A Phillips Curve for the Euro Area

Laurence Ball and Sandeep Mazumder*

September 2019

Abstract

Two puzzles have accompanied euro area inflation in recent years: the missing disinflation following the two recessions, and the absent inflation during the subsequent recovery. Existing research puts forth a variety of explanations for these inflationary puzzles, such as a de-anchoring of inflation expectations. We instead argue that inflation in the eurozone is not particularly puzzling when one estimates a Phillips curve with survey forecasts of inflation and the OECD measure of the output gap, where our measure of core inflation is the weighted median inflation rate. Moreover, augmenting the model with the pass-through of headline-into-core inflation improves the fit of the model further still.

*Ball: Department of Economics, Johns Hopkins University, lball@jhu.edu. Mazumder: Department of Economics, Wake Forest University, mazumds@wfu.edu.

1 Introduction

The behavior of European inflation over the last decade has puzzled economists and policymakers. Researchers have suggested that the puzzle has two parts: a "missing disinflation" in the wake of the twin recessions of 2008 and 2011, and a "missing inflation" more recently as the economy has recovered (e.g. Ciccarelli and Osbat, 2017; Abdih et al., 2018). Both parts of the puzzle are apparent failures of inflation to respond to the level of economic slack in the way predicted by a conventional Phillips curve.

This experience has produced many suggestions of possible factors behind the puzzles. For example, the Eurosystem reported edited by Ciccarelli and Osbat considers a de-anchoring of inflation expectations, an increase in the persistence of shocks to the inflation rate, non-linearity or time-variation in the effects of slack on inflation, and changes in world commodity prices and exchange rates. Recently, top ECB officials have suggested that the behavior of inflation has been influenced by structural changes in the economy such as digitalization and globalization (Draghi, 2019) and the growth of the service sector (Coueré, 2019).

This paper argues that European inflation behavior is not as puzzling or complex as recent discussions suggest. A simple Phillips curve captures most of the movements in inflation over the twenty years that the Euro has existed.

Like many researchers, we examine a measure of *core* inflation that strips out the effects of large relative price changes on headline inflation. We do not, however, focus on the most common measure of core inflation, the inflation rate excluding food and energy prices; instead, our preferred measure is the weighted median of industry inflation rates, a concept developed by Bryan and Cecchetti (1994) that we also use in recent work on U.S. inflation (e.g. Ball and Mazumder, 2019). In both the U.S. and Europe, weighted median inflation is less volatile and easier to explain than the conventional core measure. Section 2 of this paper examines the evolution of the two core inflation series since 1999.

We have two main specifications of the Phillips curve. The first, presented in Section 3, relates quarterly core inflation to expected inflation, as measured by five-year forecasts of the Survey of Professional Forecasters, and the Euro area output gap, as measured by the OECD. Expected inflation does not vary much over our sample period, so our equation is close to a relationship between core inflation and the output gap alone. This equation explains a large fraction of the movements in weighted median inflation ($\overline{R}^2$=0.64), but substantial residuals remain, including inflation rates that are higher than the predicted levels over 2010-2013 and lower since then.

Our second specification, examined in Section 4, adds another variable to the equation for core inflation: the deviation between headline and core inflation over the current and previous three quarters. Such a variable is not common in core-inflation equations, but including it captures the idea that movements in headline inflation can be passed through into core inflation through wage adjustment and the costs of intermediate goods. This idea appears in some previous research on European inflation (e.g. Peersman and Van Robays, 2009) and some ECB discussions (e.g., ECB, 2017).

Adding this new variable to the Phillips curve produces a substantial improvement in fit ($\overline{R}^2$=0.76), and it largely eliminates the perceived puzzles about inflation since the Great Recession. Based on our specification, there is no missing disinflation, and only a modest amount of missing inflation that arises in 2017-2018.

Section 5 extends our analysis to the United States with the goal of comparing European and U.S. inflation behavior. We find similarities but also differences; in particular, we do not find any pass through of headline into core inflation in the United States, a difference that could be explained by differences in wage determination.

Section 6 concludes the paper with suggestions for future research.

2 Core Inflation in the Euro Area

This paper seeks to explain the quarterly behavior of core inflation, as measured by the weighted median inflation rate. Here we examine the behavior of this variable over 1999-2018 to see the inflation movements that we need to explain. We also compare weighted median inflation to the inflation rate excluding food and energy prices (XFE inflation), which is the version of core inflation most commonly used in research and in discussions of ECB policy.

Our previous work (e.g. Ball and Mazumder, 2019) discusses the theoretical and empirical case for measuring core inflation with the weighted median. At a practical level, while the XFE inflation rate filters out the transitory effects of large changes in food and energy prices, the weighted median filters out large price changes in all industries, producing a less volatile measure of underlying inflation.

We construct a weighted median inflation rate from the inflation rates of 94 industries that make up the HICP price index for the Euro area, with weights equal to industries' relative importance in the index. The industry data come from the IMF. We use data on the HICP inflation rate excluding food and energy from the ECB.

Figure 1 compares the median and XFE inflation rates at three frequencies: monthly, quarterly, and a four-quarter moving average. We can see that monthly XFE inflation is quite a bit more volatile than monthly median inflation: the standard deviation of changes in inflation is 0.93 for XFE and 0.51 for median. This difference diminishes at higher levels of time aggregation, but the standard deviation is still larger for XFE (by a factor of 1.22 for quarterly data and 1.09 for 4-quarter averages).

One period in which XFE inflation is especially volatile is the last two years of the sample, 2017-2018. During those years, annualized monthly rates of XFE inflation vary over a four-percentage-point range, from -1.3% to 2.7%, while median inflation rates vary by only one percentage point, from 0.8% to 1.8%.

To understand why the XFE series can be volatile, consider the month of October 2017, when XFE inflation is -1.3% and median inflation is +1.3%. For that month, Figure 2 shows a histogram of industry price changes within the XFE index. Each bar in the graph represents an interval of 5 percentage points in annualized inflation rates and shows the total weights of the industries in that range. We see a large tail of price decreases that skews the histogram to the left and pulls down XFE inflation. Eight industries with total weights of about 7% have annualized inflation rates below -10%, including education (-31%) and transport insurance (-26%); by contrast, no industry has an inflation rate above +10%. In some other months, the opposite pattern arises: large price increases in certain industries skew the distribution of price changes to the right and cause an upward spike in XFE inflation. Median inflation is more stable than XFE inflation because the effects of large industry price changes are filtered out.

The two core inflation rates also differ somewhat in their average levels: for 1999-2018, average annual inflation is 1.71% for median and 1.45% for XFE. The average of headline HICP inflation is also 1.71%, the same to two decimal places as the average of the median. This result means that large price changes that shift headline inflation relative to the median are roughly balanced over time between positive and negative changes. The fact that XFE inflation is lower on average than headline inflation means that the relative price of food and energy has risen over time.

In Figure 1C, which shows four-quarter averages of core inflation, we also show the two recessions that have occurred in the Euro area (as dated by the CEPR): the double dip in 2008Q1-2009Q2 and 2011Q3-2013Q1 associated with the global financial crisis and the European debt crisis. Much of the recent discussion of European inflation concerns the responses or lack thereof to these recessions and to the recovery since 2013, which are often characterized as a "missing disinflation" and a "missing inflation."

The Figure suggests that reports of a missing disinflation are a bit misleading, because the first recession led to a substantial fall in inflation: the four-quarter average of median inflation fell from 2.9% in 2008Q3 to 1.1% to 2010Q2. What may be more surprising is that inflation rebounded sharply between the two recessions, reaching 2.2% in 2011Q4, so the net decrease from the inflation peak to late-2011 was modest. The second recession led to a second fall in inflation, to 0.9% in 2015Q1.

Over the last five years, there has been a missing inflation in the sense that inflation has persistently fallen short of the ECB's target of "below but close to 2%." Inflation has risen somewhat, however, as the economy has recovered: 4-quarter median inflation rose from 0.9% in 2016Q2 to 1.5% in 2018Q4. Notice that 4-quarter XFE inflation was only 1.2% in 2018Q4, so the common focus on this core inflation measure has magnified the apparent missing-inflation puzzle.

In the rest of this paper, we ask how well the ups and downs of inflation in Figure 1 can be explained by simple Phillips curves.

3 The Basic Phillips Curve

Here we examine a simple version of a textbook Phillips curve and find that it explains a large fraction of the fluctuations in core inflation in Figure 1, especially when core inflation is measured by weighted median inflation.

Specification and Data

We examine a simple version of Milton Friedman's (1968) Phillips curve, in which the quarterly core inflation rate is determined by expected inflation and the level of slack in the economy. Here, we measure slack with the gap between output and potential output, as estimated by the OECD. In an Appendix, we consider the robustness of our results to using the other common measure of slack, the deviation of unemployment from its natural rate.

Our Phillips curve is:

$$\pi_t = \pi_t^e + \alpha\overline{(y - y^*)}_{t-1} + \epsilon_t, \tag{1}$$

where π is the core inflation rate, π^e is expected inflation, and $\overline{(y - y^*)}_{t-1}$ is the log difference between the four-quarter averages of actual and potential output from $t-4$ through $t-1$. In assuming that quarterly inflation depends on slack over four quarters, we follow previous research on the U.S. Phillips curve.[1]

We compare results with the core inflation rate measured with the weighted median and with XFE inflation. For expected inflation, we use 5-year forecasts of inflation from the European Survey of Professional Forecasters. These forecasts have been highly stable over our twenty-year sample: they always lie in a range from 1.8 to 2.0. Therefore, in practice our specification is close to one in which π^e is constant and the output gap is the only variable explaining movements in core inflation.

Our data on actual output come from Eurostat, and we derive potential output from actual output and OECD estimates of the output gap. The actual output data are quarterly but the output gap estimates and implied levels of potential are available only at the annual frequency. We must adjust the annual series for potential output for use in our quarterly regressions. This task is made easier by the fact that the output-gap variable in our equation is a four-quarter moving average. We interpret our estimate of potential in a year as a four-quarter average of potential through the fourth quarter of the year; take logs of these fourth-quarter observations; and then linearly interpolate to estimate the log of four-quarter averages of potential in quarters 1, 2, and 3. We subtract the resulting quarterly series from the log of the four-quarter average of actual output to obtain the four-quarter output gap

[1] See, for example, Stock and Watson (2009) and Ball and Mazumder (2019). The exact timing of lags varies in previous work, and it is not clear what is best. We have also estimated a version of equation (1) in which the slack variable is the average of the output gap from $t-3$ through t. That changes improves the fit of the equation modestly: the $\overline{R}^2$ rises from 0.64 to 0.67. However, when we add the deviation of headline from median inflation to the Phillips curve in the next section, the fit is better with the gap from $t-4$ through $t-1$ ($\overline{R}^2$=0.76) than with the gap from $t-3$ through t ($\overline{R}^2$=0.73). We have only considered four quarter averages of the output gap, not separate quarterly lags, because our method of constructing y^* from the OECD's annual series only allows recovery of the averages.

series in our regressions.

Estimates for Alternative Core-Inflation Measures

Table 1 shows estimates of equation (1) with core inflation measured with median inflation and with XFE inflation. In each case, we present results with and without a constant term added to the equation. The theory underlying the Phillips curve implies that the constant should be zero: when the output gap is zero, inflation should equal expected inflation. Therefore, one test of the theory is whether the estimated constant is zero.

When the dependent variable is median inflation, the fit of the Phillips curve is good. The $\overline{R}^2$ is 0.64 with no constant included, and when the constant is added it is small and statistically insignificant. The coefficient on the output term (with no constant) is 0.23: a one percentage point increase in the average output gap over the previous four quarters raises the inflation rate by a bit less than one quarter of a percentage point.

When the dependent variable is XFE inflation, the coefficients on the output gap are similar to those for median inflation. But the fit of the equation is substantially worse: the $\overline{R}^2$ is 0.21 without a constant and 0.46 with a constant. This deterioration reflects the relatively large transitory fluctuations in XFE inflation shown in Figure 1, which are not explained by the output gap. When a constant term is included, it is significant with an estimated value of -0.32. This result implies that XFE inflation falls short of expected inflation by an economically meaningful amount when output is at potential.

Fitted Values for Median Inflation

How well does our Phillips curve explain the inflation movements that have puzzled observers? To help answer this question, Figure 3 compares the path of median inflation to the fitted values from our estimated equation for that variable, with no constant. Figure 3A shows the results for quarterly data, and Figure 3B shows smoother series created by taking four-quarter moving averages of actual and fitted inflation rates. The Figure confirms the fact, indicated by the $\overline{R}^2$ of our regression, that the equation captures most of the broad movements in the median inflation rate.

The differences between actual and fitted values show, however, that the puzzles about inflation since the Great Recession are not fully resolved. The fitted values capture the actual fall in inflation from 2008Q3 to 2010Q2 fairly well, but fail to explain most of the rise in inflation over 2010Q2-2011Q4. This pattern produces a significant amount of missing disinflation in the sense that actual four-quarter inflation in 2011Q4 (2.16%) exceeds the level predicted by our equation (1.62%)

Consistent with suggestions of a more recent missing inflation, actual four-quarter inflation is lower than the fitted values after 2014Q3. However, the differences between the two series are modest, peaking at 0.37 percentage points in 2016Q2, when actual inflation is 0.86% and the fitted value is 1.23%. In 2018Q4, the end of the sample, actual four-quarter inflation is 1.52% and the fitted value is 1.76%. Inflation rates after 2014 are well below the ECB's "close-to-2-percent" inflation target, but only part of this shortfall is a puzzle. Part of it is explained by a negative output gap over the period, albeit one that diminishes to near zero at the end of the sample.

4 Pass-Through from Headline to Core Inflation

Here we show that the fit of our basic Phillips curve can be improved substantially by adding one more variable to the equation: a four-quarter moving average of the deviation of headline inflation from median inflation. With this variable included, the apparent puzzles about inflation behavior since 2008 are mostly eliminated.

<u>Motivation</u>

Our modification of the Phillips curve is motivated by a look at the data in Figure 4. That Figure shows the two series from Figure 3B-the four-quarter averages of median inflation and of the fitted values from our basic Phillips curve–and adds the four-quarter average of headline inflation. The graph strongly suggests a relationship between fluctuations in headline inflation and the residuals in our basic Phillips curve, which means that headline inflation helps explain movements in median inflation that are not explained by the output gap.

In particular, the major ups and downs in headline inflation since 2008 seem to pull median inflation in the same direction relative to the predictions of the basic Phillips curve. Headline inflation goes through a cycle in which it is below median inflation over 2008Q4-2009Q4, then rises above median over 2010Q1-2013Q1, then falls below it again over 2013Q2-2016Q4. In each of these periods, median inflation deviates from the fitted values in the Figure in the direction of headline inflation.

Econometric studies of core inflation generally ignore the behavior of headline inflation. They implicitly assume that the core and non-core parts of inflation are determined independently, the first by a Phillips curve and the second by relative price changes. However, in ECB policy discussions, it is common to suggest that non-core movements in headline inflation-especially movements caused by oil-price changes-feed into core inflation with some lag. For example, for much of 2017, the discussions of inflation in the ECB's Economic Bulletins predicted that past decreases in oil prices would dampen core inflation going forward. In the second half of 2018, after oil prices rose, the Bulletins predicted that increases in core inflation would follow. (See, for example, the discussions of the inflation outlook in the Bulletins for September 2017 and September 2018.)

These discussions of headline/core interactions are based on the conventional definition of core as inflation excluding food and energy prices, not median inflation. But this distinction, which is important in much of this paper, is not essential for the basic point here.

Although XFE and median inflation behave differently, movements in both are small relative to movements in headline inflation, meaning that the deviations of headline from core are not very sensitive to the measurement of core. This reflects the fact that food and energy (especially energy) are most often the industries with large price changes that influence headline inflation but are filtered out by the median.

Specifically, letting π^h denote headline inflation and π without a superscript denote core, the correlation between $\pi^h - \pi$ with π measured by median and $\pi^h - \pi$ with π measured by XFE is 0.94. Notice that when π is measured by XFE, $\pi^h - \pi$ is approximately proportional to the change in the relative price of food and energy. In Figure 4, the down/up/down cycle of headline inflation from 2008Q4 to 2016Q4 is explained mainly by a similar cycle in energy prices.

Why might non-core movements in headline inflation feed into core inflation? The ECB's December 2014 Bulletin discusses this question for the case of headline movements caused by oil-price changes. The Bulletin offers a straightforward explanation based on the chain of production. Increases in oil prices lead to increased costs of producing goods in non-energy industries that use oil as an input, which cause price increases in those industries and thereby contribute to core inflation. This process occurs over time as the effects of oil price increases move through production chains.

There is another channel through which headline inflation movements can affect core inflation: wage adjustment. This idea–again, focusing primarily on the effects of oil prices–is suggested in the classic work of Bruno and Sachs (1985) and appears in some studies since then, such as Peersman and Van Robays (2009). These authors suggest that the wage-setting process is such that nominal wages adjust to headline-inflation shocks, protecting workers' real wages from the shocks. Increases in nominal wages are then passed through, at least in part, into core inflation.

Both Bruno and Sachs and Peersman and Van Robays attribute the response of nominal wages to inflation to features of European labor markets, including formal or informal indexation arrangements and labor unions with the strength to protect real wages. These authors contrast the experience of Europe to that of the U.S., which lacks the labor market institutions that protect real wages from headline-inflation shocks. We return to this comparison when we discuss the U.S. Phillips curve below.

A Modified Phillips Curve

We modify the basic Phillips curve, equation (1), by adding a term involving $\pi^h - \pi$, where π is core inflation as measured by the weighted median. This term captures deviations of headline inflation from core caused by large price changes in any industry; in practice, as we have discussed, it is heavily influenced by oil price changes.

Many past studies include variables similar to $\pi^h - \pi$ in Phillips curves for headline inflation. For example, Gordon's (1982) classic Phillips curve specification includes the change in the relative price of food and energy; again, this is roughly equivalent to including $\pi^h - \pi$ when π is measured by XFE inflation. Ball and Mankiw's (1995) Phillips curve includes a measure of large price changes in all industries that is closely related to this paper's $\pi^h - \pi$. As discussed above, however, researchers do not usually include such variables in equations for *core* inflation.[2]

Based on the idea that it takes time for headline inflation shocks to influence core inflation, we include lags of $\pi^h - \pi$ in our modified Phillips curve. Our preferred specification is

$$\pi_t = \pi_t^e + \alpha\overline{(y - y^*)}_{t-1} + \beta\overline{(\pi^h - \pi)}_t + \epsilon_t, \tag{2}$$

This equation assumes that quarterly core inflation is influenced by the average of $\pi^h - \pi$ from $t - 3$ through t. We have also experimented with specifications that allow different

[2] Ball and Mankiw measure the contribution to aggregate inflation of industries with relative prices that rise or fall by more than a cutoff of X percent (with X set to 10 or 25). This variable is positive when large price increases skew the distribution of price changes to the right, and negative when large price decreases skew the distribution to the left.

coefficients on the current $\pi^h - \pi$ and its three lags, with results consistent with our preferred specification: the contemporaneous term is significant; the three lags are jointly significant; and we cannot reject the restriction that the four terms have equal coefficients, so only their average matters.

Table 2 presents estimates of equation (2), with and without a constant term. The output gap remains highly significant and its coefficient is close to the one in the basic Phillips curve in Table 1. The new $\overline{\pi^h - \pi}$ term is also highly significant, and including it raises the $\overline{R}^2$ from 0.64 to 0.76. The coefficient on the new term is 0.34: if headline inflation over four quarters has exceeded core inflation by one percentage point, that raises current core inflation by about three tenths of a point. As in the basic Phillips curve, when the constant term is included it is insignificant.

Fitted Values for the Modified Phillips Curve

Figure 5 shows fitted and actual values of median inflation for the modified Phillips curve, both quarterly and four-quarter averages. We see that the fit of the equation improves markedly compared to the basic Phillips curve. For four-quarter averages, the fall in inflation from 2008Q3 to 2010Q2 matches the fitted values from the equation almost exactly. The equation also captures almost all of the rise in inflation from 2010Q3 to 2011Q4.

The missing inflation puzzle of recent years is attenuated. For four-quarter averages, actual inflation falls persistently below the fitted values only in 2017Q2, compared to 2014Q3 for the basic Phillips curve. The gap between fitted and actual four-quarter inflation peaks at 0.29 points in 2017Q4, when the fitted value is 1.52 and actual inflation is 1.23. The gap in the last quarter of the sample, 2018Q4, is 0.28 (fitted value of 1.80 and actual inflation of 1.52).

5 Comparison to the United States

We have developed a simple Phillips curve that captures the behavior of median inflation in Europe. Does our specification capture features of inflation behavior that are specific to Europe, or more general?

To address this question, we estimate the same Phillips curves for the United States that we have estimated for Europe. We consider both equation (1) with only the output gap and equation (2) that includes our term for headline-inflation shocks. As a measure of median inflation, we construct a quarterly series from the monthly series for weighted median CPI inflation published by the Federal Reserve Bank of Cleveland. Using annual OECD estimates of potential output in the U.S., we construct a four-quarter moving average of the output gap using the same approach as we used for Europe. We measure expected inflation with ten-year forecasts from the U.S. Survey of Professional Forecasters. The sample period is 1986-2018, which is based on the availability of the potential output series.

Table 3 presents estimates of the U.S. Phillips curve. For the basic Phillips curve (1), the results are similar to those for Europe. The output gap is highly significant, and when a constant term is included it is not significant. The output gap coefficients are somewhat smaller than those for Europe (for the equation without a constant, the coefficient is 0.17 for the U.S. and 0.23 for Europe), but the differences are not statistically significant.

In contrast, when the $\pi^h - \pi$ term is added to the Phillips curve, the U.S. results diverge sharply from the European results. There is no evidence that the headline-inflation shocks captured by the $\pi^h - \pi$ term push core inflation in the same direction. The estimated coefficient on the variable is slightly negative and far from statistically significant ($t < 0.4$), and including the term reduces the $\overline{R}^2$.

Recall that one motivation for including the $\pi^h - \pi$ term in the European Phillips curve is that nominal wages respond to headline-inflation shocks. Bruno and Sachs and Peersman and Van Robays suggest that this phenomenon is the result of labor-market institutions that exist in Europe but not the United States. Our findings that $\pi^h - \pi$ influences core inflation in Europe but not the U.S. is consistent with this line of thinking.

Our results are also consistent with discussions of U.S. inflation by Fed policymakers and economists. In contrast to the European case, we have *not* found instances in recent years in which Fed sources suggest that headline inflation movements will feed into core inflation. Instead, when Fed officials discuss inflation developments, they often emphasize that deviations of headline from core inflation are transitory, and say or imply that these shocks are not relevant for the future path of core inflation. For example, Janet Yellen (2014) said:

> Inflation has continued to run below the Committee's 2 percent objective, and the recent sizable declines in oil prices will likely hold down overall inflation in the near term. But as the effects of these oil price declines and other transitory factors dissipate and as resource utilization continues to rise, the Committee expects inflation to move gradually back toward its objective.

Our results suggest that this Fed view is correct for the United States, just as the European view of pass-through from headline inflation is correct for Europe.

6 Conclusion

The analysis in this paper leads us to several conclusions. First, in Europe (as in the United States), the weighted median inflation rate is a more stable measure of core inflation, and one that is easier to understand, than the inflation rate excluding food and energy. Second, most of the ups and downs in median inflation since 1999 are explained by a simple Phillips curve that includes expected inflation and the output gap. Finally, the fit of the Phillips curve is even better if it includes a term capturing current and past deviations of headline inflation from median inflation.

Our preferred specification fits the data well enough that the much-discussed puzzles about inflation since 2008-the missing disinflation and the missing inflation-are mostly resolved. The main exception is that median inflation is below the levels predicted by our equation by modest amounts over 2017-2018. Going forward, we will see whether this deviation persists, which would suggest that new thinking is needed about the behavior of European inflation.

Another topic for future research is the usefulness of our Phillips curves for real-time forecasting of inflation. The series for the output gap used in this paper is constructed by the OECD based on data through the present. A critical issue for inflation forecasting is whether the output gap can be reliably estimated in real time.

Appendix: A Phillips Curve with Unemployment Gaps

Empirical studies of the Phillips curve sometimes measure economic slack with the gap between output and potential output, and sometimes with the gap between unemployment and its natural rate. The main text of this paper uses the output gap and this Appendix uses the unemployment gap. The basic Phillips curve becomes:

$$\pi_t = \pi_t^e + \alpha\overline{(u - u^*)}_{t-1} + \epsilon_t, \tag{A1}$$

We use annual OECD estimates of the natural rate u^*. We use this series and quarterly unemployment data to construct the four-quarter moving average of $u - u^*$ that appears in (A1); our method parallels our construction of four-quarter averages of $y - y^*$ in our main analysis.

Table A1 presents estimates of the basic Phillips curve (A1), and of that equation augmented with our headline-inflation-shock variable, $\overline{(\pi^h - \pi)}_t$. We omit constant terms (which are insignificant if we include them). Broadly, the results tell the same story as our main results with slack measured by the output gap. The coefficient on the unemployment gap is negative and highly significant, confirming the basic Phillips-curve tradeoff between economic slack and inflation.

However, the fit of the equations is not as good as when slack is measured with the output gap. Without the $\pi^h - \pi$ term, the $\overline{R}^2$ is 0.64 for the Phillips curve with the output gap, and only 0.53 with the unemployment gap. With the $\pi^h - \pi$ term, the $\overline{R}^2$s are 0.76 with the output gap and 0.67 with the unemployment gap.

Figure A1 shows four-quarter moving averages of actual median inflation and the fitted values from the Phillips curve with slack measured by the unemployment gap, and $\overline{(\pi^h - \pi)}_t$ included. We can compare this Figure to Figure 5B, which is the same except that slack is measured by the output gap. We see that the fit deteriorates with the unemployment gap in part because the fitted values of median inflation fall by substantially less than actual median inflation over 2009Q1-2011Q2, in the aftermath of the global financial crisis. These results suggest there was an "excessive disinflation" compared to the predictions of the Phillips curve. As shown in Figure 5B, with the output gap in the Phillips curve, the fitted values over the same period match actual inflation very well.

These results reflect the fact that the series for $\overline{(u-u^*)}$ rises only modestly above zero after the global financial crisis, reaching 0.42 in 2009Q4. This level of slack is much lower than the level following the European debt crisis, when $\overline{(u-u^*)}$ rises to 3.00 in 2013Q4. By contrast, the $\overline{(y-y^*)}$ series implies almost as much slack in the first episode (the variable falls to -3.36 in 2009Q4) as in the second (the variable falls to -3.81 in 2013Q4). The Phillips curve is better able to explain the substantial fall in inflation in both episodes when slack is measured with the output gap.

References

Abdih, Y., L. Lin, and A. Paret (2018): "Understanding Euro Area Inflation Dynamics: Why So Low for So Long?" IMF Working Papers 18/188, International Monetary Fund.

Ball, L. and N. Mankiw (1995): "Relative-Price Changes as Aggregate Supply Shocks," *The Quarterly Journal of Economics*, 110, 161–93.

Ball, L. and S. Mazumder (2019): "The Nonpuzzling Behavior of Median Inflation," NBER Working Papers 25512, National Bureau of Economic Research, Inc.

Bruno, M. and J. Sachs (1985): *Economics of worldwide stagflation*, Harvard University Press.

Bryan, M. and S. Cecchetti (1994): "Measuring Core Inflation," in *Monetary Policy*, NBER, NBER Chapters.

Ciccarelli, M. and C. Osbat (2017): "Low inflation in the euro area: Causes and consequences," Occasional Paper Series 181, European Central Bank.

Cœuré, B. (2019): "The rise of services and the transmission of monetary policy," *Speech at the 21st Geneva Conference on the World Economy*.

Draghi, M. (2019): "Twenty Years of the European Central Bank's Monetary Policy," *Speech at the ECB Forum on Central Banking, Sintra*.

ECB (2017): "Monthly Bulletin: Economic and Monetary Developments," Monthly Bulletin 8/2017, European Central Bank.

Friedman, M. (1968): "The Role of Monetary Policy," *American Economic Review*, 58, 1–17.

Gordon, R. (1982): "Inflation, Flexible Exchange Rates, and the Natural Rate of Unemployment," in *Workers, Jobs, and Inflation*, ed. by M. Baily, The Brookings Institution, 89–158.

Peersman, G. and I. Van Robays (2009): "Oil and the Euro area economy," *Economic Policy*, 24, 603–651.

Stock, J. and M. Watson (2009): "Phillips Curve Inflation Forecasts," in *Understanding Inflation and the Implications for Monetary Policy*, ed. by J. Fuhrer, Y. Kodrzycki, J. Little, and G. Olivei, Cambridge: MIT Press, 99–202.

Yellen, J. (2014): "Transcript of Chair Yellens FOMC Press Conference," *FOMC Press Conference*.

Table 1: Euro Area Basic Phillips Curve, 1999Q1-2018Q4

$$\pi_t = \pi_t^e + \alpha\overline{(y - y^*)}_{t-1} + \epsilon_t$$

	Median Inflation		XFE Inflation	
$Constant$		-0.052		-0.320
		(0.060)		(0.086)
α	0.228	0.221	0.238	0.194
	(0.024)	(0.025)	(0.037)	(0.031)
$\overline{R}^2$	0.643	0.646	0.207	0.459
$S.E. of Reg.$	0.345	0.343	0.540	0.447

Note: OLS with robust (HAC) standard errors is used (standard errors in parentheses). π_t is represents core inflation, and π^e is the ECB's SPF mean point forecast of 5-year ahead inflation (1999Q2-4 and 2000Q2-Q4 are linearly interpolated due to missing data). The gap variable uses OECD data on the output gap to derive estimates for $\overline{y}^*$.

Table 2: Euro Area Phillips Curve with Price Shock, Median Inflation, 1999Q1-2018Q4

$$\pi_t = \pi_t^e + \alpha\overline{(y - y^*)}_{t-1} + \beta\overline{(\pi^h - \pi)}_t + \epsilon_t$$

$Constant$		-0.066
		(0.047)
α	0.209	0.200
	(0.017)	(0.019)
β	0.341	0.349
	(0.065)	(0.062)
$\overline{R}^2$	0.755	0.764
$S.E. of Reg.$	0.286	0.280

Note: OLS with robust (HAC) standard errors is used (standard errors in parentheses). π_t is represents core inflation, π_t^H is headline inflation, and π^e is the ECB's SPF mean point forecast of 5-year ahead inflation (1999Q2-4 and 2000Q2-Q4 are linearly interpolated due to missing data). The gap variable uses OECD data on the output gap to derive estimates for $\overline{y}^*$.

Table 3: U.S. Phillips Curve, 1986Q1-2018Q4

$$\pi_t = \pi_t^e + \alpha\overline{(y - y^*)}_{t-1} + \beta\overline{(\pi^h - \pi)}_t + \epsilon_t$$

	Without Price Shocks		With Price Shocks	
$Constant$		0.040		0.036
		(0.065)		(0.065)
α	0.169	0.177	0.172	0.178
	(0.049)	(0.047)	(0.049)	(0.047)
β			-0.030	-0.026
			(0.080)	(0.080)
$\overline{R}^2$	0.665	0.664	0.664	0.662
$S.E.of Reg.$	0.506	0.507	0.507	0.508

Note: OLS with robust (HAC) standard errors is used (standard errors in parentheses). The gap variable uses OECD data on the output gap to derive estimates for $\overline{y}^*$.

Table A1: Euro Area Phillips Curve with Unemployment, 1999Q1-2018Q4

$$\pi_t = \pi_t^e + \alpha\overline{(u - u^*)}_{t-1} + \beta\overline{(\pi^h - \pi)}_t + \epsilon_t$$

	Without Price Shocks	With Price Shocks
α	-0.312	-0.286
	(0.037)	(0.027)
β		0.378
		(0.083)
$\overline{R}^2$	0.527	0.666
$S.E.of Reg.$	0.397	0.333

Note: OLS with robust (HAC) standard errors is used (standard errors in parentheses). π_t is represents core inflation, and π^e is the ECB's SPF mean point forecast of 5-year ahead inflation (1999Q2-4 and 2000Q2-Q4 are linearly interpolated due to missing data). The gap variable uses OECD data on the unemployment gap to derive estimates for $\overline{u}^*$. *Chow* shows p-value for stability test with break date of 2009Q1.

Figure 1: Euro Area XFE and Median Inflation

(a) Monthly

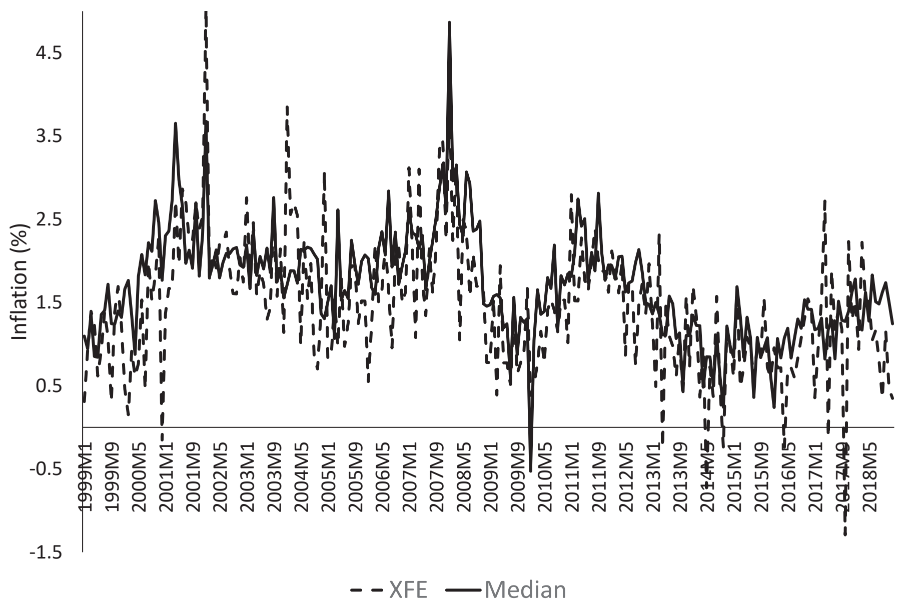

(b) Quarterly

(c) 4-Quarter Moving Average

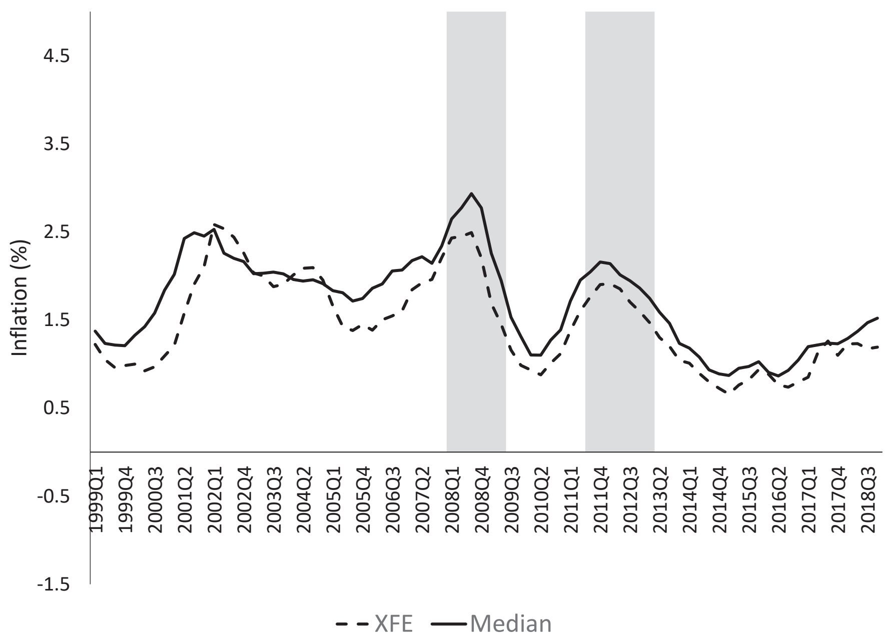

Figure 2: Histogram of Industry Price Changes in October 2017

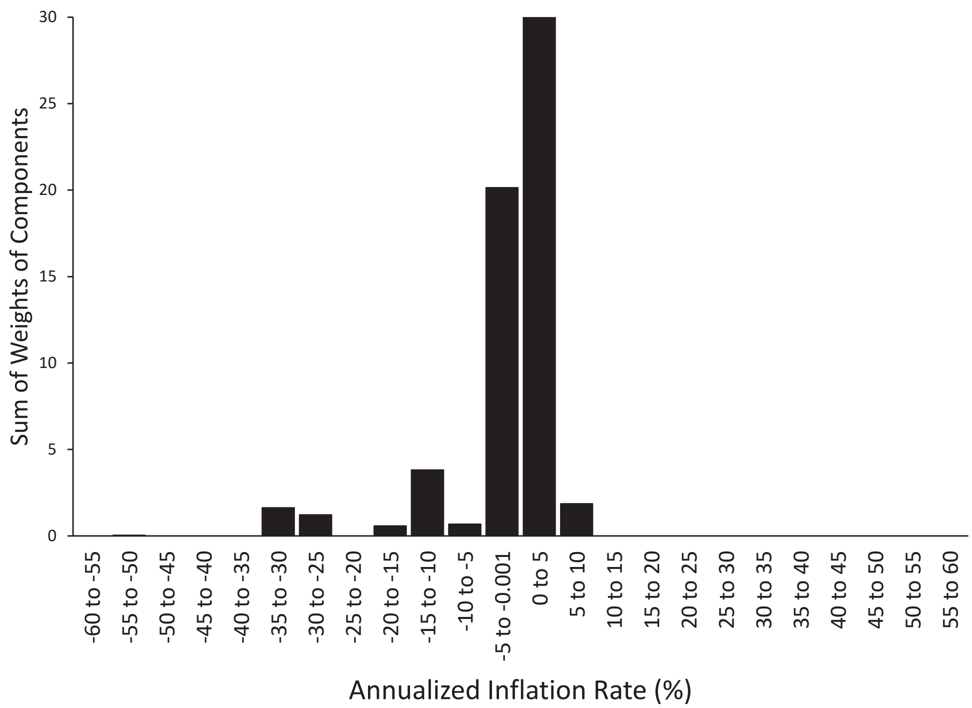

Note: The vertical axis is cut off at 30–the sum of industry weights in the 0 to 5% inflation range is 70. Food and energy industries are excluded.

Figure 3: Actual and Fitted Values from Euro Area Basic Phillips Curve (Table 1, column 1)

(a) Quarterly

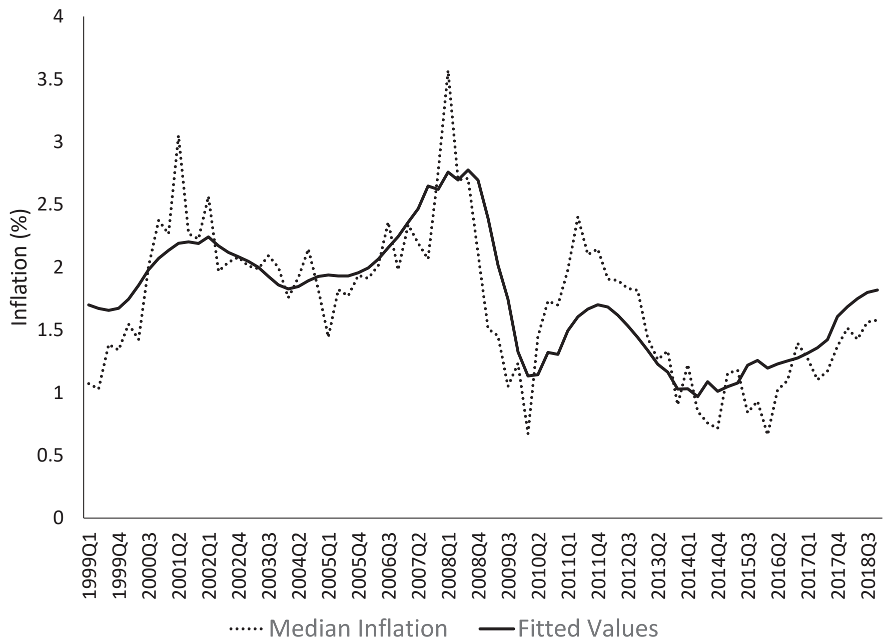

(b) 4-Quarter Moving Average

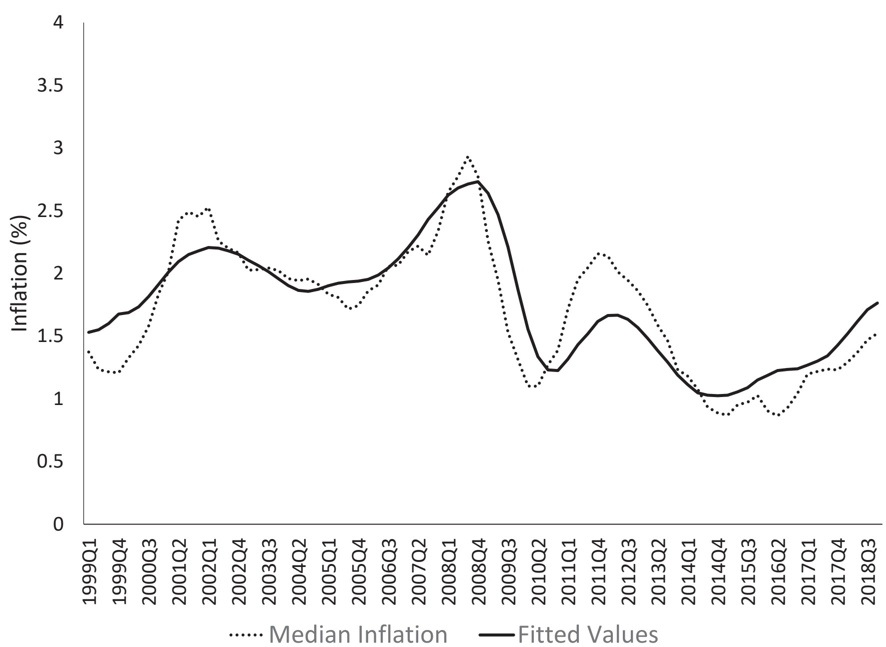

Figure 4: Euro Area 4-Quarter Averages of Median Inflation, Fitted Values from Basic Phillips Curve, and Headline Inflation

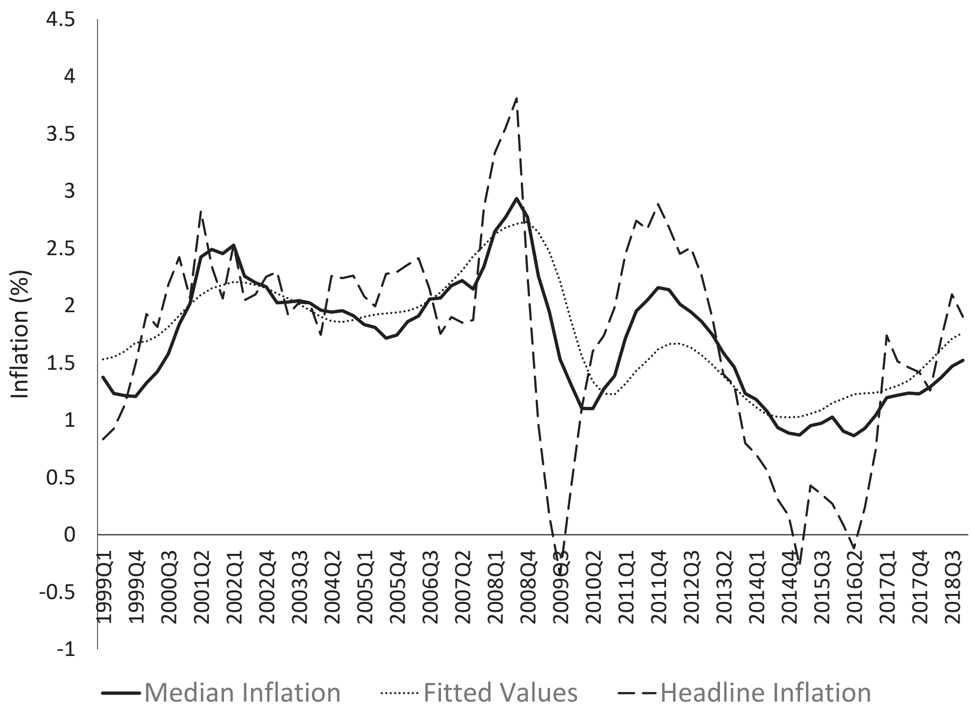

Figure 5: Actual and Fitted Values from Euro Area Phillips Curve with Price Shock (Table 2, column 1)

(a) Quarterly

(b) 4-Quarter Moving Average

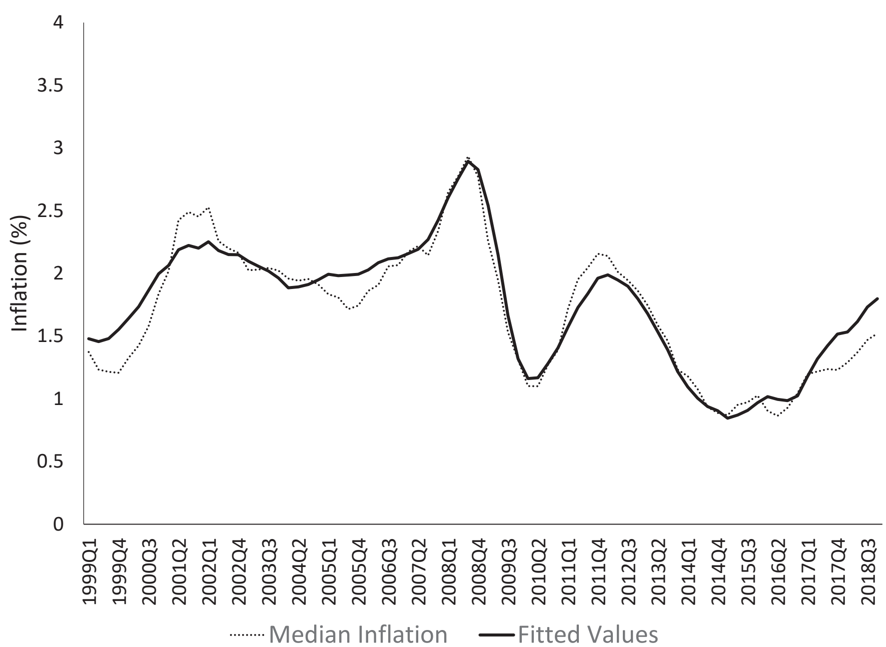

Figure A1: Actual and Fitted Values from Euro Area Phillips Curve with Unemployment Gap and Price Shock

(a) Quarterly

(b) 4-Quarter Moving Average