EUROPEAN CENTRAL BANK

EUROSYSTEM

Working Paper Series

Laura Moretti, Luca Onorante,
Shayan Zakipour Saber

Phillips curves in the euro area

No 2295 / July 2019

Abstract

We perform a robust estimation of the Phillips curve in the euro area using a battery of 630 theory-driven models. We extend the existing literature by adding model specifications, taking into account the uncertainty in the measurement of variables and testing for potential non-linearities and structural changes. Using Dynamic Model Averaging, we identify the most important determinants of inflation over the sample. We then forecast core inflation 12 quarters ahead and present its probability distribution. We compare the distribution of forecasts performed in recent years, and we assess, in a probabilistic manner, the convergence towards a sustainable path of inflation.

Keywords: Phillips curves, Dynamic Model Averaging, non linearities, structural changes, density forecast

JEL Codes: C30, E52, F41, E32.

Non-technical Summary

In this paper, we evaluate the importance of the Phillips curve, the standard theoretical and empirical benchmark stating a relationship between real activity and inflation, after the recent financial crisis. The paper provides four main contributions. First, we confirm the existence of a Phillips curve in the euro area, using simple models and testing different cycle indicators.

Second, we proceed to its robust estimation, extending the existing literature. To account for model specification uncertainty, we estimate, using Dynamic Model Averaging, a battery of 630 models. Using inclusion probabilities, we identify the main determinants of core inflation over the sample and we confirm that the main drivers of core inflation change between the first and the second dip of the recession. The first dip is characterized by a stronger role of external variables, while the second by domestic factors. Another robust finding is that expectations are the single most important determinant of core inflation in the sample.

Third, we estimate the slope of the Phillips curve and we expand the battery to 1260 models to test for non-linearities. We provide evidence in favour of using unemployment-based measures of the cycle and we conclude that the slope of the Phillips curve is sizeable and statistically significant. However, we do not find any evidence of non-linearities.

Finally, using the battery of 630 models, we forecast the probability distribution of HICP inflation excluding energy and unprocessed food three-year ahead for different samples ending in 2016Q1, 2017Q1, and 2018Q1. At each point in time, the distribution accounts for shocks, parameter and model uncertainty. We find an increasing, although still moderate, probability of core inflation to converge towards its long-term average, compatible with headline inflation reaching the objective.

We conclude that the Phillips curve is still a valid policy instrument once it is robustly estimated.

1 Introduction

The Great Recession and the following sovereign crisis have determined a long period of low inflation and put the monetary policy of the ECB to a severe test. The link between inflation and the economic cycle seems to have weakened after 2008, and inflation has proven difficult to forecast and to explain ex post. Headline inflation partially recovered after 2017, mostly driven by energy prices, but core inflation remains subdued, despite the strong positive contribution of unconventional monetary policy.

This paper evaluates the importance of the Phillips curve, the standard theoretical and empirical benchmark stating a relationship between real activity and inflation, after the recent financial crisis. We perform a robust estimation of the Phillips curve in the euro area and compute the main determinants of core inflation over time. Furthermore, accounting for parameter and model uncertainty, we forecast the distribution of core inflation in the medium term.

Our paper provides four main contributions. First, we show that, using simple models in the spirit of Giannone et al. (2014), the Phillips curve can explain the inflation dynamics also in the aftermath of the crisis.

Second, we extend the existing literature and estimate the curve robustly, considering different specifications (see for example Ciccarelli and Osbat, 2017, among others). We estimate a battery of 630 models using *Dynamic Model Averaging* (DMA), an econometric technique that allows the model used (and the regressors) to change over time.[1] In line with our Bayesian approach, we compare different specifications on the basis of their out-of-sample forecasting performance (See Jarocinski and Lenza, 2016). Given the uncertainty over the level of slack in the economy, we test traditional measures of output gap and unemployment, but also U6, a broad measure of unemployment. Following previous studies highlighting the importance of global factors[2], we add external factors (including import price deflator, real effective exchange rate and world industrial production). Moreover, we compare survey and market-based measures of inflation expectations. The use of model averaging techniques will allow us to identify the most important determinants of core inflation over the sample. We confirm the findings in Bobeica and Jarocinski (2019) that external factors were more relevant during the first trough, but after 2012 domestic

[1] See Raftery, Karny and Ettler (2010) for further details and Koop and Korobilis (2012) for an application to inflation forecasting using US data.

[2] See Ferroni and Mojon (2016), Forbes (2018), among others.

factors (proxied by labour market indicators) have increased their influence. Moreover, our analysis emphasizes the role of inflation expectations as determinant of core inflation over the sample and their overall contribution to inflation movements.

Third, we estimate the slope of the Phillips curve and we expand the battery to 1260 models to test for non-linearities after the Great Recession. We show that the slope of the Phillips curve is sizeable and statistically significant, in particular when using unemployment-related measures of slack. However, we do not find any evidence of non-linearities in the Phillips curve.

Finally, we forecast, using our battery of 630 models, the *distribution* of HICP inflation excluding energy and unprocessed food (HICPx) three-year ahead in different moments (2016Q1, 2017Q1 and 2018Q1). Taking advantage of our Bayesian approach, we assess the probability of core inflation to converge to its long-term average.

The rest of the paper is structured as follows. Section 2 shows, using a simple model that the Phillips curve relationship holds in the euro area after the recent financial crisis. Section 3 introduces the methodology, and Section 4 presents the data used. Section 5 discusses the results and 6 concludes.

2 Is there a Phillips curve in the Euro Area?

The Phillips Curve, a backbone of macroeconomics stating the relationship between inflation and economic slack, has been at the center of the recent policy and academic debate. During the stable inflation environment of the so-called 'Great Moderation' many economists have argued that the relation was not holding any longer. A vast literature analyses whether it was due to luck, i.e the absence of major economic shocks, or good macroeconomic policies, in particular monetary.[3]

However, Giannone et al. (2014) document the re-emergence of the Phillips curve in the euro area during the Great Recession, and Stock and Watson (2008) suggest that some forms of non-linearities make the Phillips curve stronger when deviations of unemployment from its natural level are large. Nevertheless, the persistent low inflation in the presence of a closing output gap has led to a renewed debate about the usefulness of the curve as a policy instrument both in the United States (Ball and Mazumder, 2011, and Coibion and Gorodnichenko, 2015) and in the Euro Area (Bobeica and Jarocinski, 2019

[3] See Dotsey et al. (2017), Atkenson and Ohanian (2001) and Stock and Watson (2007) for a discussion on the usefulness of Phillips curves for forecasting.

and Blanchard et al., 2015). Recently, Berson et al. (2018) show that the slope of the Phillips curve has remained stable and significantly different from zero in a sample of G7 countries, although it has flattened out since the 1980s.

The existence of the Phillips curve has important policy implications in the current juncture. The absence of a systematic relation between slack (i.e. output gap or unemployment) and inflation would imply that demand-side policies are not very effective on prices. Instead if the Phillips curve holds, demand policies and the ECB monetary policy in particular have a stronger effect on prices. Additionally, the closure of the economic gap should naturally push up inflation towards the ECB's objective.

In this section, we use a simple exercise to confirm the existence of the Phillips curve in the euro area after the crisis. We estimate a series of simple bivariate-BVARs, each containing core inflation and one measure of real activity over the pre-crisis sample 1999q1-2007q4. We set the lag order to 4 and use a loose Minnesota prior.[4] We use three output gap measures (i.e. from the OECD, the European Commission, and the IMF WEO) and three unemployment measures (i.e. rate, gap, and U6[5], the broad indicator of unemployment). We focus on core inflation in order to abstract from the fluctuations in food and energy prices, which would require a high number of controls. We then perform conditional forecasts over the period 2008Q1-2018Q1 using Waggoner and Zha (1999) and Blake and Mumtaz (2012) to obtain a distribution of the forecasts conditional on the actual path of each slack measure.

Our results indicate that the Phillips curve has remained relevant after the Great Recession. Figure 1 displays the actual path of core inflation (dashed black line), and the median forecasts of core inflation conditional on the different real activity measures (solid lines). It is worth noticing that the difference in the predictions only depends on the measure of real activity used. Although in the immediate aftermath of the crisis all measures tend to over predict core inflation, not surprisingly since the disinflation was mainly driven by external factors, afterwards unemployment measures are able to explain a considerable portion of the variance in core inflation. Looking at recent times, most measures tend to follow the actual path of core inflation quite closely, confirming the existence of the Phillips curve.

[4]The results are robust to alternative lag orders of 1 and 8. The tightness of the prior follows Blake and Mumtaz (2012) with $(\lambda_1, \lambda_3, \lambda_3, \lambda_4)$ are set to (0.2, 0.5, 1, 10^5).

[5]U6 takes into account also of discouraged workers in the computation of unemployment rate. See Section 4 for more details.

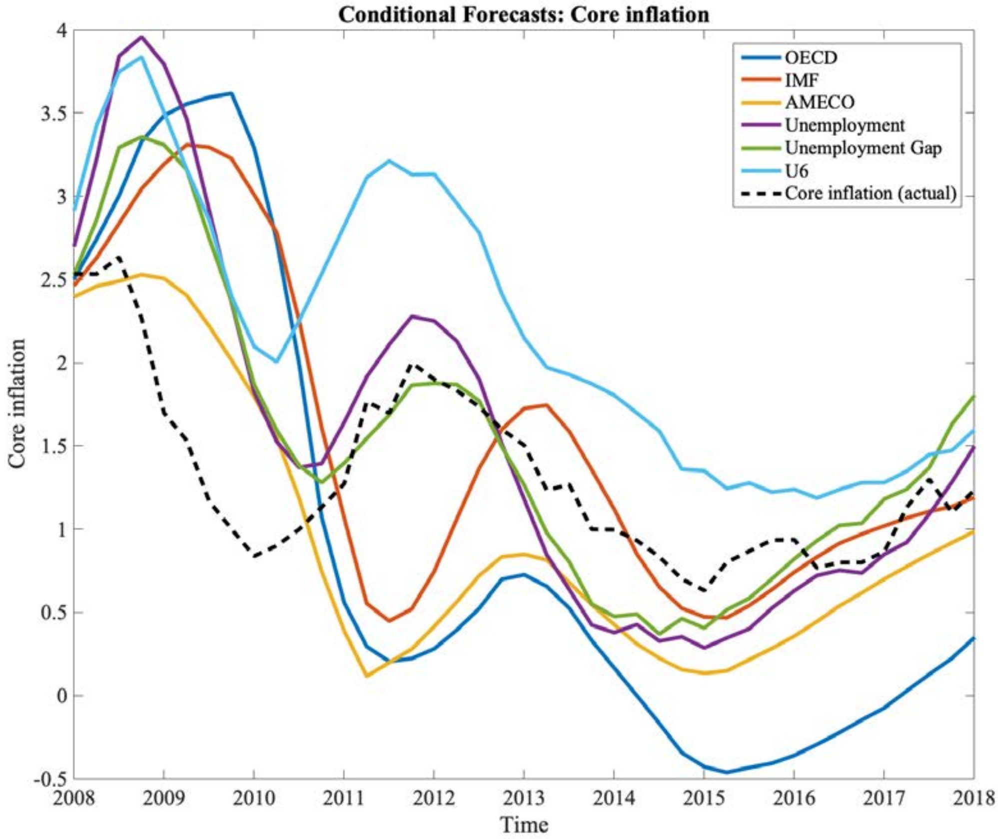

Fig. 1: Conditional forecasts for several indicators of economic conditions.

Figure 2 compares the conditional forecasts using the two measures with the lowest root mean square of the median forecasts, i.e. the best unemployment-rate predictor and best output-gap predictor, and including the one-standard deviation intervals. In particular, the unemployment rate is a valid indicator in real-time, as opposed to gap measures, which in addition are regularly revised, and does not suffer from mis-measurement issues. Moreover, despite the small uncertainty bands, since 2011 core inflation has been broadly within the one-standard deviation bands. Overall, the Phillips curve works well in real time when adopting a specification that uses unemployment as a measure of the cycle.[6]

[6]It is important to emphasize that this is a pseudo-real time exercise, in other words we do not use real time data but the latest (revised) vintages, thus giving an advantage to output gap measures. Instead unemployment figures are not revised and readily available, and only U6 is available with more than one quarter delay.

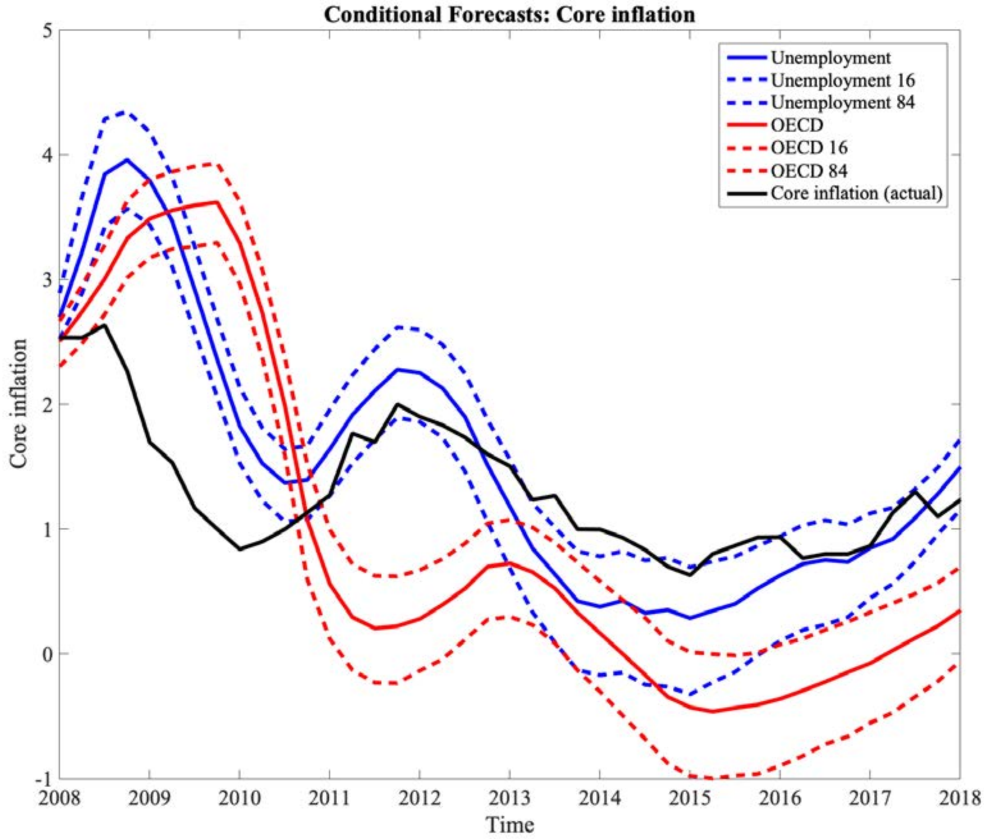

Fig. 2: Forecast distribution for best indicators of economic conditions.

Our preliminary exploration suggests the following conclusions. First, even a very simple Phillips curve model can partially explain the movements in the immediate aftermath of the crisis, as suggested by Giannone et al. (2014) for headline inflation. Second, the model does not fully capture the dynamics of core inflation during the first euro area disinflation due to the absence of global variables in the specification. However, it explains them well during the second trough, driven mainly by weak domestic activity. Finally, the different paths in conditional forecasts show that the choice of the economic activity indicator is important, as suggested in a recent contribution by Lenza and Jarocinski (2016).

Although a simple model can demonstrate the presence of the Phillips curve after the Great Recession, there is a substantial model uncertainty surrounding the Phillips curve not only regarding the specification of the model, but also the measure of slack to be used.[7] Therefore, we propose the use of a methodology that allows us to estimate the curve robustly, taking into account the uncertainty of model specification, to compute the main determinants of core inflation over time, and to forecast core inflation over the medium term.

3 Econometric framework

In this section, we discuss the methodology used to estimate the battery of Phillips curves, to identify the main determinants of core inflation over the sample, and to forecast core inflation over the medium term.

The uncertainty in both the model specification and the regressors to include is the main issue in estimating the Phillips curve. This makes it impossible to relay on a single model for policy assessments. One solution is to resort to Bayesian model averaging (BMA), an econometric technique that allows the researcher to be agnostic about the specification, to estimate a large battery of models and to average them based on their forecasting accuracy. BMA has the advantage of using parsimonious models that yield more stable estimates because fewer degrees of freedom are used in individual models. Furthermore, it allows the identification of important regressors, making the results more informative and easier to interpret.

However, the relevant model for forecasting, as well as the coefficients on the predictors, might change over time, for example the set of predictors during the first trough might be different from the ones in the second (see Bobeica and Jarocinski, 2019). Therefore, we use Dynamic Model Averaging (DMA), which allows the weights on each forecasting model and the coefficients of the predictors to change over time, thus dealing with potential structural changes. DMA was first proposed by Raftery et al. (2010) and allows the weights used in the model averaging to change overtime, making it possible to uncover the role and the importance of different regressors in the sample.

DMA has affinity with other model selection techniques. For example, both DMA and thick modelling consider a multiplicity of predictive models rather than a single one. Thick modelling, however, considers or discards models on the basis of simple heuristics (e.g. keeps a percentage of best performing models) rather than on the complete predictive likelihood. Additionally, it is not dynamic, and if the specification of the true model changes over time (e.g. possibly before and after the crisis), thick modelling produces invalid in-

[7] See ECB (2014).

ference. DMA also shares important features with the recursive modelling approach of, e.g., Pesaran and Timmermann (1995), where the forecasting model changes on the basis of some measure of past performance. BMA techniques, however, have the advantage of accounting for the within-model uncertainty. Moreover, the Bayesian literature has shown that it is possible to improve on the performance of even the best model by averaging on the basis of the complete forecasting distribution at the price of a higher computational complexity.

In the next section we present the methodology in more details.

3.1 Model uncertainty: Dynamic Model Averaging

DMA is developed in Raftery, Karny and Ettler (2010) and used in Koop and Korobilis (2012). The reader is referred to these papers for complete details. The dynamic aspect of DMA arises because it allows for a different model to hold at each period in time. We assume a population p_k of K models

$$p_k\left(y_t|y^{t-1}\right), k = 1..K \tag{3.1}$$

where $y^s = (y_1, .., y_s)'$ is the past information up to time s and $p_k\left(y_t|y^{t-1}\right)$ is the predictive density for model k at time t. We estimate our battery of models and evaluate them on the basis of their out-of-sample properties (on predictive density). Let $q_{t|s,j} = \Pr\left(k = j|y^s\right)$ be the probability that model j holds at time t given information through time s. DMA is a recursive algorithm which allows for the calculation of $q_{t|t,j}$ and $q_{t|t-1,j}$ for $j = 1, .., K$. Once calculated, weights $q_{t|t-1,j}$ can be used when forecasting y_t given information through time $t-1$. They can also be used to compute the "inclusion probability" of a variable or a set of models, i.e. the probability (and the importance) of these models relative to the complete set of K models. When estimating coefficients or impulse responses, $q_{t|t,j}$ can be used to carry out model averaging in a time-varying fashion.

To see how the weights are calculated, note that the predictive density appears in the model updating equation of:

$$q_{t|t,s} = \frac{q_{t|t-1,s}\, p_k\left(y_t|y^{t-1}\right)}{\sum_{l=1}^{K} q_{t|t-1,l}\, p_l\left(y_t|y^{t-1}\right)}. \tag{3.2}$$

If we knew $q_{t|t-1,s}$ then, starting with an initial $q_{0|0,s}$ (in our case set to equal weights, $q_{0|0,s} = 1/N$) we would be able to recursively calculate the key elements of DMA: $q_{t|t,j}$

and $q_{t|t-1,j}$ for $j = 1, .., K$. Raftery et al. (2010) provide this missing link by using the approximation:

$$q_{t|t-1,s} = \frac{q^{\alpha}_{t-1|t-1,s}}{\sum_{l=1}^{K} q^{\alpha}_{t-1|t-1,l}}. \tag{3.3}$$

A detailed justification of this approximation is given in Raftery et al. (2010). Suffice it to note here that it implies:

$$q_{t|t-1,s} \propto \left[q_{t-1|t-2,s} p_s \left(y_{t-1}|y^{t-2}\right)\right]^{\alpha} \tag{3.4}$$

$$= \prod_{i=1}^{t-1} [p_s(y_{t-i}|y^{t-i-1})]^{\alpha^i} \tag{3.5}$$

The previous equation emphasizes that a model j receives more weight at time t if it fit well in the recent past (fit is measured by the predictive likelihood, $p_j \left(y_{t-i}|y^{t-i-1}\right)$). The interpretation of "recent past" is controlled by the forgetting factor, α. Thus, if $\alpha = 0.99$ (our benchmark value and also the value used by Raftery et al., 2010), forecast performance five years ago receives 80% as much weight as forecast performance last period (when using quarterly data). If $\alpha = 0.95$, then forecast performance five years ago receives only about 35% as much weight.

In our short data set, the potential advantages of DMA are clear. We can include models featuring a large number of explanatory variables. However, if these are overfitted their predictive performance will be low and DMA will attach more weight to more parsimonious models, thus lessening the problems caused by the curse of dimensionality while keeping all candidate models.[8] Since we do not have to worry for misspecification, as misspecified models will be selected out by the DMA algorithm, we use for each model a weakly informative prior, where the regression parameters are centered around zero with unitary variance and the variance of the residuals is set to a very large value. Furthermore, DMA allows for model change. It can capture cases where certain explanatory variables or models frameworks are important in certain periods, but not in others. Given our application covers the time period since the introduction of the euro, allowing for such change may be important.

We note also that, in the past, DMA has been used in the context of time-varying parameter (TVP) models where the coefficients evolve following a random walk, i.e. as

[8]See Koop and Korobilis (2009a) for evidence that DMA can effectively find very parsimonious models.

$\beta_{it} = \beta_{i,t-1} + u_t$. Our set of models instead includes models with fixed parameters estimated recursively. Time varying parameter models were also tested, but usually underperformed out of sample (probably due to the overfitting in the relatively short sample) and were eliminated in the final DMAs. This result is consistent with Koop and Korobilis (2012), who found that allowing models to switch over time has greater empirical benefit than allowing coefficients to evolve in a TVP fashion.

4 Model Specification and Data

In the literature, there are different specifications of the Phillips Curve that include different variables and use different functional forms. An all-encompassing model would not be possible to estimate, given the high number of potential regressors, and the estimates would be meaningless as the presence of several combination of variables (e.g. the simultaneous presence of three different cycle indicators) would lead to multicollinearity. To account for model uncertainty, we adopt a robust approach and consider 630 specifications of the Phillips curve.

Following Stock and Watson's (2008) recommendation to use parsimonious models, we focus on univariate specifications, some of which include a substantial number of estimated parameters, and we compare their out-of-sample explanatory power.

The dependent variable is core inflation, defined as the year-on-year percentage change in the HICP excluding energy and unprocessed food, to abstract from more volatile components.[9] We consider a high number of regressors and divide them into four groups: real activity, inflation expectations, labour market indicators, and global indicators. Each specification of the Phillips Curve always includes an autoregressive component of inflation, one lagged real activity variable and a permutation of at most one lagged variable from each of the remaining groups.[10] Figure 3 presents a summary of the variables used. Our dataset is quarterly and spans the period from 2001Q3 to 2018Q1.[11]

The real activity group includes three measures of output gap (from OECD, the European Commission, and the WEO), the unemployment rate, the unemployment gap, i.e. the difference between unemployment rate and the NAIRU, and U6, a broader measure of

[9]Please note that the ECB refers to HICP excluding energy, food and tobacco as the main measure of underlying inflation. All or our regressions included either one or four lags of core inflation. However, the performance of the four-lags specifications was low, and we dropped them form the set of specifications.

[10]We include only one lag of inflation, but the results are robust to the inclusion of 4 lags.

[11]We extrapolate data for the 5years-in-5years swap rates using the French data.

unemployment. U6 is constructed by expressing the number of unemployed and underemployed, together with the potential additional labour force (i.e. the estimates of those available but not seeking work and those seeking work but not available), as a percentage of the extended labour force, i.e. the sum of the active labour force, which is employed plus unemployed, and the potential additional labour force.[12] Labour market indicators include (nominal) compensation per employee and total unit labour cost, a measure of labour cost adjusted for productivity.

Extending the previous literature (e.g. Coibion and Gorodnichenko, 2015), we consider both survey and market-based indicators of inflation expectations. We include three survey measures: the Survey Professional Forecasters (SPF) one year ahead, two years ahead and five years ahead; and three measures of market-based inflation expectations: the one year ahead inflation swap rate, the 1 year-in-1 year forward inflation linked swap (ILS) rate and the 5 year-in-5 year forward ILS rate. The market-based measures are chosen to match broadly the horizon of the survey measures, and to include the 5 year-in-5 year ILS rate, a benchmark measure of medium to long-term inflation expectations for central banks (see for example Draghi, 2014a).

Finally, building on previous literature that recognizes the importance of international factors in determining inflation (see Ciccarelli and Mojon 2010, Ferroni and Mojon 2016, Forbes, 2018, among others), we include external factors such as the real effective exchange rate (REER), the world industrial production, the import price deflator, and the contemporaneous price of oil. On the one hand, Auer et al. (2017) show the growing importance of global factors due to expansion of global value chains. On the other hand, Mikolajun and Lodge (2016) argue that, with the exception of commodity prices, there is little reason to include global factors into traditional reduced form Phillips curves. In our model, we include both oil prices and world industrial production as indicators of global factors, and there is no need to select one or the other because the model will weigh the more relevant determinant over the sample.

5 Empirical results

The models have initially equal weights ($q_{0|0,s} = 1/N$). We use the evolution of weights in the DMA to assess the role of each group of variables in explaining core inflation. Figure 4 reports the sum of the inclusion probabilities of the variables belonging to each group,

[12] See ECB, 2017, Box 3 for further details.

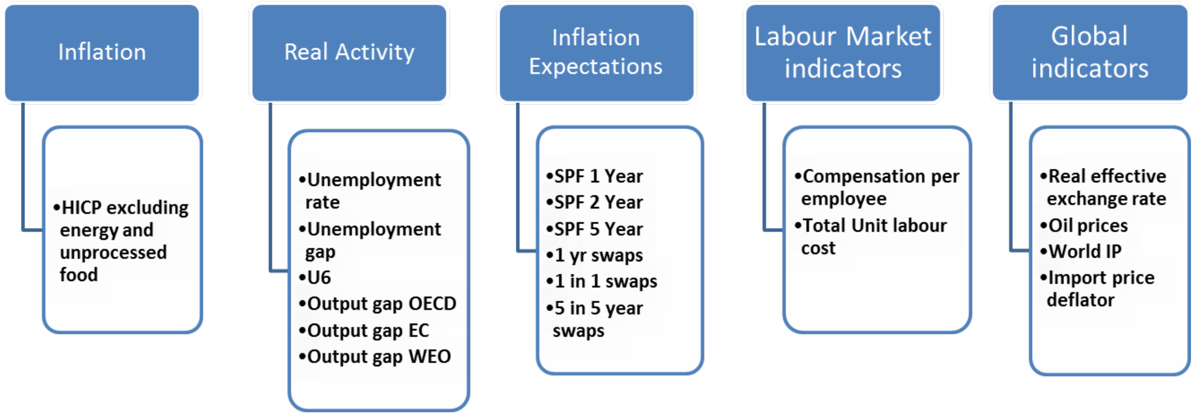

Fig. 3: Variables included in the model.

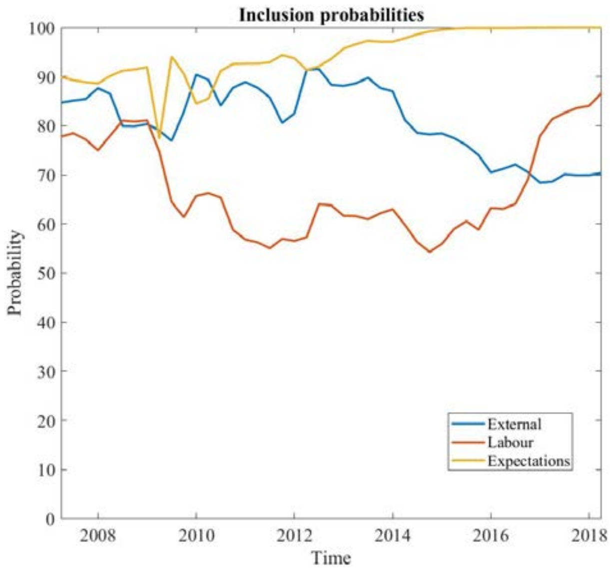

Fig. 4: Inclusion probability of the different groups of variables.

thereby showing the relative importance over time of inflation expectations (yellow line), external factors (blue line), and labour market, or domestic, indicators (red line).

First, it is important to notice that external variables are initially important, but their weight starts declining after 2012. Second, the role of domestic factors (proxied by the labour market indicators), although important before the crisis, picks up again after 2014. In fact, both the 2008 downturn and the sovereign crisis, with output away from potential, drastically reduced the importance of labour markets as a determinant of inflation. Although this might be compatible with a non-linear Phillips curve, where cost-push dynamics do not matter if the economy is far away from full capacity utilization, we do not find any evidence of non-linearities, as discussed in more detail in Section 5.4.

Finally, inflation expectations are the most important determinant of core inflation during (almost) the whole sample, second to external factors only briefly around 2010. In an environment in which the slope of the Phillips curve is significant but small, the anchoring of inflation expectations is crucial and the credibility of monetary policy is even more relevant. Jordá et al. (2019) reaches a similar conclusion for the US.

We now turn to a more detailed analysis for the single variables.

5.1 Cycle indicators

A well-know issue with the Phillips curve is that most measures of slack are based on output, which is available with considerable lags and is subject to substantial revisions. GDP in real time can be very different from the re-estimated one, and this makes the use of the curve in real time problematic. Additionally, gaps are also statistical artifacts: when used in estimation, they may bias the coefficient of slack towards zero, suggesting a spurious irrelevance of the curve.

To achieve robustness in estimation, we include in our specification several indicators of output gap and unemployment.[13] We compare cycle indicators by imposing that one of them must be present in each specification and we report the inclusion probabilities. It should be noted that we do not use real time indicators, but the last available vintage of output gaps, thereby giving an advantage to gaps over the more timely (and hardly revised) unemployment indicators. The only exception is U6, the broader measure of unemployment, that is available only with more than a quarter delay.

Figure 5 presents the inclusion probabilities of each measure of the cycle. Since we always include one measure of the cycle, these probabilities always sum to one. We find that different measures are better predictor in subsamples. The output gaps estimated by the OECD and the IMF have a higher performance at the beginning of the crisis (i.e. 2008-2009), while the output gap estimated by the European Commission remains relevant until 2010. Thereafter, unemployment measures become more important: unemployment gap (i.e. the difference between unemployment and NAIRU) is prominent in the aftermath of the sovereign crisis, followed by unemployment rate and U6. In particular, the inclusion

[13] Unemployment is hardly revised. For output gaps, we chose the last available vintage, because the focus is on the assessment of the Phillips curve ex-post rather than on forecasting.

probability of U6 has been increasing steadily since 2012, and it has become slightly larger than unemployment in the last few years.

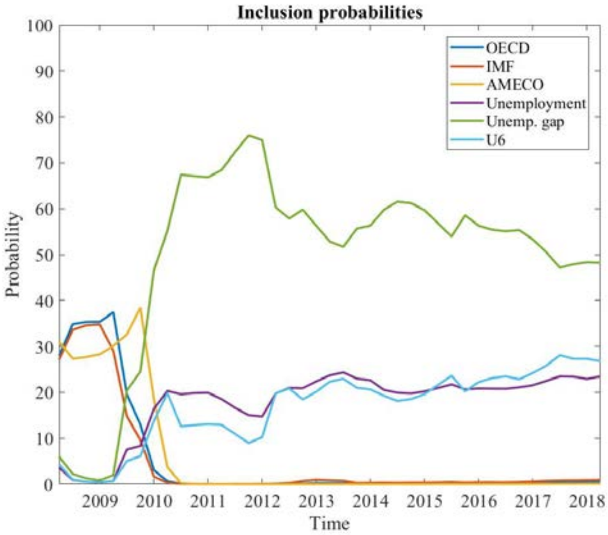

Fig. 5: Inclusion probability of different cycle variables.

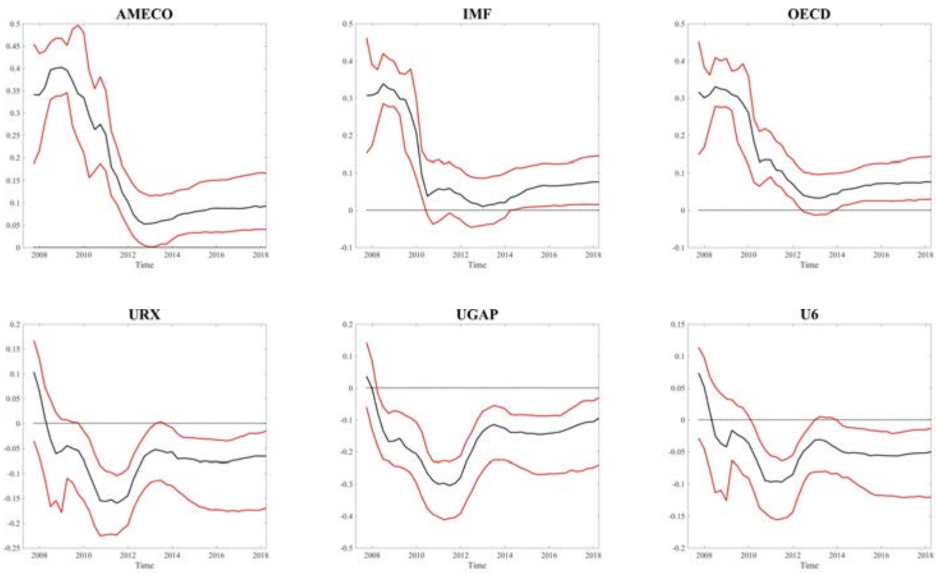

Notes: Recursive estimation across models of gap variables (one standard deviation confidence bands)

Fig. 6: Cycle variables - coefficients over time.

We turn now to the analysis of the slope of the Philips curve using recursive estimates for each of the cycle indicators considered (see Figure 6). When estimated with gap indicators, the Phillips curve appears to have flattened, a common puzzle in the literature. However, Figure 5 shows that gap indicators are associated with models having low weights in the model averaging, due to their poor predictive power. Gaps do not perform very well because they are filtered (and therefore noisier) and frequently revised measures of the economic cycle. The second row shows instead the unemployment-related predictors, for which the slope of the curve is sizeable and statistically significant. Overall, our robust regressions confirm the preliminary results in Section 2 and suggest that the Phillips curve is best measured when using unemployment-related variables of the cycle in estimation.

5.2 Labour market indicators

Stagnating labour markets have been indicated as a possible cause of the persistence of low inflation. These considerations translate into some specific formulations of the Phillips curve. We first calculate the weight of labour market-related curves (those including wages), then the contribution of these two variables to overall inflation and its evolution over time. Figure 7 shows the inclusion probabilities for unit labour cost and compensation per employee. The results show that the importance of compensation per employee has remained stable over the sample, while the weight of unit labour cost has increased after 2012.

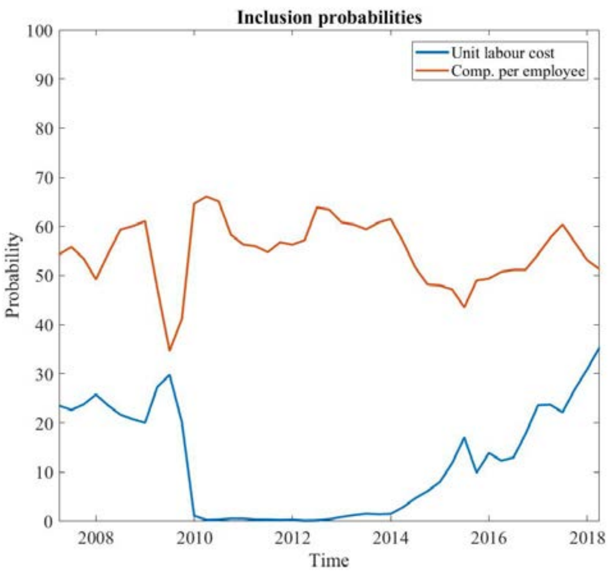

Fig. 7: Inclusion probability of different labour variables.

5.3 Measures of inflation expectations

In this section, we turn to the most important determinant of core inflation in our sample. Inflation expectations move the Phillips curve and can constitute a serious concern for the monetary authority. In fact, lower expectations determine permanently lower inflation, other factors being equal. It is then not surprising that in the current juncture a considerable attention has been given to this subject, both in theoretical models (Locarno et al., 2017) and in empirical analysis.

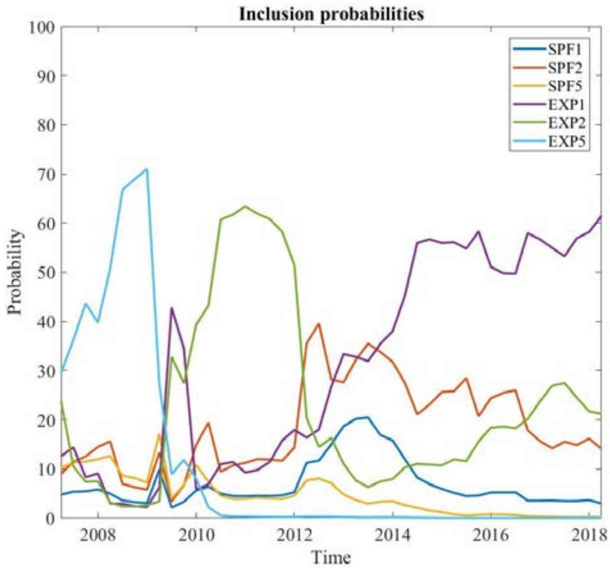

Fig. 8: Inclusion probability of different measures of inflation expectations

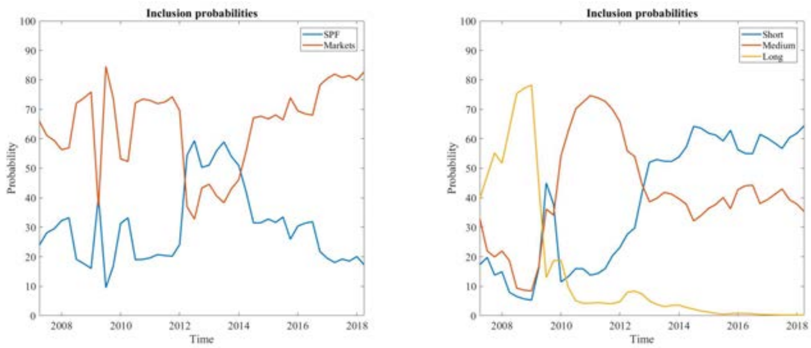

Fig. 9: Inclusion probability of inflation expectations: survey vs market-based measures (left), short vs medium vs long term expectations (right).

Inflation expectations can be empirically proxied using either market-based or survey measures, both with advantages and drawbacks. Market-based measures of inflation expectations are considered an important indicator of the credibility of monetary policy (Draghi, 2014b and Yellen, 2015). They are truthfully revealed, as agents disclose them by putting their money, but include risk premia and may be affected by market inefficiencies. Survey expectations in principle do not suffer from biases, but they tend to react slowly (Bowles et al., 2007, for the ECB's Survey of Professional Forecasters). Moreover, both represent the expectations of a small subset of the economic agents. In this paper, we use both survey and market-based measures of inflation expectations with analogous temporal profiles and we compare the role of expectations by type and time horizon.

Figure 8 presents the contribution of each variable and we can make three main observations. First, the 5 year-in-5 year ILS rate (5y5y) is the most relevant indicator during the financial crisis, but its weight declines sharply afterwards. Second, there is not a clear dominant indicator, with the 5y5y peaking between 2008 and 2010 and the 1 year-in-1 year ILS rate (1y1y) between 2010 and 2012, while after 2014 the one year ahead inflation swap rate becomes increasingly predominant.

When computing the inclusion probabilities aggregating all the survey measures and the market-based indicators, we observe that market-based measures react faster and remain relevant until 2012 and, after a short period in which survey measures are more relevant, again after 2014 (see Figure 9, left).

To analyze further the role of inflation expectations, we compare the importance of the inflation expectations measures by horizon. Figure 9 (right) presents the inclusion probabilities for short (i.e. SPF one year ahead and one year ahead inflation swap rate), medium (i.e. SPF two years ahead, and 1y1y), and long (SPF 5 years ahead, and 5y5y) indicators of inflation expectations. It is worth noticing that the long-term indicator, driven by the 5y5y, has a higher weight in the aftermath of the financial crisis, but, after the sovereign crisis (2010-2012), the medium term indicator gains importance, while the short-term index becomes the most relevant after 2013. The growing weight attributed to shorter term expectations suggests that the low levels of past inflation may be affecting the current developments.

5.4 Is there a non-linear Phillips curve?

It has been argued that the Phillips Curves may present non-linearities, whereby persistent supply side constraints would imply a stronger upward impact on core inflation. In this section, we test whether the Phillips curve is non-linear in the output gap measure.

We perform several tests of non-linearities in the curve. First, we add to our set of 630 base specifications the same models including the squared terms of the cycle (quadratic specifications), therefore estimating a total of 1260 models. If the Phillips curve is non-linear in the gap, we expect the formulations including quadratic terms to have a greater predictive likelihood than the corresponding linear specifications. We report two metrics: the relative importance ("inclusion probability") of models including the non-linear term; and the estimated coefficients (and their evolution over time) of the non-linear terms.

Figure 10 suggests that specifications that include squared output gaps appear to receive only little support from the data. Models with quadratic terms perform worse out-of-sample and therefore have lower weights in the final averaging. Starting from a prior of 50%, the non-linear models have in recent years an ex-post weight between 20% and 40%.

Moreover this result is confirmed by the fact that the coefficients of the quadratic terms are non significant, especially in recent years, and vary in sign. As shown in Figure 11, the non-linear part is not relevant for any of the six cycle indicators, with the partial exception of the European Commission output gap. Therefore, the Phillips curve does not show strong signs of non-linearities.

In the rest of the paper we restore the initial pool of 630 linear models.

5.5 Determinants of past core inflation - contributions

After presenting the determinants of core inflation, in this subsection we discuss the contribution of the different groups of variables, using optimal model weights, to the dynamics of core inflation (see Figure 12).[14]

[14]It should be noted that we run a pseudo-real time exercise, using the final vintages of the output gap. This introduces some uncertainty in determining the gap component of the decomposition. On the one hand, a complete real-time exercise might imply even lower inclusion probabilities and a smaller contribution of the gap in the inflation decomposition. On the other hand, the necessity of interpolating the output gaps provided by different institutions at an annual frequency may lad to an underestimation of the gap effect.

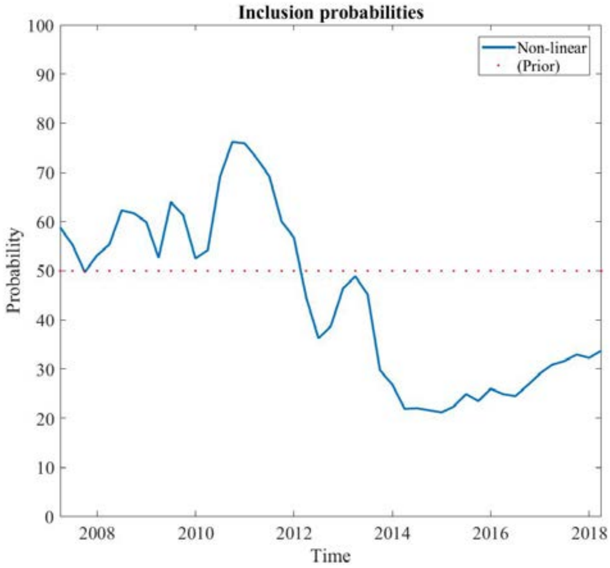

Fig. 10: Inclusion probability of non-linear specifications (prior: 50%).

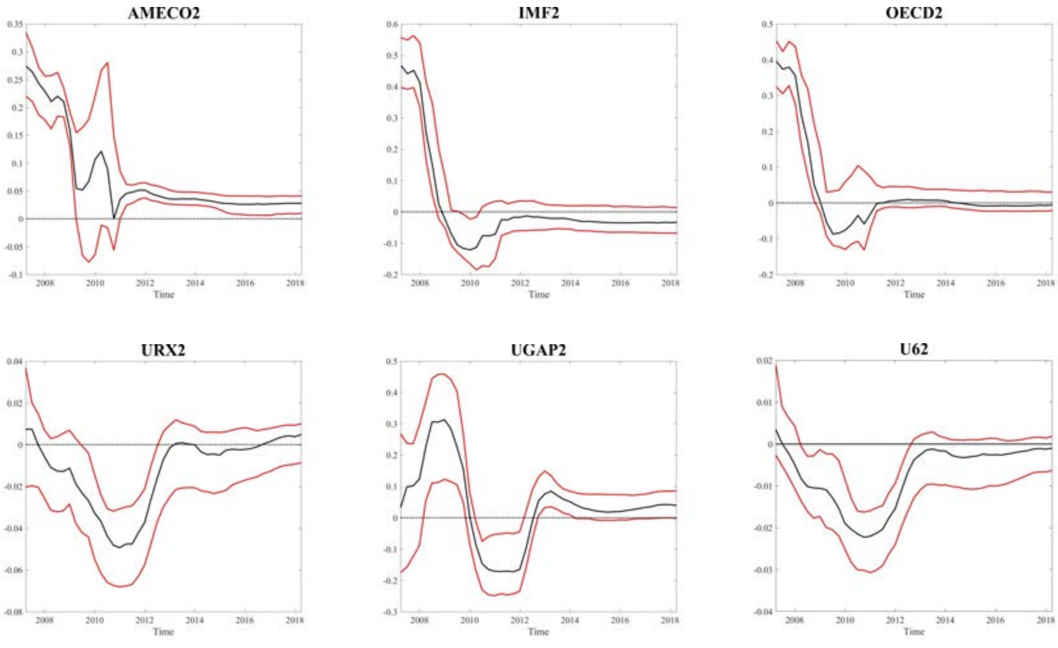

Notes: Recursive estimation across models of squared gap variables (bands are one standard deviation)

Fig. 11: Nonlinear curves - coefficients of quadratic terms over time.

The decomposition emphasizes the role of expectations and their reaction to different polices, and it could be divided in three periods. In the first, before 2008, expectations remained well anchored and provided a relatively minor contribution to the inflation peak.

In the second period, between the collapse of Lehman Brothers and the (euro area) sovereign crisis (i.e. from 2008 to the end of 2014), conventional and non-conventional policies were used jointly. The ECB lowered its key policy rate to an unprecedented level of 0.05%.[15] At the same time, to respond to the increased demand for liquidity and to reduce the risk of financial disruptions, the main refinancing operations have been conducted as fixed rate tender procedures with full allotment since October 2009. Additional measures included the targeted longer term refinancing operations (TLTROs), aimed at improving bank lending to the euro-area non-financial private sector (excluding loans to households for house purchase), the expansion of the list of marketable assets accepted as collateral in Eurosystem credit operations, the asset backed security (ABS) purchase programme and the covered bond purchase programme. In this phase, the contribution of inflation expectations was increasingly negative.

However, with the implementation of further unconventional monetary policy measures, the trend has reversed. The ECB Governing Council cut progressively interest rates[16] and announced the expanded asset purchase programme (APP) in January 2015 with combined monthly asset purchases of 60 billion.[17] Over time, APP has been recalibrated five times and contributed to the expansion of the ECB balance sheet. During this phase, the negative contribution of inflation expectations declined decidedly, supporting the progress towards the convergence of core inflation to its long-term average.

Overall, the chart shows that the ECB policies have been successful in the years following the Great Recession and inflation expectations have normalized progressively, contributing to avoid deflation. At the same time, for inflation to increase further, higher contributions should be expected from labour market wages and the closure of the output gap.

[15]The Governing Council of the ECB started decreasing interest rates in October 2008 reaching 1% in May 2009, but increased by 25bp in both April and July 2011, reverting the decisions in November and December 2011. The Governing Council decreased MRO further by 25bp in three meetings, in July 2012, in May and in November 2013, while keeping DFR at 0.0%. Furthermore, the Governing Council cut the MRO by 10bp both in June and in September 2014 to 0.05%, while driving the DFR into negative territory (-0.20%).

[16]The DFR was further cut by 10bp in December 2015 and in March 2016 to -0.40%, when MRO was also reduced by 5bp and reached 0.00%.

[17]The Governing Council recalibrated APP in December 2015 (extension until March 2017 and reinvestment of principal payments), in March 2016 (expansion of monthly asset purchases from 60 to 80 billion), in December 2016 (extension until December 2017). The Programme was further recalibrated in October 2017 (extension until September 2018 with a monthly pace of purchases of 30 billion starting from January 2018), and in June 2018 (extension until December 2018, with a monthly pace of 15 billion from October 2018).

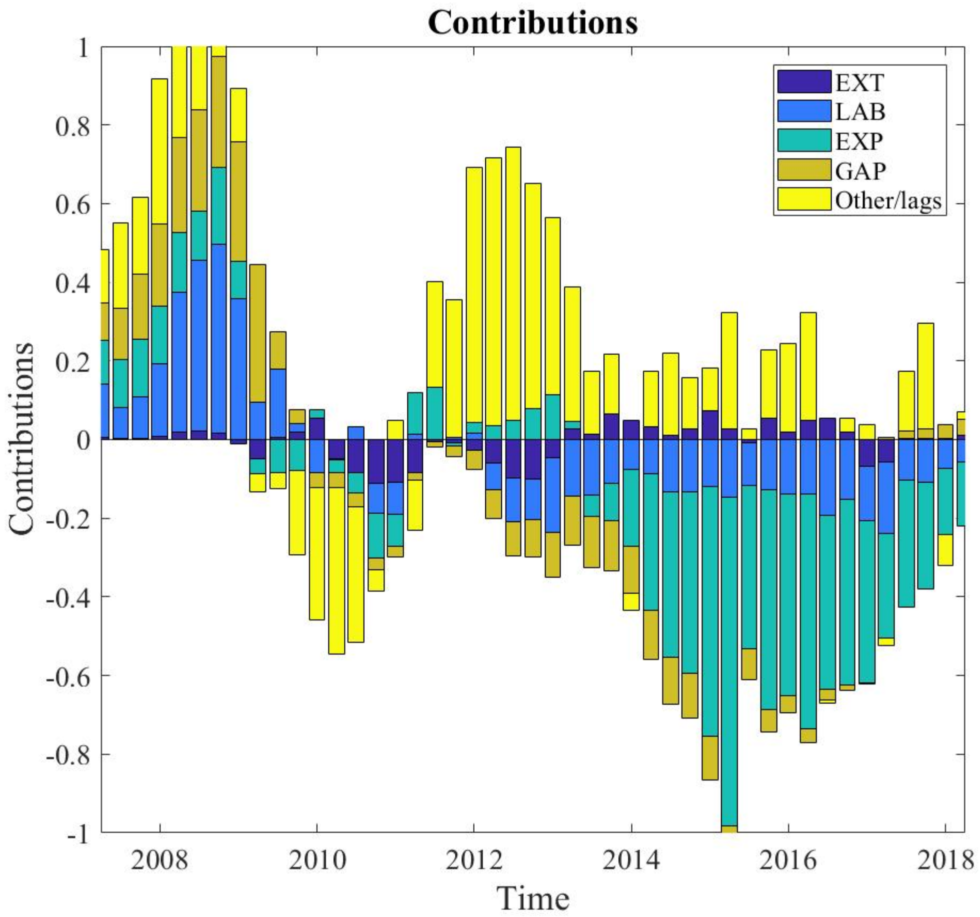

Fig. 12: Historical decomposition of core inflation.

5.6 Estimation and Forecast

Probabilistic forecasts are an important tool for policymakers who want to keep the risk of an adverse outcome (disinflation) to an acceptable level or to reach a goal (sustainable target) with an acceptable probability. In particular, the assessment towards a sustained path of inflation "requires three conditions to be in place. [...] The second is confidence: we need to be sure that this upward adjustment in inflation has a sufficiently high *probability* of being realised" (Draghi, 2018).[18]

In this section, we present the forecast of HICP inflation excluding energy and unprocessed food (HICPx) three-year ahead using the battery of 630 linear Phillips curve

[18]Italics is ours.

models.[19] We also compute the (optimally weighted) probabilistic distribution of forecast inflation and we propose to use it to assess, in probabilistic terms, the convergence of core inflation towards its long-term average.

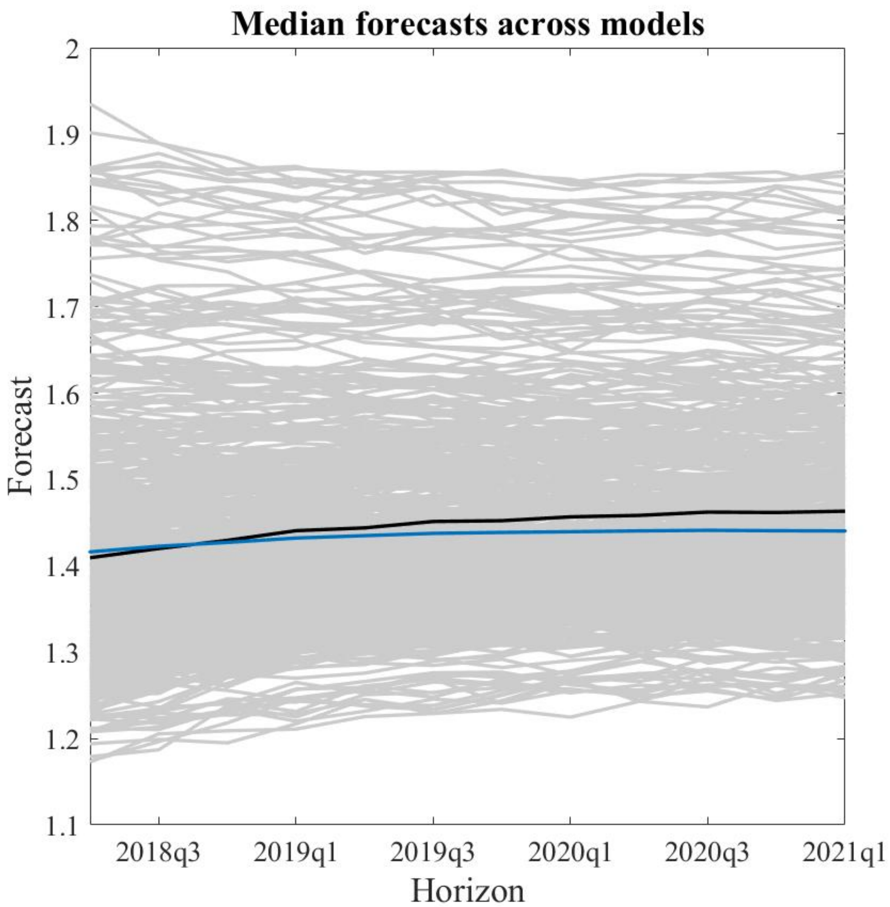

Fig. 13: Median forecasts of core inflation three years ahead using the battery of 630 models. All forecasts (gray), mean (blue line), DMA (black line).

Figure 13 presents the median unconditional forecasts three years ahead of all the models (gray lines), the simple average across models (blue line) and the DMA with the optimal weights (black line).[20] The robust forecast using DMA produces slightly lower inflation forecasts than the unweighted average of the models in the short run and is equiv-

[19]This section extends the work of Banbura and Mirza, mimeo, 2013. Differently from them, we use predictive likelihood as an evaluation criterion for different models and concentrate on the distribution rather than on the optimal averaging.

[20]We use the weights computed in the last period of the sample, and all forecasts are conditional to simple AR1 processes for the exogenous variables.

alent at longer horizons. The gray lines represent the median forecasts of each model, and show that there is a strong variation across single forecasts. This highlights the importance of taking into account the uncertainty within and between models rather than guessing a specification, as some reasonable specifications can be quite misleading. The forecasts using model averaging show that core inflation is likely to remain broadly at the current levels.

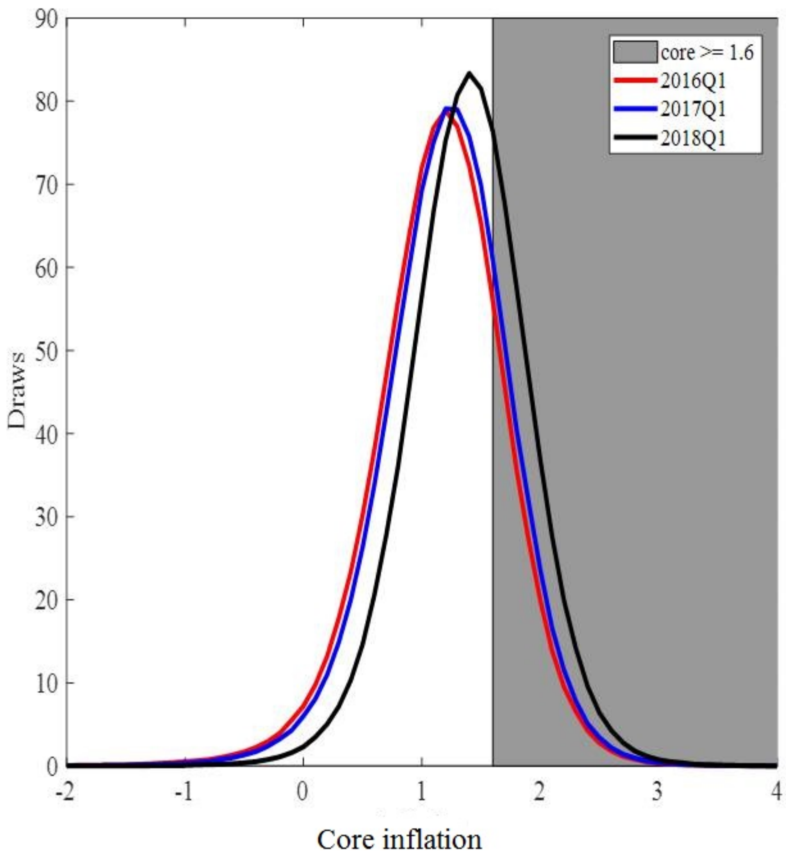

Fig. 14: Probability distribution of HICP excl. food and energy three years ahead using data ending in 2016Q1 (red), 2017Q1 (blue), and 2018Q1 (black).

We then move to the probabilistic assessment, based on the whole distribution of the forecast. Figure 14 compares the distribution forecast of HICPx inflation three-year ahead using samples ending in 2016Q1 (red), 2017Q1 (blue), and 2018Q1 (black). We derive the forecast distribution of each model in the battery and then, using the DMA optimal weights, we compute the distribution forecast for HICPx that takes into account the uncer-

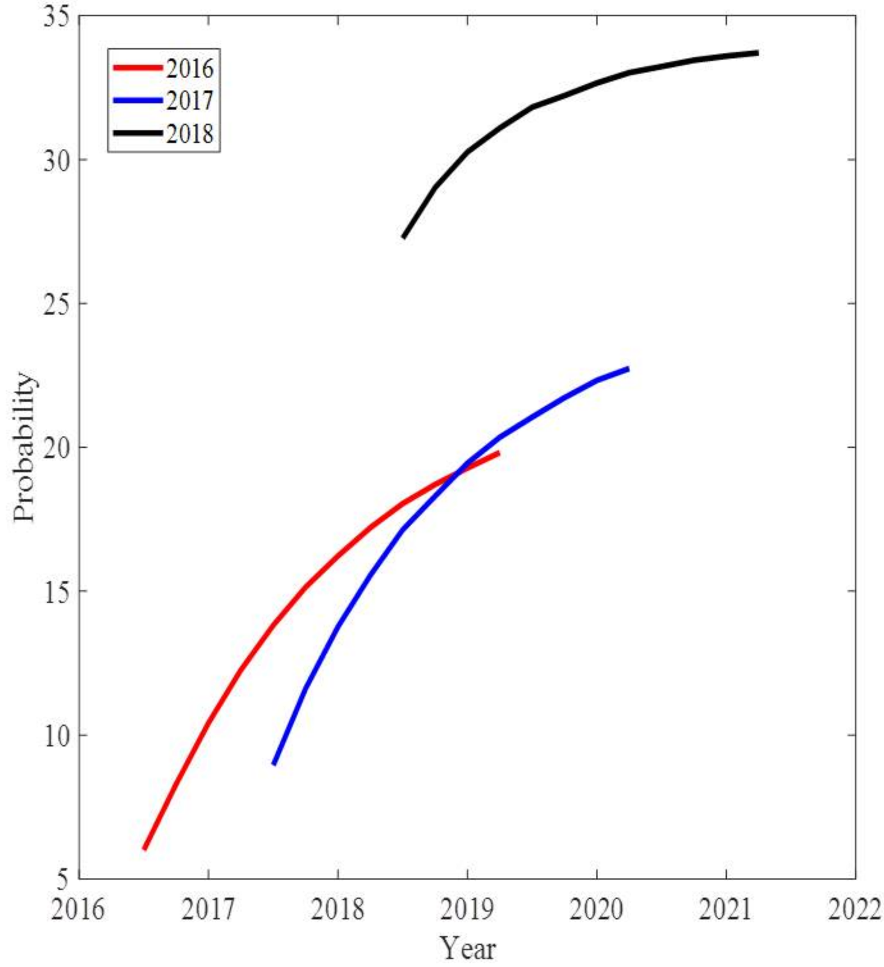

Fig. 15: Probability of HICP excl. food and energy of reaching 1.6% three years ahead using data ending in 2016Q1 (red), 2017Q1 (blue), and 2018Q1 (black).

tainty within and between models. We show the progressive upward shift in the expected probability distribution for HICPx and that the risk of deflation is now negligible.

Our regressions focus on core inflation, which is easier to forecast due to the absence of the more volatile components, such as unprocessed food and energy. However, the ECB defines its target in terms of headline (or HICP) inflation. In order to use our results as a guideline for policy makers, we need to derive a relationship between headline and core inflation. In the period 1998-2008, headline inflation was on average just below 2%, while HICPx inflation averaged 1.7%. Therefore, assuming a constant difference between headline and core inflation, the objective of sustainable inflation would be reached if headline inflation is at 1.9%, or HICP excluding energy and unprocessed food at 1.6%.

Figure 15 compares the probability of HICPx reaching at least 1.6 percent, the long-term average of core inflation using the three samples of data. The red line, for example, is constructed with data available until 2016Q1 and shows that the probability of core inflation to converge to its long-term average was very low in the short run, and increased towards 20% three years ahead. The black line repeats the exercise with data available up to 2018Q1, and shows that the probability is still relatively low, but it has considerably improved over time and it is close to 35% three years ahead.

6 Conclusions

In this paper we reassess the Phillips curve after the crisis, emphasizing robust estimation and density forecasts for policy purposes.

First, we confirm, using simple models and testing different cycle indicators, the existence of a Phillips curve in the euro area after the recent financial crisis.

Second, we proceed to its robust estimation. To account for model specification uncertainty, we estimate, using Dynamic Model Averaging, a battery of 630 models and we identify the main determinants of core inflation over the sample. Using inclusion probabilities, we confirm that the main drivers of core inflation change between the first and the second dip of the recession. The first dip is characterized by a stronger role of external variables, while the second by domestic factors. Another robust finding is that expectations are the single most important determinant of core inflation in the sample. Moreover, we estimate the contributions of the different groups of variables in our battery and we show how the negative effects of inflation expectations have considerably receded after the implementation of unconventional monetary policy, in particular following APP rounds.

Third, we estimate the slope of the Phillips curve and show that is relatively small. Furthermore, we extend our suite of models to 1260 to test for non-linearities and we do not find any signs of non-linearity in the price Phillips curve.

Finally, we argue that the convergence to the inflation objective should be assessed in a probabilistic context. We use the 630 estimated models to forecast HICPx three-year ahead and we project the probability distribution for different samples ending in 2016Q1, 2017Q1, and 2018Q1. At each point in time the distribution accounts for shocks, parameter and model uncertainty. We find an increasing, although still moderate, probability of core inflation reaching its long-term average, compatible with headline inflation reaching the objective.

Overall, we conclude that the Phillips Curve is still a valid policy instrument once it is robustly estimated.

References

Atkenson, Andrew and Lee E. Ohanian, 2001. "Are Phillips curves useful for forecasting inflation?", Federal Reserve Bank of Minneapolis Quarterly Review 2511.

Auer, Raphael, Claudio Borio and Andrew Filardo, 2017. "The globalisation of inflation: the growing importance of global value chains," BIS Working Papers No. 602.

Ball, Laurence and Sandeep Mazumder, 2011. "Inflation Dynamics and the Great Recession", emphBrookings Papers on Economic Activity, Spring.

Banbura, M. and H Mirza, 2013. "Forecasting euro area inflation with the Phillips curve", —manuscript, European Central Bank.

Blanchard, Olivier, 2016. "The Phillips Curve: Back to the '60s?", *American Economic Review, Papers and Proceedings*, Vol. 106(5): 3134

Berson, Clémence, Louis de Charsonville, Pavel Diev, Violaine Faubert, Laurent Ferrara, Sophie Guilloux-Nefussi, Yannick Kalantzis, Antoine Lalliard, Julien Matheron, and Matteo Mogliani, 2018. "Does the Phillips curve still exit?", Rue de la Banque no. 56, Banque de France.

Bhatnagar, Sanjana, Anne-Katherine Cormier, Kristina Hess, Patrisha de Leon-Manlagnit, Elise Martin, Vikram Rai, Renaud St-Cyr and Subrata Sarker, 2017. "Low Inflation in Advanced Economies: Facts and Drivers", Bank of Canada Staff Analytical Note N. 2017-16.

Blake, Andrew and Haroon Mumtaz, 2017. "Applied Bayesian econometrics for central bankers", Center for Central Banking Studies Technical Book.

Bobeica, Elena and Marek Jarociński, 2019. "Missing Disinflation and Missing Inflation: A VAR Perspective," *International Journal of Central Banking*, vol. 15(1): 199-232. .

Bowles, Carlos, Roberta Friz, Veronique Genre, Geoff Kenny, Aidan Meyler, Tuomas Rautanen, 2007. "The ECB survey of professional forecasters (SPF) A review after eight years experience", *Occasional Paper Series* 59, European Central Bank.

Casarin, Roberto, Stefano Grassi, Francesco Ravazzolo and Herman K. van Dijk, 2016. "Dynamic Predictive Density Combinations for Large Data Sets in Economics and Finance", *Tinbergen Institute Discussion Papers* 15-084/III, Tinbergen Institute, 2017.

Ciccarelli, Matteo and Benôit Mojon, 2010. "Global Inflation", *Review of Economics and Statistics*, Vol. 92(3): 524-535.

Ciccarelli, Matteo and Chiara Osbat (eds), 2017. "Low inflation in the euro area: Causes and consequences", ECB Occasional Working Paper No. 181.

Coibion, Olivier and Yuriy Gorodnichenko, 2015. "Is the Phillips Curve Alive and Well after All? Inflation Expectations and the Missing Disinflation". *American Economic Journal: Macroeconomics*, 7 (1): 197-232.

Dotsey, Michael, Shigeru Fujita, and Tom Stark, 2017. "Do Phillips Curves Conditionally Help to Forecast Inflation?", Federal Reserve Bank of Philadelphia, Working Paper No. 17-26.

Draghi, Mario, 2014a. "Unemployment in the euro area", Speech delivered at the Annual central bank symposium in Jackson Hole, 22 August.

Draghi, Mario, 2014b. "Monetary policy in the euro area", Opening keynote speech delivered at the Frankfurt European Banking Congress Frankfurt am Main, 21 November.

Draghi, Mario, 2018. "Monetary Policy in the Euro Area", Speech delivered at The ECB and Its Watchers XIX Conference organized by the Institute for Monetary and Financial Stability, Frankfurt, 14 March.

European Central Bank, 2014. "The Phillips curve relationship in the euro area", *ECB Monthly Bulletin*, July.

European Central Bank, 2017. *Economic Bulletin*, Issue 3.

Ferroni, Filippo and Benôit Mojon, 2016. "Domestic and Global Inflation", Mimeo.

Giannone, Domenico, Michele Lenza, Daphne Momferatou, and Luca Onorante, 2014. "Short-term inflation projections: A Bayesian vector autoregressive approach," *International Journal of Forecasting*, Elsevier, vol. 30(3), pages 635-644.

Jarociński, Marek and Lenza, 2016. "An inflation-predicting measure of the output gap in the euro area," ECB Working Papers Series NO. 1966.

Jordá, Óscar, Chitra Marti, Fernanda Nechio, and Eric Tallman, 2019. "Inflation: Stress-Testing the Phillips Curve", FRBSF Economic Letter, 2019-05.

Koop, Gary and Dimitris Korobilis, 2012. "Forecasting Inflation using Dynamic Model Averaging,"*International Economic Review*, Vol. 53(3): 867-886.

Koop, Gary and Luca Onorante, 2012. "Estimating Phillips curves in turbulent times using the ECB's Survey of Professional Forecasters", ECB Working Paper Series No. 1422.

Locarno, Alberto, Davide Delle Monache, Fabio Busetti, and Andrea Gerali, 2017. "Trust, but verify. De-anchoring of inflation expectations under learning and heterogeneity", *Working Paper Series* 1994, European Central Bank.

Mikolajun, Irena and David Lodge, 2016. "Advanced economy inflation: the role of global factors", ECB Working Paper Series No. 1948.

Pesaran, M. Hashem and Allan Timmermann, 1995. "Predictability of Stock Returns: Robustness and Economic Significance", *Journal of Finance*, Vol 50(4): 1201-1228.

Raftery, Adrian E., Miroslav Kárný, and Pavel Ettler, 2010. "Online Prediction Under Model Uncertainty via Dynamic Model Averaging: Application to a Cold Rolling Mill", *Technometrics* 52(1):52-66.

Stock, James and Mark Watson, 2007. “Why has U.S. inflation become harder to forecast?”, *Journal of Money, Credit and Banking*, 39(1):3-34.

Stock, James and Mark Watson, 2008. “Phillips curve inflation forecasts,” NBER Working Paper no. 14322.

Yellen, Janet L., 2015. “Inflation dynamics and monetary policy,” Speech delivered at the Philip Gamble Memorial Lecture, University of Massachusetts, Amherst, Amherst, Massachusetts.

Waggoner, Daniel F. and Tao Zha, 1999. “Conditional Forecasts In Dynamic Multivariate Models”, *The Review of Economics and Statistics*, MIT Press, vol. 81(4): 639-651.

Acknowledgements

For helpful comments and suggestions we would like to thank Philip Lane, Elena Bobeica, Mario Porqueddu, Tobias Linzert, an anonymous referee and participants to the presentations at the ECB, the Central Bank of Ireland, the International Conference on Computational and Financial Econometrics (London), the Workshop on Economic Forecasting (Vienna), the Irish Economic Association annual conference (Dublin), the Annual International Conference on Macroeconomic Analysis and International Finance (Crete), and the International Association for Applied Econometrics Annual Conference (Montr´eal). The views expressed in this paper represent those of the authors only and not necessarily those of the European Central Bank, the Central Bank of Ireland or the Eurosystem. All errors are our own.

Laura Moretti
European Central Bank, Frankfurt am Main, Germany; email: laura.moretti@ecb.europa.eu

Luca Onorante
European Central Bank, Frankfurt am Main, Germany; email: luca.onorante@ecb.europa.eu

Shayan Zakipour Saber
Central Bank of Ireland, Dublin, Ireland; email: shayan.zakipour.saber@centralbank.ie

Postal address 60640 Frankfurt am Main, Germany
Telephone +49 69 1344 0
Website www.ecb.europa.eu

This paper can be downloaded without charge from www.ecb.europa.eu, from the Social Science Research Network electronic library or from RePEc: Research Papers in Economics. Information on all of the papers published in the ECB Working Paper Series can be found on the ECB’s website.

PDF ISBN 978-92-899-3557-9 ISSN 1725-2806 doi:10.2866/221205 QB-AR-19-076-EN-N