INFLATION AND ACTIVITY

TWO EXPLORATIONS, AND THEIR MONETARY POLICY IMPLICATIONS[1]

BY OLIVIER BLANCHARD
EUGENIO CERUTTI
LAWRENCE SUMMERS

ABSTRACT.

1 INTRODUCTION

We focus in this paper on two empirical issues triggered by the Great Financial crisis. First, in most advanced countries, output remains far below the pre-recession trend, leading researchers to revisit the issue of hysteresis. Second, while inflation has decreased, it has decreased less than was anticipated (an outcome referred to as the ``missing disinflation''), leading researchers to revisit the relation between inflation and activity.

Clearly, if confirmed, either the presence of hysteresis or the deterioration of the relation between inflation and activity would have major implications for monetary policy and stabilization policy more generally. In the first case, it would imply that the cost of output shortfalls is much higher than typically assumed. In the second case, the lack of a reliable relation between inflation and activity, be it output or unemployment gaps, would require a major rethinking of the inflation targeting architecture.

With this motivation in mind, we have a broad look at the evidence. First, we revisit the hysteresis hypothesis, defined as the hypothesis that recessions may have permanent effects on the level of output relative to trend. Second, we revisit the evidence on the strength of the relation between the unemployment gap and inflation, the Phillips curve.

We do this by looking at output, unemployment, and inflation over the course of roughly 50 years for 23 advanced economies. We draw the following conclusions:

We find that a high proportion of recessions (about two-thirds) are followed by lower output relative to the pre-recession trend. Even more surprisingly, in about two-thirds of those cases, the recession is followed not just by lower output, but by lower output growth relative to the pre-recession output trend. That is as time passes following recessions, the gap between output and projected output on the basis of pre-recession trend increases.

If these correlations were causal, they would suggest important hysteresis effects and even "superhysteresis" effects. (to use Larry Ball's term for impacts of recessions on growth rates). Correlation however does not imply causality. The causality may indeed run from the recession to lower output later, and hysteresis or superhysteresis may indeed be at work. The correlation may however reflect instead common third factors: Supply shocks, such as an increase in oil

[1] Preliminary, May 18, 2015. We thank Larry Ball and Sandeep Mazumder for comments and help, as well as Yangfan Sun and Daniel Rivera for excellent research assistance. We thank Zeno Enders, Stephan Danninger, Chris Erceg, for comments. Our paper builds very much on Martin and Wilson (2013), and the IMF WEO April 2013, chapter 3.

prices, or a financial crisis, may be behind both the initial recession and lower output later. Or the correlation may reflect reverse causality: The anticipation of lower output or lower growth in the future may lead to a decrease in consumption and investment spending, and, as a result, to a recession today.

This leads us to look at recessions associated with different shocks. We find that, indeed, recessions associated with either oil price increases or with financial crises are more likely to be followed by lower output later. But we find that recessions plausibly triggered by demand shocks, for example recessions associated with decreases rather than increases in inflation, are also often followed by lower output or even lower output growth. Even in the case of recessions associated with intentional disinflations, which probably represent the purest case of demand shocks we can identify in the sample, we find that still roughly half of those recessions are associated with lower output later and that a significant fraction of these are associated with lower output growth.

We draw two tentative conclusions. The high proportion of recessions followed by lower output or lower output growth, together with a number of other observations from other studies, leads us to conclude that the correlation between recessions and subsequent poor economic performance reflects to a significant extent reverse causality: The realization that growth prospects are lower than was previously assumed naturally leads to both a recession and subsequent poor performance. But the finding that recessions plausibly triggered by demand shocks are also often followed later by lower output, or even, in some cases, lower output growth, suggests that hysteresis, and perhaps even superhysteresis may indeed also be at work. Both conclusions have important, but very different, implications for monetary policy, to which we come back later.

Turning to the Phillips curve relation, we start by estimating, for each country, a benchmark relation between inflation, long term inflation expectations, lagged inflation, and a measure of the unemployment gap. The specification allows for the natural rate to change over time, and for the coefficients to evolve over time. We confirm that the coefficient on long term expected inflation has steadily increased over time. This explains in large part why we have not observed a deflation spiral, despite the presence of sustained large unemployment gaps. But we also find clear evidence that the effect of the unemployment gap on inflation has steadily decreased over time, with all the decrease taking place before the crisis. Indeed, in the benchmark specification, the coefficient on the unemployment gap is often insignificant. We explore a number of variations on the benchmark specification, allowing for hysteresis, or replacing unemployment by short term unemployment. In most specifications and in most countries, the effect of the unemployment gap remains small, and at best marginally significant.

In the last section, we explore the implications of our findings for monetary policy.

The findings of the first section have ambiguous implications for monetary policy. To the extent that a large number of recessions are due to anticipations of coming lower underlying growth later, this implies that estimates of potential output, based on the assumption of an unchanged underlying trend, may be too optimistic, and lead to too strong a policy response to movements in output. However, to the extent that recessions have hysteretic or super hysteretic effects, then the cost of allowing downward movements in output in response to shifts in demand increases, implying that a stronger response to output gaps is desirable.

The findings of the second section yield a puzzle and a potential challenge for inflation targeting. At the same time as inflation expectations have become more anchored, the ability of central banks to affect inflation through the unemployment gap, as measured by the size and

reliability of the coefficient, has declined. Put another way, the faith in the ability of central banks to achieve their target has increased, while the ability of central banks to actually achieve it has decreased. The worry is a clear one, namely whether the faith will remain, and if it does not, what may happen to inflation and to monetary policy in the future.

2 HYSTERESIS OR REVERSE CAUSALITY?

The issue of hysteresis in output and unemployment surfaced in Europe in the 1980s (Blanchard and Summers 1986), and never got settled. It eventually lost centre stage. The crisis has brought it again to the fore. The reason is not hard to see, and is shown in Figure 1. The figure shows the evolution of the United States and the Euro area's output since 2000. Its visually striking implication is that, after the crisis, output appears to be evolving on a lower path, perhaps even a lower growth path, especially in the Euro area.

Figure 1: Advanced Economies Real GDP (Index, 2000Q1=100).

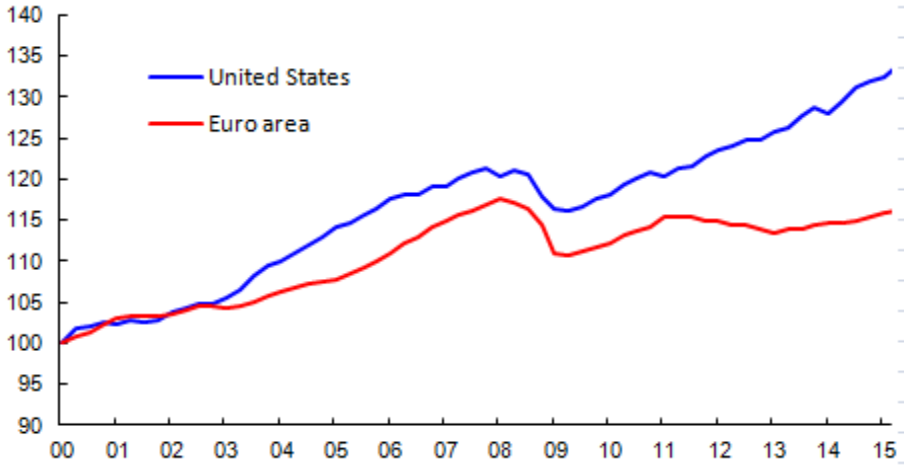

Some researchers (Ball 2014) have taken this as evidence of hysteresis. But correlation does not imply causality. One can plausibly argue that the lower path is due to institutional changes in response to the crisis, such as tougher capital requirements, or changes in bank business models, a form of institutional hysteresis. But one can also plausibly argue that the sharp decline in output at the start and the later lower growth path are due to the same underlying cause, namely the crisis of the financial system, manifesting itself through an acute effect at the start, and a more chronic effect thereafter. As a matter of logic, one could even, although less plausibly in this case, argue that the recession was partly due to the anticipation of lower growth to come.

This discussion is what leads us to look at a much larger set of recessions, over many countries and many years, and proceed in two steps. First, by establishing stylized facts and correlations: How often have recessions been followed by lower output relative to trend, or even by a lower trend? Second, by attempting to control for the cause of the recession, and focusing on those recessions which were more likely to be caused by demand rather than by supply factors, thus where causality was more likely to run from the recession to subsequent developments.

To do so, we look at the evidence from 23 advanced countries, using quarterly data starting in 1960 (or whenever data starts being available). In doing so, we build on the work of Martin and

Wilson (2013),[2] our contribution is in using a slightly different methodology, and looking at the relation conditional on different types of shocks.

We rely on a non parametric method, focused on recessions rather than on general fluctuations.[3]

The methodology requires defining both recessions and trends:

We follow the literature in defining recessions using the methodology of Harding and Pagan (2002). Roughly speaking, the method identifies peaks and troughs as local maxima and minima in the log level real GDP series, and, with some exceptions, defines recessions as times between a peak and a trough.[4] Estimating pre-recession trends and identifying both their position and their slope is more challenging:

- The first issue is how to take into account that the economy may have been in a boom, and thus above trend, before the recession started. We explore two alternatives. The first is to exclude the two years before the recession from the computation of the trend, and to base the start of the estimated trend at the value of log real GDP two years before the recession. The second recognizes one of the lessons of the crisis, namely that the economy may be on an unsustainable path even if output growth does not appear unusually high, but financial imbalances are building up which must eventually lead to an adjustment and to lower growth. Empirically, we use a rule in which, to estimate the trend and the starting point of the trend extrapolation, we exclude at least the last two years before the recession and possibly more years if they are characterized by unusual credit growth. To define unusual credit growth, we rely on the episodes identified by Dell'Ariccia et al (2012), which are based on an annual growth rate of the credit to GDP ratio exceeding 10% and a deviation from a credit to GDP trend greater than 1.5 times its standard deviation.

- The second issue is the length of time used to estimate the pre-recession trend. We explore two alternatives, one in which the trend is estimated over 4 years (so, in the absence of a credit boom, over t-8 to t-24, where t is measured in quarters). This allows for a flexible trend, but makes the estimated trend quite sensitive to what may have in effect been cyclical fluctuations. The other is thus to estimate the trend over 10 years, so, in the absence of a credit boom, over t-8 to t-48. The potential shortcoming is the symmetrical risk that this may not capture recent changes in the underlying trend. [5]

This gives four different combinations, and we derive results in each country for each of the four combinations. The figures giving actual log real GDP, recession dates, and estimated trends, are actually worth looking at one by one, and are given in a web appendix. Figures 2 and 3, which are based on a 4-year trend anchored two years before the recession, give a flavour of these graphs. In the figures, the black dash lines in each case give the one-standard deviation band associated with uncertainty about the value of the estimated trend coefficient.

[2] Martin and Wilson build in turn on Cerra and Saxena (2008).
[3] There will be an appendix discussing what can and cannot be learned from time series methods, from standard ARIMAs to Campbell-Mankiw and Blanchard-Quah decompositions, and the pros and cons of parametric versus non parametric methods in this context.
[4] Following Harding and Pagan (2002), we set to 2 quarters the number of observations on both sides over which local minima and maxima are computed, to 2 quarters the minimum duration in every contraction or expansion phase, and to 5 quarters the minimum duration between two peaks and two troughs.
[5] To state the obvious: two-sided filters, such as an HP filter, cannot be used for these purposes, as the behavior of output after the recession would affect the estimated trend before the recession. By construction, output would return to the constructed trend, thus negating any level or growth effect of recessions.

Figure 2: United States –Evolution of log real GDP and Extrapolated Trends

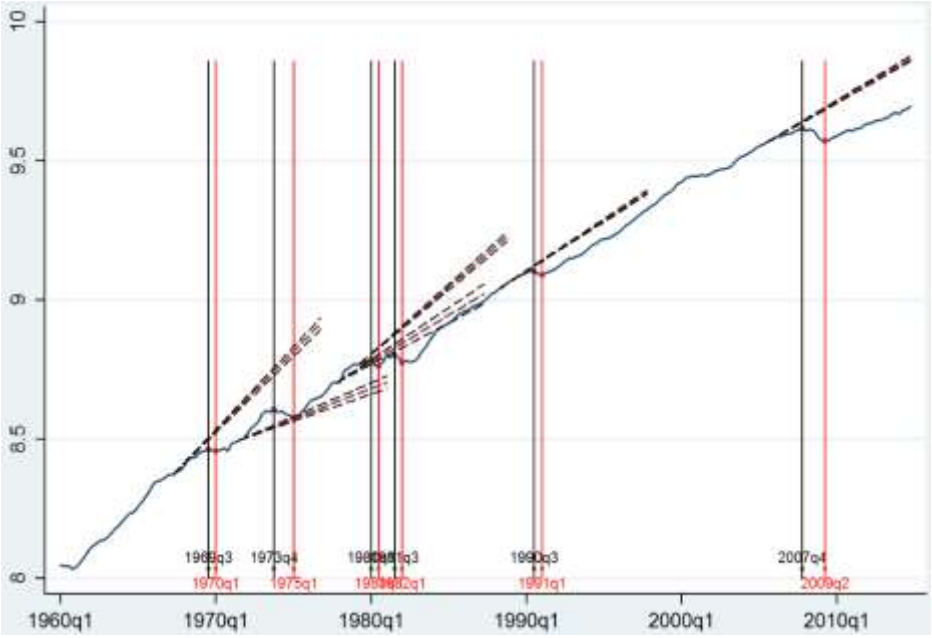

Figure 2 shows the evolution of the United States. While the decrease in output relative to trend is most striking in the case of the Great Financial Crisis, some of the other recessions appear to be associated with a lower level of output relative to trend.

Figure 3 shows the evolution of Portugal, and is representative of the evolutions of a number of European countries. All but one of the recessions since 1960 appear to be associated not only with a lower level of output relative to trend, but even with a subsequent decrease in trend growth, and thus increasing gaps between actual output and past trend.

Figure 3: Portugal –Evolution of log real GDP and Extrapolated Trends

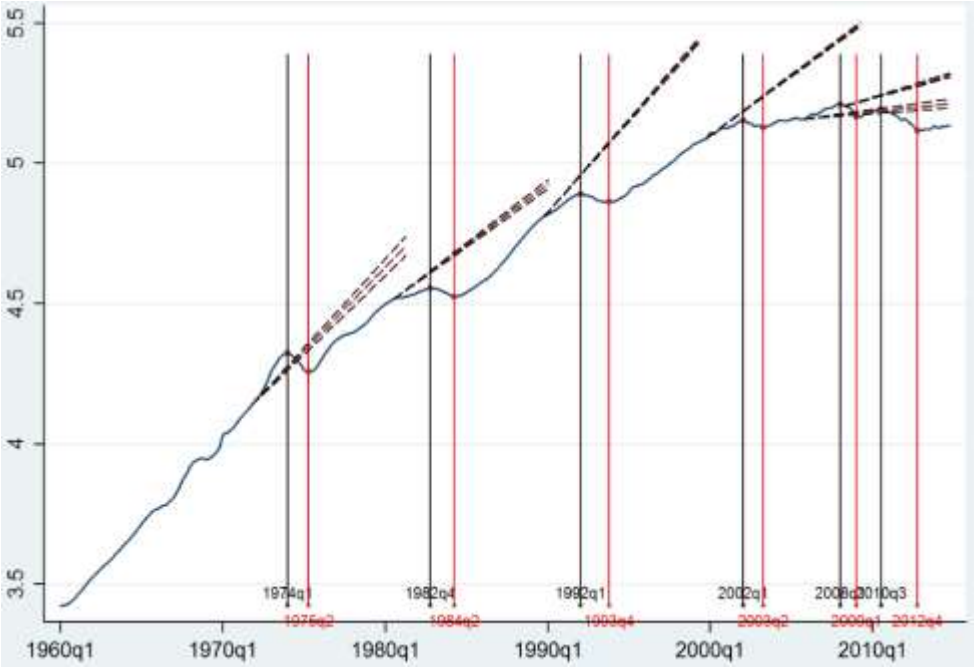

A comprehensive set of statistics is given in the top half of Table 1.

The two sets of columns give the results corresponding to the two ways of computing the pre-recession time trends, over 4 years or over 10 years respectively. The two lines correspond to the two ways of computing and removing the pre-recession years (leaving out the two years before the recession, or leaving out more years if there is evidence of a credit boom).

For each of the two time trend treatments, the table has four columns. The first gives the proportion of recessions where the gap between the output level and the past trend is not significantly different from zero after 3-7 years after the recession. The second gives the proportion where the gap is significantly different from zero. The last two columns decompose this proportion between the proportion of recessions where the gap is significantly different from zero but stable, and those where the gap is not only significantly different but also increasing.

The results, shown in the first two rows, are very similar for all four combinations. They show that in only 36-40 percent of the recessions, the recession was not followed by a sustained gap between the actual series and the estimated trend.[6] Equivalently, in 60-64 percent of the cases, the recession was followed by a significant output gap. In 40 to 46 percent of the cases, the gap between the output level and the past trend was not only significantly different from zero but also increasing.

As a robustness test, we calculate the gap and trend using log real GDP per capita (calculated as GDP over population of 16 to 64 years old) for the same recession periods. The results, shown in the next two rows of Table 1, are roughly similar. In 56 to 57 percent of the recessions, the recession was followed by a sustained output gap. The proportion of cases where the output gap was increasing runs between 36 and 40 percent.

Table 1: Analysis of the differences between output level and trend across recessions

GDP series used	Trend Extrapolation Starting Point	Trend Calculation: 4 year window				Trend Calculation: 10 year window			
		Episodes with NO sustained gap	Episodes with sustained gap	of which:		Episodes with NO sustained gap	Episodes with sustained gap	of which:	
				Increasing over time	Stable over time			Increasing over time	Stable over time
Log Real GDP	Benchmark: 2 Years Before	36%	64%	46%	18%	39%	61%	40%	20%
	Adjusted for Credit Booms	39%	61%	43%	17%	40%	60%	40%	20%
Log Real per capita GDP	Benchmark: 2 Years Before	43%	57%	40%	17%	43%	57%	36%	20%
	Adjusted for Credit Booms	44%	56%	39%	17%	43%	57%	37%	20%

Note: A total of 122 recession episodes are included in the analysis. Recession episodes during the 1960s are not included due to lack of data for estimating trends. Similarly recessions after 2010 are not included due to lack of enough observations.

We have also performed a number of visual robustness checks (is the increasing output gap in years 3 to 7 due to an outlier, to another recession?). Our conclusion is that, in 80% of the cases classified as "increasing output gap'', the increase was indeed unambiguous. This suggests that at least 30% of all recessions are followed by lower output growth later.

Focusing on those recessions followed by either a stable or an increasing output gap, we can think of three potential explanations:

[6] The total number of the recessions reported in the tables is 122 because we are not including in the sample the ones that started in the 1960s (where we lack data for the 10 year trend estimations) and the ones that started in the 2010s where we do not have enough post-crisis data. "Adjusted'' means adjusted for credit booms in the computation of the trend, as explained earlier

The first is indeed hysteresis: Recessions have lasting effects and are indeed the cause of the lower output later. A number of mechanisms have been adduced which might generate such effects. In the labor market, the recession and the associated high unemployment may lead some workers either to drop out permanently, or to become unemployable. Prolonged unemployment may lead to a change in labor market institutions, which in turn affects the natural rate later (these were the hypotheses explored by Blanchard and Summers in the 1980s to explain the increase in unemployment in Europe). Firms may invest less, leading to a lower capital stock for some time (although presumably not forever). Firms may do less R&D, leading to a permanently lower productivity level than would have been the case absent the recession. The recession may lead to lower job creation and job destruction, and thus lower reallocation and productivity growth, which is not made up later. It is fair however to say that none of these hypotheses has been conclusively shown to be empirically important.

And it is also fair to say that it is more difficult to think of mechanisms through which the recession leads to lower output growth later, to "super hysteresis". Permanently lower output growth requires permanently lower total factor productivity growth; the recession would have to lead to changes in behaviour or in institutions which lead to permanently lower R&D or to permanently lower reallocation. These may range from increased legal or self imposed restrictions on risk taking by financial institutions, to changes in taxation discouraging entrepreneurship. While these mechanisms may sometimes be at work, the proportion of cases where the output gap is increasing seems too high for this to be a general explanation.

The second is that supply shocks may be behind both the recession and the lower output later. For example, if real wages are sticky in the short run, an increase in oil prices may lead to a sharp initial recession, and, unless long run labor supply is fully inelastic, lower employment and lower output later. A financial crisis may lead to worries about liquidity and a collapse of financial intermediation in the short run; long run effects of changes in bank behaviour, or bank regulation, in the form of higher capital ratios for example, may lead to less risky but also less efficient intermediation, and lower output later. One might even argue that less efficient intermediation may decrease the efficiency of the reallocation process and generate not only lower output, but even lower growth.

The third is that the correlation reflects reverse causality: An exogenous decrease in underlying potential growth leads households to reduce consumption and firms to reduce investment, leading to an initial recession[7]. A variation on this theme is that it may take time for households and firms to realize that underlying growth has started, so that the decrease in productivity may start before the recession. Two intriguing facts support this hypothesis. A fact documented by Robert Gordon (2003), in which productivity declines at the end of the expansion.[8] A fact documented by Paul Beaudry and co-authors (2014), in which firms appear to over accumulate capital during expansions. Both are what you would expect if firms and households took some time to realize that productivity growth had actually slowed down.

One way to make some progress is to differentiate between recessions due to different factors, and see how the outcomes differ.

[7] A model along these lines is presented and estimated in Blanchard, L'Huillier and Lorenzoni (2015). The model however assumes that the news is bad news about the level of productivity, not bad news about the growth rate.

[8] Robert Gordon however offers a different interpretation of the fact. He argues that the decrease in productivity during the boom is due to over optimistic expectations by firms, which hire too many workers. He sees the recession as correcting this over hiring, and thus correcting the decrease in productivity. This however would not explain why productivity growth remains permanently lower after the recession.

In the first breakdown, we separate out those recessions associated with either financial crises, or oil price increases, and others. The motivation is straightforward: In both cases, the supply side factors behind the recession may also be behind lower output later.

The results of the financial crisis breakdowns are shown in Table 2. They are similar across the different specifications for each type of classification. They show, as one might expect, that recessions associated with financial crises, as defined in Laeven and Valencia (2013), are more likely to show a subsequent output gap, 71% on average across specifications, compared to 56% in the absence of a financial crisis. In 58% of the cases, recessions associated with financial crises are followed by an increasing output gap.

Table 2: Recessions with/without financial crises

GDP series used/ Trend Extrapolation	Scenario	Trend Calculation: 4 year window				Trend Calculation: 10 year window			
		Episodes with NO sustained gap	Episodes with sustained gap	of which:		Episodes with NO sustained gap	Episodes with sustained gap	of which:	
				Increasing over time	Stable over time			Increasing over time	Stable over time
Log Real GDP / Benchmark	With financial crisis	30%	70%	57%	13%	22%	78%	70%	9%
	Without financial crisis	37%	63%	43%	19%	43%	57%	33%	23%
Log Real GDP / Adjusted	With financial crisis	35%	65%	52%	13%	22%	78%	70%	9%
	Without financial crisis	40%	60%	41%	18%	44%	56%	33%	22%
Log Real per capita GDP / Benchmark	With financial crisis	39%	61%	48%	13%	22%	78%	61%	17%
	Without financial crisis	43%	57%	38%	18%	48%	52%	30%	21%
Log Real percapita GDP / Adjusted	With financial crisis	39%	61%	48%	13%	22%	78%	61%	17%
	Without financial crisis	45%	55%	36%	18%	48%	52%	31%	20%

Note: A total of 122 recession episodes are included in the analysis, of which 23 happened together with financial crisis (based on Laeven and Valencia 2012 definition of systemic financial crisis).

The results based on oil price changes are shown in Table 3. Recessions linked to oil price increases are more likely to show a subsequent output gap, 86% on average, compared to 54% in the rest of the cases. In 71% of the cases, recessions associated with an increase in the price of oil are followed by an increasing output gap.

Table 3: Recessions with/without oil price increases

GDP series used/ Trend Extrapolation	Scenario	Trend Calculation: 4 year window				Trend Calculation: 10 year window			
		Episodes with NO sustained gap	Episodes with sustained gap	of which:		Episodes with NO sustained gap	Episodes with sustained gap	of which:	
				Increasing over time	Stable over time			Increasing over time	Stable over time
Log Real GDP / Benchmark	With oil price increases	22%	78%	72%	6%	0%	100%	72%	28%
	Without oil price increases	38%	62%	41%	20%	46%	54%	35%	19%
Log Real GDP / Adjusted	With oil price increases	22%	78%	72%	6%	0%	100%	72%	28%
	Without oil price increases	42%	58%	38%	19%	47%	53%	35%	18%
Log Real per capita GDP / Benchmark	With oil price increases	28%	72%	67%	6%	6%	94%	72%	22%
	Without oil price increases	45%	55%	36%	19%	50%	50%	30%	20%
Log Real percapita GDP / Adjusted	With oil price increases	28%	72%	67%	6%	6%	94%	72%	22%
	Without oil price increases	47%	53%	34%	19%	50%	50%	31%	19%

Note: A total of 122 recession episodes are included in the analysis, of which 18 concided with oil prices increases (mostly during the 1970s).

In the second breakdown, we separate out those recessions associated with an increase in inflation and those associated with a decrease in inflation.[9] The motivation is also straightforward: The first set is more likely to be associated with supply shocks, which may

[9] We classify as recessions with increasing inflations those ones for which the average inflation during the year before the start of the recession is below the average inflation during the recession. Recessions with declining inflation capture the rest.

have an effect lasting for some time after the recession. The second set is more likely to be associated with demand shocks, which are less likely to be associated with those after effects.

The results are presented in Table 4. The results conform to priors, but less so than we expected. Recessions associated with increasing inflation are more likely to show subsequent lower output, with a frequency of 60%, compared to 57% for those associated with decreasing inflation. But, another way of reading the table is that, even for those recessions associated with decreasing inflation (and thus more likely to be due to demand shocks), the proportion of recessions followed by lower output is still 57% (with the large majority of those due to an increasing gap over time rather than just a larger but stable gap !).

Table 4: Recessions with/without increasing inflation

GDP series used/ Trend Extrapolation	Scenario	Trend Calculation: 4 year window				Trend Calculation: 10 year window			
		Episodes with NO sustained gap	Episodes with sustained gap	of which:		Episodes with NO sustained gap	Episodes with sustained gap	of which:	
				Increasing over time	Stable over time			Increasing over time	Stable over time
Log Real GDP / Benchmark	With increasing inflation	38%	62%	44%	18%	38%	62%	43%	19%
	With declining inflation	33%	67%	49%	18%	42%	58%	36%	22%
Log Real GDP / Adjusted	With increasing inflation	42%	58%	42%	17%	39%	61%	43%	18%
	With declining inflation	36%	64%	47%	18%	42%	58%	36%	22%
Log Real per capita GDP / Benchmark	With increasing inflation	43%	57%	40%	17%	39%	61%	38%	23%
	With declining inflation	42%	58%	40%	18%	51%	49%	33%	16%
Log Real percapita GDP / Adjusted	With increasing inflation	44%	56%	39%	17%	39%	61%	39%	22%
	With declining inflation	44%	56%	38%	18%	51%	49%	33%	16%

Note: A total of 122 recession episodes are included in the analysis, of which 77 happened with increasing inflation (the average inflation during the year before the start of the recession was below the average inflation during the recession).

"Demand shocks" comprise many different types of shocks, some of which can have lasting effects on potential output. The cleanest demand shocks we can think of are the episodes of intentional disinflations, which happened mostly in the 1980s. We identify recessions associated with intentional disinflations as those recessions characterized by a large increase in nominal interest rates, followed by a subsequent disinflation. We identify 28 such recessions. Table 5 shows the breakdown for recessions with and without intentional disinflations.

As expected, recessions associated with intentional disinflations are less likely to show a subsequent output gap, 42% on average, compared to 64% for others. But, again, the results can be read as saying that even those recessions are followed by lower output in 42% of the cases. In 21% of cases, they appear to be actually followed not only by lower output, but by lower output growth. After visual examination, at least 15% or so of the cases show a clearly increasing output gap. This proportion is small, but it is still such as to suggest that, sometimes, recessions may give rise to some form of ``superhysteresis''.

Table 5: Recessions with/without intentional disinflations

GDP series used/ Trend Extrapolation	Scenario	Trend Calculation: 4 year window				Trend Calculation: 10 year window			
		Episodes with NO sustained gap	Episodes with sustained gap	of which:		Episodes with NO sustained gap	Episodes with sustained gap	of which:	
				Increasing over time	Stable over time			Increasing over time	Stable over time
Log Real GDP / Benchmark	With intentional disinflation	50%	50%	32%	18%	57%	43%	14%	29%
	Without intentional disinflation	32%	68%	50%	18%	34%	66%	48%	18%
Log Real GDP / Adjusted	With intentional disinflation	50%	50%	32%	18%	57%	43%	14%	29%
	Without intentional disinflation	36%	64%	47%	17%	35%	65%	48%	17%
Log Real per capita GDP / Benchmark	With intentional disinflation	61%	39%	21%	18%	64%	36%	14%	21%
	Without intentional disinflation	37%	63%	46%	17%	37%	63%	43%	20%
Log Real percapita GDP / Adjusted	With intentional disinflation	61%	39%	21%	18%	64%	36%	14%	21%
	Without intentional disinflation	39%	61%	44%	17%	37%	63%	44%	19%

Note: A total of 122 recession episodes are included in the analysis, of which 28 were classified as intentional disinflation periods given that they were followed by important decreases in inflation and also accompanied by large increases in the policy rate.

To summarize: A surprisingly high proportion of recessions are followed by lower output relative to the pre-recession trend. Even more surprisingly, a large proportion of those are followed not only by lower output, but also by lower growth relative to the pre-recession trend. These proportions are larger for recessions associated with supply shocks, but not overwhelmingly so. Even for recessions plausibly induced by intentional disinflations, the proportion of recessions followed by an output gap remains high, around 42%. And the fact that some non negligible proportion of recessions due to intentional disinflations is followed by lower output growth is intriguing, and possibly suggestive of superhysteresis.

From these findings, we draw two tentative conclusions. First, to the extent that we do not have very convincing hysteresis explanations for lower growth following a recession, we conclude that the causality runs the other way, that many recessions are probably caused by the perception by households and firms of a slowdown in underlying growth, rather than the other way around. Second, to the extent that even recessions associated with intentional disinflations are associated with lower output relative to the pre recession trend, hysteresis may well be present. Both conclusions have important implications for monetary policy that we develop in the last section.

3. UNEMPLOYMENT AND INFLATION

As the crisis unfolded and GDP declined, most economists expected inflation to decrease sharply, with some forecasting a deflation spiral, along the lines of what had been observed in the Great Depression. As Figure 4, which plots inflation in the United States, the euro area, UK, and Japan since 2007, shows, inflation indeed declined, and in some countries, has now turned into deflation, but deflation has remained limited.

Much of the reason clearly comes from the changes in the way people and firms form expectations of inflation. As has been documented by many, the shift to inflation targeting and stable inflation for the two decades preceding the crisis have led forecasts of future inflation to put less weight on past inflation, and more weight on the perceived target of the central bank. This in turn has led to a shift from an “accelerationist Phillips curve”, in which the unemployment gap or the output gap led to a change in inflation, to something closer to a “level Phillips curve”, in which the gap is associated with a level of inflation.

Figure 4: Advanced Economies CPI Headline Inflation (percent, year over year)

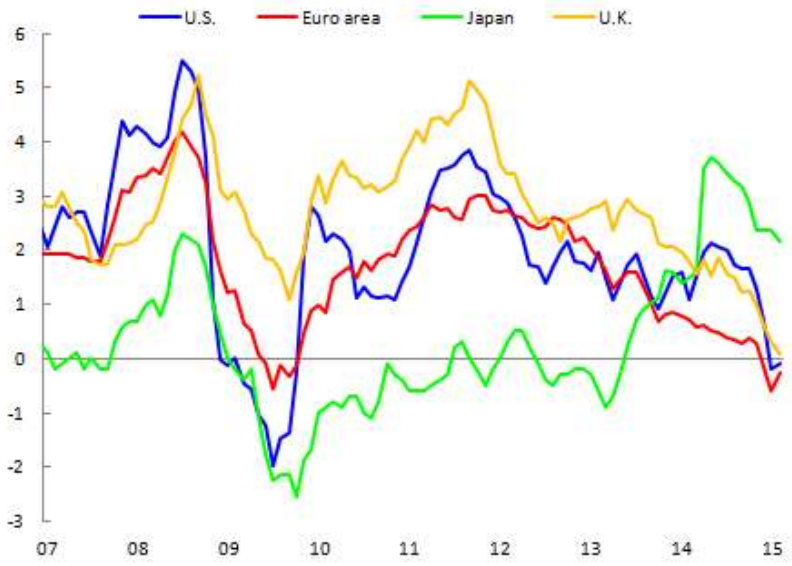

Source: IMF, *Global Data Source*.

The empirical evidence suggests however that more has been at work, namely that, controlling for expected inflation, the effect of the unemployment gap (i.e. the distance between the actual and natural unemployment rates) on inflation has steadily diminished over time. This was in particular the conclusion from the IMF April 2013 WEO (Chapter 3) study, which we extend here.[10]

The study, which was based on data from 20 countries since 1960 showed the results of estimation of the following relation (see Matheson and Stavrev (2013) for more specification and estimation details):

$$\pi_t = \theta_t(u_t - u_t^*) + \lambda_t \pi^e{}_t + (1-\lambda_t)\pi_{t-1} + \mu_t \pi_{mt} + \varepsilon_t \qquad (1)$$

where π_t is headline CPI inflation, u_t is the unemployment rate, u_t^* is the natural rate, $\pi^e{}_t$ is long term inflation expectations, π_{t-1} is one quarter lag of headline inflation (measured year-over-year), and π_{mt} is import price inflation relative to headline inflation (measured as deviations from average).

The parameters λ_t *(the coefficient reflecting the stability of inflation expectations),* θ_t *(the slope of the Phillips curve),* and μ_t (*the coefficient reflecting the importance of import-price inflation*), as well as the natural rate, which is unobservable, are assumed to follow constrained random walks (θ_t and $\mu_t \geq 0$, and $0 \leq \lambda_t \leq 1$).[11]

The system is estimated separately for each country, by maximum likelihood, using a non linear Kalman filter.

Figure 5 shows median estimates for λ_t and θ_t, the two coefficients we focus on, together with the interquartile range of estimates across countries. Figure 6 shows estimates for the US and Germany, two countries which are representative of other countries, together with one-standard deviation bands. (The results for other countries are presented in the web appendix.)

Figure 5 confirms the two conclusions of the earlier IMF study:

Since the mid 1970s, short-run inflation expectations have become more stable (λ_t has increased), and

The slope of the Phillips curve (θ_t) has flattened over time. Much of the decrease took place from the mid 1970s to the mid 1990s. The coefficient does not appear to have decreased further during the crisis.

What figure 5 does not show however is that, for most countries, the coefficient θ_t today is not only small, but statistically insignificant. This can be seen for example in Figure 6. In both the US and Germany, the one-standard-deviation band reaches the horizontal axis some time in the mid 1990s, and remains there thereafter (the estimated coefficient is constrained to be non negative).

[10] Capturing in a more poetic way the argument in the previous paragraph, the title of the study was called ``The dog which did not bark".

[11] Further details about the specification of the equation are given in appendix 1 of the IMF chapter

Figure 5: Median estimates (across countries)

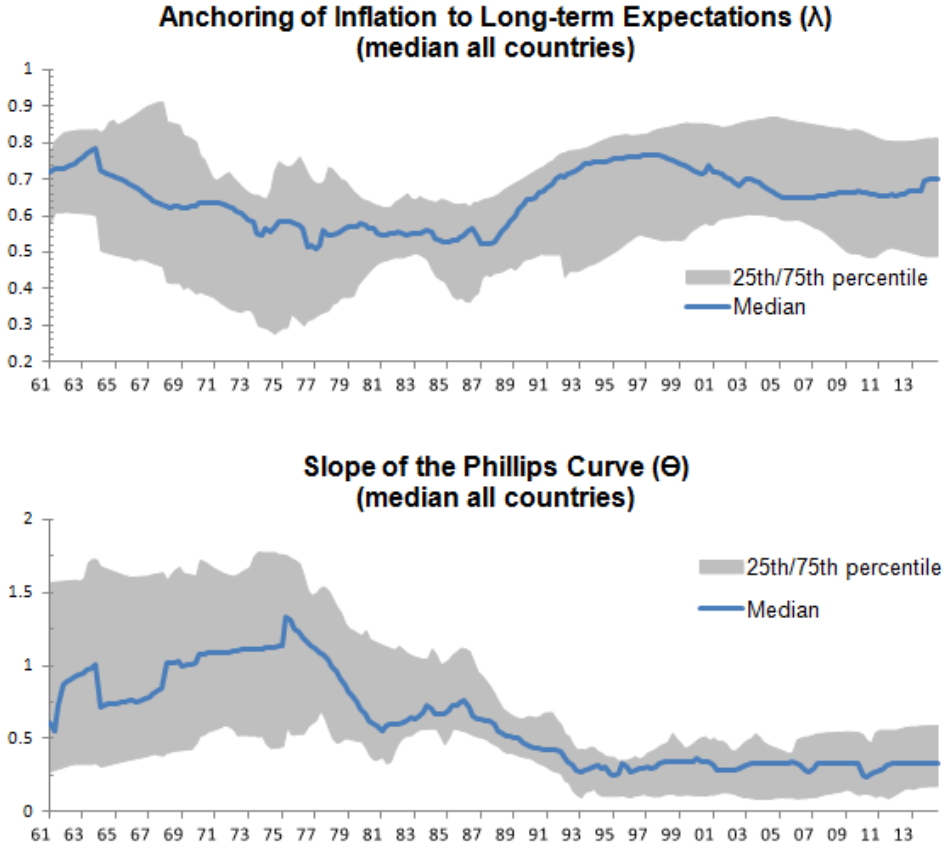

Figure 6: Estimates for Germany and the UK

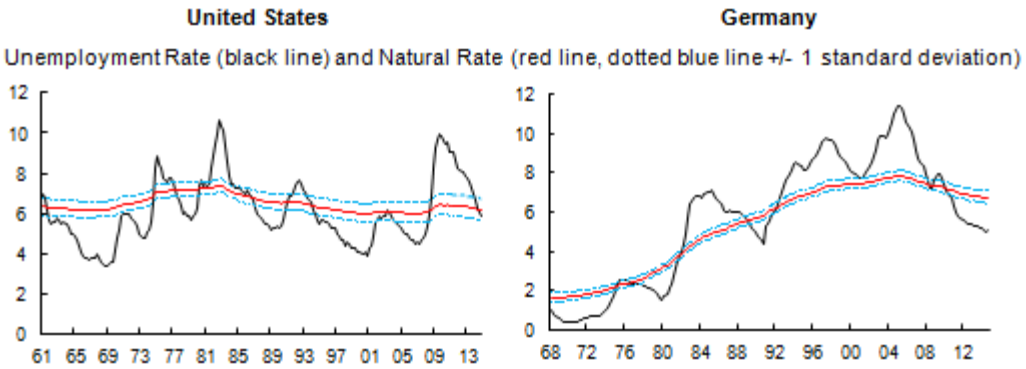

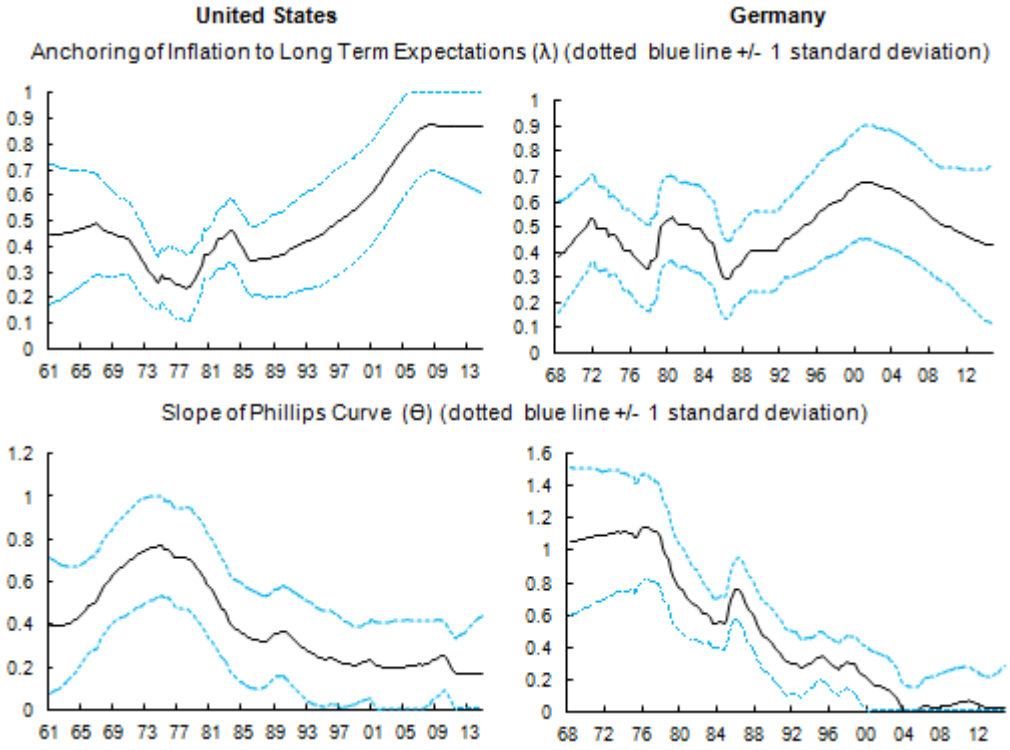

Table 6 gives the estimated coefficient θ_t for each of the 20 countries for three dates, 1985, 2000, and 2014. For most countries, θ_t has steadily declined from the mid-1980s on, from a median value (across countries) of about 0.7 to about 0.3 now. For 16 out of the 20 countries, the coefficient is no longer significantly different from zero (using a t-statistic of 1.3). This was not the case in 1985 when all but 2 coefficients were statistically significant.

Table 6: Coefficients on the unemployment gap, in 1985, 2000, 2014.

Country	Date	Coef	Stdev	Country	Date	Coef	Stdev
United States	1985	0.34	0.23	Switzerland	1985	3.32	1.60
United States	2000	0.23	0.18	Switzerland	2000	0.52	1.05
United States	2014	0.16	0.28	Switzerland	2014	0.96	1.24
Japan	1985	4.97	1.67	Sweden	1985	0.73	0.71
Japan	2000	0.78	0.98	Sweden	2000	0.61	0.57
Japan	2014	3.45	2.98	Sweden	2014	0.76	0.94
Germany	1985	0.72	0.18	Belgium	1985	0.67	0.22
Germany	2000	0.17	0.20	Belgium	2000	0.51	0.42
Germany	2014	0.03	0.25	Belgium	2014	0.56	0.76
United Kingdom	1985	0.80	0.37	Norway	1985	0.84	0.38
United Kingdom	2000	0.02	0.46	Norway	2000	0.43	0.39
United Kingdom	2014	0.24	0.94	Norway	2014	0.56	0.64
France	1985	1.11	0.21	Austria	1985	0.59	0.27
France	2000	0.38	0.33	Austria	2000	0.58	0.27
France	2014	0.65	0.42	Austria	2014	0.57	0.28
Italy	1985	1.31	0.34	Denmark	1985	0.58	0.29
Italy	2000	0.04	0.42	Denmark	2000	0.11	0.30
Italy	2014	0.40	0.37	Denmark	2014	0.22	0.39
Canada	1985	0.56	0.21	Ireland	1985	0.72	0.30
Canada	2000	0.35	0.26	Ireland	2000	0.29	0.23
Canada	2014	0.08	0.47	Ireland	2014	0.24	0.40
Australia	1985	0.27	0.31	Greece	1985	-	-
Australia	2000	0.99	0.50	Greece	2000	0.15	0.11
Australia	2014	0.08	0.79	Greece	2014	0.15	0.11
Spain	1985	0.39	0.11	Portugal	1985	1.88	1.12
Spain	2000	0.11	0.14	Portugal	2000	0.38	1.06
Spain	2014	0.18	0.10	Portugal	2014	0.22	0.91
Netherlands	1985	0.33	0.13	New Zealand	1985	1.07	0.59
Netherlands	2000	0.34	0.13	New Zealand	2000	0.10	0.76
Netherlands	2014	0.33	0.14	New Zealand	2014	0.98	1.22

There are however good reasons to explore the evidence further. The appropriate specification of the Phillips curve, if there is indeed one, is far from settled. Importantly, the issue we dealt with in the previous section is highly relevant. To the extent that potential output moves with or because of movements in actual output, the empirical measures of the output gap, and (perhaps a lesser extent) of the unemployment gap may not be the right ones. To take an extreme and implausible case, maybe the lack of further deflation during the crisis is largely explained by the fact that the unemployment gap has been much smaller than assumed in the above equation. While the specification allows the natural rate to move over time, it forces a specific type of movement, namely a random walk with small variance of innovations.

Again, there is no simple solution to these issues, and so, as in the previous section, we explore a few alternative specifications.

Under the first alternative, we allow the natural rate to depend partly on the past actual rate, reflecting a crude form of hysteresis. More specifically, we assume:

$$u_t^* = bu_{t-1}^* + (1-b)u_{t-1} + \eta_t \qquad (2)$$

where we estimate b for each country, constraining b to satisfy $0.9 \leq b \leq 1$.

Figure 7 presents the evolution of the median coefficients and interquartile ranges. As the natural rate follows the actual rate more closely, the unemployment gap is correspondingly smaller, and the coefficient on the unemployment gap (θ_t) correspondingly larger. But the declining evolution of this coefficient is similar as before if we focus on the interquartiles, and a bit less dramatic in relation with the median country that went from a value of about 1 in the mid 1970s to about 0.3 today.

Figure 8 shows the evolutions of the coefficients for the US and Germany, which are again representative of other countries. The estimation give a value of b equal to the lower bound (0.9) for the US and 0.91 for Germany, and the estimated natural unemployment rate now follows more closely the actual rate. The coefficient on the unemployment gap steadily decreases over time, becoming insignificantly different from zero at the end of the sample. More generally, for 15 out of the 20 countries, the estimated coefficient for 2014 is not significantly different from zero.

Figure 7. Median estimates, allowing for hysteresis

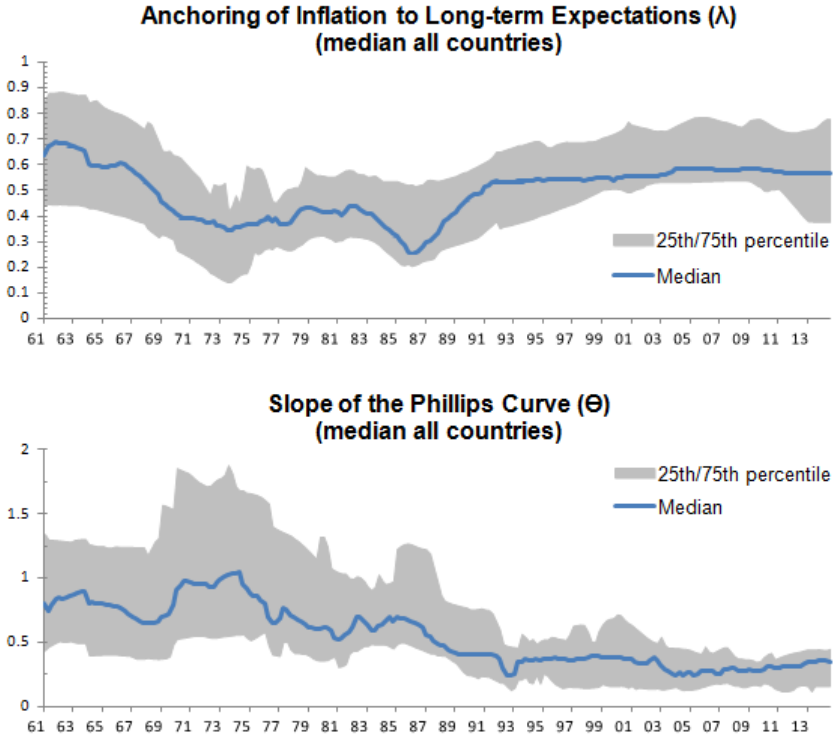

Figure 8: Specification with Hysteresis, Germany and UK

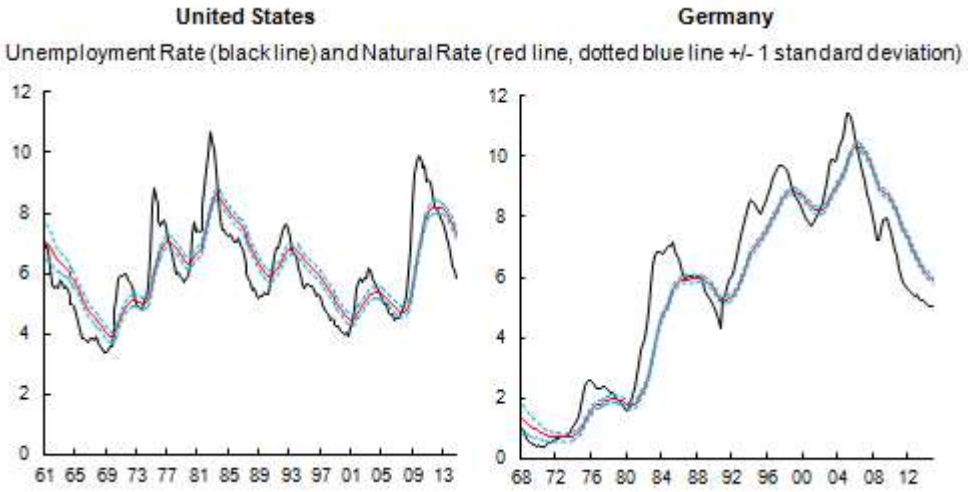

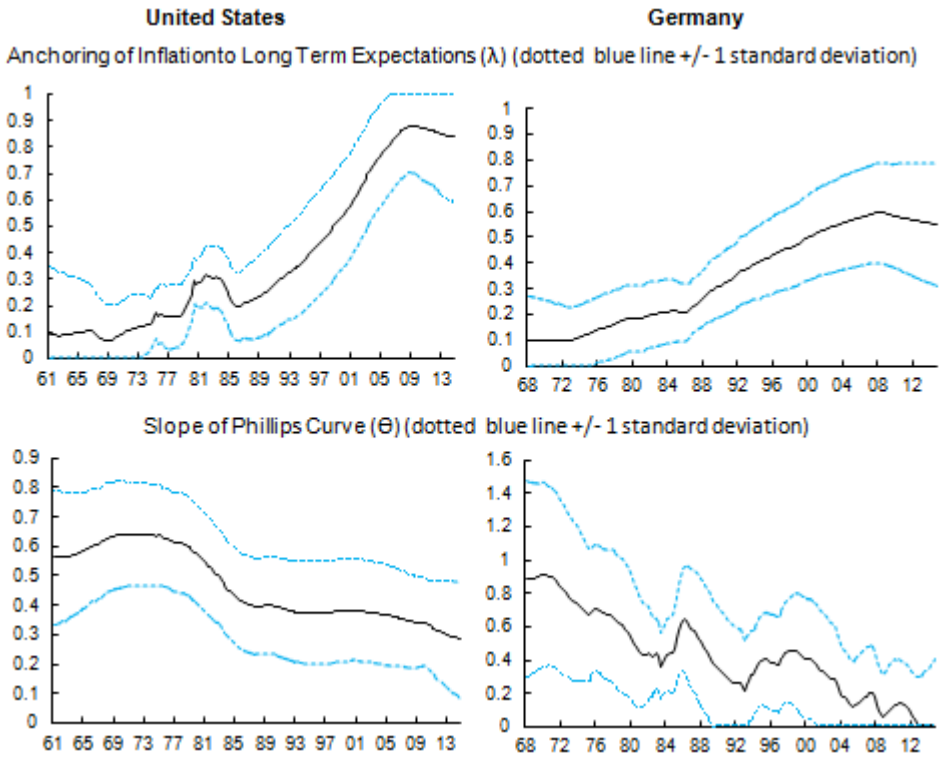

Under the second alternative, we replace the unemployment rate with the short term unemployment rate, another crude way to capture the notion that the long term unemployed matter less in the determination of wages, and in turn the determination of prices. Theoretical arguments for why the long term unemployed matter less in wage setting have been explored in a number of papers, and range from ranking by firms of applicants by unemployment duration, to loss of skills or morale, making the long term unemployed less employable, and less relevant to wage formation. A specification including only short term unemployment is surely too strong, but is a useful starting point (Ball and Mazumder [2014] have argued that such a specification works well in the US.).

The results are presented in Figure 9 for the cross-country summary, and in Figure 10 for the US and Germany. The period of estimation is shorter, due to lack of available data. Figure 9 suggests a stronger increase in the anchoring of expectations than in the benchmark, and less of a decrease in the slope. This smaller decrease reflects more an heterogeneity of evolutions than a general decrease. Figure 10 shows that, for Germany for example, the slope of the Phillips curve still steadily decreases, being insignificant for most of the sample. For the US, Figure 10 shows that the slope of the Phillips curve also steadily decreases but it is marginally significant at the end of the sample. In contrast, in countries such as Italy or Spain, the coefficient increases, becoming marginally significant at the end of the sample. Overall, the coefficient remains insignificant today in 13 out of the 20 countries

Figure 9: Median estimates. Specification with Short-term Unemployment

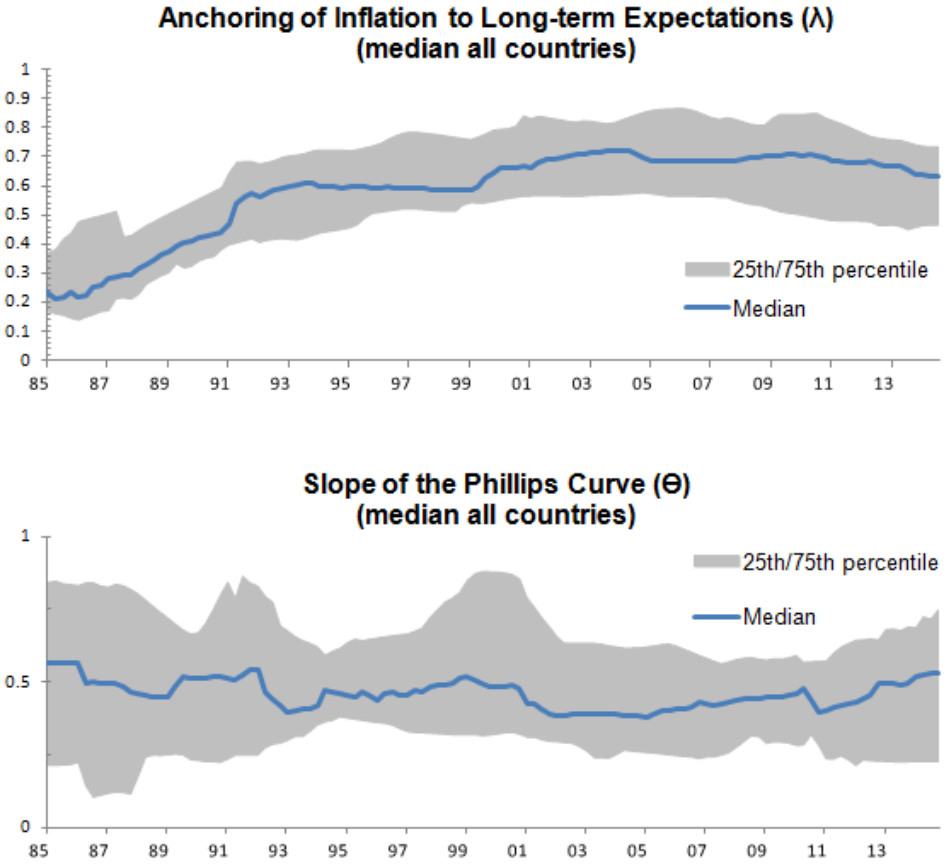

Figure 10: Specification with Short-term Unemployment, UK and Germany

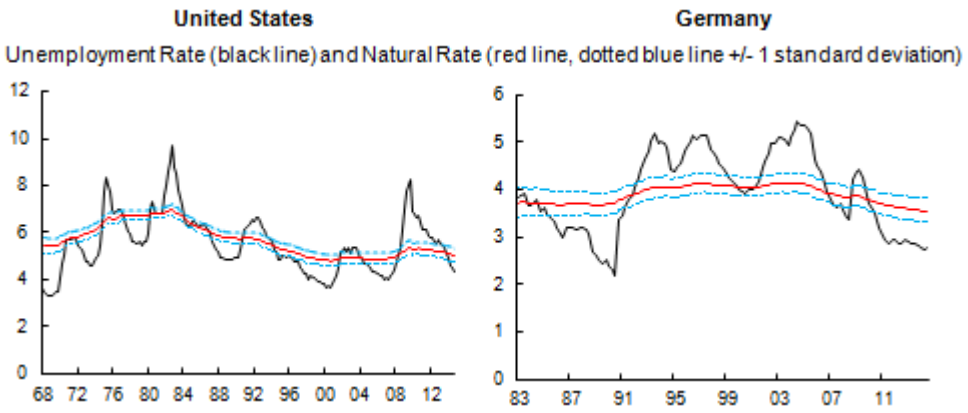

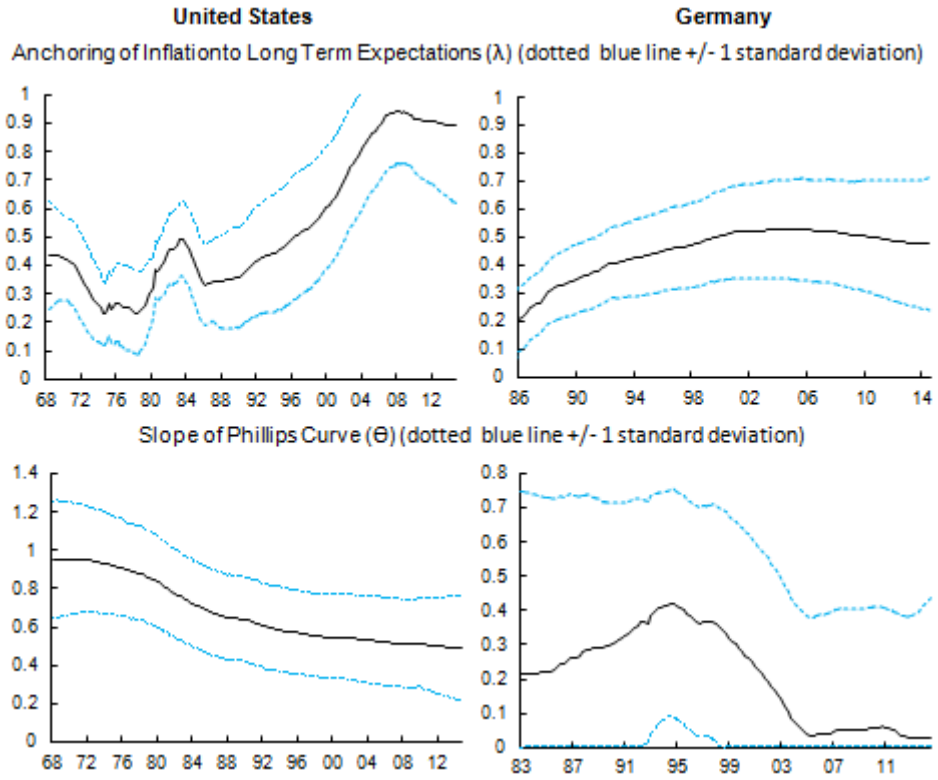

To summarize: It is clear that that the slope of the Phillips curve has decreased over time in most countries. Most of the decline has taken place from the mid 1970s to the early 1990s. The coefficient has not decreased further during the crisis. (Thus, it does not look likely that the lower coefficient comes from the zero lower bound on wage growth, which has only been binding only for the past few years).

In many cases, we find that the coefficient is not significantly different from zero. That particular conclusion again comes with a number of caveats.

The very nature of the Kalman filter may lead to large estimated standard deviations; in a number of cases, a much rougher econometric approach based on fixed coefficients over 10-year rolling samples, yields smaller standard deviations, and higher significance. For example, as of 2014, 6 countries have t-statistics above 2 and 12 countries have t-statistics above 1.

A common specification, while it imposes discipline, may not do justice to the specificities of each country. Indeed, for a number of countries in the sample, we are aware of specifications and choices of inflation and gap variables, which yield more significant estimates (although they typically also find a decrease in the slope of the Phillips curve over time) (for example Broadbent (2014) for the UK, Ball and Mazumder (2014) and Coibion and Gorodnishenko (2015) for the US, Andrle et al (2013) for the euro area). In allowing for different specifications across countries, the line between flexibility, and potential data mining, is a fuzzy one.

To draw strong policy implications, one would want to identify why the slope has indeed become smaller over time, whether it comes from wage or from price behaviour, whether it comes from changes in the structure of wage bargaining, or in the pricing behaviour of firms in the product market. This would go far beyond what we have done.

3. IMPLICATIONS FOR MONETARY POLICY

Based on the conclusions from our empirical work, what are the implications for monetary policy? A full answer would require more certainty than we have about the conclusions, and, as we indicated earlier, an understanding of the underlying mechanisms behind hysteresis if indeed present, or the decrease in the slope of the Phillips curve, both of which remain rather mysterious. We still feel we can draw the following conclusions:

The findings of the first section have potentially conflicting implications:

To the extent that recessions are caused by an underlying decrease in growth, there is the risk of overestimating potential output during and after the recession, and by implication of overestimating the output gap. This may in turn lead to too strong a response of monetary policy to output movements during and after the recession. (It is indeed often the case that estimates of output gaps associated with recessions are revised down ex post.) This is illustrated in Figure 11 below. Suppose that after time t, potential growth decreases, and that it takes a while for firms and households to realize it. For some time, growth will continue at close to the old trend, until the adjustment of expectations leads to a recession. If, in real time, the central bank constructs the output gap under the assumption that the underlying trend has not changed, the negative output gap will be measured by the sum of the orange area and the right blue area in the picture, whereas the true negative output gap will be given by the right blue area only. Only over time, will it become clear what the correct output gaps (blue areas) were and what monetary policy should have been.

Figure 11. Decreases in growth, recessions, and output gaps.

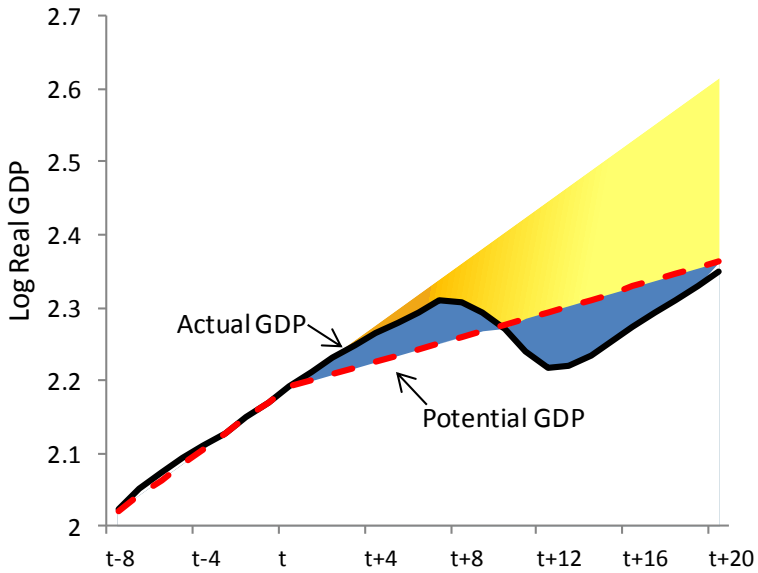

To the extent however that hysteresis is present, it implies that deviations of output from its optimal level are much longer lasting and thus more costly than usually assumed. The implication is straightforward, namely that monetary policy should react more strongly to output

movements, relative to inflation.[12] It also implies that output gaps, as opposed to the distance of output from its optimal level, may give the wrong signals for the conduct of monetary policy.

The findings of the second section potentially raise major issues for the conduct of monetary policy. If the output gap only has a small effect on inflation, this suggests that stabilizing inflation would require very large movements in the output gap. If, in addition, the effect is not only small but also uncertain, even large movements in the output gap may not succeed in controlling inflation. This suggests that monetary policy should focus on stabilizing the output gap rather than inflation.[13] Indeed, in the limit, if the output gap does not affect inflation, monetary policy should focus only on the output gap. The obvious question in this case is what determines inflation. As is clear also from our results, what appears to determine inflation at this point, is largely inflation expectations, themselves anchored to the central bank target level of inflation.

But here lies the puzzle and the challenge: Put starkly, what we have observed is an increase confidence in the central bank meeting its inflation target, while at the same time, the ability of the central bank to achieve that target has steadily decreased. Why should people trust the central bank to achieve its target, why should inflation remain anchored? And, if it doesn't, what does imply for monetary policy in the future?

REFERENCES

Andrle, M., Bruha, J., and Solmaz, S., (2013) "Inflation and output comovement in the Euro area: Love at second sight*", International Monetary Fund*, working paper 13/192

Ball, L. (2014), "Long-term damage from the Great Recession in OECD countries", NBER Working Paper, No. 20185, May.

Ball, L. and Mazumder, S., (2014), ``A Phillips curve with anchored expectations and short term unemployment'', NBER working paper 20715, November

Beaudry P., Galizia, D. and F Portier (2014), "Reconciling Hayek's and Keynes' views of recessions", CEPR Discussion Paper 9966, and also NBER Working Paper 20101.

Blanchard, O. L'Huillier J.P. , and G. Lorenzoni, (2013), "News, Noise, and Fluctuations: An Empirical Exploration", *American Economic Review*, 103:7, pp. 3045-3070.

Blanchard, O. and L. Summers, (1986), "Hysteresis and European Unemployment," In Stanley Fischer (ed.), NBER Macroeconomics Annual. MIT Press, September, pp. 15-77

Broadbent, B. (2014), "Unemployment and the conduct of monetary policy in the UK", Jackson Hole, August.

Cerra, V. and Saxena, S. (2008), "Growth Dynamics: The Myth of Economic Recovery," *American Economic Review*, 98:1, pp. 439-457.

[12] . Fiscal implications, when monetary policy is constrained by the zero lower bound, were explored by De Long and Summers (2012).

[13] One caveat: This argument is based on the assumption that the change in the slope coefficient does not affect the weights of inflation and the output gap in the objective function. In formal NK models, this may not be the case (see for example Woodford 2003, chapter 6.

Coibion, O. and Gorodnishenko, Y., (2015) , “Is the Phillips curve alive and well after all? Inflation expectations and the missing disinflation’’, *American Economic Journal*, 7:1, pp. 197-232

Dell’Ariccia, G., Igan, D., Laeven, L. and Tong, H., (2012), Policies for Macrofinancial Stability:How to Deal with Credit Booms*,” IMF Staff Discussion Note* 12/06.

De Long, B. , and L. Summers, (2012), “Fiscal Policy in a Depressed Economy”, BPEA, Spring, pp 233-274

Gordon, R., (2003), “Exploding productivity growth: Context, causes, and implications”, *Brookings Papers on Economic Activity*, (2), pp. 207-298.

Harding, D. and Pagan, A., (2002), "Dissecting the cycle: a methodological investigation," *Journal of Monetary Economics*, 49(2), pp.365-381.

IMF World Economic Outlook (2013), Chapter 3, “The dog that didn’t bark: Has inflation been muzzled or was it just sleeping?”, April

Laeven, L. and F. Valencia., (2013), “Systemic Banking Crises Database," *IMF Economic Review*, 61(2), pp. 225-270.

Martin and Wilson (2013), “Potential Output and Recessions: Are We Fooling Ourselves?” manuscript , Board of Governors.

Matheson, T. and Stavrev, E., (2013), “The Great Recession and the inflation puzzle,” *Economics Letters*, 120(3), pp. 468-472.

Woodford, M., (2003), “Interest and prices”, Princeton University Press