Benoît Cœuré: Scars or scratches? Hysteresis in the euro area

Speech by Mr Benoît Cœuré, Member of the Executive Board of the European Central Bank, at the International Center for Monetary and Banking Studies, Geneva, 19 May 2017.

* * *

Accompanying slides.

One of the most interesting questions for policymakers to have emerged from the crisis is the strength of links between the demand and supply sides of the economy.[1] The classical view – that only cyclical policies influence the former, and structural policies the latter – has been challenged in two ways: first, by the observation that long periods of weak demand can lead to rising structural unemployment and a permanently lower capital stock – the so-called hysteresis effects. And second, by corollary suggestion that stronger demand might be able to reverse such effects.

These insights are not entirely new, of course. In 1986, Blanchard and Summers' seminal "insider-outsider" model laid bare how recessions might lead to higher underlying "structural" unemployment rates even after shocks have fully dissipated.

The question of hysteresis, however, received renewed interest as we have gone through the steepest collapse in demand since the Second World War, which has exposed our economies to significantly higher risks of hysteresis than in the past. And, as a consequence, policymakers are asking themselves if this should change the way they implement monetary policy. Should they, as usual, raise rates when there are signs that price pressures are building, so as to be ahead of the curve? Or should they stay their hand, in the expectation that what appears to be structural unemployment might actually be cyclical?

This is the debate that I would like to enter into today, from the perspective of the euro area – that is, to ask whether we have indeed seen significant hysteresis effects, and whether this warrants changing how we think about our monetary policy decisions.

My tentative answer today is that we have not yet seen hysteresis in its typical form – we have seen scratches, not scars. And though this more subtle type of hysteresis may warrant policy accommodation over a longer period, this is solely with the aim of achieving our price stability mandate, not preventing a rise in structural unemployment per se. Indeed, if those two aims were ever in conflict, our choice would be absolutely clear.

Evidence of hysteresis in the euro area

Let me start with a chart similar to the one Larry Summers showed for the US economy in his seminal 2014 speech.[2] You can see that the level of potential output in the euro area today is much lower than what we had expected before the outbreak of the crisis, assuming it would have continued to grow at its 2007 growth rate.

By this measure, the output gap has largely closed – not because the economy has been doing so well but because potential GDP has been revised significantly.

These effects have been taken by many as evidence for hysteresis.[3] But the discussion is arguably more complex than this. For sure, potential output estimates are chronically unreliable and often subject to substantial revisions *ex post*. Faced with the despair of failure, we often tend to underestimate our ability to overcome crises – a phenomenon I will return to in a minute.

Moreover, recessions often tend to have a permanent effect on the *level* of potential output. It is much less clear, however, whether the crisis will also leave a permanent footprint on the trend

potential *growth rate* – something Laurence Ball called “super-hysteresis” effects.

In fact, on the right-hand side of the first slide you can see that potential output growth had been slowing in the euro area well before the crisis, mainly reflecting a long-term slowdown in total factor productivity (TFP).

So the question I would like to ask today is not so much why potential output growth appears to be decelerating in the long run, or whether we are entering a period of secular stagnation – the jury is still out on this. It is whether the crisis, and the persistent shortfall in aggregate demand, has *accelerated*, or not, the decline in potential growth.

The first point to consider is whether the crisis has had a significant effect on labour participation – that is, the size of the total labour force. Hysteresis effects might be caused by unemployed workers becoming discouraged and discontinuing their search for employment, leading them to become detached from the workforce and lose their skills over time. Although this would reduce unemployment, it would lead to a permanent reduction in labour supply and growth potential.

So is there any evidence of this in the euro area?

The overall picture suggests not: the euro area has not been affected by strong adverse participation effects, as was seen for example in the United States – where the unemployment rate would have been significantly higher during the crisis had it not been for workers dropping out of the labour force. You can see this on the left-hand side of slide 2.

For the euro area, this is not a statistical artefact due to pension reforms or other possibly distorting measures: participation is growing virtually within all age groups.[4] Workers have on the whole remained attached to the labour market and, today, headcount employment has virtually regained its pre-crisis level. This is shown on the right-hand side. The same is true of the euro area employment rate, which expresses headcount employment as a share of the population aged 15 to 74.

This in itself provides strong evidence against participation-related hysteresis in the euro area.

But even though labour force participation has held up relatively well in the euro area, the second issue we need to look at is the *composition* of the unemployed.

It might be the case that, given the depth and severity of the crisis, some workers have remained in the labour force but have been without a job for so long that they have become virtually unemployable – that is, that they have become structurally unemployed. This would also have a negative effect on growth potential.

The natural starting point to assess the magnitude of such effects is to consider estimates of the NAIRU.

What we see here is that, in certain countries where unemployment rose sharply and where there was a strong sectoral dimension to the crisis, estimates of structural unemployment have indeed risen significantly. Current real-time estimates from a variety of international institutions, for example, suggest that Spanish structural unemployment increased by over three percentage points in the aftermath of the crisis, to around 17%.

At face value, this would suggest that virtually all of the unemployment that we currently observe in Spain is of a structural nature – and that the crisis has left permanent and irreversible scars.

And although current estimates do not suggest a notable increase in structural unemployment for the euro area as a whole, this was not true just a few years ago. For example, on the third slide you can see that in 2013 the European Commission estimated that structural unemployment for

the euro area as a whole was likely to increase from pre-crisis levels of around 8.5% to 11.6%.[5]

Today, Commission staff believe that structural unemployment in the euro area is closer to 9% – compared with an actual unemployment rate of 9.5% in March 2017.

In other words, without the significant downward revisions of recent years, we would currently face a *negative* unemployment gap in the euro area. And even according to the Commission's latest estimates, policy-relevant labour market slack would have vanished, by and large.

You can see what I'm getting at.

If structural unemployment had indeed risen during the crisis to the extent suggested, then we would have expected inflationary pressures to emerge much more rapidly, since the unemployment gap would have been closing faster than otherwise.

But this is not what we see in the euro area – or, in fact, in many other economies.

Despite a rapid fall in the unemployment rate, wages have remained stubbornly low. Annual growth in compensation per employee hovered around 1.2% since mid-2014 and only increased to 1.5% at the end of last year – substantially below its historical average of 2.1%.[6] And measures of core inflation – in its simplest form headline inflation excluding volatile items such as food and energy – have also remained at depressed levels of below 1% for most of the past three years.

This is confirmed by Phillips curve estimates. As you can see from the right-hand side chart on slide 3, ECB in-house estimates using a variety of slack and inflation expectations measures suggest that core inflation today is well below what historical regularities would suggest.

Of course, a number of factors – such as changes in the slope, weak productivity growth or second-round effects of past low inflation – might be at work here. But in the current context this is unlikely to be the whole story.

For one, there is considerable uncertainty around the slope of the Phillips curve. As we heard this morning, some estimates show that the curve has flattened in recent years. ECB research, by contrast, found evidence that inflation has become more responsive to slack lately, but with notable heterogeneity across euro area Member States.[7] Blanchard, Cerutti and Summers have not found much change at all during the crisis.[8]

Second, our Phillips curve estimates show that even once we control for past low inflation and weak productivity growth, there is still a considerable part that cannot be explained by the model.

These analyses therefore do offer one useful conclusion: that slack may be larger and that structural unemployment may perhaps not have risen by as much as some estimates suggest. This also means that policymakers should be careful in attaching too much weight to notoriously unreliable and hard-to-estimate measures such as the NAIRU. Indeed, already 20 years ago Staiger, Stock and Watson[9] pointed out that the most striking feature of NAIRU-based inflation estimates is their lack of precision, not to mention model uncertainty as evidenced by the right-hand chart on slide 3.

So what, then, is causing the sluggish performance of inflation?

Here I want to argue that there are more subtle forms of hysteresis in the economy that are altering how we typically understand labour market slack. These are the scratches rather than the scars.

A different perspective on labour market slack

One of the features of the recent employment recovery is that many new jobs have been of a lower "quality" compared with those seen before the crisis. While employment ratios have rebounded strongly, employment of full-time workers with permanent contracts has in fact declined.

The latest EU Labour Force Survey (LFS) data suggest that "core" employment – that is, full-time positions on open-ended permanent contracts – now account for just over two thirds of all employees. You can see on slide 4 that this is the lowest share since the survey was first conducted. The LFS also suggests that of the net employment created since the crisis, around a third has been for workers on temporary contracts, and around a quarter part-time.

Of course, a fraction of those seeking part-time employment do so as a way to balance work and other commitments. But, for others, the choice for part-time work is involuntary. Currently there are more than seven million underemployed part-time workers across the euro area – those who would like to work more hours but cannot – an increase of around one million since the crisis.

This may have important implications for inflation dynamics.

While the main objective for workers in permanent work is typically higher wages, those in temporary or part-time positions may pursue objectives other than wage increases, such as full-time employment or an increase in hours worked.

Or if they do seek higher wages, they may be less likely to benefit from union representation and thus have less bargaining power. An increase in temporary or part-time employment as a share of total employment may therefore lead to less pressure on wage inflation.

Recent estimates by ECB staff – which you can see on the right-hand side of slide 4 – suggest that broader measures of labour market slack, including the unemployed, underemployed and those marginally attached to the labour force currently affect around 18% of the extended euro area labour force.[10]

This broader measure of slack also appears to be less sensitive to GDP growth than traditional measures: it has declined rather less than the reductions in the unemployment rate over the course of the recovery – a broader version of Okun's, law if you will.

And our analysis suggests that this might help explain why underlying inflation has been so sluggish recently.

Rerunning our Phillips curve models using this broader measure of slack, or equivalently, a measure of "core" employment only, leads to the prediction errors for inflation becoming measurably smaller. Look at how the gap between actual and predicted inflation largely disappears on slide 5, in particular for the recent history.

All this essentially means that it may take longer for inflation to gain steam and wage pressure might only start to rise meaningfully once adjustments in the "intensive margin" take hold – that is, once those who want to work more hours also succeed in doing so and once those who are still willing to work, but not currently counted as unemployed, are reabsorbed. Indeed, hours worked per employee have not recovered since the double-dip recession.

So what is the link between these developments and hysteresis?

To be sure, the increasing prevalence of part-time and temporary work reflects several factors – including, in part, important structural reforms in many euro area economies that have helped strengthen the resilience of the euro area.

Yet, the compositional effects we are seeing in the labour market may also reflect more crisis-related factors, in particular uncertainty about the strength and durability of the recovery and pessimism about the long-term impact of the downturn on growth.

As Blanchard and co-authors have recently postulated, a revision to future potential growth prospects may cause firms to cut back on investment plans, causing muted growth expectations to become self-fulfilling.[11]

On the right-hand side of slide 5 you can see Blanchard and Co.'s analysis replicated for the euro area. And we see signs that such effects might have also been at play in the euro area: there is a positive correlation between potential growth revisions and investment error forecasts. This might suggest that investment, and therefore demand, is influenced by long-term animal spirits – changes in expectations concerning the steady-state path of the economy.

But the logic may also run in reverse and create a vicious feedback loop: this would be the case if the pessimism-induced period of low demand ultimately affected long-run potential growth.

I would call this an "anxiety trap" that could ultimately lead to the more subtle forms of hysteresis I mentioned earlier.

The reason is that if the effect of low expectations and high uncertainty is to quell animal spirits, large demand shocks can have long-lasting effects: the increase in risk aversion slows down capital accumulation, thereby reducing productivity growth, while uncertainty encourages firms to enter into more flexible forms of employment that shed labour market risks onto workers.

This then feeds back into wage growth in three ways: (i) through lower labour productivity, which depresses real wages, (ii) through a higher level of effective labour market slack and (iii), in an environment of high uncertainty, through employees and trade unions valuing job security more than higher wages – in particular after recessions where wage growth would likely have been negative for a while had it not been for downward nominal wage rigidity.

Put differently, what we could be seeing today is a hysteresis in the euro area economy that has an effect on wages quite unlike that of "traditional" labour market hysteresis.

The issue might not be so much that fewer "insiders" may be able to push aggregate wages up – which was a key insight of Blanchard and Summers' earlier work. It is rather that hysteresis may be affecting labour markets through employers' risk perceptions and subsequent hiring and investment choices, which have so far had the contrary effect of pushing wages *down*.

But while this form of hysteresis might leave some workers working less than they would like, they remain – by and large – attached to the labour market. And by retaining their professional network and skills, they are less likely to become marginalised and, on aggregate, structural unemployment is less likely to become entrenched.

This is what I mean when I say that the crisis may have caused scratches, but not necessarily scars.

Still, while the labour market has proven surprisingly resilient to the repercussions of the double-dip recession, the perils of hysteresis – of real scars – remain: you can see on slide 6 that long-term unemployment as a share of total unemployment has increased from levels around 45% before the outbreak of the crisis to 50% today and has barely declined in recent years. In other words, one in every two of the more than 15 million currently unemployed face the risk of becoming increasingly detached from the labour market.

Some of the long-term unemployed will be absorbed back into the labour force as the recovery continues, and some of the marginally attached will also rejoin. But expecting the economic cycle

to fully resolve these wider problems would be misguided.

This is particularly worrisome for younger generations, for whom extended spells of unemployment at the start of their working lives may lead to lifelong "scarring" effects – including a higher propensity for further spells of unemployment, which are likely to constrain labour market choices and earnings throughout their working lives.[12] In short, they risk becoming a "lost generation".

Concerns regarding labour market mismatches – effectively precluding the reabsorption of those displaced from some permanently downsized sectors – also remain. The euro area Beveridge curve – as you can see on the right-hand side of slide 6 – looks to have shifted outward during the crisis. This is consistent with a lower degree of matching efficiency post-crisis, with the unemployment level still remaining high despite rising labour demand over the duration of the recovery.

Hysteresis effects and monetary policy

So what is the role of monetary policy as regards hysteresis effects?

Though our primary mandate is price stability and not employment, in conditions of weak aggregate demand there is typically a "divine coincidence" between those two variables, since unemployment rises while inflation falls after adverse shocks hit the economy. So, only by reducing labour market slack can we create the conditions for wage and price pressures to emerge and for inflation to return to our preferred aim within our definition of price stability.

This is largely what we have seen in the euro area. Since the adoption of the ECB's credit easing package in June 2014, around five million jobs have been created in the euro area. Model-based evidence indicates that our policy has contributed meaningfully to this increase: the current recovery has, by and large, been driven by monetary policy, with additional temporary support from the fall in oil prices.

So, forceful policy action in the face of risks to our price stability mandate has measurably reduced the risks of hysteresis by creating the financial conditions necessary for growth and employment to flourish.

What is more, our actions have also defeated a potentially very nasty form of hysteresis, namely a destabilisation of inflation expectations. This could have had hysteresis-like effects by shifting the Phillips curve inwards and permanently lowering the inflation rate, towards which the economy gravitates once all shocks have dissipated. Fortunately, survey-based measures of medium-term inflation expectations have consistently remained close to our definition of price stability.

But the concrete anatomy of labour market dynamics and the possible effects of hysteresis arguably become more important in informing monetary policy decisions at later stages of the easing cycle or at turning points.

With the recovery in the euro area now fully underway, policymakers must judge at what point they expect inflationary pressures to emerge in a sustainable manner that would warrant a gradual withdrawal of policy accommodation.

Conditions in the labour market are, of course, an important part of this discussion. Wage developments are a key driver for service price inflation and, thus, for headline inflation. The amount of slack we expect to see in the economy, and in the labour market in particular, therefore has an important bearing on our stance.

For example, were we to ignore the facts that I just presented, namely that labour market slack

may be larger than is suggested by headline unemployment measures, we could run the risk of tightening policy prematurely. Choking off growth too quickly would lead to us failing on our medium-term inflation mandate and would needlessly keep people out of work.

Similarly, the changes in the composition of labour contracts that I mentioned earlier are likely to imply that inflationary pressures might take longer than usual to build up. We need to factor this information into our policy deliberations.

At the same time, we need to monitor carefully developments in long-term unemployment, labour force participation rates and hours worked with a view to informing our assessment on the likely evolution of structural unemployment.

Should the share of long-term unemployed stabilise at one point, with no convincing signs of further likely downward adjustments, the risks of inflationary pressures starting to emerge at higher levels of unemployment would increase.

At this point – when scratches indeed become scars – the level of slack relevant for determining medium-term price stability would decline. Policymakers would therefore need to start tightening policy earlier to remain faithful to their primary mandate.

As I mentioned in my introduction, however, there are some who argue that such a response would be premature, since if hysteresis effects can be caused by lack of demand, then they can equally be reversed by excess demand. In a high-pressure economy, for example, firms might be encouraged to dig deeper into the labour market to find recruits, thereby reactivating and retraining some of those workers who were previously thought to be unemployable. Glenn Rudebusch and John Williams therefore conclude that optimal monetary policy would allow inflation to overshoot its target.[13]

I have some doubts about this view, mainly for two reasons. First, and most importantly, unlike in the US, allowing inflation to overshoot our price stability definition for the sole purpose of bringing more people back into employment would not be consistent with our primary mandate. The "divine coincidence" applies only insofar as employment developments are consistent with our definition of price stability. Beyond that there is no room for trade-offs.

Second, there are doubts as to the likely efficacy of such a policy. We have seen some evidence in the United States that the participation rate can become partially endogenous to demand, as discouraged workers return to job-seeking in a strengthening labour market.

But we have seen no such evidence in the euro area. Our participation rate has continued to increase over time. So the question in the euro area is much more whether, or not, long-term unemployment is likely to be amenable to continued policy accommodation.

The evidence raises at least doubts. For example, although many claim that monetary policy was too accommodative during the boom years of 2005 and 2006 – as also evidenced by policy rates often being lower than what standard Taylor rules would have predicted – the share of long-term unemployed in the euro area has remained fairly constant over this period.

So, even if the desired effects of accommodative policy lead to more investment and employment, it remains questionable whether employers will respond by taking on chronically unemployed or whether they would instead rely on "insiders" whose pay has been stagnant or, alternatively, on fresh university graduates. A recent study analysing the experience of the Spanish labour market confirms this view: over time job finding rates become less and less responsive to improvements in the aggregate labour market.[14]

I would even go one step further, namely to question whether allowing inflation to overshoot would help address the more subtle forms of hysteresis in the economy linked to uncertainty and

investment.

I have no doubt that investment would have been much lower today were it not for the impulse generated by our policies. But central banks cannot address all of the underlying causes of weak capital formation.

For example, if it is true that investment is being affected to some extent by weak long-term growth expectations, then monetary policy cannot be the whole answer – unless one believes in a very hard form of “reverse hysteresis”.

The truth is that should weak long-term growth expectations indeed become self-fulfilling, then we would also see signs of hysteresis in the natural rate of interest, which economists frequently call "r-star". A lower r-star, in turn, would push the interest rate level at which policy becomes expansionary even lower – feeding concerns about the ability of monetary policy to address advanced cases of hysteresis.

We don't want the central bank itself to fall into the “anxiety trap”.

The answer therefore has to address the structural causes of low trend growth and long-term unemployment through policies supporting productivity and minimising both dualities in the labour market and social exclusion.

Other scratches also need early and effective treatment to avoid becoming scars. For example, insofar as investment is being affected by post-crisis legacies, such as the debt overhang in parts of the private and public sector – the famous Fisher-Minsky-Koo channel of debt deflation – history suggests that effective deleveraging nearly always involves a combination of growing out of debt and writing down debt – which is to say, structural policies are also key.[15]

And action also at the EU level is crucial. Efforts to complete the Single Market, for example, would incentivise more firms to invest and grow as they have a larger market to serve and exploit economies of scale. Progress with the Capital Markets Union (CMU) would likewise help to raise productivity by facilitating capital reallocation.

Conclusion

So, to conclude, the evidence suggests that hysteresis as it is often understood – meaning rising structural unemployment – has not yet materialised meaningfully in our economy. We see more scratches than we see scars. And unlike “classical” hysteresis, these scratches are currently reducing price pressures, warranting our currently very accommodative monetary policy stance.

We can be confident that this stance will have the desired effect. The ongoing broadening of the recovery, together with the gradual disappearance of downside risks, increases the chances that those who want to work more will be able to do so. This, in turn, will create the necessary price and wage pressures for a return of inflation towards levels closer to 2%.

But should we reach a point where the path of inflation is expected to be self-sustaining, but long-term unemployment remains high, there should be no doubt as to how I would decide regarding our policy stance. Monetary policy cannot “run the economy hot” as insurance against labour market risks – that is, the risk that long-term unemployment becomes structural unemployment. This would be neither desirable in view of our primary mandate, nor efficient given the euro area circumstances. It will be up to other policymakers to tackle this issue – as they must.

---

[1] I would like to thank V. Jarvis and E. Lis for their contributions to this speech. I remain solely responsible for the opinions contained herein.

[2] See Summers, L. (2014), “U.S. Economic Prospects: Secular Stagnation, Hysteresis, and theZero Lower

Bound", Keynote Address at the NABE Policy Conference, 24 February 2014.

[3] See, e.g., Ball, L. (2014), "Long-Term Damage from the Great Recession in OECD Countries", NBER Working Paper No. 20185.

[4] Participation rates have fallen for the under-25s and remain some way below their pre-crisis levels – in part reflecting higher enrolment rates, as young people exercise their "outside option" of staying in education in the face of less advantageous labour markets.

[5] The level expected to prevail in 2015.

[6] The picture changes little when hourly wage growth is considered, with annual rates ranging from 1.0% to 1.3% over the course of the recovery. Yet, it picked up more strongly in the last two quarters of 2016, standing at 1.8% in the last quarter of 2016.

[7] See Ciccarelli, M. and C. Osbat (2017), "Low inflation in the euro area: Causes and consequences", ECB Occasional Paper No 181.

[8] See Blanchard, O., E. Cerutti and L. Summers (2015), "Inflation and activity – Two Explorations and their Monetary Policy Implications", IMF Working Paper No 230.

[9] Staiger, D., J. Stock and M. Watson (1997), "The NAIRU, Unemployment and Monetary Policy", Journal of Economic Perspectives, 11(1): 33-49.

[10] See the recent Economic Bulletin box entitled "Assessing labour market slack", Economic Bulletin, Issue 3, ECB, 2017. The figure is computed by expressing the numbers of unemployed and underemployed, together with estimates of those available for work but not seeking it, and those seeking work but not available for it (the "potential additional labour force"), as a percentage of the extended labour force (i.e., the employed and the unemployed, who comprise the active labour force, plus the potential additional labour force).

[11] See Blanchard, O., G. Lorenzoni and J.-P. L'Huillier (2017), "Short-Run Effects of Lower Productivity Growth: A Twist on the Secular Stagnation Hypothesis", NBER Working Paper No 23160, February 2017.

[12] See Bell, D. and D. Blanchflower (2011), "Young people and the Great Recession", Oxford Review of Economic Policy, 27, pp. 241–267.

[13] Rudebusch, G. and J. Williams (2016), "A wedge in the dual mandate: Monetary policy and long-term unemployment", Journal of Macroeconomics, Volume 47, Part A, March 2016, Pages 5–18.

[14] Bentolila, S., J. Ignacio García-Pérez and M. Jansen (2017), "Are the Spanish Long-Term Unemployed Unemployable?", CESIFO Working paper No. 6338.

[15] For example, streamlined insolvency frameworks, better judicial efficiency and supervisory action against non-performing loans.