a limited number of positions. Deliberation in decision-making was achieved by restricting the number of representatives, having them convene and discuss regularly and agree on a set of proposals. The Council of Five Hundred sits slap bang in the middle of the optimal range of people suggested by Wagner and Vinaimont.

The Greek Empire of course did not last. Nor did its model of deliberative democracy for decision-making. In the period since, two alternative models of representative decision-making have risen to prominence. Interestingly, these seek to strike a rather different balance between diversity and deliberation, both relative to the Athenian model and especially relative to each other.

One is *executive democracy*. Under this model, an elected or appointed set of representatives take decisions on society’s behalf, often through a Parliament or Congress. This model underpins modern democracies in a great many countries; it has grown in popularity for much of the past 200 years. Such has been its success, many people these days would define democracy in terms of this model.

For most of history, this was not the case. Aristotle was not alone in seeing an executive democratic model as an “elected oligarchy”.⁴² At the time many executive democracies were being put in place, their Founding Fathers (they were all still men) were often deeply suspicious of Athenian-style democracy. John Adams, second President of the United States, captured it thus: “Remember democracy never lasts long. It soon wastes, exhausts and murders itself”.⁴³ The Greek Empire lasted only 350 years before exhausting itself.

⁴⁰ Hennig (2017), Van Reybrouck (2016).

⁴¹ Gruen (2018), Fiskhin *et al* (2008).

⁴² *Politics*, Aristotle.

⁴³ From John Adams to John Taylor, 17 December 1814.

While Adams was wrong about the durability of executive democracy, suspicions about whether it properly represents all slices of society have persisted. Indeed, these suspicions have risen recently. At the same time as the steady ascent in the numbers of countries adopting the executive democratic model, there has been a steady descent in public trust in the representatives of executive democratic institutions.

At the heart of those trust problems are concerns about whether decision-making is sufficiently reflective of wider society. By and large, the public accept that the executive democratic model is a reasonable vehicle for discussion and *deliberation* in decision-making. Indeed, the public can watch those deliberations on TV. But there is a concern among the public about whether decision-making and decision-makers are sufficiently *diverse*, whether they properly reflect and represent societies' views and needs.

This is not an accusation that could be levelled at the alternative model of representative decision-making – so-called *direct democracy*. This puts voice and vote firmly in the hands of public on decisions with an important bearing on wider society. This model has been used to varying degrees across various countries at various times. In some countries, public referenda are reserved for decisions whose societal impact is especially large and long-lasting. I am sure we can all think of examples.

This model is recognised as having both pros and cons. It draws in a much wider range of expert and non-expert opinion, so achieves a potentially greater degree of representativeness and diversity than alternatives. This very breadth, however, makes it difficult to have a deliberative dialogue in which participants have the time to consider carefully the facts, competing hypotheses and possible solutions. This risks the collective view being less well-informed, and potentially more polarised, than would be ideal.

With degrees of discontent about both the executive and direct democratic models, there has recently been a rekindling of interest in the Athenian, or deliberative democracy, model. This model, at least in principle, offers greater diversity than the executive democratic model. It also offers greater scope for debate and deliberation than the direct democratic model. It is not a third way, as it was in fact the first. But it is a middle way. It is in some ways the institutional incarnation of the “Insight Curve”.

Interestingly, a number of countries have over recent years begun experimenting with deliberative democratic processes to debate and resolve some of societies' most complex problems. Citizens' panels or juries have been used to tackle such varied issues as voting reform (in Canada), gay rights (Ireland), reconciliation between indigenous and non-indigenous communities (Australia), nuclear power (South Korea), membership of the euro (Denmark) and environmental pollution (United States).⁴⁴

Here in Estonia, you have experience of using people's assemblies to debate key societal issues. In 2013, a people's assembly (*Rahvakogu*) was used to debate Estonian electoral laws. Last year, a people's assembly was used to deliberate on the future of ageing and pension reform. People's juries have also been

⁴⁴ See <https://cdd.stanford.edu/deliberative-polling-timeline/>

used to debate regional issues, including water transportation in the Emajõgi River region of the country. Estonia has been a relative hotbed of activity on deliberative decision-making in practice.

Each of these experiments used slightly differently approaches. This has provided useful case law on which to decide what works and when. One interesting finding is that citizens' panels can be effective not just in getting people to engage on issues of national or regional importance, but in providing an effective forum for people to change their minds, and converge in their opinions, on these issues. They can reduce polarisation. These benefits lie at the heart of the deliberative democratic process.

A second lesson is that technology is a potential game-changer for deliberative decision-making. This relies on a distributed, modular structure. Historically, structures such as the size of the *boule* placed physical constraints on the scope for information-making and debate. Digital structures loosen those constraints, allowing longer reach and greater connectivity at greater speed. Some countries have begun using digital technologies to support their deliberative processes, including here in Estonia.

The benefits of doing so are partly practical, with greater numbers of people able to connect and communicate than ever previously. The public assemblies of Ancient Greece brought together people in their thousands. Today, technology could do so for millions of people. There are also, however, behavioural benefits. The "trust shift" among the public means credibility and understanding is more often these days built through distributed means and direct connections. A digital infrastructure can help those ends meet.

It could be argued that these country-level experiments are just that – experiments. The deliberative model has not been tried and tested in anger for the full gamut of societal decisions. As put, that is true. Yet there is at least one domain of public policy where this model has been tried and tested (both literally and metaphorically) for a wide range of societal decisions over many centuries: the judicial system.

Many judicial systems draw on both expert (judges) and non-expert (juries) opinion, like the Athenians. Non-experts are drawn randomly from the population, as in Athens. There is a clear process for deliberation among participants, as in the *boule*. This system is the nerve-centre of decision-making on a defined set of issues, as in Ancient Greece. The English judicial system has already lasted for almost 1,000 years, more than twice as long as the Greek Empire.

The Next Revolution

Against this backdrop, what concrete steps might central banks take to build an improved infrastructure for engagement with the general public? Can it build on others' experience in engaging wider society, in ways which improve the public's and policymakers' understanding of the economy? Let me offer some suggestions, which are motivated by recent initiatives at the Bank of England aimed at doing just that.

(a) *Minding Your Language*

Historically, the language of central banks has often been pitched at a level which has made it difficult for most non-experts to understand. This has contributed to the public's understanding and trust deficits. Slowly but surely, that is beginning to change. Central bank's external communications are being reoriented towards a wider, non-specialist audience. Let me give a couple of examples.

The first is so-called "forward guidance" about monetary policy.⁴⁵ How you judge its success hinges crucially on what you consider to have been its objective and audience. If you are a financial market participant whose job it is to price the future path of interest rates, the most useful type of policy guidance is precise and time-specific. By contrast, if you are a company or household considering whether to spend, a general idea of the direction and destination of interest rates is likely to be sufficient.

The critics of forward guidance have tended to be financial market participants for whom it lacks sufficient precision and time-specificity. But forward guidance simply cannot be that precise; it depends on the path of the economy. And pricing interest rates in financial markets was not the main purpose of forward guidance in the first place. Its purpose was to support spending in the economy by households and companies. That calls for shorter and simpler guidance, focussed on the broad direction and destination of interest rates.

The MPC first used the words "limited and gradual" in 2014 when describing the likely future course of interest rates rises. It was short and simple – three little words. It was guidance aimed squarely at households and companies, offering them an indication (but no promise) of the broad direction and destination of interest rates. Surveys suggest this message was understood by a majority of companies and some households.⁴⁶ By allaying fears about too-rapid a rise in rates, it is likely to have encouraged spending and supported the economy.

When the MPC did come to raise interest rates, in November 2017 and again in August 2018, it is interesting to see how well these were understood by companies and households. Around three-quarters of households, and around 90% of companies, had expected a rise in rates within the year ahead in surveys held immediately prior to the decisions. Simple, directional forward guidance on monetary policy appears, for this non-expert audience who comprise most of the spending in the economy, to have been effective.

A second monetary policy example comes from the Bank's quarterly *Inflation Report*. This was first produced over 25 years ago, as part of the first wave of the transparency revolution. It is a long and technical document, with language requiring at least 13-14 years of education to understand.⁴⁷ The complexity of its language meant the *Report* was probably only accessible to around 10% of the population.

⁴⁵ Bank of England (2013) provides a useful summary of concepts and issues relating to forward guidance; see Carney (2018) for a more recent discussion.

⁴⁶ Carney (2018).

⁴⁷ Haldane (2017a).

In November 2017, the Bank began publishing two simpler, “layered” versions of its *Report*. Layer 1 is the single sentence/single graphic version, pitched at a level which is accessible to around two-thirds of the population. Layer 2 is the single page/simple graphics version, accessible to over 40%. Measured in web hits and downloads, these layered communications appear to have broadened notably the reach of the MPC’s monetary policy messages, in particular among a non-expert audience.

A couple of recent research studies have used experimental methods to look more closely at the impact of the Bank’s simpler communications on public understanding and trust.⁴⁸ For example, in a randomised control trial of over 2,000 members of the general public, a recent study found that use of the visual summaries increased the public’s understanding of the *Inflation Report*’s messages by around 25%. This is a very substantial gain in public understanding for a fairly modest change in format.

This study went, however, one step further and asked whether further improvements in understanding were possible. The format they considered personalised the language, making it more relatable to people’s everyday lives – for example, by greater use of the first and second person (us/you), greater use of commonly-used words (“rising prices” rather than “inflation”) and with more engaging, interactive graphics. Doing so increased the gain in public understanding by 40%.

There are limits to what simplification, personalisation and visualisation can achieve. The same study assessed how these changes in format affected the general public’s trust in the Bank of England’s policies. These gains were modest. Perhaps that should come as no surprise. Understanding is a necessary but not sufficient condition for trust-building. And trust tends to be earned slowly from repeated experience.

Nonetheless, these studies point to an important role for simpler, personalised communications in lowering the twin deficits. There probably further for this agenda to run. For example, there is no reason in principle why the same methods could not be applied to other central bank publications, covering the financial system and banknotes. The Bank has already made some strides in this direction through its *Financial Stability Report and Knowledge Bank*, aimed at improving public understanding of the economy and finance.

Looking ahead, in my view greater personalisation of central bank communications holds promise. Central banks often affect people’s lives in ways which are lagged and indirect and hence difficult for them to observe. Earlier this year, I constructed “personalised” scorecards, which set out the impact of the Bank’s monetary policy easing actions after the global financial crisis on the income, wealth and well-being of citizens across the UK.⁴⁹

This recognised that the impact of monetary policy on people’s lives is often quite different depending on their personal circumstances – old versus young, borrower versus saver, home-owner versus renter, North

⁴⁸ Haldane and McMahon (2018) and Bholat *et al* (2018).

⁴⁹ Haldane (2018a).

versus South. Casting the impact of monetary policy in terms of averages or aggregates tends not to resonate with people's personal experience. Spelling out the benefits in personalised terms can potentially increase understanding and awareness of central banks' actions.

There is already a precedent for this in another arm of UK policymaking – tax policy. Taxpayers in the UK receive a personal scorecard, telling them how their tax payments have funded spending on various public goods at the macro-economic level. This personalisation has been found to have improved the public's understanding and acceptance of tax policies.⁵⁰

(b) *Public Education*

A second way of improving public understanding is through improved public education. Many central banks have made efforts in this direction over a number of years, including through school visits, schools competitions and schools materials, including at the Bank of England and Bank of Estonia. That does not yet appear to have satiated demand from pupils and their parents, almost 80% of who say they think economics should form part of the curriculum.⁵¹

In part in response, earlier this year the Bank of England launched a new education strategy. There are a number of elements. One is a new set of schools competitions, aimed at reaching a broader range of students. A second is a schools ambassador programme, with Bank officials giving 200 talks in schools across the UK this year. We are on schedule to hit our target.

The most ambitious part of our education strategy has been to develop classroom materials on the economy and financial system for use in schools as part of the curriculum for pupils aged 11-16. These go by the name econoME. As the name suggests, and like our recent public communications efforts, these materials describe the economy and finance in terms which are personal and relatable to young people's lives. Why does the economy and finance matter to me? And how do my decisions in turn affect the economy?

We set ourselves a target of reaching 400 state schools with econoME during the course of this year. So far since launch in April, over 1,000 schools have downloaded the materials, conceivably covering around 90,000 pupils.⁵² This demonstrates the potential pent-up demand for school materials on economic and financial issues. The Bank's aim is to reach close to 1,500 schools by mid-2020.

We are currently considering where next to take our education initiatives. A promising avenue would be a younger age range, say 7-11 year olds. This will require different materials and possibly a different approach. A more challenging market still would be adults, perhaps linked to existing initiatives to boost

⁵⁰ Haldane (2018a) and Barnes *et al* (2017).

⁵¹ ING-Economics Network Survey of the Public's Understanding of Economics (2017).

⁵² Assuming 30 students per class and 3 classes per school.

economic and financial literacy in the public at large. The scale of the public understanding deficit means central bank efforts can only be part of the solution but, as public institutions, an important part.

(c) *A Distributed Architecture*

Central banks historically have tended to be, well, centralised. But in a world where understanding of the economy is often localised and trust is built on a distributed basis, a centralised infrastructure may be far from ideal. What is needed instead is a *distributed* architecture for engagement. In response, central banks have over time changed their structures and further change seems likely.

The Bank of England has had a network of Agents around the UK for around 90 years. Today, the Bank's twelve regional agencies cover all corners of the UK, with an extensive network that generates around 9,000 meetings with company contacts each year. This is a valuable source of intelligence on the economy and financial system. Indeed, I would say it is increasingly valuable. It provides colour and context to accompany our data and models. It offers folk wisdom, the wisdom of the company crowd.

But companies are only one crowd among many. Companies are an important lens on the economy, but far from being the only one. Reflecting that, the Bank has recently sought stronger connections with some other crowds, crowds with a different lens on the economy with which it has perhaps had less contact historically – trades unions, charities, community organisations and the like.

For more than a year, I have augmented those Agents' efforts with my own "Townhall" meetings across the UK.⁵³ By design, these seek out the views of a wider range of citizens, in a wider range of locations on a wider range of issues around the economy and financial system than has been the case in the past.

These events have established for me that there is wisdom in these crowds when gauging some of the key issues shaping the economy and financial system. The lived experience of everyday people making everyday decisions is what makes the economy tick (and sometimes tock). Tapping that lived experience can add to the Bank's understanding of the economy and help it when setting policy to keep the economy stable. There is wisdom in crowds of citizens, as there is in crowds of companies.⁵⁴

The Bank is now taking a further step towards putting in place an infrastructure to formalise its engagement with citizens. Earlier this year, in response to a recommendation from the RSA, we announced that the Bank would be setting-up a set of Citizens' Panels across the UK.⁵⁵ In the remainder of this year, we will be running trials of these citizen panels to help us decide how best to organise them, before rolling them out systematically, UK-wide, next year.

⁵³ See <https://www.bankofengland.co.uk/outreach>

⁵⁴ Taylor (2018) also discusses the benefits of deliberative democracy in the RSA's 2018 Annual Lecture.

⁵⁵ Haldane (2018b).

There are a number of key design questions that need to be answered. How do you ensure as wide, diverse and representative a set of citizens as possible? How do you engage citizens who might not understand or trust the Bank? How much time needs to be set aside to enable the deliberative process to be fruitful? And how many people should you engage to get the balance right between diversity and deliberation? We are working with a range of partners, including the RSA, to come up with answers.

The plan would be to set up one citizen panel for each of the Bank's 12 regional agencies. With roughly 25 citizens per panel, this would give us a non-expert panel of around 300 people. This is edging up towards the mid-point of the Wagner/Vinaimont range and just shy of the Athenian *boule*. By using digital technology, we would aim to connect these citizen panels to allow national, as well as regional, deliberation.

This is deliberative democracy in practice. As with the judicial system, it will involve a blend of expert and non-experts. As with juries and judges, there will be a structured process for deliberation. As in the legal system, the role of the non-experts (citizen panels) is to help establish the facts and of the experts (MPC) to set policy. What's good enough for decisions on incarceration is good enough for decisions on inflation.

Conclusion

Central banks, like all public institutions, were created by the public to serve the public. Despite a communications revolution, public understanding of central banks remains low and public trust in central banks has fallen. Improving both calls for a second revolution, focussed on educating and learning from the general public. The Age of Innocence may be over, but the Era of Engagement has only just begun.

This new era may mean 21st century central banks becoming more disburbed and personalised in their communications, creating new infrastructure to interact with the public and extract their folk wisdom. Thomas Jefferson said democracy is something you can only learn by doing. The same is true of central banking. For the Bank of Estonia entering its second century, and for the Bank of England in its fourth, there is much still to learn and even more to do.

References

- Akerlof, G and Shiller, R (2009)**, *Animal Spirits: How Human Psychology Drives the Economy and Why It Matters for Global Capitalism*, Princeton University Press.
- Bailey, M, Cao, R, Kuchler, T and Stroebel, J (2016)**, 'Social Networks and Housing Markets', *NBER Working Paper*, No. 22258.
- Baker, S, Bloom, N and Davis, S (2016)**, 'Measuring Economic Policy Uncertainty', *Quarterly Journal of Economics*, Vol. 131, No. 4, pp. 1593-1636.
- Bank of England (2013)**, 'Monetary policy trade-offs and forward guidance', August 2013.
- Barnes, L, Feller, A, Haselswerdt, J and Porter, E (2018)**, 'Information, Knowledge, and Attitudes: An Evaluation of the Taxpayer Receipt', *The Journal of Politics*, Vol. 80, No. 2.
- Barro, G and Gordon, D (1983)**, 'Rules, discretion and reputation in a model of monetary policy', *Journal of Monetary Economics*, Vol. 12, No. 1, pp. 101-121.
- Berger, H, Ehrmann, M and Fratzscher M (2011)**, 'Monetary Policy in the Media', *Journal of Money, Credit and Banking*, Vol. 43, No. 4, pp. 689-709.
- Bholat, D, Broughton, N, Parker, A, Ter Meer, J and Walczak, E (2018)**, 'Enhancing central bank communications with behavioural insights', *Bank of England Staff Working Paper*, No. 750.
- Bloom, N (2009)**, 'The Impact of Uncertainty Shocks', *Econometrica*, Vol. 77, No. 3, pp. 623-685.
- Bloom, N, Bond, S and Van Reenen, J (2007)**, 'Uncertainty and Investment Dynamics', *Review of Economic Studies*, Vol. 74, pp. 391-415.
- Botsman, R (2017)**, *Who Can You Trust? How Technology Brought Us Together – and Why It Could Drive Us Apart*, Portfolio Penguin.
- Brunner, K (1981)**, 'The Art of Central Banking', *Centre for Research in Government Policy and Business Working Paper*, No. 81-6.
- Carney, M (2018)**, 'Guidance, Contingencies and Brexit', speech available at <https://www.bankofengland.co.uk/speech/2018/mark-carney-society-of-professional-economists-annual-dinner>
- Earwicker, R (2016)**, 'The impact of problem debt on health – a literature review', *EU Joint Action Programme on Health Inequalities*.
- FCA (2017)**, 'Understanding the financial lives of UK adults', *Financial Conduct Authority*.
- Fishkin, J, Luskin, R, Panaretos, J, Siu, A and Xekaliki, E (2008)**, 'Returning Deliberative Democracy to Athens: Deliberative Polling for Candidate Selection', paper presented at meeting of the *American Political Science Association*.
- Giles, C and Atkins, R (2008)**, 'Doubt cast on Bank's inflation credentials', article in *Financial Times*, 19 May 2008.
- Gruen, N (2018)**, 'An Ancient Greek idea could foil Brexit's democratic tragedy', article in *The Guardian*, 16 January 2018.
- Gürkaynak, R, Sack, B and Swanson, E (2005)**, 'Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements', *International Journal of Central Banking*, Vol. 1, No.1, pp. 55-93.

Gürkaynak, R, Levin, A and Swanson, E (2010), 'Does Inflation Targeting Anchor Long-Run Inflation Expectations? Evidence from the U.S., UK and Sweden', *Journal of the European Economic Association*, Vol. 8, No. 6, pp. 1208-1242.

Haldane, A (2016), 'The Sneetches', speech available at <https://www.bankofengland.co.uk/speech/2016/the-sneetches>

Haldane, A (2017a), 'A Little More Conversation, A Little Less Action', speech available at <https://www.bankofengland.co.uk/speech/2017/a-little-more-conversation-a-little-less-action>

Haldane, A (2017b), 'Everyday Economics', speech available at <https://www.bankofengland.co.uk/speech/2017/andy-haldane-speech-during-regional-visit>

Haldane, A (2018a), 'How Monetary Policy Affects Your GDP', speech available at <https://www.bankofengland.co.uk/speech/2018/andy-haldane-david-finch-public-lecture-melbourne>

Haldane, A (2018b), 'Climbing the Public Engagement Ladder', speech available at <https://www.bankofengland.co.uk/speech/2018/andy-haldane-royal-society>

Haldane, A and McMahon, M (2018), 'Central Bank Communications and the General Public', *AEA Papers and Proceedings*, Vol. 108, pp. 578-83.

Hennig, B (2017), *The End of Politicians: Time for a Real Democracy*, Unbound Digital.

Hong, L and Page, S (2004), 'Groups of diverse problem solvers can outperform groups of high-ability problem solvers', *Proceedings of the National Academy of Sciences of the United States of America*, Vol. 101, No. 46, pp. 16385-16389.

Janis, I (1972), *Victims of Groupthink*, Houghton, Mifflin.

Kahneman, D (2012), *Thinking, Fast and Slow*, Penguin.

King, M (2003), speech available at <https://www.bankofengland.co.uk/-/media/boe/files/speech/2003/east-midlands-development-agency-dinner>

Kydland, F and Prescott, E (1977), 'Rules Rather than Discretion: The Inconsistency of Optimal Plans', *Journal of Political Economy*, Vol. 85, No. 3, pp. 473-492.

Kynaston, D (2017), *Till Time's Last Sand: A History of the Bank of England 1694-2013*, Bloomsbury Publishing.

Mackay, C (1841), *Extraordinary Popular Delusions and the Madness of Crowds*, Richard Bentley.

Mankodi, A and Pike, T (2018), 'Can central bankers become Superforecasters?', *Bank Underground*, 12 March 2018.

Nelson, E (2005), 'The Great Inflation of the 1970s: What Really Happened?', *Advances in Macroeconomics*, Vol. 3.

Newman, N, Fletcher, R, Levy, D and Nielsen, R (2016), 'Reuters Institute Digital News Report 2016', *Reuters Institute for the Study of Journalism*.

Page, S (2008), *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*, Princeton University Press.

Papacharissi, Z (2002), 'The virtual sphere: the internet as a public sphere', *New Media & Society*, Vol. 4, No. 1, pp. 9-27.

Richardson, T, Elliott, P, Roberts, R and Jansen, M (2017), 'A Longitudinal Study of Financial Difficulties and Mental Health in a National Sample of British Undergraduate Students', *Community Mental Health Journal*, Vol. 53(3), pp. 344-352.

Roubini, N and Mihm, S (2011), *Crisis Economics: A Crash Course in the Future of Finance*, Penguin.

Shafik, M (2017), 'In experts we trust?', speech available at <https://www.bankofengland.co.uk/speech/2017/in-experts-we-trust>

Shiller, R (2017), 'Narrative Economics', *American Economic Review*, Vol 107, No. 4, pp. 967-1004.

Stock, J and Watson, M (2003), 'Has the Business Cycle Changed and Why?', *NBER Macroeconomics Annual*, Vol. 17.

Surowiecki, J (2004), *The Wisdom of Crowds: Why the Many Are Smarter Than the Few and How Collective Wisdom Shapes Business, Economies, Societies and Nations*, Doubleday Books.

Taylor, M (2018), 'What Democracy Needs Now', *Royal Society for the encouragement of Arts, Manufactures and Commerce*, RSA Chief Executive's Lecture 2018.

Tetlock, P and Gardner, D (2016), *Superforecasting: The Art and Science of Prediction*, Broadway Books.

Tucker, P (2018), *Unelected Power: The Quest for Legitimacy in Central Banking and the Regulatory State*, Princeton University Press.

Tuckett, D (2011), *Minding the Markets: An Emotional Finance View of Financial Instability*, Palgrave.

Tuckett, D and Nikolic, M (2017), 'The role of conviction and narrative in decision-making under radical uncertainty', *Theory & Psychology*, Vol. 27, No. 4, pp. 501-523.

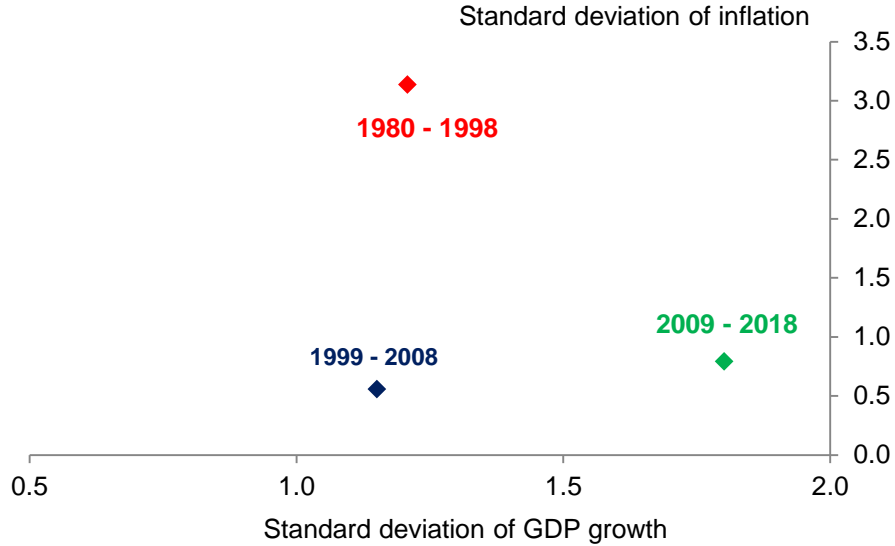
Turner, M and Pratkanis, A (1998), 'Twenty-Five Years of Groupthink Theory and Research: Lessons from the Evaluation of a Theory', *Organizational Behavior and Human decision Processes*, Vol. 73, Nos. 2/3, pp. 105-115.

Wagner, C and Vinaimont, T (2010), 'Evaluating The Wisdom of Crowds', *Issues in Information Systems*, Vol. 11, No. 1, pp. 724-732.

Van Reybrouck, D (2016), *Against Elections: The Case for Democracy*, Bodley Head.

Annex

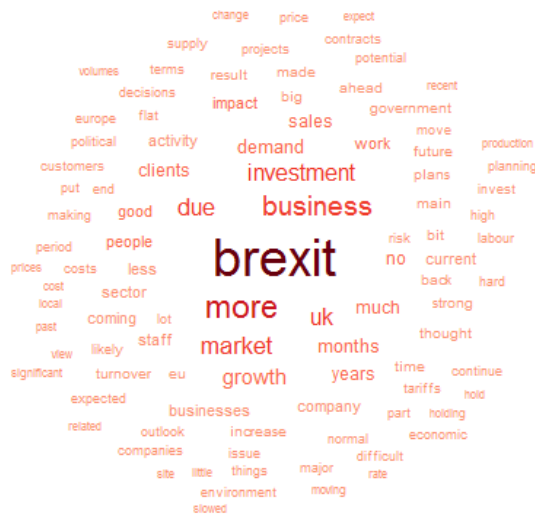
Chart 1: Advanced economy volatility of growth and inflation



Sources: IMF and Bank calculations.

Notes: Calculations use April 2018 IMF World Economic Outlook, with data at annual frequency. 2018 assumed to be data.

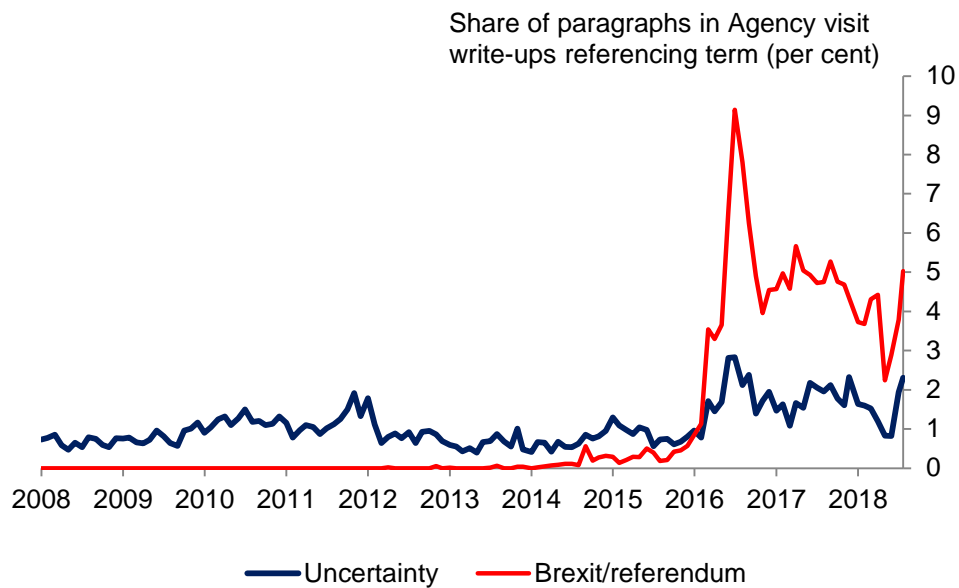
Figure 1: Word cloud for terms associated with 'uncertainty' in Agents' reports



Sources: Bank of England


Notes: Refers to Agents' reports compiled between 1 January and late-August 2018.

Chart 2: Frequency of ‘uncertainty’ and ‘Brexit’ appearing in Agents’ reports



Sources: Bank of England.

Chart 3: Inflation forecast errors of different groups since 2008 (RMSE)



Sources: Bank of England/TNS Inflation Attitudes Survey, Barclays Basix, YouGov/CitiGroup, HMT, CBI and Bank calculations.

Notes: Forecast errors since 2008Q2. Some measures start a few quarters later than 2008Q2, or have missing observations. Individual measures are included when there are no more than three missing observations. Financial market measures are adjusted for RPI-CPI wedge; and the Bank/TNS household measure is adjusted for the perceptions wedge.

** Except for household measures, where we need to use five-year-ahead expectations.

~ Simple average of all measures available in a given quarter. This will include measures that become available part way through the sample.

Figure 2: Illustrative diagram of 'wisdom of crowds'

