

Phillips curves

Preliminary version - do not quote

Laura Moretti², Luca Onorante ^{*1}, and Shayan Zakipour-Saber²

¹European Central Bank

²Central Bank of Ireland

Abstract

We perform a robust estimation of the European Phillips curve, taking into account the different specifications existing in the literature, the uncertainty in the measurement of variables and the potential non linearities and structural changes deriving from the Great Recession. We extend the existing literature by adding model specifications, accounting for different measures of inflation expectations and external factors, and choosing different specifications on the basis of their out-of-sample forecasting performance. Additionally, we attempt at identifying the most important determinants of inflation over sample.

Keywords: Phillips curves, non linearities, structural changes, density forecast

JEL Codes: RENEW C30, E52, F41, E32.

*Corresponding author: E-mail: Luca.Onorante@ecb.int. Any views expressed in this paper represent those of the authors only and not necessarily of the Central Bank of Ireland, the European Central Bank or the Eurosystem.

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. Inflation has proven difficult to forecast and to explain ex post. It remains subdued, despite the strong positive contribution of the unconventional monetary policies implemented.

The recent signs of recovery in the euro area economy have raised the question of when a prudent policymaker can be sufficiently sure about the reversion of inflation to the target (and away from the danger of deflation) to avoid a premature normalization in monetary policy. This requires not only a forecast, but also an assessment of the uncertainty surrounding it, including the uncertainty about the appropriate model to use.

In this paper we focus on the robust estimation of the Phillips curve and we use our results to calculate distributions of inflation when the economy reverts to full potential. We proceed in several steps.

First, we show with a simple model (in the spirit of GLOM, XXXX) that a Phillips curve model can account for the dynamic of inflation in the aftermath of the crisis in the euro area. Second, we proceed to a robust estimation of the curve, taking into account the different specifications existing in the literature, the uncertainty in the measurement of variables and the potential non linearities and structural changes deriving from the Great Recession. This part extends the existing literature Bobeica & Jaroncinski (2016) (LIFT) by adding model specifications, accounting for different measures of inflation expectations and external factors, and choosing different specifications on the basis of their out-of-sample forecasting performance (See Lenza Jaroc). The robustness to models specification is particularly important when using the model for policy decisions because the conclusions might depend, to some extent, on the model chosen. This issue with few exceptions [see Norway, ravazzolo] is not properly accounted for in the (policy) literature. Third, we attempt at identifying

the most important determinants of inflation over sample. Finally, we explore the consequences of our findings for monetary policy.

We conduct our analysis using Dynamic Model Averaging (or DMA, see Raftery, Karny and Ettler, 2010) and we test the role of each group of variables over time to evaluate the components that dominate the determination of inflation over time (see Koop and Korobilis, 2012 for an application to US data).

Our main findings can be summarized as follows. First,

The remainder of the paper is structured as follows. We first show using a simple model (in the spirit of GLOM) that the Phillips curve relationship holds in the euro area post crisis. We overall conclude that there is scope for action of monetary policy, but the objective of inflation close but below 2 percent will also crucially depend on the (largely exogenous developments of food and energy prices, and most crucially on expectations.

2 Is there a Phillips curve in the Euro Area?

The Philips 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. In fact, its existence and its possible use as a practical policy tool remains a important topic of discussion. During the stable inflation environment of the so-called 'Great Moderation' many economists started to consider the curve as 'dead', and a vast literature analyses whether this was more due to luck, i.e the absence of major economic shocks, or good macroeconomic policies, in particular monetary. Moreover, the break-down of the relationship made its use less reliable for forecasting.

However, during the Great Recession Giannone et al. (2014) document the re-emergence of the Phillips curve relationship in the euro area. Moreover, Stock and Watson (2008) suggest, as a possible reason, 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 Jarocinski 2016, Wauters (2016), Blanchard et al (2015), where the decline in unemployment from the peak of 12% to under 10% has not been accompanied by corresponding increase in inflation, which remains subdued, especially core inflation.

The existence of the Phillips curve has important policy implications in the current juncture. The absence of a systematic relation between output or unemployment and inflation would imply that demand-side policies are not very effective on prices. If the Phillips curve holds, demand policies and the ECB monetary policy in particular are likely to be successful, and additionally the closure of the economic gap should naturally push up inflation towards the ECB's target.

While the assessment of the presence of a Phillips curve relation can only be done ex-post (see Bobeica and Jarocinski, 2017), the Phillips curve is useful also for forecasting inflation (Stock and Watson). However, the principle of prudence in taking policy decisions, especially when coming out of a deep recession, requires taking into account the presence of model uncertainty. In fact, the conclusion on the takeoff of inflation could depend, to some extent, on the model chosen.

In this section we present an exercise supporting the existence of a Phillips curve in the euro area after the crisis. This result is important because it contradicts the idea that the Phillips curve has (again) flattened.

Our approach aims at being simple and robust to the choice of the real activity variable. We estimate a number of BVARs, each containing core inflation and one measure of real activity, over the pre-crisis sample 1999q1-2007q4. We test output gap measures (from OECD, EC, WEO) and unemployment measures (rate and gap). We then perform conditional forecasts over the period 2008Q1-2017Q1 using Wag-

goner & Zha (1999) and Blake & Mumtaz (2012) to obtain a distribution of the forecasts. We focus our attention on core inflation in order to abstract from the very strong fluctuations in the food and energy prices, which would require a high number of controls as shown in Bobeica and Jarocinski (BJ, 2016).

Our results indicate that the Phillips curve has been important at least since the great recession. Chart 1 displays the actual path of core inflation (dashed black line), and the median conditional forecasts of core inflation given the different real activity measures (colored solid lines). It is worth noticing the difference in the predictions depending on the measure of real activity used. Although in the immediate aftermath of the crisis all measures tend to over predict inflation. However, looking at recent times, at least three of the measures (i.e. output gap produced by the European Commission, the unemployment rate, and the unemployment gap) actually follow the actual path quite closely, pointing towards the existence of the Phillips curve relationship.

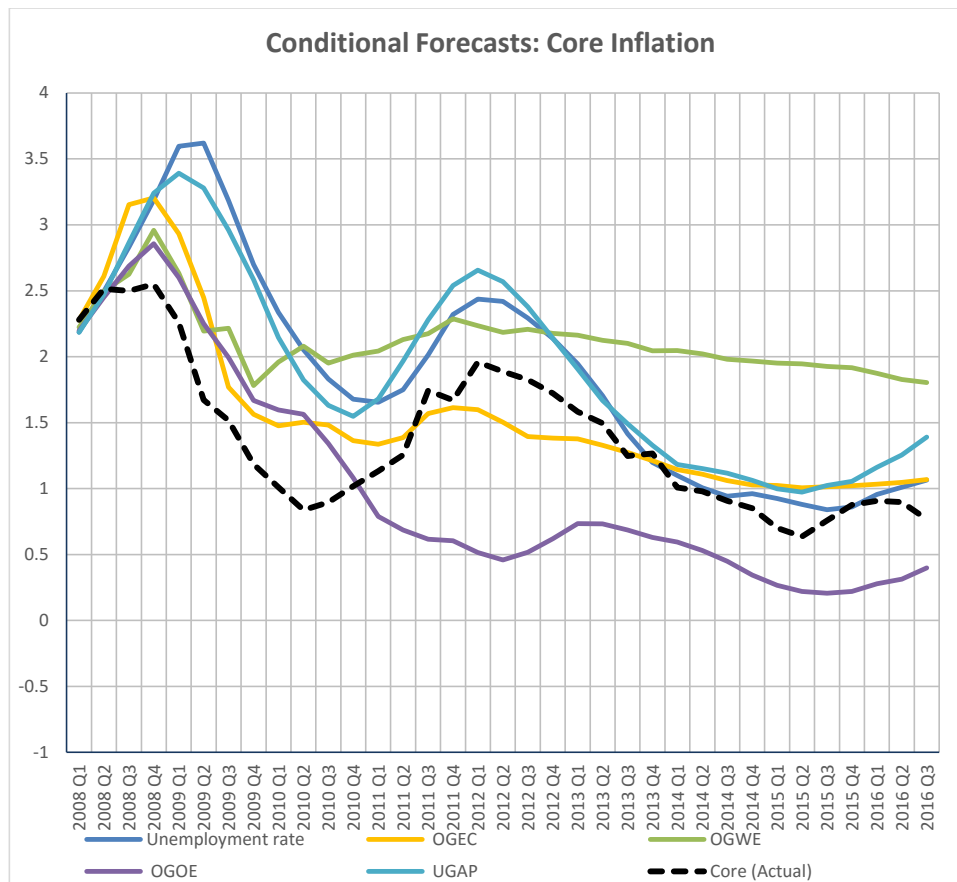


Fig. 1: Conditional forecasts for several indicators of economic conditions - mode.

Chart refFig:Phil2 compares the forecasts using unemployment rate and output gap from the European Commission, i.e. the two measures that produce the lowest root mean square of the median forecasts, including the one-standard deviation intervals. In particular, the unemployment rate is a valid indicator in real-time, as opposed to gap measures, which are regularly revised, and does not suffer from mis-measurement issues. Moreover, despite the small uncertainty bands, from 2013 core inflation always lies within the one-standard deviation bands. Put simply, the Phillips curve works well in real time when adopting a specification that uses unemployment as a measure of the cycle.




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

We reach the following conclusions. First, even a very simple Phillips curve can partially explain the movements in the immediate aftermath of the crisis, as suggested by GLOM for headline inflation. Our models cannot capture the whole inflation dynamics mainly due to the extreme nature of the event and because they do not include the effects of global variables, which are key to understanding the first euro area disinflation (BJ). Second, weak domestic real activity is key to understanding the second disinflation, therefore the same simple model explains well the movements of core inflation after the sovereign crisis. Finally, the choice of the economic activity indicator is important, as suggested in a recent contribution by Lenza and Jarocinski (2016).

Following these conclusions, we aim at pinning down a robust estimation of the curve.

3 Econometric framework

The curse of dimensionality and the uncertainty about the regressors to be used make the identification and the use of a single model impossible. We resort therefore to Bayesian techniques of model averaging (BMA). Advantages of BMA include the possibility of using parsimonious models, which yield more stable estimates, because fewer degrees of freedom are used in individual models. Also, BMA can identify important regressors, making the results more informative and easier to interpret. It accounts for model uncertainty, and can be used as a tool to select the best indicator to measure a concept, for example choosing between different measures of slack in a Phillips curve. Finally, it can also be used to account for uncertainty about model structure beyond variable selection (e.g. linear versus nonlinear models, univariate versus multivariate models, fixed versus time-varying parameters). We use Dynamic Model Averaging (DMA), which additionally allows models to change over time, thus also dealing with structural changes. DMA was first proposed by Raftery et al.(2010)

and allows the weights used in the model averaging to change overtime, allowing to uncover the role and importance of different regressors over time.

3.1 Model uncertainty: Dynamic Model Averaging

DMA is developed in Raftery, Karny and Ettler (2010) and used in Koop and Korobilis (2011) and the reader is referred to these papers for complete details. The dynamic aspect of DMA arises since it allows for a different model to hold at each different time period. We assume a population p_k of K models

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

where $y^s = (y_1, \dots, y_s)'$ is the past information up to time s and $p_k (y_t|y^{t-1})$ 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 $q_{t|s,j} = \Pr (k = j|y^s)$ 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 $j = 1, \dots, K$. 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, that is 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:

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

If we knew $q_{t|t-1,s}$ then, starting with $q_{0|0,s}$ we could recursively calculate the key elements of DMA: $q_{t|t,j}$ and $q_{t|t-1,j}$ for $j = 1, \dots, K$. Raftery et al. (2010) provide this missing link by using the approximation:

$$q_{t|t-1,s} = \frac{q_{t-1|t-1,s}^\alpha}{\sum_{l=1}^K q_{t-1|t-1,l}^\alpha}. \quad (3.3)$$

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

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

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

The previous result clarifies that a model j will receive more weight at time t if it has fit well in the recent past (fit is measured by the predictive likelihood, $p_j(y_{t-i}|y^{t-i-1})$). 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. These considerations suggest that we focus on the interval $\alpha \in (0.95, 0.99)$.

In our short data set, the potential advantages of DMA are clear. We can include models featuring a large number of explanatory variables, but if these are overfitted their predictive density 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.¹ Furthermore, DMA allows for model change. It can capture cases where certain explanatory variables or models frameworks are im-

¹See Koop and Korobilis (2009a) for evidence that DMA can effectively find very parsimonious models.

portant in certain periods, but not in others. Given our application, which covers the time from the introduction of the euro and the outcomes the recent financial crisis, allowing for such chance may be important.

We note also that, in the past, DMA has been done in the context of time-varying parameter (TVP) models where the coefficients evolve as $\beta_{it} = \beta_{i,t-1} + u_t$. Our set of models include time varying parameters as well as fixed parameter models. It turns out that the time varying models usually under perform out of sample (probably due to overfitting) and have low weights in the final DMAs. This result is consistent with Koop and Korobilis (2011), who found that allowing for models to switch over time is of greater empirical benefit than allowing for coefficients to evolve in a TVP fashion.

3.2 Model weights

In this section, we previous the weight of the different models over our sample. Figure 3 shows their evolution with some specifications dominating depending on the period. AS it is customary in DMA, a relatively reduced set of models dominates at each point in time.

4 Model Specification and Data

In the literature, there are different specifications of the Phillips Curve, including different variables and using different functional forms. An all-encompassing model would not be possible to estimate, given the high number of potential variables and the estimates would be meaningless anyway, as the presence of several combination of variables (e.g. the presence of three different cycle indicators) would lead to multicollinearity. Additionally, we wish to consider all the Phillips curves of interest to account for model uncertainty. To solve these issues, we adopt a robust approach and we include a large number of model formulations.




Fig. 3: Evolution of model weights over time.

Following the recommendation of Stock and Watson to use parsimonious models, we focus on univariate specifications. Some of our specifications will include in any case a substantial number of estimated parameters, and we will compare their out-of-sample explanatory power with the one of smaller models.

The dependent variable is core inflation, defined as the year-on-year percentage change in the HICP excluding energy and unprocessed food. We focus our attention on core inflation to abstract from more volatile components, but we repeat the analysis also using HICP as dependent variable and we discuss the different results.

We consider a high number of regressors and divide them into four groups: real activity, inflation expectations, labour market indicators, and global indicators (see Figure 4). Each specification of the Phillips Curve includes always one real activity

variable and a permutation of at most one variables from each group. The real activity group includes three measures of output gap (from OECD, the European Commission, and the WEO), the unemployment rate and the unemployment gap, i.e. the difference between unemployment rate and the NAIRU. Labor market indicators include compensation per hour and total unit labor cost.

Contrary to the previous literature, we consider both survey and market-based indicators. We include three survey measures, i.e. Survey Professional Forecasters (SPF) one year ahead, two years ahead and five years ahead, and the 5 year-in-5 year inflation swap rate as the market-based measure. In particular, we select the five-year in five-year forward inflation swap rate that is the benchmark measure of medium to long-term inflation expectations for central banks (see for example Draghi, 2014).

Finally, building on the literature that recognizes the role of international factors in determining inflation (Ciccarelli and Mojon 2010, Ferroni and Mojon 2016, Bobeica and Ciccarelli), we include external factors such as the real effective exchange rate (REER), one year ahead oil futures, the world industrial production, and the import price deflator. In fact, previous literature has highlighted the importance of such global factors in the determination of inflation (see Ciccarelli and Mojon, 2010, Ferroni and Mojon, 2016, and Bobeica and Jarocinski, 2016). Figure 4 presents a summary of the variables used. Our dataset is quarterly and spans the period from 2001Q1 to 2017Q2.

5 Empirical results

5.1 Importance of different variables over time

We use the estimated weights in the forecasting using the DMA to test the role of each group of variables. Figure 5 presents the relative importance over time of inflation

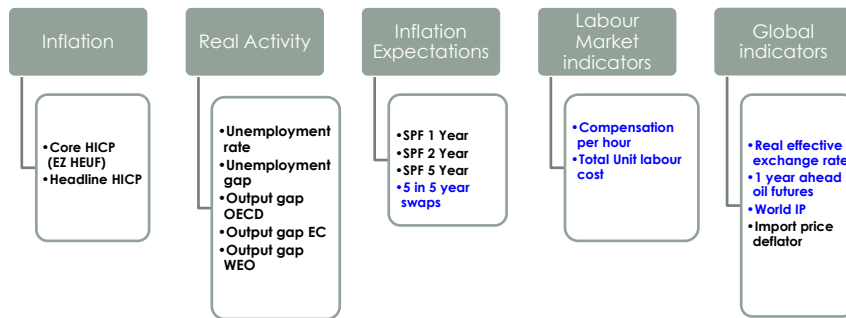


Fig. 4: Variables included in the model.

expectations (in yellow), external factors (in blu), and labor market indicators (in red).

It is important to notice that, as soon as the equal prior weights are updated, external variables become the most important determinants of core inflation and remain relevant up to the second dip in the euro area recession (i.e. the beginning of the sovereign crisis), when their importance drops. Afterwards, external variables regain their importance up to 2014 when their relevance starts to wane. Wage and labor market dynamics are important before the crisis and their role is increasing towards the end of the sample. Both the 2008 downturn and the sovereign crisis, with output away from potential, drastically reduced the role of labor markets as determinants of inflation; this evidence is compatible with a non linear Phillips curve, where cost-push

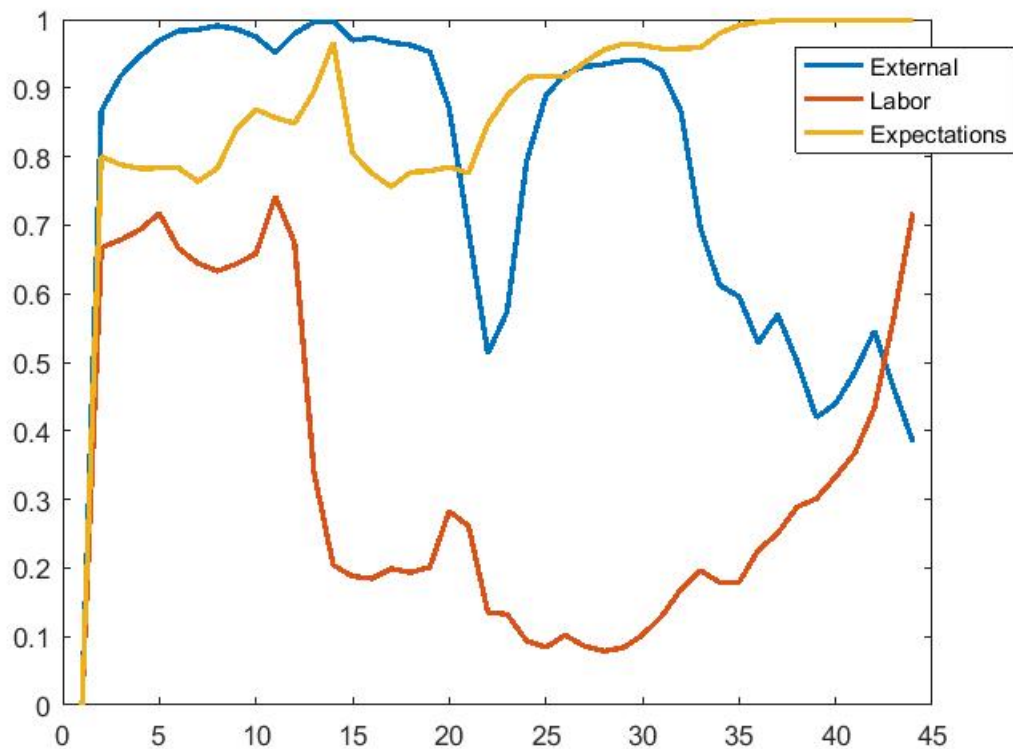


Fig. 5: Inclusion probability of different sets.

dynamics do not matter if the economy is far away from full capacity utilization. [TO ADD This evidence is also consistent with Bobeica and Jarocinski]. Finally, expectations are an important driver of inflation, second to external factors in the initial period of the sample, but becoming more and more relevant over time.

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

5.2 Cycle indicators and real-time use of the Phillips curve

A well-know issue in using the Phillips curve is that most measures of slack are based on output, a quantity which is available with considerable lags and it is subject to substantial revision. In fact, 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 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.

We compare cycle indicators by imposing that one of them must be always 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, for example, unemployment. We find that different measures are better predictor in subsamples. Figure 6 presents the results, where the solid line is the inclusion probability for each measure of the cycle. The output gap estimated by the OECD has a higher performance in the “great moderation” times that characterized the first years of the EMU, but becomes severely misleading (and its importance drops to practically zero) during the crisis, to recover progressively thereafter. The alternative gap provided by the European Commission dominates the second subsample because it estimates a sizeable output gap, going in the direction of the 6 per cent gap indicated in Lenza and Jarocinski as the best predictor of inflation.

As in Lenza et Jarocinski, this performance is only valid *ex post* and does not help in real-time. Specifications including unemployment, however, are different in nature. We already mentioned in our preliminary evidence the potential advantages of this measure of slack. Unemployment is slower in reacting, and does not compare to the best *ex-post* output gap indicator in tranquil times (although it performs similarly to most others). However, unemployment becomes a better indicator when it is most needed: during and post the crisis. This is hardly surprising. Unless the economy is in tranquil times that allow the gaps to be estimated with minor revisions, unemploy-

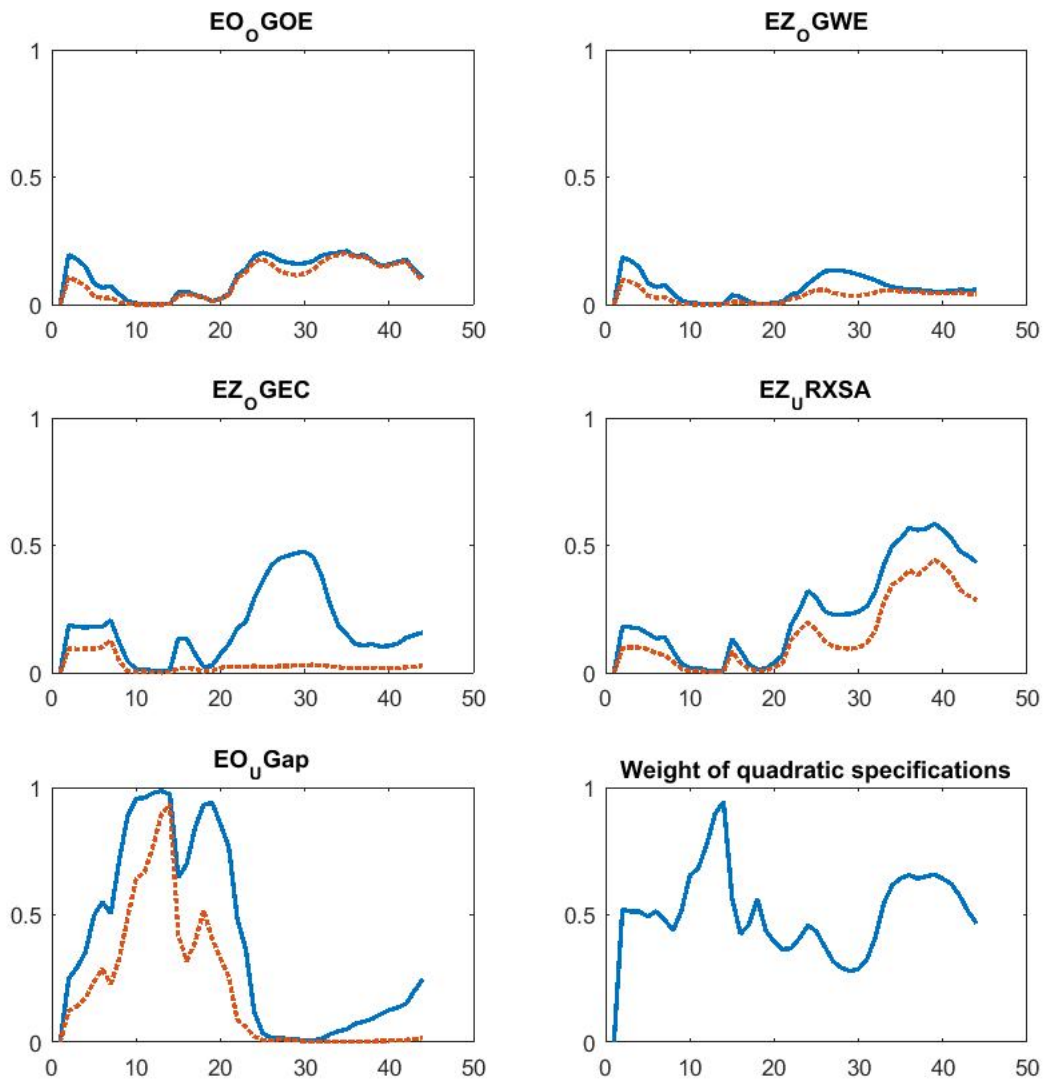


Fig. 6: Inclusion probability of different cycle variables

ment is a better indicator because it is immediately observed and hardly revised. This result is not new for the Euro Area (***) cit. KO)

5.3 We do not need a nonlinear Phillips curve

In Figure 6 we present also the inclusion probability of the squared gap measures (dotted lines). We perform several tests of non linearities in the curve. First, we add to our set of 300 the same models, including squared terms of the cycle (quadratic specifications). If the Phillips curve is non linear in the gap, we expect those to have a greater predictive likelihood than the corresponding linear specifications. The last panel shows that this is not the case: the overall weight of quadratic specifications remains at about 50 percent, indicating that the quadratic terms do improve our models in any way. This result is confirmed by the fact that the coefficients of the quadratic terms are non significant, and vary in sign.

We test the hypothesis that the coefficients change with the crisis by adding to the set of linear models the same specifications in the context of a smooth transition model, where the coefficients are allowed to change once. These models perform badly, do not identify a clear transition period (we expected them to identify at least the beginning of the crisis) and the coefficients are not statistically significant before and after the transition. Their bad out-of-sample performance skews their weights towards zero.

The overall evidence suggests that non linear or time varying version of the curve do not outperform linear specifications. In the rest of the paper we go back to the initial pool of 300 linear models.

5.4 Labour market as transmission mechanism

MORE CITATION NEEDED HERE. Stagnating labor 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 (the relative importance) of labor market-related curves (those including wages and/or unemployment). We then calculate the contribution of these two variables to overall

inflation and its evolution over time. Figure 7 shows the inclusion probabilities for the unit labor cost and the compensation per employee. The results show that their importance declined rapidly after the initial years of the euro, but it increased after the crisis. In particular, the inclusion probability of the compensation per employee reached nearly 70 percent.

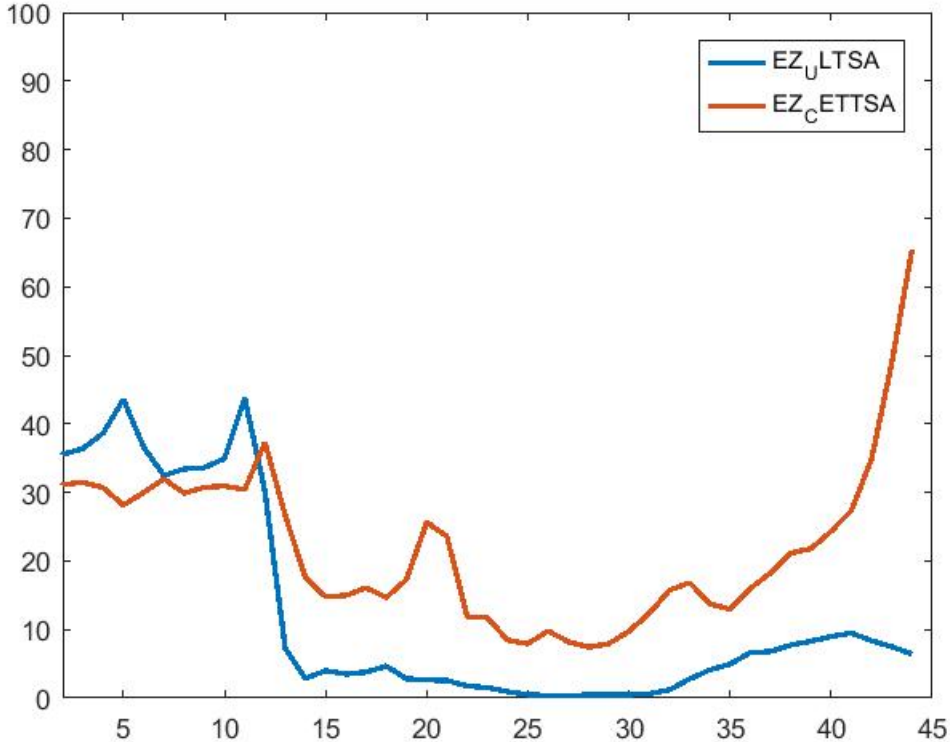


Fig. 7: Inclusion probability of different labour variables

5.5 Are expectations deanchored?

The de-anchoring of inflation expectations moves the Philips curve and can constitute a serious concern for the monetary authority. Lower expectations permanently lower inflation, other factors being equal. It is then not surprising that in the current junc-

ture a considerable attention has been given to this subject, both in theoretical models (Gerali) and in empirical analysis. To measure inflation expectations it is possible to use both market-based and survey measures. They all presents some advantages and some drawbacks. Market-based expectations are considered an important indicator of the credibility of monetary policy (Draghi, 2014). 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 do not suffer from these bias, but they tend to react slowly (Kenny et al). Moreover, both represent a subset of the economic agents. We first compare the different measures on the basis of predictive power. Figure 13 shows the inclusion probability of each of the variables. Conditional on being present in the model, the contribution of each variable is depicted in figure X.

5.6 Contributions

Finally, in Figure 9 we present the contribution of different groups of variables to the dynamics of core inflation. The first picture is similar to many that are based on a single model and/or in sample decomposition, and is produced by averaging models with equal weights. It may appear that expectations have become progressively disanchored, despite ECB unconventional monetary policy interventions, and their contribution is currently negative. This might be the reason why we do not see inflation. However, equal weights are far from optimal.

However, when using optimal model weights, the resulting decomposition tells us a very different story, as shown in Figure 10. First, observe that before the crisis, inflation could have been higher given wages and positive cycle (and as of 2006 external), but anchored expectations were at work and contributing negatively, helping to keep inflation close to the target as advocated by ECB. Then, after the 2008 crisis, after an initial sudden decrease, expectations contributed positively as ECB introduced APP

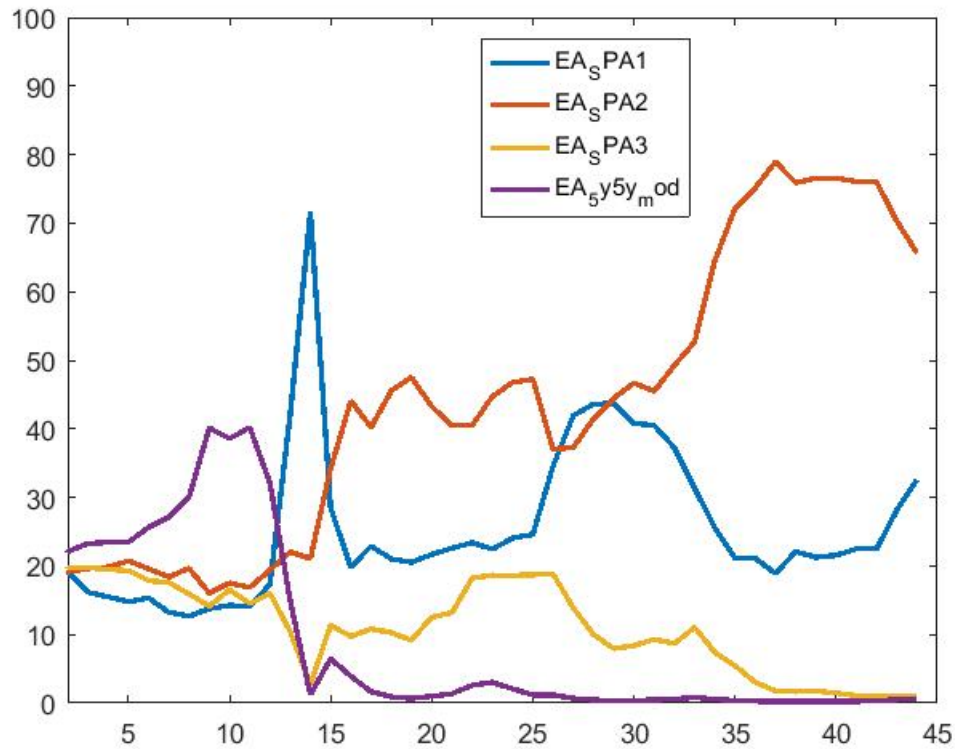


Fig. 8: Inclusion probability of different measures of inflation expectations

and Forward Guidance. In other words, the chart tells us that the ECB policies have been highly successful in the years following the Great Recession and expectations (yellow area) have contributed to maintain inflation above zero. At the same time the effect on expectations of APP and Forward Guidance has been diminishing in time, and if inflation has to increase further, higher contributions should be expected from labor market wages and closure of the gap.

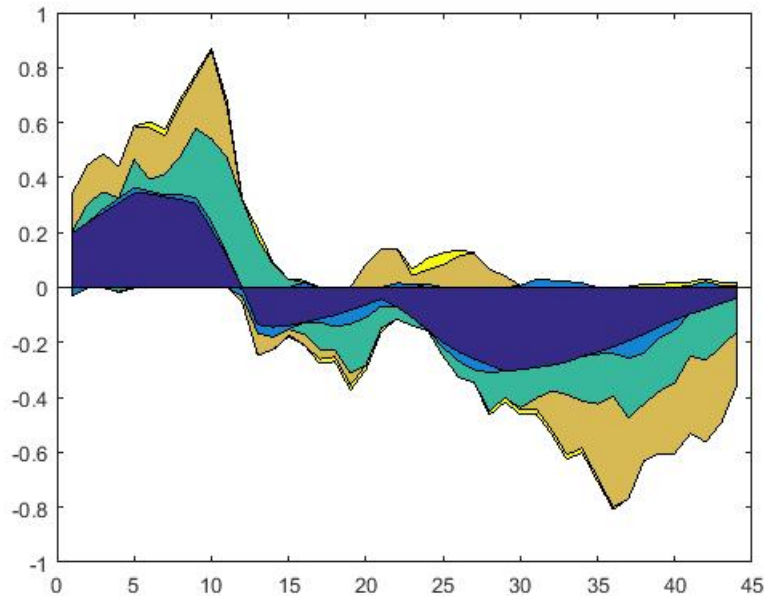


Fig. 9: Dynamic decomposition - CONSTANT WEIGHTS - blue = cycle, light blue = external, gree = labour, brown = expect, yellow = qe

5.7 Estimation and Forecast

As discussed in the Introduction, despite the ongoing economic recovery, headline and core inflation rates in advanced economies and in the euro area in particular remain low. Headline inflation has been affected downwards by the large decline in oil prices, but the same phenomenon has been observed for core or underlying rates of inflation. We have argued that this does not imply that the Phillips curve should not be considered a valid tool.

In this section, we present and discuss the results of the robust forecasting exercise using the battery of Phillips curve models and we discuss the distribution of future inflation and the probability of reaching the ECB target of inflation of below but close to 2 percent. Then, we assess the future developments in core inflation in a

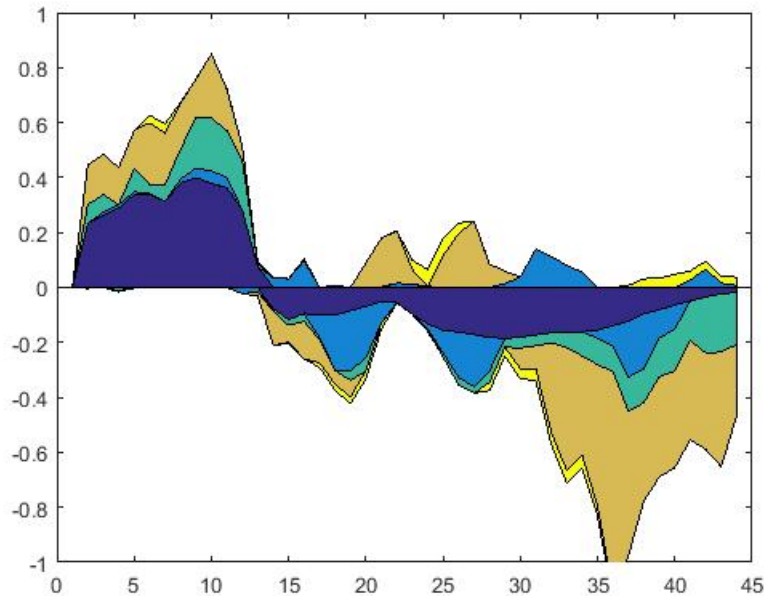


Fig. 10: Dynamic decomposition - DMA DYNAMIC WEIGHTS

probabilistic fashion. In fact, the return of inflation to the ECB target should be evaluated in terms of probabilities.

Figure 13 shows the 12-quarter ahead forecasts of HICP inflation excluding energy and unprocessed food using the battery of 300 models. We report the unconditional forecast of all the models and we highlight the simple average (the dotted line) and the DMA using the optimal weights (continuous line). We use the weights of the last period, and all forecasts are conditional to simple AR1 processes for the exogenous variables. We notice that the robust estimations present moderate differences depending on the specification and on the indicators adopted. The forecasts using our specifications show that inflation will remain subdued, below the ECB target of below but close to 2 percent. In particular, the DMA is lower than the unweighted average, and points to a very contained inflation in the medium term.




Fig. 11: Battery of models - all forecasts (gray), mean (dots), BMA (continuous line)

Figure ?? shows our results taking into account the uncertainty within and between models. The black line indicates the (combined) probability density function for inflation three years ahead (i.e. in the second quarter of 2019). The combination involves densities from all 300 different models via Bayesian Model Averaging weights. The shadowed area indicates the probability of core inflation to be above 1.6 percent. The results show that, even taking into account within and between model uncertainty, deflation is now a negligible outcome. However, the probability mass to the right of 1.6 percent (i.e. the probability that core inflation reaches 1.6 percent) is still relatively low and assessed at 30 percent. Whether this probability is sufficient depends on the loss function of the policymaker. However, these considerations re-

quire a complete distribution of the forecast constructed over a robust set of models. [Sargent Hansen 08 highlighted this, in robust control monopol.] Probabilistic forecasts are an important tool, given that the goal of the policymaker includes keeping the risk of an adverse outcome (disinflation) to an acceptable level or to reach a goal (sustainable target) with an acceptable probability.

We analyzed here the probabilistic distribution of core inflation three-year ahead. 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 relate them to the headline inflation. In order to do that, we look at the inflation developments before the financial crisis and we derive the relationship between headline and core inflation. In the period 1998-2008, headline inflation was on average just below 2 percent, while core inflation averaged 1.71 per cent. Therefore, the objective of sustainable inflation would be reached if headline inflation is at 1.9 percent, or core at 1.6. According to our previous results, the probability of reaching the target of inflation remains quite low.

The second point requires computing the complete distribution of inflation forecast 12-quarter ahead (up to 2019Q2). 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.

Overall, we conclude that the Curve is still a valid policy instrument once it is correctly estimated and the uncertainty relative to the cycle indicators and available specifications is resolved. Highlight that it is the whole distribution that matters to assess risk. Including model uncertainty. Highlight that this is at current policies, as we only implicitly account for them, and it also includes current expectations of future policies.




Fig. 12: Battery of models - all forecasts (gray), mean (dots), BMA (continuous line)

6 Closing remarks

[to be added]

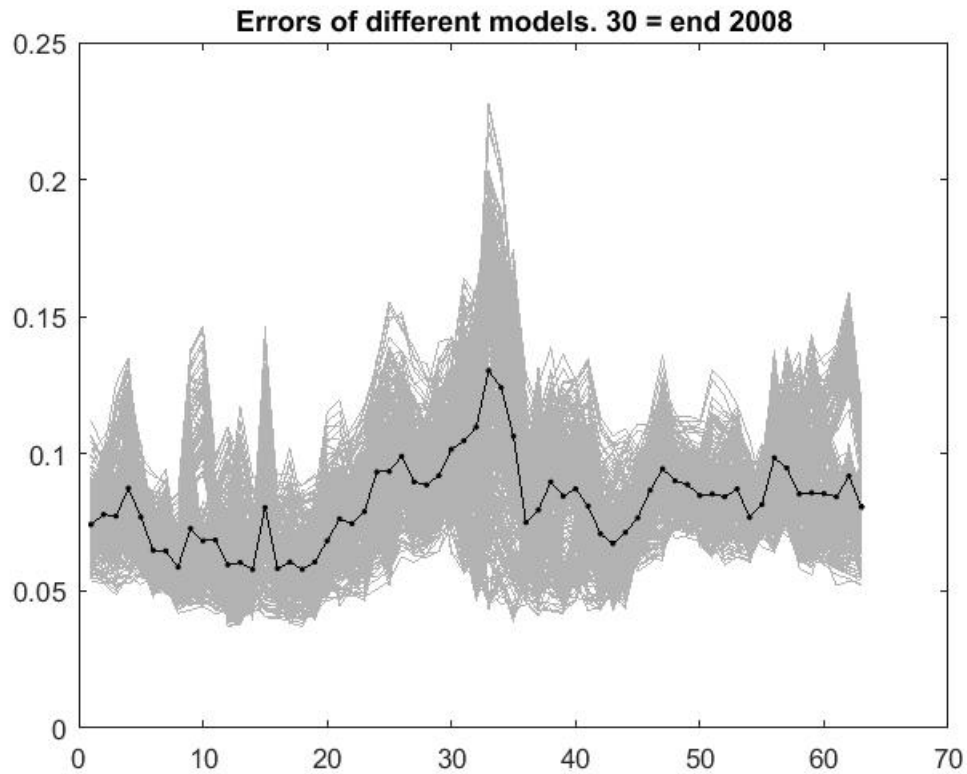


Fig. 13: Battery of models - all forecasts (gray), mean (dots), BMA (continuous line)

References

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