

NBER WORKING PAPER SERIES

INFLATION AND ACTIVITY – TWO EXPLORATIONS AND THEIR MONETARY
POLICY IMPLICATIONS

Olivier Blanchard
Eugenio Cerutti
Lawrence Summers

Working Paper 21726
<http://www.nber.org/papers/w21726>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2015

Extended version of the paper presented at the ECB Forum on Central Banking in Sintra, Portugal on “Inflation and unemployment in Europe” on May 22, 2015. We thank Larry Ball and Sandeep Mazumder for comments and help, as well as Yangfan Sun and Daniel Rivera for excellent research assistance. Comments by Larry Ball, who was our discussant, led to substantial changes in the second part of the paper. We also thank Zeno Enders, Stephan Danninger, Chris Erceg, Jaewoo Lee, and other IMF colleagues for comments. Our paper builds on Martin and Wilson (2013) and IMF (April 2013) Chapter 3. This paper was also published as IMF WP/15/230. The online appendix is available at: <http://www.imf.org/external/pubs/ft/wp/2015/wp15230app.pdf>. The views expressed herein are those of the authors and do not necessarily reflect the views of the IMF, its board of directors, or the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2015 by Olivier Blanchard, Eugenio Cerutti, and Lawrence Summers. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Inflation and Activity – Two Explorations and their Monetary Policy Implications
Olivier Blanchard, Eugenio Cerutti, and Lawrence Summers
NBER Working Paper No. 21726
November 2015
JEL No. E31,E32,E50

ABSTRACT

We explore two issues triggered by the crisis. First, in most advanced countries, output remains far below the pre-recession trend, suggesting hysteresis. Second, while inflation has decreased, it has decreased less than anticipated, suggesting a breakdown of the relation between inflation and activity. To examine the first, we look at 122 recessions over the past 50 years in 23 countries. We find that a high proportion of them have been followed by lower output or even lower growth. To examine the second, we estimate a Phillips curve relation over the past 50 years for 20 countries. We find that the effect of unemployment on inflation, for given expected inflation, decreased until the early 1990s, but has remained roughly stable since then. We draw implications of our findings for monetary policy.

Olivier Blanchard
Peterson Institute for International Economics
1750 Massachusetts Avenue, NW
Washington, DC 20036
and NBER
oblanchard@piie.com

Lawrence Summers
Harvard Kennedy School of Government
79 JFK Street
Cambridge, MA 02138
and NBER
lhs@harvard.edu

Eugenio Cerutti
Research Department
International Monetary Fund
700 19th Street N.W.
Washington, DC 20431
Ecerutti@imf.org

Content

Abstract	1
I. Introduction	3
II. Aftereffects of Recessions: Hysteresis?	5
A. Measuring the Effects of Recessions: A Non-Parametric Approach	6
B. What could be behind the results? Hysteresis and alternative hypotheses	12
C. Controlling by Type of Recession	13
D. Relative Output Gaps	16
E. Main Takeaways	17
III. Does Unemployment Affect Inflation?	18
A. Empirical Estimations of Phillips Curves: A Time-Varying Approach	18
B. Empirical Estimations of Phillips Curves: A Simpler Approach	22
IV. Implications For Monetary Policy	24
Tables	
1. Analysis of the differences between output level and trend across recessions	11
2. Recessions with/without financial crises	14
3. Recessions with/without oil price increases	14
4. Recessions with/without increasing inflation	15
5. Recessions with/without intentional disinflations	16
6. Slope of the Phillips Curve	23
Figures	
1. Advanced Economies Real GDP	5
2. Actual GDP and Extrapolated Trends	7
3. United States –Evolution of log real GDP and Extrapolated Trends	8
4. Portugal –Evolution of log real GDP and Extrapolated Trends	9
5. Histogram of Average Adjusted Output Gaps	10
6. Histogram of Average Relative Output Gaps	17
7. Advanced Economies CPI Headline Inflation	18
8. Median estimates	20
9. Estimates for Germany and the UK	21
10. Decreases in growth, recessions, and output gaps	26
References	27

I. INTRODUCTION

We explore 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, and, more generally, the effect of recessions on output. 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 for stabilisation 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 these motivations 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 and draw the conclusions set out below.

We find that a high proportion of recessions, about two-thirds, are followed by lower output relative to the pre-recession trend even after the economy has recovered. Perhaps more surprisingly, in about one-half 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 the pre-recession trend increases.

If these correlations are causal, they suggest important hysteresis effects and even “super-hysteresis” effects (the term used by Laurence Ball (2014) for the impact of a recession on the growth rate rather than just the level of output). Correlation however does not imply causality. The causality may indeed run from the recession to lower output later, and hysteresis or super-hysteresis may be at work. However, the correlation may instead reflect common third factors. Supply shocks, such as an increase in oil prices, or a financial crisis, may be behind both the initial recession and lower output later. Alternatively, 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 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 nearly two-thirds are associated with lower output later and that a significant fraction of those are associated with lower output growth.

We draw two tentative conclusions. It is likely that, in many cases, the correlation between recessions and subsequent poor economic performance reflects 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 intentional disinflations are also often followed later by lower output, or even, in some cases, lower output growth, suggests that hysteresis, and perhaps even super-hysteresis may indeed also be at work. Both conclusions have radically different, but important, implications for monetary policy, to which we shall come back below.

Turning to the Phillips curve relation, we start by estimating, for each country, a relation between inflation, expected and lagged inflation, and a measure of the unemployment gap. The specification allows for both the natural rate and the coefficients to evolve over time. We confirm that the coefficient on long-term expected inflation (as opposed to the coefficient on lagged 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 substantially decreased since the 1970s. Most of the decrease, however, took place before the early 1990s. Since then, the coefficient appears to have been stable, and, in most cases, significant; indeed it does not appear to have decreased during the crisis.

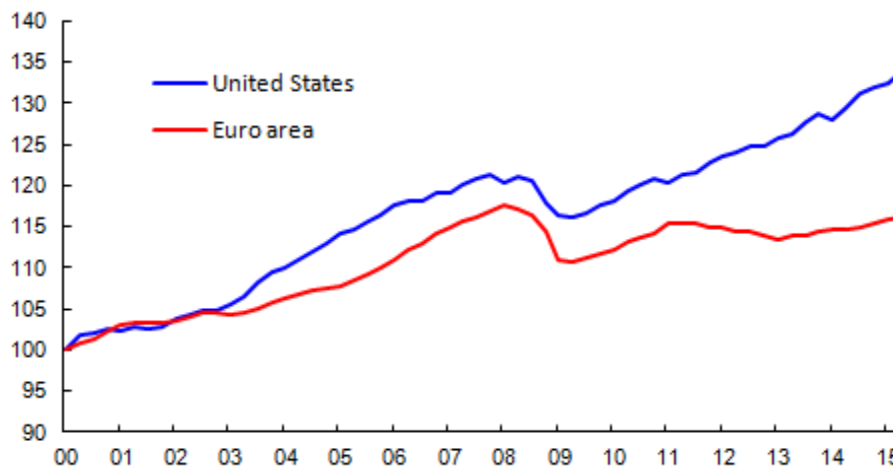
Finally, in the last section, we explore the implications of our findings for monetary policy. The findings of the second section have opposite implications for monetary policy depending on their interpretation. To the extent that recessions are due to the perception or anticipation of lower underlying growth, 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 hysteresis or super-hysteresis effects, then the cost of allowing downward movements in output in response to shifts in demand increases implies that a stronger response to output gaps is desirable.

The findings of the third section yield less dramatic conclusions. To the extent that the coefficient on the unemployment gap, while small, remains significant, the implication is that, within an inflation targeting framework, the interest rate rule should put more weight on the output gap relative to inflation. A more general conclusion is that this small coefficient reinforces the case for a dual mandate: stabilizing inflation may require very large changes in the unemployment gap, and lead to large welfare losses.

II. AFTEREFFECTS OF RECESSIONS: HYSTERESIS?

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 the issue again to the fore. The reason is not hard to see, and is shown in Figure 1. The Figure shows the evolution of output in the United States and in the euro area 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. One can indeed 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. However, correlation does not imply causality. 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 argue, although less plausibly in this case, that the recession was partly owing

to the anticipation of lower growth to come.²

These considerations have led us to look at a much larger set of recessions, over many countries and many years and to 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 that 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. The rest of the section presents our results.

A. Measuring the Effects of Recessions: A Non-Parametric Approach

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).³ Our contribution is in using a slightly different methodology, looking at the effect of recessions conditional on different types of shocks, and in the interpretation of the results. We rely on a non-parametric method, focused on recessions rather than on fluctuations more generally.

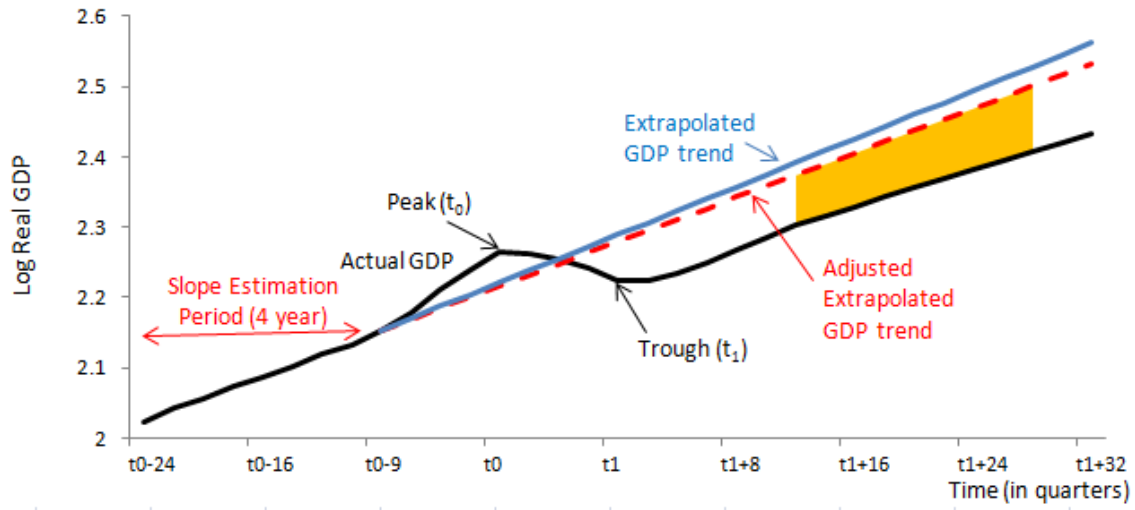
Identifying Recessions

We define 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, for example t_0 and t_1 in Figure 2 below.⁴ We keep all the recessions for which we have sufficient data before and after the recession to estimate pre-recession trends and post-recession gaps. This forces us to ignore the recessions of the 1960s and those of the 2010s, and leaves 122 recessions.

² “Less plausibly in this case”, because, as far as we know, nobody mentioned such an underlying decrease in growth at the start of the Great Financial Crisis. More recently, however, research has concluded that, at least in the United States, there was indeed a slowdown in productivity starting a few years before the crisis (Fernald 2014).

³ Martin and Wilson build in turn on Cerra and Saxena (2008), which, using an autoregressive projection model and consensus surveys, documented the behaviour of output following financial crises and civil wars.

⁴ Following Harding and Pagan (2002), we set to two quarters the number of observations on both sides over which local minima and maxima are computed, to two quarters the minimum duration in every contraction or expansion phase, and to five quarters the minimum duration between two peaks and two troughs.

Figure 2: Actual GDP and Extrapolated Trends.

Estimating Pre-recession Trends

The first issue is how to take into account the fact that the economy may have been in a boom, and thus above trend, before the recession started. We explore two alternatives. The first, similar to Martin and Wilson (2013), 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 alternative 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 choose the starting point of the trend extrapolation, we exclude at least the two years prior to 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 Crowe et al. (2011), 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. In only eleven recessions does this additional constraint bind, so that, for most recessions, the trend is anchored at a point in time corresponding to two years before the recession.

The second issue is the length of time to be used in estimating that trend. We explore again two alternatives, one in which the (log linear) trend is estimated over four years (so, in the absence of a credit boom, over t_0-9 to t_0-24 , as we measure time in quarters, in Figure 2). 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 ten

years, so, in the absence of a credit boom, over t_0-9 to t_0-48 . The potential shortcoming is the symmetrical risk that this may not capture recent changes in the underlying trend.⁵

The various alternatives (the two choices for the start of the pre-recession trend, and two periods of estimation for this trend) give us four different combinations. We derive results in each country for each of the four combinations. The graphs giving actual log real GDP, recession dates, and estimated trends, are given in a web appendix. Figures 3 and 4, which are based on a four-year trend anchored two years before the recession, give a flavor 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.

Figure 3 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 also to be associated with a lower level of output relative to trend.

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

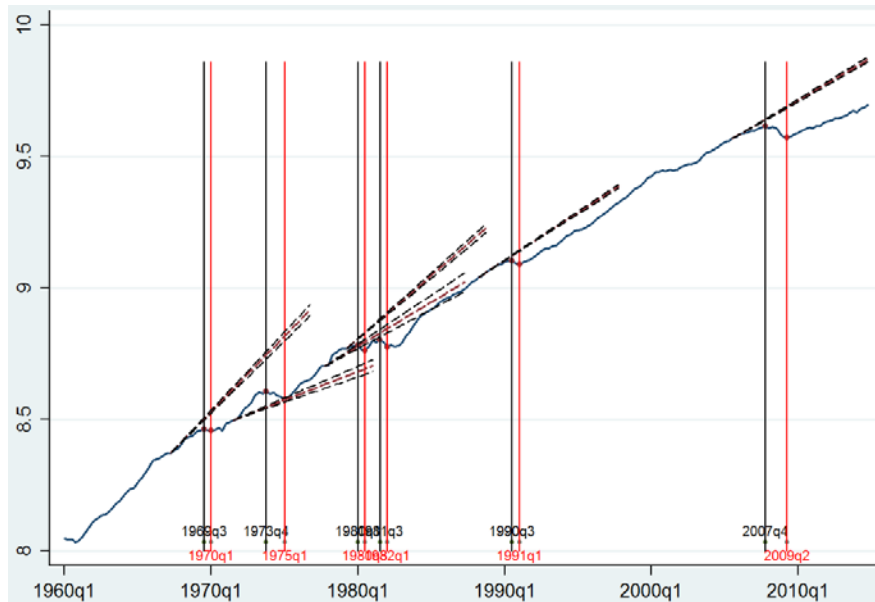
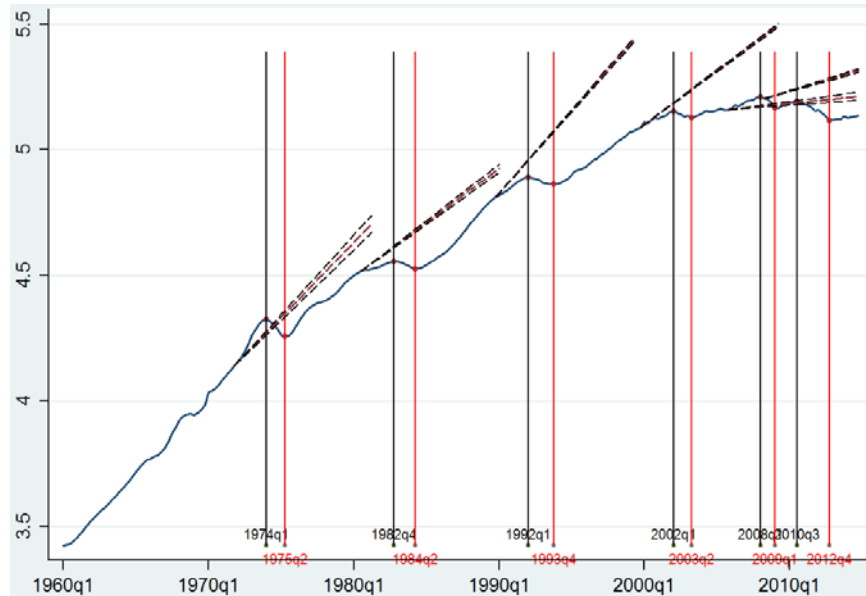


Figure 4 shows the evolution of Portugal and is representative of the evolutions of a number of European countries. All but one of the recessions since 1960 appear to be associated not

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.

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 4: Portugal –Evolution of log real GDP and Extrapolated Trends



To reduce the probability of finding a spurious post-recession gap, we made two additional adjustments.

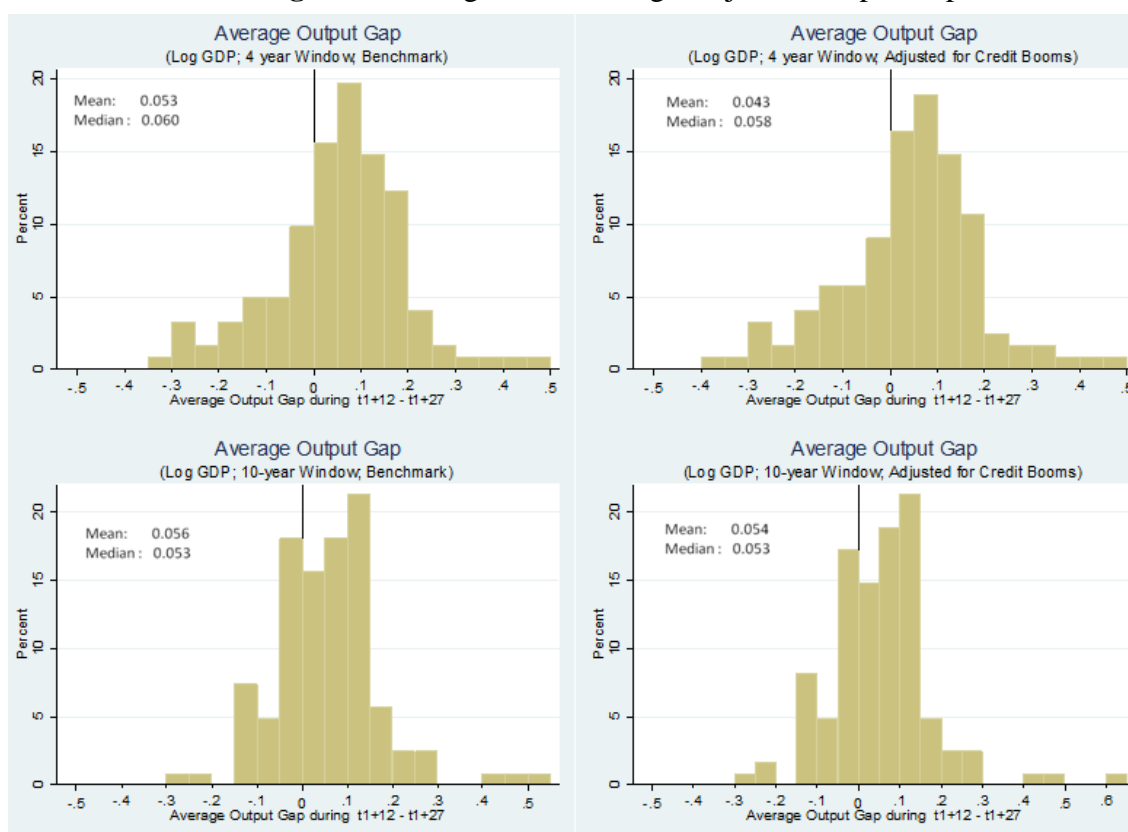
First, given the uncertainty about the exact value of the estimated time trend, we adjust down the slope of the trend using the estimated coefficient minus one standard deviation (e.g. the adjusted extrapolated trend in Figure 2 or the lower bands in Figures 3 and 4).

Second, and more empirically importantly, one needs to take into account the fact that, since the 1960s, most countries have experienced declining growth. Surely, much of this decline reflects a secular decrease in growth, which has nothing to do with the potential after-effects of recessions. This implies that, on average, log linear time trends estimated on earlier data will tend to overpredict later output, and, thus, will generate spurious gaps between later output and the time trend estimated on earlier data. To account for this, we first regress log GDP on linear and quadratic time trends over the whole sample. Then, for each recession, and for each post-recession quarter, we adjust down the estimated pre-recession trend by a factor equal to the coefficient on the quadratic term times the square of the difference between the mid-range of the trend estimation period (t_0-16 and t_0-28 for the four and ten

year windows respectively) and the relevant quarter.⁶ This adjustment turns out to be empirically significant, and to substantially reduce the number of cases where we find output to be below the adjusted pre-recession trend.

We define output gaps as the difference between the adjusted pre-recession trend and actual log GDP. We define post-recession output gaps as the average output gaps from three to seven years after the recession, thus, in terms of quarters, from t_1+12 to t_1+27 (the shaded area in Figure 2). Figure 5 plots the resulting distributions of the output gaps for each of our four different sets of assumptions. In all cases, the means and the medians of the distribution are positive. Thus, on average, output is lower than the pre-recession trend.

Figure 5: Histogram of Average Adjusted Output Gaps



Finally, to summarize results here and below, we use the following classifications. If the average gap during t_1+12 to t_1+27 is non-positive, we classify the recession as having no sustained gap. If the average gap is positive, we classify the recession as having a sustained

⁶ To understand the adjustment, note that, if the coefficient on the quadratic term is c , the derivative of the growth rate relative to time is equal to $2c$, and thus, after n quarters, log output will be lower relative to a linear trend by $2c(0+1+\dots+(n-1)) = cn^2$

gap. To see whether the recession is associated with a lower level or a lower growth rate of output, we regress each gap from t_1+12 to t_1+27 on a constant and a time trend. If the estimated time trend is positive (i.e. if the gap is increasing) and statistically significant at the 1% level, we classify the recession as having not only a sustained gap, but an increasing gap.

Based on this methodology and these classifications, the results are shown in the first two lines of Table 1.

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	29%	71%	31%	40%	32%	68%	34%	34%
	Adjusted for Credit Booms	31%	69%	31%	38%	33%	67%	34%	34%
Log Real GDP per working-age pop.	Benchmark: 2 Years Before	27%	73%	30%	43%	33%	67%	33%	34%
	Adjusted for Credit Booms	31%	69%	30%	39%	34%	66%	34%	33%

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.

The two sets of columns give the results corresponding to the two ways of computing the pre-recession time trends, over four years or over ten years respectively. The two lines correspond to the two ways of anchoring the pre-recession trend (leaving out the two years before the recession, or leaving out more quarters 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 with no sustained gap (according to the definition above). The second gives the proportion with a sustained gap. The last two columns give the proportion of recessions with a stable gap, and the proportion of recessions with an increasing gap.

The results are very similar for all four combinations, and are the main findings from this section. Taking the average over the four combinations, they show that in only 31% of cases, the recession was not followed by a sustained gap. Equivalently, in 69% of the cases, the recession was followed by a sustained gap. And, in 47% of these sustained-gap cases (33% of all the cases), the recession was followed by an increasing gap.

As a robustness test, we calculated the gap using log real GDP per working-age population (calculated as GDP over population of 16 to 64 years old) for the same recession periods. The results, shown in the bottom two rows of Table 1, are very similar. In only 31% of the recessions, the recession was not followed by a sustained gap. Equivalently, 69% of the recessions were followed by a sustained output gap, with 46% of recessions followed by an increasing output gap.⁷

⁷ In the literature, results by Haltmaier (2012) – using a methodology based on HP filters – and Martin and Wilson (2013) also suggest that recessions (in general) lead, in many cases, to a lower level of output. Studies

(continued...)

We have also performed a number of visual robustness checks (i.e. whether or not the increasing output gap in years three to seven was as a result of an outlier or another recession). Our conclusion is that, in roughly 80% of the cases classified as “increasing output gap”, the increase was indeed unambiguous. This suggests that at least 30% of all recessions were followed by lower output growth later.

B. What could be behind the results? Hysteresis and alternative hypotheses

Focusing on the recessions followed by either a stable or increasing output gap, we can think of three potential explanations for the previous subsection results:

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 suggested 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 labour 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 (e.g., due to weaker prospects, debt overhang, etc.), leading to a lower capital stock for some time (although presumably not forever). Firms may do less research and development during the recession, leading to a permanently lower productivity level than would have been the case without the recession.⁸ 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.

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, i.e. to “super-hysteresis”. Permanently lower output growth requires permanently lower total factor productivity growth; the recession would have to lead to changes in behavior or in institutions which lead to permanently lower research and development 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.

focusing on deep recessions, such as Cerra and Saxena (2008), Reinhart and Rogoff (2009) and IMF WEO (2009), among many others, also highlight highly persistent effects on the level of output.

⁸ Such a mechanism is explored by Comin and Gertler (2006).

The second explanation 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 behavior, 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.⁹

The third explanation is that the correlation instead 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.¹⁰ A variation on this theme is that it may take time for households and firms to realize that underlying growth has started to decline, so that the decrease in productivity may start before the recession. Two intriguing facts support this hypothesis. An old fact documented by Robert Gordon in 1980 and revisited by him more recently (2003), in which productivity declines during the end of the expansion (i.e. before the recession starts).¹¹ A fact documented by Beaudry et al. (2014), in which firms appear to over-accumulate capital during periods of expansion. Both are what one would expect if firms and households took some time to realize that productivity growth had actually slowed down.

C. Controlling by Type of Recession

In order to make some progress in differentiating among the factors presented above, we focus on how the outcomes differ when we control for the cause of the recessions.

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

⁹ For example, Ennis and Keister (2003) have illustrated in a theoretical model how a higher probability of banks runs can reduce the stock of capital and output as well as the long run growth rate. Empirical evidence in this regard has been provided by Ramirez (2009) in his analysis of the 1893 US financial crisis.

¹⁰ A model along these lines is presented and estimated in Blanchard, L'Huillier and Lorenzoni (2013). The model, however, assumes that the news is bad news about the level of productivity, not bad news about the growth rate.

¹¹ Robert Gordon, however, offers a different interpretation of the fact. He argues that the decrease in productivity during the boom is due to overoptimistic 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.

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 sustained output gap, 83% on average across specifications, compared to 66% in the absence of a financial crisis. In 35% of all 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	17%	83%	35%	48%	13%	87%	48%	39%
	Without financial crisis	31%	69%	30%	38%	36%	64%	31%	32%
Log Real GDP / Adjusted	With financial crisis	17%	83%	39%	43%	13%	87%	48%	39%
	Without financial crisis	34%	66%	29%	36%	37%	63%	30%	32%
Log Real GDP working- age pop / Benchmark	With financial crisis	22%	78%	22%	57%	13%	87%	30%	57%
	Without financial crisis	28%	72%	32%	39%	37%	63%	33%	29%
Log Real GDP working- age pop / Adjusted	With financial crisis	26%	74%	22%	52%	13%	87%	35%	52%
	Without financial crisis	32%	68%	31%	36%	38%	62%	33%	28%

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 increases are shown in Table 3. Recessions linked to oil price increases are more likely to show a sustained gap, 90% on average, compared to 65% in the rest of the cases. In 76% of all the cases, recessions associated with an increase in the price of oil are followed by an increasing 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	17%	83%	72%	11%	0%	100%	78%	22%
	Without oil price increases	31%	69%	24%	45%	38%	63%	27%	36%
Log Real GDP / Adjusted	With oil price increases	17%	83%	72%	11%	0%	100%	78%	22%
	Without oil price increases	34%	66%	24%	42%	38%	62%	26%	36%
Log Real GDP working- age pop / Benchmark	With oil price increases	22%	78%	72%	6%	0%	100%	83%	17%
	Without oil price increases	28%	72%	23%	49%	38%	62%	24%	38%
Log Real GDP working- age pop / Adjusted	With oil price increases	22%	78%	72%	6%	0%	100%	83%	17%
	Without oil price increases	33%	67%	22%	45%	39%	61%	25%	36%

Note: A total of 122 recession episodes are included in the analysis, of which 18 coincided 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.¹² The motivation is again

¹² We classify as recessions with increasing inflation those 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.

straightforward: the first set is more likely to be associated with supply shocks, which may 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 distinction is far from tight. While financial crises should be thought as a supply shock, they may also lead, as they did in the recent crisis, to a very large initial decrease in demand, leading to a decrease in output larger than the decrease in natural output, and thus to a decrease in both output and inflation. It remains the case that recessions associated with an increase in inflation are more likely to come from supply shocks, recessions associated with a decrease in inflation to come from demand shocks.)

The results are presented in Table 4. The results conform to priors, but less strongly than in the previous two tables. Recessions associated with increasing inflation are more likely to show a sustained gap, with a frequency of 72%, compared with 63% for those associated with decreasing inflation. 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 63% (with the large majority of those associated with an increasing gap over time rather than just a 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	29%	71%	32%	39%	26%	74%	38%	36%
	With declining inflation	29%	71%	29%	42%	42%	58%	29%	29%
Log Real GDP / Adjusted	With increasing inflation	31%	69%	31%	38%	27%	73%	36%	36%
	With declining inflation	31%	69%	31%	38%	42%	58%	29%	29%
Log Real GDP working- age pop / Benchmark	With increasing inflation	23%	77%	31%	45%	29%	71%	36%	35%
	With declining inflation	33%	67%	29%	38%	40%	60%	27%	33%
Log Real GDP working- age pop / Adjusted	With increasing inflation	29%	71%	30%	42%	30%	70%	36%	34%
	With declining inflation	36%	64%	29%	36%	40%	60%	29%	31%

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 shock, 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. So, as a final step, 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 (see online appendix). Table 5 shows the breakdown for recessions with and without intentional disinflations.

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	Without intentional disinflation	28%	72%	34%	38%	29%	71%	39%	32%
	With intentional disinflation	32%	68%	21%	46%	43%	57%	18%	39%
Log Real GDP / Adjusted	Without intentional disinflation	31%	69%	34%	35%	30%	70%	38%	32%
	With intentional disinflation	32%	68%	21%	46%	43%	57%	18%	39%
Log Real GDP working- age pop / Benchmark	Without intentional disinflation	27%	73%	34%	39%	30%	70%	36%	34%
	With intentional disinflation	29%	71%	18%	54%	43%	57%	21%	36%
Log Real GDP working- age pop / Adjusted	Without intentional disinflation	32%	68%	33%	35%	31%	69%	37%	32%
	With intentional disinflation	29%	71%	18%	54%	43%	57%	21%	36%

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.

Recessions associated with intentional disinflations are somewhat less likely to show a sustained gap, 63% on average, compared to 70% for others. But the difference is small, and again, the results can be read as saying that even recessions associated with intentional disinflations are followed by a sustained gap in 63% of the cases. In 20% of all the cases, they appear to be actually followed not only by lower output, but by lower output growth.

D. Relative Output Gaps

We have so far looked at the output gap, independent of the size of the initial recession. Under a simple hysteresis hypothesis, one would expect the ratio of the output gap to the initial output fall during the recession, call it the relative output gap, to be between 0 and 1. With this motivation in mind, Figure 6 presents results for the relative output gap. To avoid dividing by a number close to zero (as some recessions are associated with very small decreases in output), we limit ourselves to 68 large recessions where the fall in output was larger than 2 percent.¹³

¹³ We redid the computations presented earlier for this smaller sample of recessions. In general, the results are rather similar. The proportion of recessions followed by an output gap is slightly larger (75% versus 69% for the whole sample).

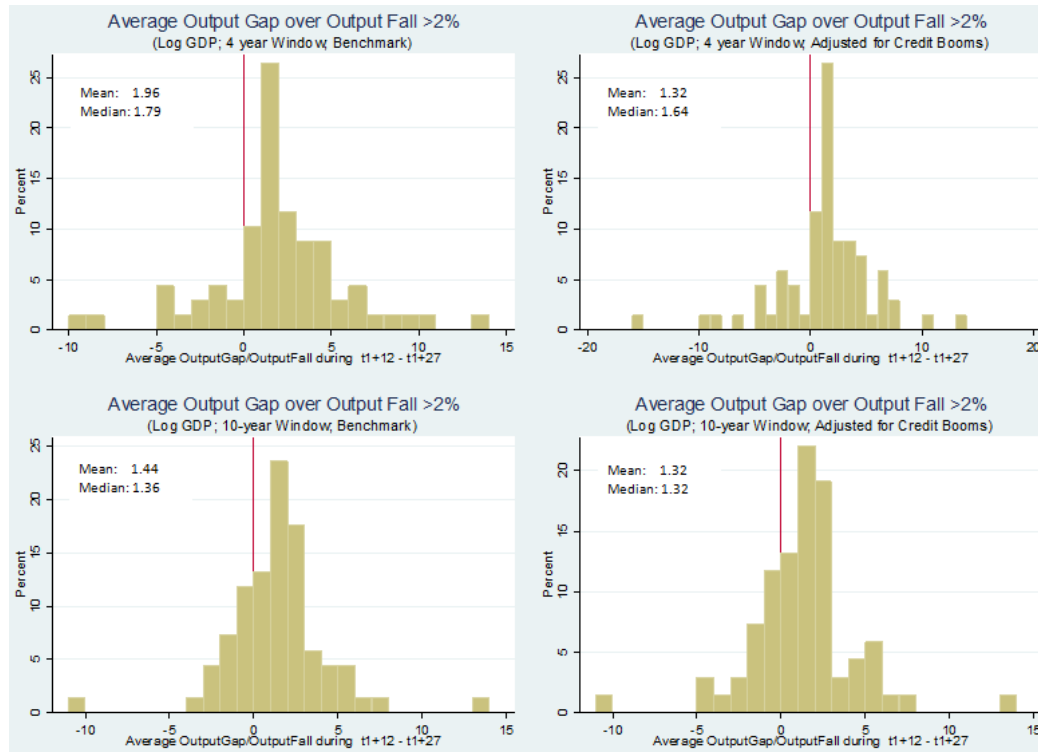
Figure 6: Histogram of Average Relative Output Gaps

Figure 6 plots the resulting distributions of relative output gaps, for each of our four different sets of assumptions. In all cases, the means and the medians of the distribution are not only positive but between 1 and 2, so larger than would be expected under a simple hysteresis hypothesis.

E. Main Takeaways

A surprisingly high proportion – two-thirds – of recessions are followed by lower output relative to the pre-recession trend. Even more surprisingly, almost one-half of those are followed not only by lower output, but also by lower output growth relative to the pre-recession trend. These proportions are larger for recessions associated with supply shocks. Even for recessions plausibly induced by intentional disinflations, the proportion of recessions followed by an output gap remains high, at around 63%.

From these findings, we conclude that it is likely that all three explanations are relevant. The fact that recessions identified with specific supply shocks, whether oil prices or financial crises, are more likely to be followed by sustained gaps suggests that such shocks probably explain both the recession and the subsequent lower output or output growth. The fact that almost two-thirds of the recessions associated with decreasing inflation, and thus with demand shocks as the more likely factor behind the recession, are associated with a sustained

gap points to one of the other two explanations, hysteresis or reverse causality. Finally, the fact that also nearly two-thirds of the recessions associated with intentional disinflations show a sustained gap is suggestive of hysteresis, since even demand shocks could turn out to affect supply more permanently. These conclusions have important implications for monetary policy that we shall develop in the last section.

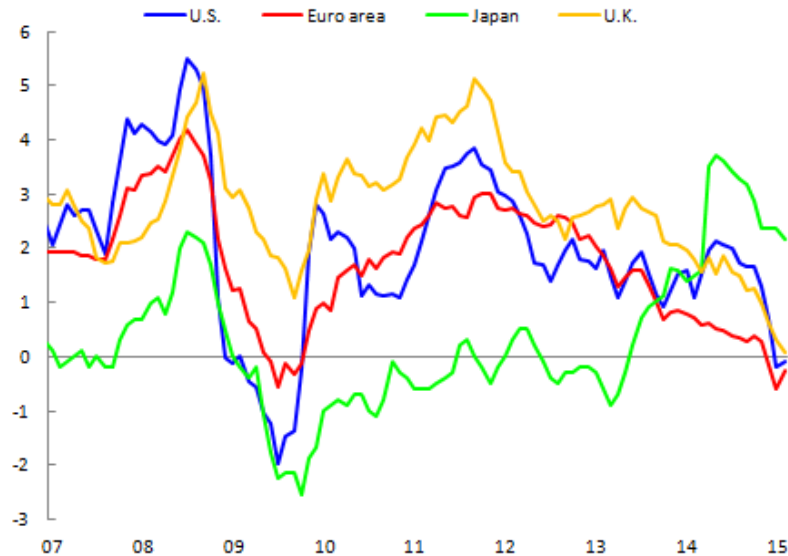
III. DOES UNEMPLOYMENT AFFECT INFLATION?

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. Figure 7, which plots inflation since 2007 in the United States, the euro area, United Kingdom and Japan, shows that inflation indeed declined, and in some countries, has now turned into deflation, but deflation has remained limited.

A. Empirical Estimations of Phillips Curves: A Time-Varying Approach

Much of the explanation 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. (An alternative interpretation is that low inflation has led to decreased inflation salience, and that less attention is paid by workers and firms to actual inflation). This in turn has led to a shift from an “accelerationist Phillips curve”, in which the unemployment gap affects the change in inflation, to something closer to a “level Phillips curve”, in which the gap is affects the level of inflation.

Figure 7: Advanced Economies CPI Headline Inflation (percent, year over year)



Source: IMF, Global Data Source.

Some 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 diminished over time. In particular, this was the conclusion of the study conducted in the IMF’s April 2013 World Economic Outlook (WEO) study, which we take as our starting point here.¹⁴

The study, which was based on quarterly 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):

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

Where π_t is headline CPI inflation (defined as quarterly inflation, annualized), u_t is the unemployment rate, u_t^* is the natural rate, π_t^e is long-term inflation expectations, π_{t-1}^* is the average of the last four quarterly inflation rates, and π_{mt} is import price inflation relative to headline inflation.

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, u_t^* , which is unobservable, are all assumed to follow constrained random walks (the constraints being θ_t and $\mu_t \geq 0$, and $0 \leq \lambda_t \leq 1$).¹⁵

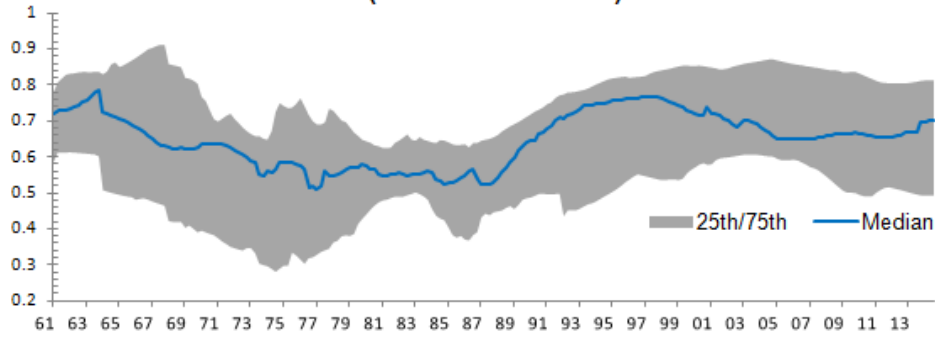
The equation is estimated separately for each country by maximum likelihood, using a non-linear Kalman filter. The results differ from those in the WEO chapter only to the extent that we use two more years of data, with the samples ending in 2014.

Figure 8 shows median estimates for λ_t and θ_t , the two coefficients we focus on, together with the interquartile range of estimates across countries.

¹⁴ 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”.

¹⁵ Further details about the specification of the equation are given in Appendix 1 of the IMF chapter.

**Figure 8: Median estimates (across countries)
Anchoring of Inflation to Long-term Expectations (λ)
(median all countries)**



**Slope of the Phillips Curve (θ)
(median all countries - shown in absolute values)**

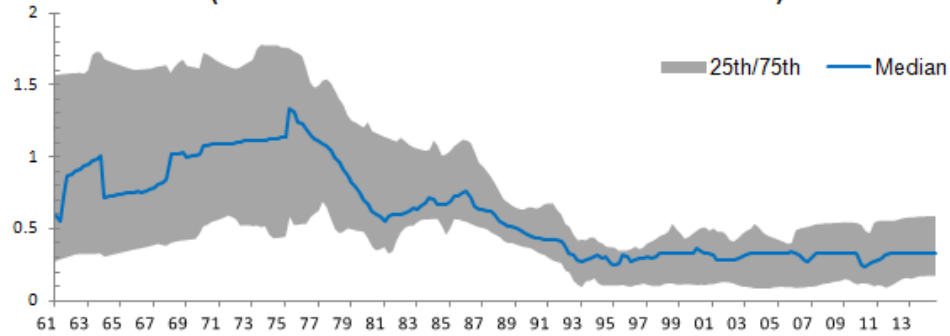
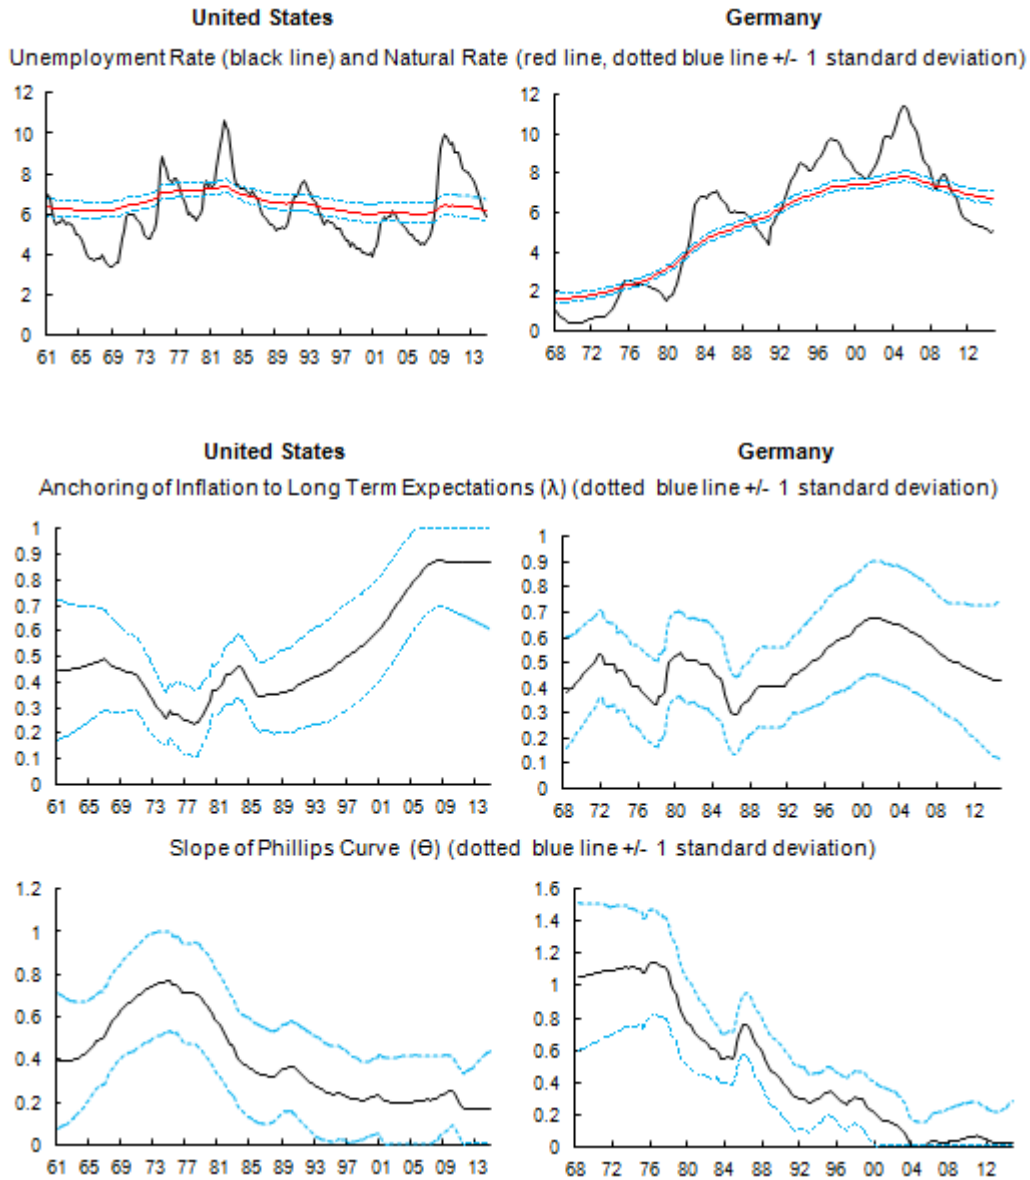


Figure 8 confirms the two main conclusions of the earlier IMF study: (i) since the mid-1970s, short-run inflation expectations have become more stable (λ_t has increased), and (ii) the slope of the Phillips curve (θ_t) has flattened over time, with nearly all of the decline taking place from the mid-1970s to the early 1990s, and the coefficient remaining roughly constant since then. It also does not show any further decrease since the beginning of the crisis.

What Figure 8 does not show, however, is that, for most countries, the coefficient θ_t today is not only small, but also statistically insignificant. This can be seen for example in Figure 9, showing estimates for the United States and Germany, two countries that are representative of other countries (the results for other countries are presented in the web appendix). In both countries, 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).

The graphs are representative of the results for the larger set of countries. While, in 1985, the coefficient θ_t was significant for all but two countries, in 2014 it is insignificant for all but four countries.¹⁶

Figure 9: Estimates for Germany and the US



¹⁶ The appropriate specification of the Phillips curve, if there is indeed one, is far from settled. The above results were robust to either allowing the natural rate to depend partly on the actual unemployment rate – reflecting a crude form of the hysteresis issue dealt in the previous section – or replacing the unemployment rate with the short-term unemployment rate.

B. Empirical Estimations of Phillips Curves: A Simpler Approach

Our discussant at the 2015 ECB Forum on Central Banking, Laurence Ball, however, expressed skepticism. He argued that the low significance of the coefficients may have come from our specification of the inflation-unemployment relation. The Kalman filter may have a hard time distinguishing how much of the change in the relation between the two came from changes in θ_t or from changes in u_t^* : for example, the lack of an apparent effect of an increase in unemployment on inflation could be equally well explained by a smaller coefficient θ_t or by an increase in u_t^* in parallel with the increase in u_t . The large standard deviations may then be the result of this poor identification. He showed that, for the United States and for the euro area, simple regressions delivered significant coefficients.

This led us to follow his lead and explore simpler and tighter specifications. Based on the observation that, in most countries, the estimated coefficient was roughly stable starting in the early 1990s, we limited the sample to the period 1990-2014. Based on the observation that, again starting in the early 1990s, inflation expectations appeared well anchored, we dropped lagged inflation as an explanatory variable, and we estimated the following relation, assuming constant rather than time-varying coefficients:

$$\pi_t = \theta(u_t - u_t^*) + \lambda\pi_t^e + \mu\pi_{mt} + \varepsilon_t \quad (2)$$

We estimated, for each country, the relation over the whole period, 1990-2014, and over the crisis period, 2007-2014. We made two different assumptions about u_t^* . In the first, we took the time series for u_t^* from the Kalman filter estimation presented earlier. In the second, we simply assumed u_t^* to be constant and equal to the average unemployment rate over the whole period (\bar{u}). Table 6 reports estimated coefficients for θ and their standard errors over both the 1990-2014 and the 2007-2014 periods, for the two specifications of the natural rate, u_t^* and \bar{u} (full results, and other coefficients, are available in the web appendix).

Table 6: Slope of the Phillips Curve

Country	Sample Period	Natural rate: U^*		Natural Rate: \bar{U}	
		Coefficient 1/	Standard Deviation	Coefficient 1/	Standard Deviation
United States	1990-2014	-0.29 ***	0.07	-0.25 ***	0.07
	2007-2014	-0.26 *	0.13	-0.24 *	0.12
Japan	1990-2014	-1.09 ***	0.17	-0.50 ***	0.09
	2007-2014	-2.37 ***	0.64	-1.54 ***	0.43
Germany	1990-2014	-0.11	0.09	-0.12	0.07
	2007-2014	0.15	0.20	0.11	0.14
United Kingdom	1990-2014	0.04	0.11	0.02	0.10
	2007-2014	-0.04	0.30	-0.04	0.28
France	1990-2014	-0.49 ***	0.14	-0.32 ***	0.08
	2007-2014	-0.63 ***	0.20	-0.52 ***	0.16
Italy	1990-2014	-0.12	0.09	-0.15 ***	0.06
	2007-2014	-0.30 ***	0.10	-0.25 ***	0.08
Canada	1990-2014	-0.32 ***	0.16	-0.21	0.13
	2007-2014	-0.52	0.45	-0.50	0.41
Australia	1990-2014	-0.68 ***	0.14	-0.49 ***	0.10
	2007-2014	-0.78 **	0.38	-0.79 **	0.36
Spain	1990-2014	-0.08 **	0.04	-0.07 ***	0.03
	2007-2014	-0.09 *	0.05	-0.07 *	0.04
Netherlands	1990-2014	-0.40 ***	0.12	-0.31 ***	0.10
	2007-2014	0.01	0.20	-0.01	0.16
Switzerland	1990-2014	-0.71 ***	0.15	-0.64 ***	0.08
	2007-2014	-0.59	0.79	-0.65	0.60
Sweden	1990-2014	-0.55 ***	0.12	-0.48 ***	0.09
	2007-2014	-1.21 ***	0.32	-1.10 ***	0.26
Belgium	1990-2014	-0.62 ***	0.20	-0.37 ***	0.14
	2007-2014	-1.13 *	0.62	-0.74 **	0.30
Norway	1990-2014	-0.06	0.20	-0.07	0.20
	2007-2014	-0.80	1.52	-0.73	1.41
Austria	1990-2014	-0.68 ***	0.25	-0.11	0.14
	2007-2014	-1.62 ***	0.42	-1.26 ***	0.36
Denmark	1990-2014	-0.17	0.12	-0.13	0.10
	2007-2014	1.72 ***	0.49	1.53 ***	0.44
Ireland	1990-2014	-0.28 ***	0.04	-0.21 ***	0.03
	2007-2014	-0.26 ***	0.08	-0.22 ***	0.07
Greece	1990-2014	-0.15 **	0.07	-0.14 ***	0.03
	2007-2014	-0.19 ***	0.06	-0.11 ***	0.04
Portugal	1990-2014	-0.23 ***	0.07	-0.17 ***	0.05
	2007-2014	-0.01	0.11	-0.02	0.08
New Zealand	1990-2014	-0.50 ***	0.09	-0.33 ***	0.11
	2007-2014	-0.60 *	0.34	-0.58 *	0.28

1/ *** indicates significance at 1 percent, ** at 5 percent, and * at 10 percent, respectively, based on robust standard errors.

The table yields the following conclusions.

The main conclusion is that the estimated coefficient, θ , is typically significant (at the 95% confidence level). When the natural rate is proxied by u_t^* , the coefficient is significant in 15 out of the 20 countries. When the natural rate is proxied by \bar{u} , the coefficient is significant in 14 out of the 20 countries. Countries where the coefficient is not significant in either specification are Germany, the United Kingdom, Norway, and Denmark.^{17 18} These results

¹⁷ In the case of the United Kingdom, there appears to be a stable and significant relation between wage inflation, expected inflation and unemployment (Broadbent, 2014). What appears to have broken down is the relation between wage inflation and price inflation.

suggest that the Kalman filter results indeed understated the significance of the coefficient. In most countries, even if the coefficients are smaller than they were earlier in time, there remains a significant relation between activity and inflation.

As in the Kalman filter results, coefficients vary widely across countries, and appear roughly inversely related to the levels and the movements in unemployment. For example, in the specification using u_t^* , the coefficient for Japan is equal to -1.08, the coefficient for the United States, is -0.29, and the coefficient for Spain is -0.08. The coefficients are typically larger when u_t^* is used. The likely explanation is that when the natural rate is allowed to move over time, it tends to move with the actual rate, reducing the unemployment gap, requiring a larger coefficient to explain the movement in inflation.

One might have expected that the reluctance of employers to impose (and of employees to accept) decreases in nominal wages (a hypothesis sometimes referred to as the “zero bound” on wage growth), combined with very low inflation, might have led to a decrease in the effect of the unemployment gap on inflation. This does not appear to be the case. The coefficients do not appear to have decreased in the recent past. Among countries with significant coefficients over the period 1990-2014, only two have much smaller and insignificant coefficients when estimated over 2007-2014, namely the Netherlands and Portugal.¹⁹ In all other cases, the coefficient is about the same or larger.

To summarize: it is clear that the slope of the Phillips curve has decreased over time in most countries. Most of the decline, however, took place from the mid-1970s to the early 1990s. Since then, the coefficient has remained roughly stable. In particular, it does not appear to have decreased during the crisis.

IV. IMPLICATIONS FOR MONETARY POLICY

Based on the conclusions from our empirical work, what are the implications for monetary policy? A full answer would require a much better understanding of the underlying mechanisms behind hysteresis, if indeed present, or behind the decrease in the slope of the Phillips curve. Nevertheless, we feel we can draw the following conclusions.

¹⁸ We do not look at the euro area as a whole. Work by Andrieu et al. (2013) also suggests the presence of a significant relation between inflation and unemployment there as well.

¹⁹ At the microeconomic level, there is substantial evidence of a binding zero lower bound on wage changes in Portugal, which might explain why the coefficient on the unemployment gap has become insignificant. The question is then why the same has not happened in other countries, which also have low wage inflation.

Depending on their interpretation, the findings of the second section have strong, but conflicting implications for monetary policy (and for macroeconomic policy in general).

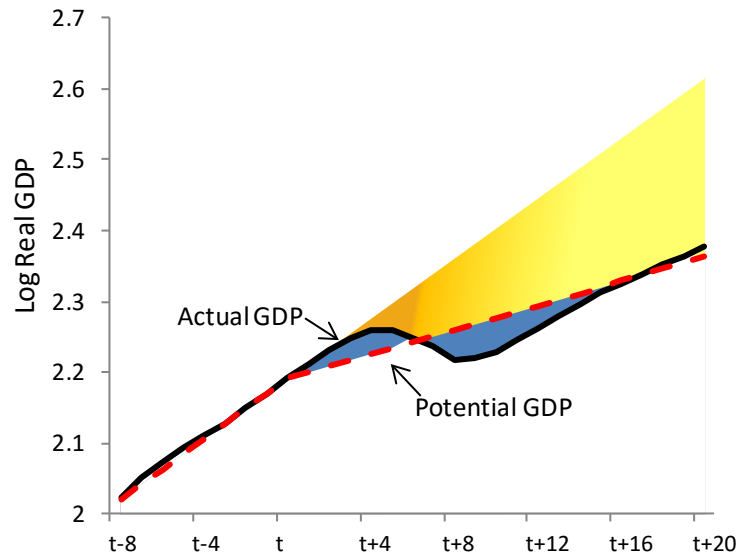
To the extent that hysteresis is present, it implies that deviations in output from its optimal level are much longer-lasting and thus more costly than usually assumed. The implication is straightforward, namely that monetary policy should react more strongly to output movements, relative to inflation. For example, by being more aggressive early on, this would reduce the increase in unemployment, and by implication, reduce the increase in the number of long term unemployed. It also implies that stabilizing inflation is definitely not the optimal policy: to the extent that an increase in actual unemployment leads to an increase in the natural rate, the unemployment gap, and by implication inflation, will give a misleading signal about the degree of underutilization of resources in the economy.²⁰

To the extent, however, that some recessions are caused by an underlying decrease in growth, there is instead the risk of overestimating potential output during and after the recession, and by implication the risk of overestimating the output gap. In turn, this may lead to too strong a response of monetary policy to output movements during and after the recession.²¹ This is illustrated in Figure 10 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, determined by demand, will continue at close to the old trend, until the adjustment of expectations leads to an adjustment in spending and 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 is 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.

The findings of the third section are less dramatic. Our initial hypothesis, which was that there might no longer be a significant relation between inflation and unemployment, is not supported by the data. While the Phillips curve coefficient is clearly lower than it was up until the early 1990s, it appears to have remained stable since then, including during the crisis, and is significant in most countries.

²⁰ We are not aware of a derivation of optimal monetary policy under hysteresis. For a beginning, see Gali (2015).

²¹ It is indeed often the case that estimates of output gaps associated with recessions are revised down ex post.

Figure 10. Decreases in growth, recessions, and output gaps.

This decrease does not, by itself, put into question the standard inflation-targeting framework, but it has implications for the optimal policy rule. To draw specific policy implications, one would need to identify why the coefficient became smaller, whether it comes, for example, from wage-setting or from price-setting behavior, whether it comes from changes in the structure of wage bargaining, or in the pricing behavior of firms in the product market.²² However, more generally, to the extent that the unemployment gap has a smaller effect on inflation, monetary policy rules should put relatively more weight on the unemployment gap relative to inflation. Trying to stabilize inflation may require very large movements in the unemployment gap.

²² More specifically, the issue is whether the factors behind the decrease in the coefficient also affect the weights of inflation and the output gaps in the welfare function. In formal NK models, this may not be the case (see, for example, Woodford 2003, Chapter 6).

References

- Adrian, T. and Shin, H.S (2010), “Financial intermediaries and monetary economics”, in Friedman, B. and Woodford, M. (eds.), *Handbook Monetary Economics*, Vol. 3, pp. 601-650.
- Andrle, M., Bruha, J. and Solmaz, S. (2013), “Inflation and output comovement in the euro area: Love at second sight”, *IMF Working Paper*, No 13/192.
- Ball, L. (2014), “Long-term damage from the Great Recession in OECD countries”, *NBER Working Paper*, No 20185, May.
- Beaudry P., Galizia, D. and Portier, F (2014), “Reconciling Hayek's and Keynes' views of recessions”, *CEPR Discussion Paper*, No 9966; *NBER Working Paper*, No 20101.
- Blanchard, O., L’Huillier, J.P. and Lorenzoni, G. (2013), “News, noise, and fluctuations: An empirical exploration”, *American Economic Review*, Vol. 103(7), December, pp. 3045-3070.
- Blanchard, O. and Summers, L. (1986), "Hysteresis and European Unemployment", in Fischer, S. (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*, Vol. 98(1), pp. 439-457.
- Coibion, O. and Gorodnishenko, Y. (2015), “Is the Phillips curve alive and well after all? Inflation expectations and the missing disinflation”, *American Economic Journal*, Vol. 7(1), pp.197-232
- Comin, D. and Gertler, M. (2006), “Medium-term business cycles”, *American Economic Review*, Vol. 96(3), pp. 523-551.
- Crowe, C., Dell’Ariccia, G., Igan, D. and Rabanal, P. (2011), “Policies for macrofinancial stability: Options to deal with real estate booms”, *IMF Staff Discussion Note*, No 11/02.
- Ennis, H. and Keister, T. (2003), “Economic growth, liquidity, and bank runs”, *Journal of Economic Theory*, Vol. 109(2), pp. 220-245.
- Fernald, F. (2014), “Productivity and potential output before, during, and after the Great Recession”, in *NBER Macroeconomics Annual 2014*, Vol. 29.
- Gali, J. (2015), “Hysteresis and the European unemployment problem revisited”, slides presented at the NBER summer institute, June.
- Gordon, R. (2003), “Exploding productivity growth: Context, causes, and implications”, *Brookings Papers on Economic Activity*, Vol. 2, pp. 207-298.
- Haltmaier, J. (2012), ““Do recessions affect potential output?””, *International Finance Discussion Paper*, No1066, Federal Reserve Board, December.
- Harding, D. and Pagan, A. (2002), "Dissecting the cycle: a methodological investigation," *Journal of Monetary Economics*, Vol. 49(2), pp. 365-381.

- IMF World Economic Outlook (2009), Chapter 3, “From recession to recovery: How soon and how strong?”, April.
- 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 Valencia, F. (2013), “Systemic banking crises database”, IMF Economic Review, Vol. 61(2), pp. 225-270.
- Martin and Wilson (2013), “Potential output and recessions: Are we fooling ourselves?”, manuscript.
- Matheson, T. and Stavrev, E. (2013), “The Great Recession and the inflation puzzle”, Economics Letters, Vol. 120(3), pp. 468-472.
- Ramirez, C. (2009), “Bank fragility, “Money under the mattress”, and long run growth: US evidence from the perfect Panic of 1893”, Journal of Banking and Finance, Vol. 33, pp. 2185-2198.
- Reinhart, C. and Rogoff, K. (2009), “The aftermath of financial crises”, American Economic Review, Vol. 99(2), pp. 466-72.
- Woodford, M. (2003), Interest and prices: Foundations of a theory of monetary policy, Princeton University Press.