Navigating with a Compass: Charting the Course of Underlying Inflation*

Nuno Lourenço,[a] João Quelhas,[a,b] and António Rua[a,b]
[a]Banco de Portugal
[b]Nova SBE

We propose a novel tool to gauge price pressures resorting to circular statistics, the inflation compass. It offers a reliable indication of inflationary pressures in the euro area since its inception. Unlike most alternative measures of underlying inflation, the inflation compass does not exclude any inflation subitems. Moreover, it is not revised, providing real-time signals about the course of underlying inflation. We also provide evidence of the usefulness of the inflation compass to forecast overall inflation, even during periods of increased turbulence, such as the COVID-19 pandemic or the recent inflation surge. Lastly, our approach can handle large-dimensional data straightforwardly.

JEL Codes: C53, C55, E31, E37.

Our compass is price stability. We will stay riveted to that compass as we fight inflation to deliver stable prices for people across Europe.

–Christine Lagarde, Brussels, September 26, 2022

Underlying inflation is not a policy target, but measures of underlying inflation can serve as a complementary cross-check of our forecasting process.

–Christine Lagarde, Frankfurt am Main, March 22, 2023

*The analyses, opinions, and conclusions expressed in this article are those of the authors and do not necessarily reflect the opinions of Banco de Portugal or the Eurosystem. Any errors and omissions are the sole responsibility of the authors. Authors' e-mails: nalourenco@bportugal.pt, jquelhas@bportugal.pt, antonio.rua@bportugal.pt.

1. Introduction

Central bank mandates vary in scope, encompassing maintenance of price stability, promotion of full employment, and fostering of economic growth, among others. Nonetheless, a prevailing characteristic shared by central banks worldwide is their unwavering attention to monitoring inflation dynamics. In the recent past, this became even more evident as inflation recorded levels not seen for decades, triggering renewed discussions among inflation targeters, e.g., the U.S. Federal Reserve, the European Central Bank (ECB), the Bank of Japan, and the Bank of England, to name a few.

Central banks' growing interest in gauging the persistence of price changes in a timely manner has reinforced the development of new methods to filter incoming data on aggregate prices. Thus, several tools have been proposed in the literature to extract relevant information from granular data, leading to the construction of different underlying inflation measures aimed at discarding transitory price movements and capturing the subjacent trend in prices. By closely monitoring these tools, policymakers can anticipate shifts or rapid reversals in the overall pace of inflation, in periods of both upward and downward price pressures.

Despite their value to guide monetary policy decisions, the lack of a consensus on a preferred instrument to track price pressures still persists. However, it is widely acknowledged that these measures ought to have desirable properties, such as accuracy in tracking overall inflationary pressures, ability to be computed in real time, lack of revisions, or forecasting performance within a regression framework. Previous work has examined some of these properties (see, inter alia, Bryan and Cecchetti 1994; Bryan, Cecchetti, and Wiggins 1997; Bakhsi and Yates 1999; Dolmas 2005). See also Clark (2001), Cogley (2002), Marques, Neves, and Sarmento (2003), Rich and Steindel (2007) or Wynne (2008) for the United States, Roger (1997) for New Zealand, Hogan, Johnson, and Laflèche (2001) for Canada, and Mankikar and Paisley (2002) for the United Kingdom. More recently, central banks have shown a renewed and heightened interest in this issue, as illustrated by the resurgence of work along this line, for example, Luciani and Trezzi (2019) or Luciani (2020) for the United States and Pincheira-Brown, Selaive, and Nolazco (2019) for Latin American countries.

In the euro area, the ECB has set as its primary objective the maintenance of price stability, which entails keeping inflation at a low, stable, and predictable level around 2 percent. Monetary policy prescriptions have thus relied on the assessment of the inflation outlook, considering incoming economic and financial data, the dynamics of inflation, and the effectiveness of monetary policy transmission. More than ever before, the recent surge in euro-area inflation, from 1 percent at the start of 2021 to 10 percent by the autumn of 2022, drew attention to the need for close monitoring of underlying inflation measures. Ehrmann et al. (2018) provide a comprehensive overview of the tools typically tracked by the ECB to monitor inflation developments. These include permanent exclusion measures such as the Harmonised Index of Consumer Prices (HICP) excluding energy and unprocessed food, the HICP excluding energy and food, and the HICP excluding energy, food, travel-related items, clothing, and footwear. Additionally, there are temporary exclusion measures, including trimmed means and the weighted median, as well as frequency exclusion measures. Examples of the latter include the Supercore and the Persistent and Common Component of Inflation (PCCI).

We depart from the existing literature on underlying inflation measures by proposing a novel tool to gauge price pressures resorting to circular statistics, the so-called inflation compass. Even though the use of circular statistics is not recurrent in economics, recent work has evidenced their usefulness in inferring the business cycle stance, as shown in Lourenço and Rua (2023). Unlike previous literature on core inflation, we take a different perspective by bringing on board the information from every subitem of the HICP basket. In particular, we retain the inflation level and the change in the inflation rate corresponding to each subitem, and convert such information to angular data. By assigning the HICP weights to this granular data, we can determine an overall direction and gain insights into the underlying price pressures in the economy. That is, starting with the HICP elementary items, we extract an angle for each subitem, and subsequently determine a weighted mean direction that summarizes the different price pressures in the consumption basket.

Such an approach leads to a compass representation, allowing us to infer whether inflation is approaching the target as well as the directional change of price pressures, i.e., whether they are accelerating or decelerating. In our application, we illustrate the compass for key episodes of high and low inflation in the euro area, showing that it provides informative signals on inflationary pressures since the inception of the monetary union. Unlike most alternative measures, the compass does not exclude any subitems of inflation, ensuring that all available information is used. On top of that, it is not subject to revisions, providing policymakers with a reliable real-time assessment. As a by-product, we take advantage of the properties of the inflation compass to shed light on the periods during which underlying inflation did not deviate from 2 percent, thus providing a characterization of the full spectrum of price pressures in the euro area.

In addition to its qualitative informational content in gauging price pressures, we provide evidence of the usefulness of the inflation compass for forecasting overall inflation up to 36 months ahead. In particular, we assess the performance (both in- and out-of-sample) of several underlying inflation measures. To address this, we adopt the usual econometric specification relating future changes in inflation to the transitory component of price changes identified by the underlying inflation measure. This regression framework not only has the advantage of being simple to interpret but is also flexible enough to incorporate alternative horizons into the analysis. Quantitatively, we find that the inflation compass outperforms the underlying inflation measures typically tracked in the euro area in a forecasting framework and that the accuracy gains are, in general, statistically significant.

As a robustness analysis, we provide evidence that the inflation compass delivers higher accuracy than the other measures, even during challenging and turbulent periods like the COVID-19 pandemic and the subsequent inflation surge. Furthermore, as our approach can handle large-dimensional data straightforwardly, we augment our initial data set, which already comprises almost 100 series, by considering two variants of data disaggregation. Firstly, we consider more detailed product-level data entailing nearly 300 HICP subitems. Secondly, we consider country-level data, as opposed to aggregate euro-area data, which enables us to consider more than 1,000 HICP series. Notably, the inflation compass is robust to these data-rich environments, as the results reaffirm its superior performance.

The remainder of the paper is organized as follows. Section 2 discusses the main building blocks behind the inflation compass representation. Section 3 describes the data. In Section 4, the empirical application is conducted. Section 5 evaluates the forecasting performance of the novel approach against other underlying inflation measures commonly tracked in the euro area and Section 6 provides a robustness analysis. Finally, Section 7 concludes.

2. The Circular Statistics Approach

As mentioned above, we depart from previous literature by addressing underlying inflation measurement from a different angle using, literally, angles. To lay out the intuition behind the approach and the main concepts around it, let us start by considering a very stylized example. For illustration purposes, suppose that inflation, π_t, follows a sinusoidal model,

$$\pi_t = \mu + A \sin(w_0 t), \ t = 1, \ldots, T, \tag{1}$$

where μ captures the mean, A denotes the amplitude and w_0 controls the frequency. In Figure 1A, we display such a process by setting, as an example, $\mu = 2$, $A = 1$ and assuming a cycle duration of eight years with $T = 200$. Naturally, we get an oscillatory behavior around the mean level. Note that, in a context of inflation targeting and central bank credibility, one would expect inflation to hover around the inflation target, with departures from time to time driven by economic conditions.[1]

When tracking inflation developments, economic agents are interested in both the level of and the change in inflation. Hence, for each point in time, one can display the inflation level against the corresponding change in the Cartesian coordinate system. In particular, the y-axis denotes the inflation level and the x-axis the change in inflation. In Figure 1B, we present for a few selected points in the time series the corresponding representation in the Cartesian coordinate system.

Note that any point in the Cartesian plane defined as (x_i, y_i) can be summarized by polar coordinates, an angle, θ_i, and a distance

[1]For a discussion on inflation cycles, see, for example, Arti et al. (1995) and Banerji and Achuthan (2019).

Figure 1. Stylized Inflation Time Series and Corresponding Representation in the Cartesian Plane

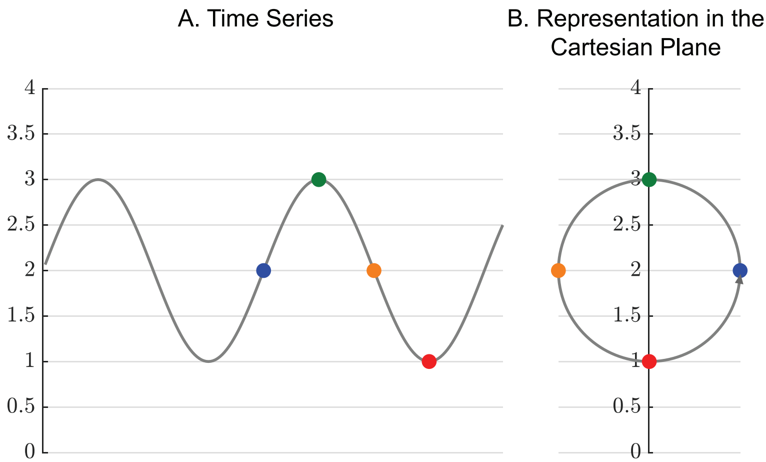

Figure 2. Angle Reading Example in the Cartesian Plane

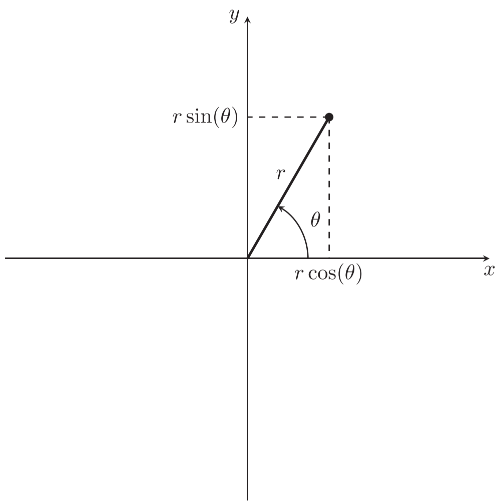

from the origin, r_i, as illustrated in Figure 2. These conversions are conducted using sine and cosine trigonometric functions, drawing on the relationships $x_i = r_i \cos\theta_i$ and $y_i = r_i \sin\theta_i$. The mathematical convention for angular measures in statistics typically uses the polar angle, measured counterclockwise from the positive x-axis (pointing east) to the line segment connecting the origin to a given point. Hence, a rotation in the opposite direction to the conventional results in a negative angle. In this case, negative angles are subtracted from 360° to convert them to the corresponding positive angles. For instance, in the case of Figure 1B, the green dot in the Cartesian plane corresponds to a 90° angle with $r = 1$.

Naturally, inflation does not have to follow the model as described in (1). In fact, given the above discussion, one can generalize and write π_t as

$$\pi_t = r_t \sin \theta_t,\ t = 1, \ldots, T, \tag{2}$$

with $\Delta\pi_t = r_t \cos \theta_t$. By merging information from the inflation level and the change in inflation, the angle θ_t is uniquely determined by the two coordinates and conveys a sense of rotation with east as the zero direction. Note that such a directional position does not depend on the distance from the origin.

Now let us assume that behind the evolution of θ_t there is an underlying unobserved direction, θ_t^u, along with an idiosyncratic error, ξ_t; that is, $\theta_t = \theta_t^u + \xi_t$. In this case, (2) becomes

$$\pi_t = r_t \sin\left(\theta_t^u + \xi_t\right),\ t = 1, \ldots, T, \tag{3}$$

with

$$\Delta\pi_t = r_t \cos\left(\theta_t^u + \xi_t\right),\ t = 1, \ldots, T. \tag{4}$$

To extract the underlying unobserved direction θ_t^u, one can, for instance, take advantage of the disaggregated inflation data. Let us consider that we have N inflation subitems. As in (2), one can write each inflation subitem i, π_{it}, as

$$\pi_{it} = r_{it} \sin \theta_{it},\ t = 1, \ldots, T \text{ and } i = 1, \ldots, N \tag{5}$$

and assume for each inflation subitem i the following:

$$\pi_{it} = r_{it} \sin\left(\theta_t^u + \xi_{it}\right),\ t = 1, \ldots, T \text{ and } i = 1, \ldots, N, \tag{6}$$

with $\theta_{it} = \theta_t^u + \xi_{it}$ and $\Delta\pi_{it} = r_{it} \cos\left(\theta_t^u + \xi_{it}\right)$.

Note that this type of formulation finds support in the way ECB officials interpret underlying inflation (see, for example, Lane 2022, 2023). Firstly, underlying inflation is seen as a latent variable that helps to look through the noise present in headline inflation data. The volatility induced by temporary idiosyncratic shocks blurs the signal about the inflationary pressures relevant for monetary policy. In fact, the assumption that inflation is driven by both signal and noise is embedded in (3).

At the same time, it is explicitly mentioned that the reaction function takes into account the dynamics of underlying inflation. This highlights the importance of monitoring the change in inflationary pressures as well. In this respect, it is recognized that the trajectory of inflationary pressures is also affected by idiosyncratic shocks. This notion is translated in (4). Hence, as both the level of and the change in inflationary pressures are relevant, the proposed approach combines the two pieces of information through a circular representation given the cyclical nature of inflation as discussed above. Such an approach enables one to extract a single indicator about underlying inflationary pressures summarized by the direction given by θ_t^u.

Finally, given the unobservable nature of underlying inflation, several measures can be devised involving different levels of statistical complexity. Such measures range from a priori exclusions of some components of headline inflation to complex statistical models that exploit the cross-sectional variability of inflation components. The suggested approach herein falls into the latter category.

Hence, at each point in time, we end up with N directions; that is, a sample of N angles, $\theta_1, \theta_2, \ldots, \theta_N$. Given this sample, one can compute the mean angle which reflects the center of mass of the distribution. The mean direction captures the underlying common direction by averaging out the idiosyncratic components. By exploiting the disaggregated inflation data, one can retrieve an underlying common direction regarding inflation developments.

In practice, although sample moments can be derived from the distribution of the angles, these require the use of methods tailored for circular data.[2] For instance, directions of $0°$ and $360°$ have the same meaning, but the sample mean computed in a linear scale

[2] An overall treatment of circular statistics can be found, for example, in Fisher (1993), Mardia and Jupp (2000), and Jammalamadaka and SenGupta (2001).

Figure 3. Inflation Compass Reading

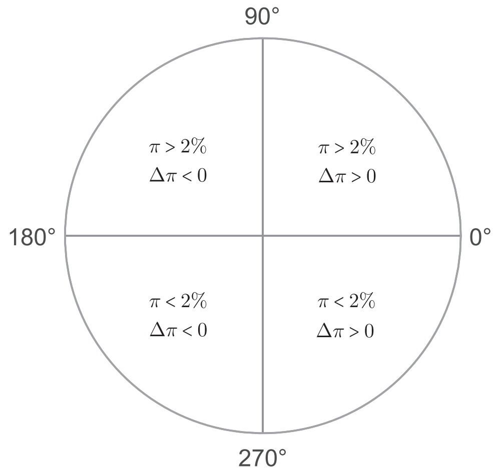

would yield 180°, thus pointing in the opposite direction.[3] Given the geometrical properties of circular data, these cannot be treated in the same way as linear data. In essence, the treatment involves transforming directions from polar coordinates to Cartesian coordinates, calculating the mean of the transformed data, and then back-transforming the outcome. Rather than working with different circles, it turns out to be natural to work with one standard circle, the unit circle, and build all circular functions from it. The unit circle can be divided into four quadrants, each of them representing a range of angles as depicted in Figure 3.

As our main interest lies in the common direction, we focus on the information conveyed by the angles and assume that the observations lie on the unit circle. This corresponds to assuming that $r_{it} = 1$ for all subitems. As illustrated with model (1), such a parameter acts as a scaling factor. Implicitly setting all these parameters to unity acts like a normalization procedure in the same spirit, for instance, as the standardization step in a factor model.

[3]Note that circular data can be measured either in degrees, in the interval between 0° and 360°, or in radians, in the interval between 0 and 2π, where the relationship between both is given by $\frac{\theta_{degrees}}{360°} = \frac{\theta_{radians}}{2\pi}$.

In our setting, we end up with an angle corresponding to every inflation subitem. In particular, the angle θ_i for item i can be determined by setting the value of x_i as the change in the inflation rate for item i, and y_i as the observed inflation rate for item i subtracted to 2 percent. That is, θ_i corresponds to the Cartesian coordinates given by $\Delta\pi_i$ and π_i minus 2 percent, respectively. The rationale for the latter normalization lies in the fact that the ECB's Governing Council, after concluding its strategy review in 2021, considered that price stability is best maintained by aiming for 2 percent inflation over the medium term. Therefore, this transformation allows for a more appealing interpretation of the directions pointing toward 0° or 180°, where inflation would be on target.[4] This implies that the counterclockwise movements around the unit circle can be interpreted as follows. In the first quadrant, lying in the upper right part of the circle, where the angle is between 0° and 90°, inflation is above the target. As we move along this quadrant, price pressures build up with inflation increasing until it attains a local maximum denoted with the 90° angle. Transitioning into the second quadrant, which spans from 90° to 180°, inflation remains above the target. However, prices start to decelerate as we progress along this quadrant until inflation reaches the target with the angle at 180°. Entering the third quadrant, which covers the range from 180° to 270°, inflation is below the target. Here, price pressures continue to decelerate as we move further along this quadrant until inflation hits a local minimum with the angle at 270°. Finally, in the fourth quadrant, spanning from 270° to 360°, prices are accelerating. As we rotate through this quadrant, inflation increases until it reaches the target again.

Given that the headline inflation rate is a weighted average of subitems whose weights reflect the relative importance of each item in the consumption basket, we compute the weighted average of the

[4]For the sake of simplicity, throughout the paper we refer to a 2 percent inflation target, even though we are aware of the previous quantitative definitions of price stability formulated by the ECB. These include the one in place after December 1998: "Price stability shall be defined as a year-on-year increase in the Harmonised Index of Consumer Prices (HICP) for the euro area of below 2%. Price stability is to be maintained over the medium term" and the ex post clarification of the same definition in 2003 to "below, but close to 2% over the medium term."

angles of those subitems. More formally, for a sample of N angles, the computation of the weighted mean angle $\overline{\theta}_\omega$ requires us to compute the rectangular coordinates

$$C_\omega = \sum_{i=1}^{N} \omega_i \cos\theta_i \tag{7}$$

and

$$S_\omega = \sum_{i=1}^{N} \omega_i \sin\theta_i \tag{8}$$

and the resulting length of the mean vector

$$R = \sqrt{C_\omega^2 + S_\omega^2}, \tag{9}$$

where ω_i corresponds to the weight of item i in the computation of the headline inflation. The weighted average angle, $\overline{\theta}_\omega$, can be determined using the cosine and sine trigonometric functions, likewise mentioned before, using

$$\cos\overline{\theta}_\omega = \frac{C_\omega}{R} \tag{10}$$

and

$$\sin\overline{\theta}_\omega = \frac{S_\omega}{R}. \tag{11}$$

To put it simply, $\overline{\theta}_\omega$ can be obtained as

$$\overline{\theta}_\omega = \mathrm{atan2}\left(\sum_{i=1}^{N} \omega_i \sin\theta_i, \sum_{i=1}^{N} \omega_i \cos\theta_i\right) \tag{12}$$

or, alternatively, it can be computed as

$$\overline{\theta}_\omega = \arg\left(\sum_{i=1}^{N} \omega_i e^{i\theta_i}\right). \tag{13}$$

The resulting $\overline{\theta}_\omega$ measure provides the overall direction underlying all inflation subitems at each point in time. Hence, the rotation of the direction given by $\overline{\theta}_\omega$ around the unit circle provides insights on the evolving underlying inflationary pressures.

Given the statistical uncertainty surrounding any measure, it is important to evaluate its magnitude for a well-grounded inference. In fact, it can be extremely useful to provide a confidence interval for the measure of interest to allow for a better assessment of the current direction.

In the case of directional data, the most used distribution is the von Mises distribution, which assumes unimodality and symmetry. It is basically the circular analogue to normal distribution. As these conditions do not always hold, the pursuit of robust methods leads to nonparametric or distribution-free techniques. Note that whereas in linear inference one can justify an assumption of normality, for instance, when dealing with means of large samples, there is no analogue rationale for invoking the von Mises distribution in directional inference. Hence, the use of bootstrap methods has been advocated for directional data, as the distributions of the statistics commonly used for inference are frequently intractable; see Mardia and Jupp (2000) for details.

In fact, bootstrap methods have proved very effective in situations where distributional assumptions are kept to a minimum or when distributional results for the statistic of interest are not available. With bootstrapping, the distribution of a statistic of interest can be assessed by resampling, i.e., sampling from the observed data, and then evaluating the statistic of interest for each of the bootstrap samples with the variability of these values taken as an estimate of the variability of the statistic over the population (see, for example, Efron and Tibshirani 1993).

In particular, to obtain a confidence interval for $\overline{\theta}_\omega$ at each point in time we proceed as follows (see Fisher (1993) for a detailed discussion). In the first stage, given a sample of N angles, we resample with replacement and generate B bootstrap samples with the same size of the original sample.[5] In the second stage, for the b^{th} bootstrap sample, we evaluate the corresponding weighted mean angle, $\overline{\theta}_\omega^b$, using (12). In the third stage, we compute the confidence

[5]In the empirical application that follows, the number of bootstrap samples is set to 10,000.

interval. To obtain the $(1-\alpha)100$ percent confidence interval, we need to compute the circular distance between the weighted mean angle for the original sample and that of the b^{th} bootstrap sample as

$$\gamma_b = \overline{\theta}_\omega^b - \overline{\theta}_\omega,\ b = \{1, \ldots, B\}, \tag{14}$$

and sort into increasing order. We then determine the quantiles $\alpha/2$ and $1-\alpha/2$ of γ_b, denoted by $\gamma_b^{\alpha/2}$ and $\gamma_b^{1-\alpha/2}$, respectively. The $(1-\alpha)100$ percent confidence interval for $\overline{\theta}_\omega$ is given by

$$CI_{\overline{\theta}_\omega}^{1-\alpha} = \left[\overline{\theta}_\omega + \gamma_b^{\alpha/2}, \overline{\theta}_\omega + \gamma_b^{1-\alpha/2}\right]. \tag{15}$$

With the above confidence interval, we can also conduct a hypothesis test with a significance level α by simply not rejecting the null hypothesis that the weighted mean is equal to a value of interest $\overline{\theta}_\omega^0$ if $\overline{\theta}_\omega^0$ lies within the confidence interval and rejecting the null hypothesis otherwise. For instance, given that directions of 0° and 180° correspond to a 2 percent inflation level, which is the ECB inflation target, we can assess if our measure $\overline{\theta}_\omega$ departs statistically from 0° and 180°. If 0° or 180° do not lie inside the confidence interval, then it signals that the underlying inflation is statistically different from the inflation target. Note that the bootstrap confidence interval is not necessarily symmetric around the point estimate, as it relies on resampling from the original data, allowing for enhanced flexibility in the characterization of the distribution.

3. Data

The HICP is the official index of consumer prices in the euro area, harmonized within member states and taken as reference for the purposes of monetary policy. It is used to assess price stability around the inflation target and price convergence required for entry into the European Monetary Union.

The empirical application that follows relies on the universe of subitems that compose the HICP in the euro area, according to the European classification of individual consumption by purpose (ECOICOP) at the four-digit level.[6] This amounts to 93 series

[6]Data are retrieved from the Eurostat database and can be accessed at https://ec.europa.eu/eurostat/web/hicp/database.

Figure 4. Measures of Underlying Inflation | Year-on-Year Percentage Change

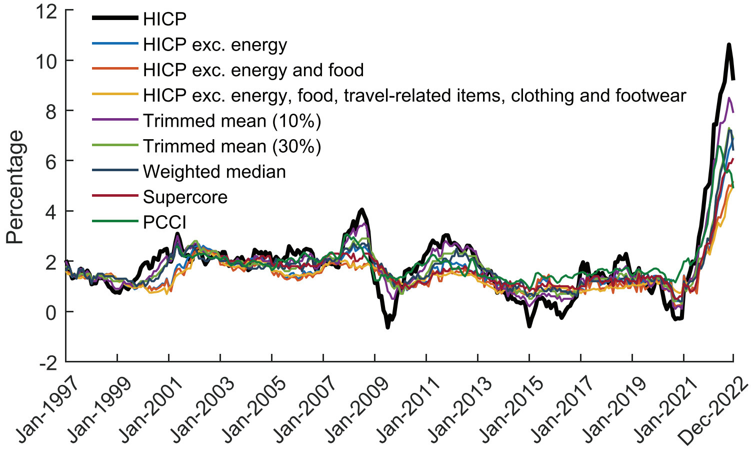

Note: The sample period for all measures starts in January 1997, except for the Supercore (January 2003) and PCCI (April 2001). The latest vintage of the PCCI is displayed (December 2022).

spanning the period from January 1996 to December 2022. Similarly, we collect data on the corresponding weights for each subitem in the reference consumption basket. Taken together, this information provides the building blocks for constructing the inflation compass. Data are neither seasonally nor working-day adjusted and we compute year-on-year rates of change for each subitem.

Figure 4 depicts several measures of underlying inflation commonly tracked in the euro-area economy. As mentioned before, these can be grouped into three broad categories: (i) permanent exclusion measures, (ii) temporary exclusion measures, and (iii) frequency exclusion measures. The idea of the first group of measures is to abstract from typically volatile components (e.g., fluctuations in oil prices that affect energy goods inflation, or atypical weather which induces strong volatility in food inflation). The second class consists of a temporary exclusion of subitems in line with a given statistical criterion. These include trimmed means (10 percent, 30 percent) or the weighted median. For example, the 10 percent trimmed mean removes 5 percent of the year-on-year rates of change from each tail of the distribution of all subitems and aggregates the remaining year-on-year rates of change by rescaling the weights. Note that the weighted median can be considered an extreme form of trimmed mean, as it trims all price changes except for the weight-based midpoint of the distribution. The third category, in turn, comprises model-based measures, built to filter out the transitory component of inflation and retain the persistent components of all subitems. Examples of these measures include the so-called Supercore and PCCI. The former relies solely on subitems that are deemed sensitive to slack, as measured by the output gap in a Phillips-curve regression framework, thus departing from purely statistical criteria (see O'Brien 2018). The latter draws on the estimation of a dynamic factor model to capture the persistent and common component of all subitems' inflation rates across 12 euro-area countries, based on ECOICOP four-digit classes, amounting to approximately 1,000 series. The estimation of the low-frequency component for each subitem entails the choice of a set of parameters, including the number of dynamic and static factors and the threshold for the minimum length of cycles allowed in the common component (see Bańbura and Bobeica (2020) for details). These measures are sourced from the ECB's Statistical Data Warehouse.

While HICP inflation has hovered around 2 percent for over 20 years in the euro area, with some periods falling below the target over the last decade, the recent surge in inflation has also been evident in all measures of underlying inflation. Although there were already some signs of easing toward the end of 2022, these remain at historically high levels.

4. The Inflation Compass for the Euro Area

In this section, we present the inflation compass for the euro area, focusing on key episodes of historically high and low inflation since the beginning of the monetary union. Figure 5 depicts such episodes, where the red (blue) dots denote the periods of high (low) inflation. Two episodes of high inflation emerge, occurring in July 2008 and October 2022, with inflation rates hovering around 4 percent and 10 percent, respectively. In contrast, in July 2009, January 2015, and September 2020, inflation in the euro area was marginally negative. Although the inflation compass for the euro area can be

Figure 5. Key Episodes of High and Low Inflation in the Euro Area | Year-on-Year Percentage Change

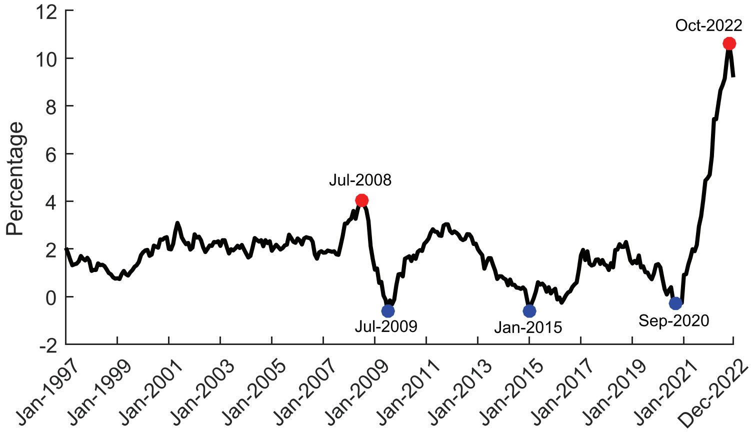

computed every month, we pick these episodes in order to illustrate the inflation compass and to confirm its usefulness during such periods.

In Figure 6, the inflation compasses for the months of high or low inflation mentioned above are displayed in the center of each row, along with the adjacent months on either side. In the compass, the needle in red denotes the weighted mean direction, whereas the arc in black around the compass depicts the corresponding 95 percent confidence interval. The compass reads as explained in detail in Section 2.

Starting with the episode of high inflation identified in July 2008, the compass needle in this month is pointing toward the north, i.e., in the 90° direction, with the confidence interval including this direction. This means that we would not reject that underlying inflation has attained a local maximum. In fact, we would reject the occurrence of an episode of high inflation in the preceding month (i.e., June 2008) as the 90° direction lies outside the confidence interval. Naturally, a wider confidence interval reflects a greater dispersion of the evolution of the subitems of the consumption basket, thus inducing higher uncertainty in the overall direction of price pressures. The rotation of the compass from June 2008 to August 2008 signals that price pressures in the euro area increased and hit a local maximum.

Figure 6. Inflation Compasses During Episodes of High and Low Inflation

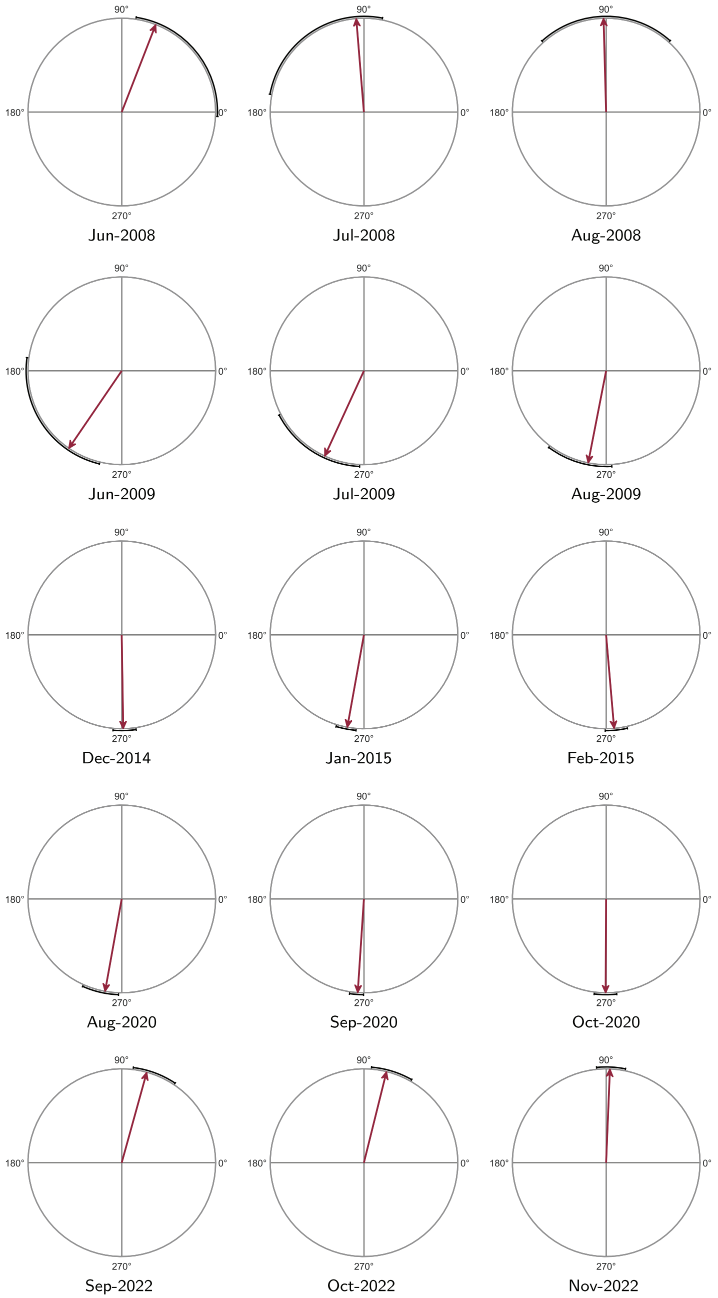

The global financial crisis was also marked by periods of (marginally) negative inflation in the euro area. The inflation compasses in the second row of Figure 6 illustrate these events. Starting in June 2009, price pressures steadily decelerated as we moved further along the third quadrant until the compass needle pointed toward south, that is, in the 270° direction. In fact, in August 2009, one cannot reject that underlying inflation has reached a local minimum. The wider confidence intervals reflect the turbulent economic landscape experienced during these times, with repercussion in the elementary items of the overall consumer price index.

The two other periods of low inflation, in early 2015 and in the autumn of 2020, are depicted in the third and fourth rows of Figure 6, respectively. We observe that the compass needles for these episodes are pointing toward the south, with little dispersion around the mean direction. The narrower confidence interval in both episodes conveys smaller uncertainty, which can be explained by the greater synchronization of the dynamics of the subitems of inflation during these periods.

After this, inflation surged worldwide. When analyzing the rotation of the compass needle between October 2020 and September 2022, we observe this buildup of inflationary pressures as the needle rotates smoothly from the south to the north. In the euro area, inflation reached double digits at the start of autumn 2022. The bottom row in the figure illustrates this. Departing from the first quadrant, price pressures gradually accelerated until underlying inflation reached a local maximum at 90°, notably in November 2022.

These results highlight that the inflation compass stands out for its appealing properties in tracking price pressures in the euro area. Unlike most existing tools, it does not exclude any subitems of inflation, whether temporarily or permanently, ensuring that all available information is used. This eliminates any ad hoc selection of subitems, thus granting a comprehensive analysis of the whole range of price data. As HICP data are not revised, the inflation compass is not subject to revisions, allowing for reliable real-time insights into price pressures. Furthermore, its transparency sets it apart from intricate underlying inflation measures like the PCCI and Supercore, providing clear and easily understandable information for policymakers and the wider public.

Figure 7. HICP Inflation | Year-on-Year Percentage Change

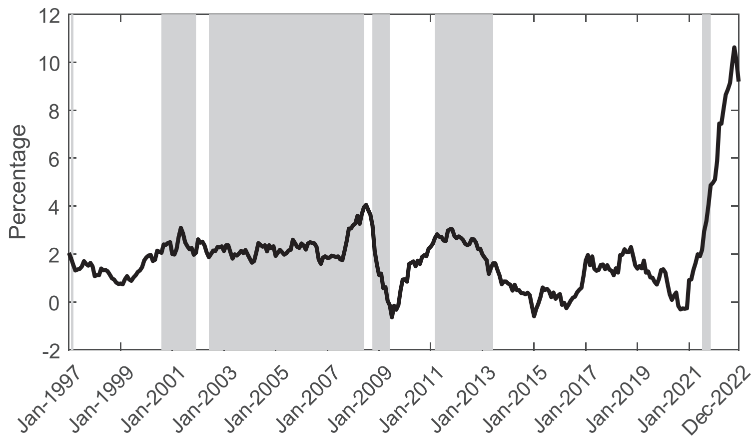

.

Note: The shaded areas denote the periods in which underlying inflation is not statistically different from the 2 percent inflation rate with a 95 percent confidence level.

For the sake of tractability and to ease presentation, the analysis so far has focused on a few episodes of high and low inflation. However, we can also take advantage of the properties of the inflation compass to characterize the price pressures since the beginning of the monetary union. In particular, drawing on the confidence interval obtained by bootstrap methods, we can evaluate at each point in time whether the underlying inflation departs from the 2 percent inflation rate, i.e., if the confidence interval encompasses the directions of 0° or 180°. Therefore, we display in Figure 7 the periods in which underlying inflation in the euro area is not statistically different from 2 percent.

The gray shaded areas map the periods where the confidence interval of the compass needle includes the 0° or 180° directions, i.e., the east or the west directions. As alluded to before, these are the directions where euro-area inflation would be on target. These results give rise to two important remarks. Firstly, the episodes addressed before lie outside the shaded areas. In fact, these were periods of high (or very high) and (marginally) negative inflation in the euro area, thus differing from 2 percent in a statistical sense. Secondly, Figure 7 highlights that the underlying trend in overall inflation in the euro area was pointing to the target for most of the time during the first decade of the 2000s and in the 2011–12 period. Overall, the 2008–11 period was marked by initially rising inflation, followed by a period of disinflationary pressures and, later, a slow return to levels around 2 percent. Moreover, euro-area inflation was persistently low from 2013 to 2019, with an average rate of 1 percent during this period, significantly below 2 percent. In fact, this low-inflation environment can be attributed to a combination of interconnected factors, including the underestimation of the economic slack, a de-anchoring of longer-term inflation expectations, and the ongoing structural trends in the economy such as globalization, digitalization or demographic trends, as advocated by Koester et al. (2021). In the aftermath of the COVID-19 pandemic and the war in Ukraine, inflation surged from 1 percent in the beginning of 2021 to double digits by the autumn of 2022, thus crossing 2 percent, as evidenced in Figure 7. Hence, this complementary analysis based on the inflation compass offers valuable insights when monitoring price pressures within the euro area, by assessing whether underlying inflation deviates from 2 percent.

5. Predictive Content of Underlying Inflation Measures

Economists and policymakers usually resort to measures of underlying inflation to predict future changes in overall inflation. As such, an empirical evaluation of the information content delivered by the inflation compass to forecast headline inflation is conducted. To this end, we compare the suggested tool in a forecasting context exercise against the underlying inflation measures commonly used in monetary policy diagnostics in the euro area. Besides its qualitative allure, we show that the inflation compass holds considerable quantitative relevance in monitoring inflation, particularly within the context of forecasting up to 36 months ahead.

The predictive content is gauged from a regression equation that relates the change in overall inflation between the current month and a future time period to the current gap between the underlying inflation measure and the overall inflation rate. Formally, we assess the ability to predict future inflation by estimating the equation

$$\pi_{t+h} - \pi_t = \alpha + \beta(\pi_t^u - \pi_t) + \epsilon_{t+h}, \tag{16}$$

where π_{t+h} represents the year-on-year headline inflation rate h months ahead; π_t denotes the year-on-year headline inflation rate in a given month; π_t^u denotes the underlying inflation measure in the same month; ϵ_t is a mean-zero random disturbance term, and t indexes time.

This regression framework has been used in several studies that evaluate the ability of underlying inflation measures to forecast inflation (see, inter alia, Clark 2001; Hogan, Johnson, and Laflèche 2001; Cogley 2002; Rich and Steindel 2007; or Bańbura and Bobeica 2020). To address the inherent subjectivity introduced by the requirement to specify an econometric model, we employ a regression framework that links future inflation changes to the transitory component of price changes identified by the underlying inflation measures. Note that the use of differences in inflation rates for both dependent and independent variables ensures stationarity and avoids issues arising from the existence of unit roots. This framing offers the advantages of both interpretability and flexibility, as it not only simplifies the readability of the exercise, but also allows for the incorporation of alternative horizons in the analysis. As Clark (2001) and Rich and Steindel (2007) argue, this framework is also aligned with the common beliefs of central bankers who consider movements in underlying inflation, by themselves, as signals of future changes in inflation.

In the case of the inflation compass, we have a circular predictor. In order to measure the association between a linear random variable, Z, and an angular random variable, Θ, the regression model for this type of association has the form

$$E(Z|\Theta = \theta) = \alpha + \gamma_1 \sin\theta + \gamma_2 \cos\theta \tag{17}$$

(see Johnson and Wehrly 1978; Fisher 1993; Mardia and Jupp 2000). This corresponds to a linear-circular regression model where the response is a linear variable and the covariate is circular. It turns out to be a simple linear regression model, as it is linear in the regression variables $\sin\theta$ and $\cos\theta$. Hence, in the case of the inflation compass, we include both the sine and cosine of $\overline{\theta}_\omega$, and run a standard linear regression with these components; that is,

$$\pi_{t+h} - \pi_t = \alpha + \gamma_1 \sin \overline{\theta}_\omega + \gamma_2 \cos \overline{\theta}_\omega - \beta \pi_t + \epsilon_{t+h}, \qquad (18)$$

which can be easily rearranged to look like (16):

$$\pi_{t+h} - \pi_t = \alpha + \beta \left(\frac{\gamma_1}{\beta} \sin \overline{\theta}_\omega + \frac{\gamma_2}{\beta} \cos \overline{\theta}_\omega - \pi_t \right) + \epsilon_{t+h}. \qquad (19)$$

We estimate the regression models for several horizons, $h = \{12, 18, 24, 30, 36\}$, to examine the medium- to long-run predictive content of the several underlying inflation measures. The emphasis on these horizons is motivated by conventional wisdom about the lags in the monetary policy transmission mechanism. In fact, many countries that implement inflation targeting typically establish a horizon at which monetary policy operates over the business cycle and central banks may effectively anchor expectations or maintain inflation at desirable levels.

5.1 In-Sample Evaluation

We start by presenting the results of the in-sample evaluation. We split the sample period roughly into two-thirds for the in-sample analysis while leaving the remaining third for the out-of-sample evaluation. In particular, we restrict the in-sample analysis to the period until December 2015, while the out-of-sample period runs from 2016 up to the end of 2022. The sample period starts in February 1997, except for the Supercore (January 2003) and the PCCI (April 2001). Table 1 reports the root mean squared error (RMSE) over the different horizons for the inflation compass relative to that of each underlying inflation measure for the same period. A reading below 1 indicates that the inflation compass has greater in-sample ability to track headline inflation than the alternative measure.

We find that the inflation compass outperforms the measures of underlying inflation commonly followed in the euro area across all horizons. The only two exceptions are the Supercore for $h = 36$ and the PCCI for $h = 12$, where the ratio is higher than 1. The inflation compass exhibits sizable improvements across various horizons when compared to the permanent and temporary exclusion measures, and to a lesser extent relative to the Supercore. In-sample analyses are commonly subject to criticism, as they tend to deviate from real-time forecasting exercises and are susceptible to overfitting.

Table 1. In-Sample Relative RMSE Across the Different Horizons

	Dependent Variable: $\pi_{t+h} - \pi_t$				
	h = 12	**h = 18**	**h = 24**	**h = 30**	**h = 36**
HICPX	0.875	0.809	0.783	0.800	0.900
HICPXX	0.952	0.890	0.842	0.851	0.945
HICPXXX	0.982	0.919	0.866	0.867	0.943
Trimmed Mean (10 percent)	0.821	0.740	0.711	0.707	0.753
Trimmed Mean (30 percent)	0.866	0.791	0.756	0.752	0.821
Weighted Median	0.897	0.828	0.800	0.808	0.881
Supercore	0.994	0.922	0.859	0.882	1.045
PCCI	1.072	0.982	0.957	0.959	0.985

Note: HICPX refers to the HICP excluding energy, HICPXX corresponds to the HICP excluding energy and food, and HICPXXX is the HICP excluding energy, food, travel-related items, clothing, and footwear.

To overcome these limitations, we proceed with an out-of-sample forecast evaluation.

5.2 Out-of-Sample Forecast Evaluation

This subsection introduces the out-of-sample forecasting exercise. For each horizon h, we start by estimating the model with data until December 2015 and produce an h-step-ahead forecast. For example, when $h = 12$, the first forecast refers to December 2016. We proceed by estimating the forecasting equation recursively, expanding the estimation window at each time by one month. This yields a sequence of h-step-ahead forecasts and forecast errors. The exercise is conducted in real time using the actual vintages of data available at each point in time.[7]

For benchmarking, we consider two additional models in the forecasting evaluation exercise. The first benchmark corresponds to a naïve forecast based on an autoregressive model, recursively estimated and with the order chosen according to the Bayesian information criterion (BIC). The second alternative model corresponds to the approach suggested by Bai and Ng (2008) for inflation forecasting. In particular, it combines factor models with a method based

[7]In particular, in the case of the PCCI, we consider the corresponding real-time vintages kindly made available by the ECB to the authors.

on penalized regressions that performs subset selection and shrinkage to target predictors, the LARS-EN (least-angle regression with elastic net) algorithm. The elastic net combines the virtues of Lasso (least absolute shrinkage and selection operator) and ridge regression, allowing for shrinkage of coefficients, elimination of regressors, and efficient selection of variables within the data set (see Zou and Hastie 2005). Hence, given a large data set, a selection of a subset of predictors targeting the variable of interest is conducted prior to factor estimation, which allows us to reduce the influence of uninformative predictors and enhance forecast accuracy. Drawing on the HICP large data set, the subitems are selected at each point in time through the LARS-EN algorithm.[8] Then, the factors are estimated from the set of selected predictors by principal components and the number of estimated factors included in the forecasting model is chosen according to a modified BIC as in Stock and Watson (1998).

To quantify the out-of-sample forecasting performance, we compute the root mean squared forecast error (RMSFE). In addition, we assess whether the forecast accuracy of the inflation compass significantly differs from that of other underlying inflation measures by employing the Clark and West (2007) test. These results are reported in Table 2.

The out-of-sample analysis shows that the inflation compass provides more accurate forecasts compared to other methods, given that the relative RMSFEs for different time horizons are below unity, indicating its superiority. Firstly, it outperforms the benchmarks for all horizons at a 1 percent statistical significance level. When comparing the inflation compass with temporary exclusion measures, it achieves an average gain of 6 percent, reaching around 10 percent for the medium-term horizons, $h = 18$ and $h = 24$, vis-à-vis the trimmed means. Compared with the model-based measures, the average gain is close to 4 percent, being higher vis-à-vis the Supercore than against the PCCI. This is very clear in the case of the horizons $h = 18$ and $h = 24$. In turn, the accuracy gains are

[8]Based on the findings by Bai and Ng (2008), we set the parameter that controls the number of predictors to be selected in Lasso to be 30 and the parameter that controls the importance of the penalty associated with the ridge regression to be 0.25. As a sensitivity analysis, we considered other alternative settings and the results are qualitatively similar.

Table 2. Out-of-Sample Relative RMSFE Across the Different Horizons

	Forecast Horizon				
	h = 12	**h = 18**	**h = 24**	**h = 30**	**h = 36**
AR	0.907***	0.882***	0.897***	0.887***	0.823***
LARS-EN-PC	0.920***	0.907***	0.916***	0.910***	0.904***
HICPX	0.958	0.929**	0.943*	0.982**	0.972*
HICPXX	0.968	0.935**	0.935*	0.969**	0.964*
HICPXXX	0.993	0.977*	0.984	0.988	0.974*
Trimmed Mean (10 percent)	0.957*	0.897**	0.901**	0.949***	0.962**
Trimmed Mean (30 percent)	0.937***	0.910**	0.924**	0.961**	0.973**
Weighted Median	0.935**	0.928**	0.948*	0.984**	0.973*
Supercore	0.930**	0.916**	0.941*	0.972**	0.954**
PCCI	0.931	0.974*	0.999	0.979**	0.962**

Note: ***, **, * denote statistical significance at 1 percent, 5 percent, and 10 percent significance levels, respectively. AR and LARS-EN-PC denote the autoregressive benchmark and the Bai and Ng (2008) approach, respectively. HICPX refers to the HICP excluding energy, HICPXX corresponds to the HICP excluding energy and food, and HICPXXX is the HICP excluding energy, food, travel-related items, clothing, and footwear. The recursive out-of-sample forecasting exercise uses as starting estimation sample the period up to December 2015.

somewhat smaller vis-à-vis the permanent exclusion measures, averaging slightly less than 4 percent. This is particularly visible against the HICPXXX, where the gains are the lowest across all the alternative measures. Nevertheless, the gains are, in general, statistically significant according to the Clark and West (2007) test procedure, which corroborates the usefulness of the proposed tool. We find that the statistical significance of the differences in performance is more marked for horizons $h = 18$ and above.

6. Robustness Analyses

6.1 The COVID-19 Pandemic and the Recent Inflation Surge

As a sensitivity analysis, we assess the forecasting performance of each measure over the last three years, a period during which the euro-area economy was severely hit by the COVID-19 pandemic and, more recently, by an inflation surge. As Chahad et al. (2022) outline, there has been a significant degradation in the precision of HICP

Table 3. Out-of-Sample Relative RMSFE During the COVID-19 Pandemic and the Recent Inflation Surge

	Forecast Horizon				
	h = 12	**h = 18**	**h = 24**	**h = 30**	**h = 36**
AR	0.935**	0.904**	0.920**	0.898**	0.819**
LARS-EN-PC	0.937**	0.915**	0.946*	0.954***	0.918***
HICPX	0.962	0.930*	0.942	0.981	0.965*
HICPXX	0.969	0.934*	0.932	0.967*	0.957**
HICPXXX	0.994	0.977	0.980	0.984	0.968*
Trimmed Mean (10 percent)	0.967	0.900**	0.905*	0.955*	0.957*
Trimmed Mean (30 percent)	0.942**	0.912**	0.926*	0.964*	0.966**
Weighted Median	0.939**	0.930**	0.948	0.984	0.966*
Supercore	0.930**	0.916*	0.940	0.970**	0.947**
PCCI	0.929	0.975	1.000	0.975*	0.955**

Note: ***, **, * denote statistical significance at 1 percent, 5 percent, and 10 percent significance levels, respectively. AR and LARS-EN-PC denote the autoregressive benchmark and the Bai and Ng (2008) approach, respectively. HICPX refers to the HICP excluding energy, HICPXX corresponds to the HICP excluding energy and food, and HICPXXX is the HICP excluding energy, food, travel-related items, clothing, and footwear. The forecast errors are computed for the period from January 2020 until December 2022.

forecasts since the outbreak of the pandemic, and particularly following the third quarter of 2021. The plunge in the accuracy of these projections is mainly ascribed to unforeseen fluctuations in energy prices. This factor, along with the consequences of the post-lockdown reopening of economies and the global supply chain bottlenecks, has spurred unprecedented surges in HICP inflation.

In the analysis that follows, we take the forecast errors from January 2020 until December 2022 for all horizons and proceed as before by computing the RMSFEs and employing the Clark and West (2007) test for equal forecast accuracy (Table 3).

The results reveal that, in general, the inflation compass delivers higher accuracy than alternative measures. Importantly, the findings are robust to the pandemic and inflation surge periods, during which monetary policy gained prominence.[9] Throughout these

[9]Note that the loss of statistical significance may be related to the lower number of observations used to perform the tests.

periods, policymakers closely monitored a set of measures, while significant interest rate hikes were implemented. Notably, these results underscore the importance of the inflation compass in tracking overall inflation developments during times of disruption or structural change. The fact that it comprises all the information available, weighting the changes in all items of the consumption basket, allows the inflation compass to capture strong and unexpected fluctuations that are not predicted by other underlying inflation measures that exclude items, such as energy, or others that smooth swift movements in headline inflation.

6.2 Bringing On Board More Data

In this subsection, we investigate whether using more disaggregated data changes the reliability of the inflation compass and its forecast accuracy. To this end, we consider two dimensions of disaggregation: (i) from ECOICOP4 to ECOICOP5 (product-level disaggregation); and (ii) from ECOICOP4 at the product level for the euro area to ECOICOP4 using data for 12 euro-area countries, i.e., we use country as opposed to aggregate data (country-level disaggregation).[10] One should note that the suggested approach is flexible enough to accommodate a time-varying composition of the input data and can handle large-dimensional data straightforwardly.

We illustrate the outcomes of these two levels of disaggregation by focusing on the two recent episodes of low and high inflation in the euro area experienced since 2020. In particular, Figure 8 reports the inflation compasses computed using the five-digit classes data encompassing nearly 300 HICP subitems and corresponding weights. We observe that the compass needles in the first row of Figure 8 are pointing toward the south, with little dispersion around the weighted mean direction as evidenced before. In addition, the bottom row in the figure reinforces the idea of price pressures gradually accelerating until inflation reached a local maximum around November 2022.

[10] The five-digit ECOICOP classes are only available from January 2017 onward and are retrieved from the Eurostat database mentioned before. The country-level data refer to the 11 founding euro-area member states, namely Austria, Belgium, Finland, France, Germany, Ireland, Italy, Luxembourg, the Netherlands, Portugal, and Spain, plus Greece, which joined the euro area in 2001.

Figure 8. Inflation Compasses Using Five-Digit Data

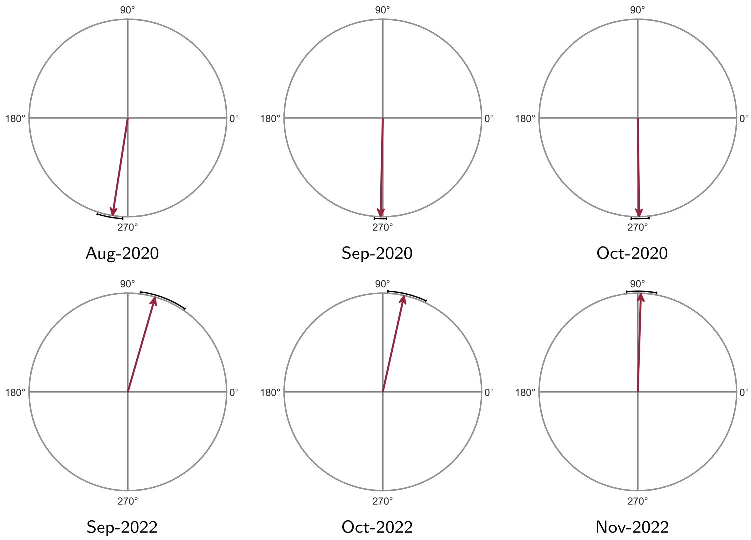

We now turn our attention to the second dimension of disaggregation leveraging the country-level data. In particular, we enlarge the data set to cover the four-digit ECOICOP classes for 12 euro-area countries as in the construction of the PCCI, amounting to more than 1,000 series in total. Thereby, we explore an even larger data set aimed to capture country-level inflationary pressures.

Figure 9 displays the inflation compasses for the two same episodes addressed above. We find that the readings of the inflation compasses during these two episodes are in line with those obtained before.

In Table 4 we compare the forecasting ability of the inflation compass based on four-digit vis-à-vis five-digit classes as well as against four-digit classes data for 12 euro-area countries.

By enlarging the data set, the forecasting performance of the inflation compass remains relatively unaffected. In fact, the inflation compass is robust to the product-level disaggregation, in the sense that using more detailed data does not enhance forecast accuracy. This is outlined by the entries nearly identical to unity in the top row of Table 4, which exhibit the ratio of RMSFE between the inflation

Figure 9. Inflation Compasses Using Four-Digit Country-Level Data

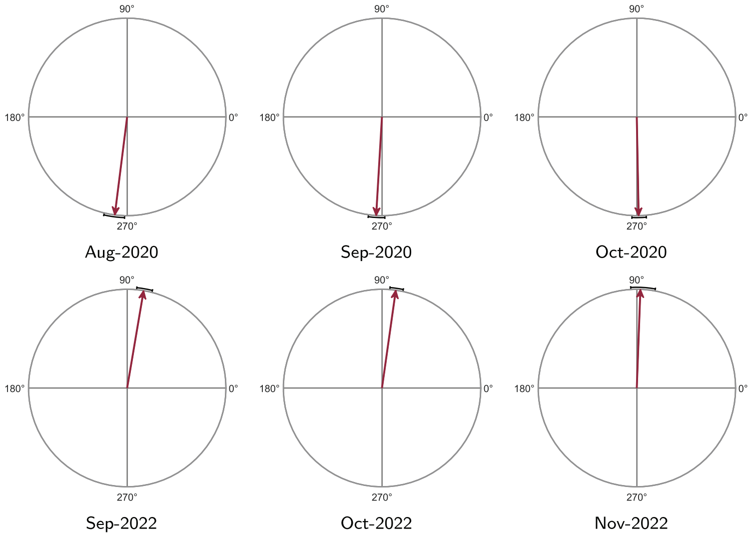

Table 4. Out-of-Sample Relative RMSFE Leveraging More Data

	Forecast Horizon				
	h = 12	**h = 18**	**h = 24**	**h = 30**	**h = 36**
Five-Digit Data Compass	0.998	1.000	1.000	1.000	0.999
Four-Digit Country-Level Data Compass	1.009	0.995	0.994	0.999	0.992

Note: The top (bottom) row corresponds to the ratio between the RMSFE of the inflation compass computed from the four-digit euro-area classes and the RMSFE using the five-digit euro-area classes (four-digit classes country-level data). The forecast errors are computed for the period from January 2020 until December 2022.

compass computed from the four-digit classes and that using the five-digit classes. In fact, these marginal differences in forecast accuracy are confirmed by the lack of statistical significance of those entries. One should bear in mind that using four-digit ECOICOP classes, as examined in Section 5, already entails considering a relatively large level of disaggregation with almost 100 HICP subitems.

Moreover, the inflation compass is robust to the country-level disaggregation, in the sense that bringing on board more data does not seem to enhance forecast accuracy, as evidenced in the bottom row of Table 4. Even though for $h = 12$ the inflation compass computed from the four-digit classes country-level data outperforms that using the four-digit classes data for the euro area, this difference is negligible, which is confirmed by the lack of statistical significance.

7. Concluding Remarks

A key feature of central bank mandates worldwide is their unwavering attention to monitoring inflation dynamics. The ECB is no exception. Monetary policy prescriptions in the euro area have typically relied on an assessment of the inflation outlook, considering incoming economic and financial data, the dynamics of inflation, and the effectiveness of monetary policy transmission.

In this regard, several techniques to filter out incoming data have been used, given the importance of gauging price pressures in a timely manner for the purposes of monetary policy responses. The idea of closely monitoring such tools is to abstract from transitory price movements and thereby produce a measure of underlying inflation.

We depart from the previous literature and propose a novel tool to gauge price pressures resorting to circular statistics, the so-called inflation compass. Given that circular statistics have not been commonly used in economics, we lay out the basic intuition while streamlining the main theoretical concepts underneath the inflation compass. In particular, we illustrate the inflation compass reading and develop statistical inference based on bootstrap methods.

Unlike most alternative measures of underlying inflation, the inflation compass does not exclude any subitems of inflation, ensuring that all available information is used. This eliminates the need for ad hoc selection of subitems and allows for a comprehensive analysis of all price data. Another feature of the inflation compass is its real-time reliability, in the sense that it is not revised backward as new data arrive, thus providing policymakers with reliable real-time insights about price pressures. Furthermore, the transparency of the inflation compass sets it apart from other, more intricate tools. In fact, the information conveyed by the inflation compass can be easily understood and interpreted, facilitating decision-making and communication.

Resorting to HICP disaggregated data for the euro area, covering almost 100 series, we illustrate the usefulness of the inflation compass for tracking underlying price pressures during key episodes of high and low inflation since the euro area was formed. We also take advantage of the properties of the inflation compass to shed light on the periods during which underlying inflation did not deviate from 2 percent, thus providing a characterization of the full spectrum of price pressures in the euro area over the last decades.

In addition to its qualitative relevance in gauging price pressures, we provide evidence of the usefulness of the inflation compass for forecasting overall inflation up to 36 months ahead. Quantitatively, the inflation compass outperforms the measures regularly monitored in the euro area in a real-time forecasting exercise. We find that the forecast accuracy gains achieved with the suggested approach are, in general, statistically significant.

Finally, we conduct a robustness analysis where we investigate the performance of the inflation compass for monitoring price pressures amid challenging periods like the COVID-19 pandemic and the recent inflation surge. Our findings confirm its dominance even in such difficult times for policymaking. Moreover, as the suggested method can handle large-dimensional data straightforwardly, we enlarge the data set either by incorporating more detailed product-level data from the consumer price index or by taking on board country-level data rather than relying on aggregate euro-area data. Either way, the main findings hold, which reinforces the robustness of the proposed procedure.

References

Arti, M. J., R. C. Bladen-Hovell, D. R. Osborn, G. Smith, and W. Zhang. 1995. "Predicting Turning Points in the UK Inflation Cycle." *Economic Journal* 105 (432): 1145–64.

Bai, J., and S. Ng. 2008. "Forecasting Economic Time Series Using Targeted Predictors." *Journal of Econometrics* 146 (2): 304–17.

Bakhsi, H., and T. Yates. 1999. "To Trim or Not to Trim? An Application of a Trimmed Mean Inflation Estimator to the United Kingdom." Working Paper No. 97, Bank of England.

Bańbura, M., and E. Bobeica. 2020. "PCCI – a Data-Rich Measure of Underlying Inflation in the Euro Area." Statistics Paper No. 38, European Central Bank.

Banerji, A., and L. Achuthan. 2019. "What the Fed Has Forgotten: The Inflation Cycle Is Not the Business Cycle." *Challenge* 62 (4): 242–52.

Bryan, M., and S. Cecchetti. 1994. "Measuring Core Inflation." In *Monetary Policy*, 195–219. Chicago: University of Chicago Press.

Bryan, M., S. Cecchetti, and R. Wiggins. 1997. "Efficient Inflation Estimation." NBER Working Paper No. 6183, September.

Chahad, M., A. Hofmann-Drahonsky, B. Meunier, A. Page, and M. Tirpák. 2022. "What Explains Recent Errors in the Inflation Projections of Eurosystem and ECB Staff?" *Economic Bulletin* (European Central Bank) 3/2022.

Clark, T. 2001. "Comparing Measures of Core Inflation." *Economic Review* (Federal Reserve Bank of Kansas City) Second Quarter.

Clark, T., and K. West. 2007. "Approximately Normal Tests for Equal Predictive Accuracy in Nested Models." *Journal of Econometrics* 138 (1): 291–311.

Cogley, T. 2002. "A Simple Adaptive Measure of Core Inflation." *Journal of Money, Credit and Banking* 34 (1): 94–113.

Dolmas, J. 2005. "Trimmed Mean PCE Inflation." Working Paper No. 56, Federal Reserve Bank of Dallas.

Efron, B., and R. J. Tibshirani. 1993. *An Introduction to the Bootstrap.* New York: Chapman and Hall.

Ehrmann, M., G. Ferrucci, M. Lenza, and D. O'Brien. 2018. "Measures of Underlying Inflation for the Euro Area." *Economic Bulletin* (European Central Bank) 4/2018.

Fisher, N. I. 1993. *Statistical Analysis of Circular Data.* Cambridge: Cambridge University Press.

Hogan, S., M. Johnson, and T. Laflèche. 2001. "Core Inflation." Technical Report No. 89, Bank of Canada.

Jammalamadaka, S. R., and A. SenGupta. 2001. *Topics in Circular Statistics.* Singapore: World Scientific.

Johnson, R. A., and T. E. Wehrly. 1978. “Some Angular-Linear Distributions and Related Regression.” *Journal of the American Statistical Association* 73 (363): 602–06.

Koester, G., E. Lis, C. Nickel, C. Osbat, and F. Smets. 2021. “Understanding Low Inflation in the Euro Area from 2013 to 2019: Cyclical and Structural Drivers.” Occasional Paper No. 280, European Central Bank.

Lane, P. R. 2022. “Inflation Diagnostics.” *ECB Blog*, November 25, European Central Bank.

———. 2023. “Underlying Inflation.” Speech delivered at Trinity College, Dublin, March 6.

Lourenço, N., and A. Rua. 2023. “Business Cycle Clocks: Time to Get Circular.” *Empirical Economics* 65: 1513–41.

Luciani, M. 2020. “Common and Idiosyncratic Inflation.” Finance and Economics Discussion Series No. 2020-024, Board of Governors of the Federal Reserve System.

Luciani, M., and R. Trezzi. 2019. “Comparing Two Measures of Core Inflation: PCE Excluding Food & Energy vs. The Trimmed Mean PCE Index.” FEDS Notes, Board of Governors of the Federal Reserve System, August 2.

Mankikar, A., and J. Paisley. 2002. “What Do Measures of Core Inflation Really Tell Us?” *Quarterly Bulletin* (Bank of England) 42 (4, Winter): 373–83.

Mardia, K. V., and P. E. Jupp. 2000. *Directional Statistics*. New York: John Wiley & Sons.

Marques, C. R., P. D. Neves, and L. M. Sarmento. 2003. “Evaluating Core Inflation Indicators.” *Economic Modelling* 20 (4): 765–75.

O’Brien, D. 2018. “The Supercore Measure of Underlying Inflation.” *Economic Bulletin* (European Central Bank) 4/2018.

Pincheira-Brown, P., J. Selaive, and J. Nolazco. 2019. “Forecasting Inflation in Latin America with Core Measures.” *International Journal of Forecasting* 35 (3): 1060–71.

Rich, R. W., and C. Steindel. 2007. “A Comparison of Measures of Core Inflation.” *Economic Policy Review* (Federal Reserve Bank of New York) 13 (3): 19–38.

Roger, S. 1997. “A Robust Measure of Core Inflation in New Zealand, 1949–96.” Working Paper No. G97/7, Reserve Bank of New Zealand.

Stock, J., and M. Watson. 1998. "Diffusion Indexes." Working Paper No. 6702, National Bureau of Economic Research.

Wynne, M. A. 2008. "Core Inflation: A Review of Some Conceptual Issues." *Review* (Federal Reserve Bank of St. Louis) 90 (3): 205–28.

Zou, H., and T. Hastie. 2005. "Regularization and Variable Selection via the Elastic Net." *Journal of the Royal Statistical Society Series B: Statistical Methodology* 67 (2): 301–20.