

Benoît Cœuré: Central banks as risk managers

Speech by Benoît Cœuré, Member of the Executive Board of the ECB, at the 53rd SEACEN Governors' Conference/ High-Level Seminar and the 37th Meeting of the SEACEN Board of Governors, Bangkok, 16 December 2017

It is a great pleasure for me to be here today.

Before I comment on this panel's topic, let me express my gratitude and satisfaction with the strategic partnership between the ECB and SEACEN, which got off to a successful start this year. Many cooperation activities, ranging from seminars on macroprudential analysis to central bank governance, have already been launched and more is being planned for 2018. I look forward to strengthening our cooperation over the coming years.

The topic of this panel deals with the implications of political risks for central banks. Given the independence of central banks and their legal separation from the political dimension, this is obviously a complex issue – and one where monetary policymakers need to tread very carefully.

For this reason, I would first like to spell out how the ECB generally incorporates different kinds of risk into its monetary policy strategy, and how this has influenced our actions over the last few years. I will argue that every central bank is to a considerable extent a risk manager, reflecting the forward-looking nature of monetary policy.

I will then explain why political risks cannot be addressed in the same way as economic risks. Central banks should not prejudge political outcomes through their actions. Rather, they should address their effects if and when they become visible in the economic and financial data that are relevant for their price stability mandates.

Monetary policy and risk management

My starting point is that monetary policy works with long and variable lags.

In the euro area, for example, the full transmission of interest rate decisions to output has been estimated to be between one and two years, and even longer for inflation.^[1] So, if we were to decide policy on the basis of past outcomes, we would always be behind the curve. Monetary policymakers therefore have to look at the economy in a forward-looking way.

To do this, we produce forecasts, on a regular basis, that indicate our

central expectations for the economy – the baseline. In principle, this should be enough to form a view on how policy should be designed today. But we all know that this would be a bad idea. Policymakers are typically poor forecasters, and central bankers are no exception.^[2] This is nothing to be ashamed of. It merely testifies to the fact that the past is often a poor predictor of the future.

You can see this quite clearly for the euro area on my first slide. We call it the “spaghetti chart”. It shows the repeated inflation forecast misses over the past few years. On each and every occasion there were good reasons to assume the economy would go the predicted way. But on each and every occasion unpredictable shocks hit our economy that made our central forecast redundant.^[3]

The implication is that we would likely have made severe policy mistakes if we had based our policy decisions entirely on our baseline.

And bear in mind that the economy can be more or less elastic to different types of shock. A tail risk, if it materialises, may cause the economy to react in a non-linear and potentially disruptive way – hyperinflation and deflation being typical examples of risks central banks want to avoid.

For all these reasons, central banks usually augment their forecasts with an assessment of the risks surrounding them. This comprises a *distribution* of risks – the range of possible outcomes and the likelihood of their happening – which, in turn, allows us to form a view on the *balance* of risks, i.e. whether they are overall tilted to the upside or downside, and on the probability of tail events.

Such risk assessments are not an exact science and there is no automatic link between them and policy decisions. But we do at times apply what Alan Greenspan famously called a “risk management” approach to monetary policy.^[4] If the balance of risks is tilted very strongly in one direction, or if the distribution of risks is especially wide, there might be a case for us to act.

For example, we might need to provide forward guidance, i.e. specifying how we would react to particular risks. Alternatively, we might need to change our policy stance pre-emptively, especially in situations where tail risks are material and it becomes cost-efficient to truncate that part of the distribution.

The ECB’s monetary policy since mid-2014 illustrates these two aspects well.

Around that time, we saw the balance of risks to the inflation outlook shift decisively downwards, while the distribution of risks widened to encompass outright deflation, as you can see by comparing the blue and red lines on my second slide. If they had materialised, those risks would have fundamentally compromised medium-term price stability, and so our strategy required us to respond – even though our central forecast at that time was for a low but positive rate of inflation in the years ahead.

We responded in two main ways.

First, we clarified our reaction function to the main risks we saw and the instruments we would use if each of those risks materialised.^[5] This sent a clear signal to observers that we were ready to respond in the case of adverse contingencies.

Then, when those contingencies arose, we followed through with our forward guidance and introduced a set of policy measures that was designed to cover the full downside distribution of risks – that is, a very accommodative policy stance to combat disinflationary forces, and an option to be even more accommodative if the situation deteriorated into outright deflation.

Thanks to these policy interventions, the distribution of risks has narrowed considerably over time – as you can see from the yellow line – and we no longer see a meaningful probability of deflation. The balance of risks has also shifted upwards as the economic recovery has gathered steam. The current economic expansion in the euro area is stronger than it has been for a decade and broader than for two decades.

This improving picture is the main reason for our recent decision to recalibrate our policy by reducing the pace of our monthly asset purchases from €60 billion to €30 billion, starting in January.

Of course, risks emanate not only from our own jurisdiction, the euro area, where we can respond with our monetary policy, but also from the rest of the world. Indeed, while the ECB's Governing Council currently sees the risks surrounding the euro area's growth outlook as broadly balanced, it sees downside risks relating primarily to global factors.

But here too we can manage risks effectively by cooperating closely with other central banks.^[6] This does not mean that we decide jointly on policy actions. It rather means that through our regular bilateral contacts, and dialogues in multilateral fora such as the IMF, the BIS and the G20, we can achieve a better understanding of global risks and their channels of propagation. And when risks do turn into shocks, this cooperation allows us to build up readiness and have the tools in place to react.

Perhaps most importantly, since 2011, the ECB has operated a permanent network of swap lines with the Bank of England, the Bank of Japan, the Federal Reserve and others, allowing all participating central banks to obtain foreign currency in the event of a liquidity squeeze. In 2013, the ECB also established a swap agreement with the People's Bank of China in recognition of its growing systemic importance as well as the rapidly growing bilateral trade and investment between the euro area and China.^[7]

Factoring in political risks

So how do we factor political risks into our decision-making?

I would argue that central banks cannot process political risks in the same way as economic risks, for two reasons.

The first relates to the degree of uncertainty that surrounds political risks.

Here it is useful to recall Frank Knight's classic distinction between risk and uncertainty.^[8] Risk is present when future events occur with measurable probability. Uncertainty arises when the likelihood of future events is indefinite or incalculable. In conditions of uncertainty, it is not possible to manage risk in the sense of quantifying a range of outcomes. Decision-making then depends on qualitative judgement.

To be sure, this is sometimes the situation central banks find themselves in when surveying the economic outlook. The economy is always characterised by both risk and uncertainty, and there are certain situations – for instance, financial crises – in which models fail and uncertainty prevails. In these cases, central banks still have to take decisions and judgement is the only basis we have.

Yet, I would venture that economic risks are, on the whole, more quantifiable than political ones, and hence more conducive to active risk management. This is because we have workable models of the economy with broadly established parameters and regularities. And even when the parameters of those models appear to change – like the Phillips curve today – they still provide us with a framework to think about those deviations and attempt to explain what we are seeing.

For politics, however, we rarely have such tools.

We may be able to gauge from opinion polls the likelihood of a political change of course happening. We may even be able to weigh up political parties' manifestos and estimate some of the economic consequences of their coming to power.

But fundamentally, we know little about how consumers and firms will react to political developments, and especially to the types of seismic political change that are macroeconomically relevant. Indeed, for such events to be considered a risk they are usually unprecedented.

This means that if we were to engage in managing political risks ex ante, most of the time we would be operating in uncertain circumstances and making judgement calls. I would question whether this could really be called risk management at all. Worse still, it would project us into the political domain on very shaky analytical foundations.

This brings me to the second reason why economic and political risks have to be treated separately, and it relates to the endogeneity between monetary policy and risks. In the economic realm, such endogeneity has been recognised as desirable and is a key reason why central banks have become much more transparent over the past two decades or so.

A clear understanding by the public of how the central bank will react to economic risks automatically reduces the likelihood of such risks materialising. For instance, if markets expect central banks to react to

adverse shocks by providing monetary accommodation, easier financial conditions will immediately follow. Such anticipation effects can increase the effectiveness of monetary policy.

For political risks, however, establishing such expectations would not be desirable. If we were to communicate that we will take decision “X” in response to political outcome “Y”, financial conditions would move as the probability of that outcome rose, and this would potentially prejudice the result. That would be controversial in the case of global political risks. For domestic ones, it would be unacceptable.

Even if the central bank had perfect foresight of the economic consequences, such a reaction function would be seen as undue interference in the political process and it could undermine the effectiveness of monetary policy, instead of increasing it.

And since our assessment would be largely based on judgement not analysis – for the reasons I mentioned – we would find ourselves being accused of political meddling. This is a position that no independent central bank would want to be in.

So when it comes to political risks, we have to be data-driven. We do not prejudge political outcomes. And we do not try to risk-manage their effects on the economy, since we can rarely predict those effects accurately – and worse, we may end up influencing political developments and thereby compromising our independence.

The only way in which we can include political risks in our policy framework is by responding to their visible impact on economic and financial conditions. This does not mean being complacent: we can and must plan for all eventualities. But we react to data, not to political events themselves.

In some ways, this is analogous to the debate about “leaning versus cleaning” of financial bubbles: faced with so much uncertainty about what constitutes a bubble, most of the time it is more efficient for central banks to use macroprudential tools to prick bubbles, or to ease policy after they burst, rather than to try and identify bubbles in advance and deflate them by hiking rates. The risk of false positives is just too high.

Two episodes in the recent history of Europe illustrate our data-dependent reaction function: the threat of a break-up of the euro area in 2012; and the threat of a country leaving the European Union in 2016, namely the United Kingdom.

In the first case, we had plenty of data showing that political risks were spilling over dangerously into the economy and financial system. Markets began pricing in redenomination risk. Financial conditions tightened significantly in some Member States. Bank lending contracted and the euro area entered a second recession. Uncertainty in the euro area, as measured by the VStoxx^[9], was on the rise – as the grey shaded area on my third slide shows.

Although at this point inflation was still being buoyed up by energy prices and indirect taxes, it was plain to see that political risks had become economic ones, and were in turn endangering the medium-term outlook for price stability. We therefore responded by launching a new monetary policy programme – Outright Monetary Transactions (OMTs) – which brought this episode of market turmoil to an end.

We did this, however, in a way that did not pre-empt political decisions. We took stock of the clear commitment of European leaders to hold our monetary union together and make it more solid by establishing a banking union. And we made the OMT programme conditional on countries participating in an assistance programme with the European Stability Mechanism.

In the case of the UK's vote to leave the EU, the situation was different, however. Various forecasts predicted severe market turbulence and macroeconomic fallout, so we had contingency plans in place for a range of outcomes. But as the slide illustrates, there were few signs of uncertainty in euro area financial markets in the run-up to the vote or after it. And, so far, there turned out to be no economic consequences with medium-term impact.

So our policy stance remained consistent with the data: unchanged. And the same logic, incidentally, can be applied to the recent political crisis in Catalonia. Though we monitored the situation very closely, we saw no changes in financial conditions or the economy that would have warranted a monetary policy shift.

Conclusion

Let me conclude.

Monetary policy is a forward-looking enterprise and policymakers always have to think in terms of risks. On several occasions in recent years the ECB has changed its monetary policy in response to emerging tail risks, even when our central forecasts for inflation painted a less alarming picture.

This can be seen as applying a risk management approach to monetary policy, in which we prioritised truncating the most dangerous tails of the distribution rather than targeting our policy at the modal point. The frequent central forecast misses we experienced suggest we were right to do so and we avoided much worse outcomes as a result.

When it comes to political risks, however, central banks cannot be risk managers, since this would bring us too close to being political actors. We can monitor political risks, and we can put in place plans for responding to them – but we can only act when the data justify such a step, and in a way that does not pre-empt political decisions.

Our actions during the crisis clearly demonstrated this reaction function.

Thank you.