What can we learn from the distributions of inflation expectations across European households?*

Andros Kourtellos† Christos Antonios Statheas‡ Marios Zachariadis§

December 6, 2025

Abstract

We study inflation dynamics using a large dataset of household inflation expectations for eleven euro-area economies from 2004:1 to 2023:11. Using a functional data approach, we characterise the full cross-sectional distribution of beliefs and show that its functional components significantly influence realized inflation. This highlights a role for the entire expectations distribution in the Phillips Curve, beyond what is implied by models based solely on consensus expectations. We document substantial heterogeneity across countries and find that the influence of these distributional components differs between low- and high-inflation environments. Moreover, empirical specifications that incorporate both past expectations of current inflation and current expectations of future inflation outperform models relying only on forward-looking expectations, particularly in turbulent periods. These findings suggest that macroeconomic models with heterogeneous agents and rational inattention provide a promising foundation for capturing an empirically relevant expectations-inflation relationship.

Keywords: inflation, inflation expectations, functional regression
JEL Classification: E31, D84

*We thank Sophocles Mavroeidis, Nicoletta Pashourtidou and Iacovos Sterghides for useful comments and suggestions, and Francesca Monti for providing us with the codes for functional regression.

†Department of Economics, University of Cyprus, 1678 Nicosia, Cyprus. E-mail: kourtellos.andros@ucy.ac.cy

‡Department of Economics, University of Cyprus, 1678 Nicosia, Cyprus. E-mail: statheas.christos-antonios@ucy.ac.cy@ucy.ac.cy

§Department of Economics, University of Cyprus, 1678 Nicosia, Cyprus. E-mail: zachariadis@ucy.ac.cy

1 Introduction

The relationship between inflation realizations and inflation expectations is central to the New Keynesian Phillips curve and is of independent interest for both policy and theory alike. Survey data provide a direct measure of expectations (e.g., Coibion and Gorodnichenko (2012)) while imposing minimal structure on expectation formation. A central empirical challenge is how to summarize these expectations (Mavroeidis et al. (2014)), given the substantial heterogeneity across households and firms, which simple measures of central tendency fail to capture. Although dispersion measures partially address this issue (Fofana et al. (2024)), other features of the cross-sectional distribution of inflation expectations could also be important but have remained relatively unexplored.

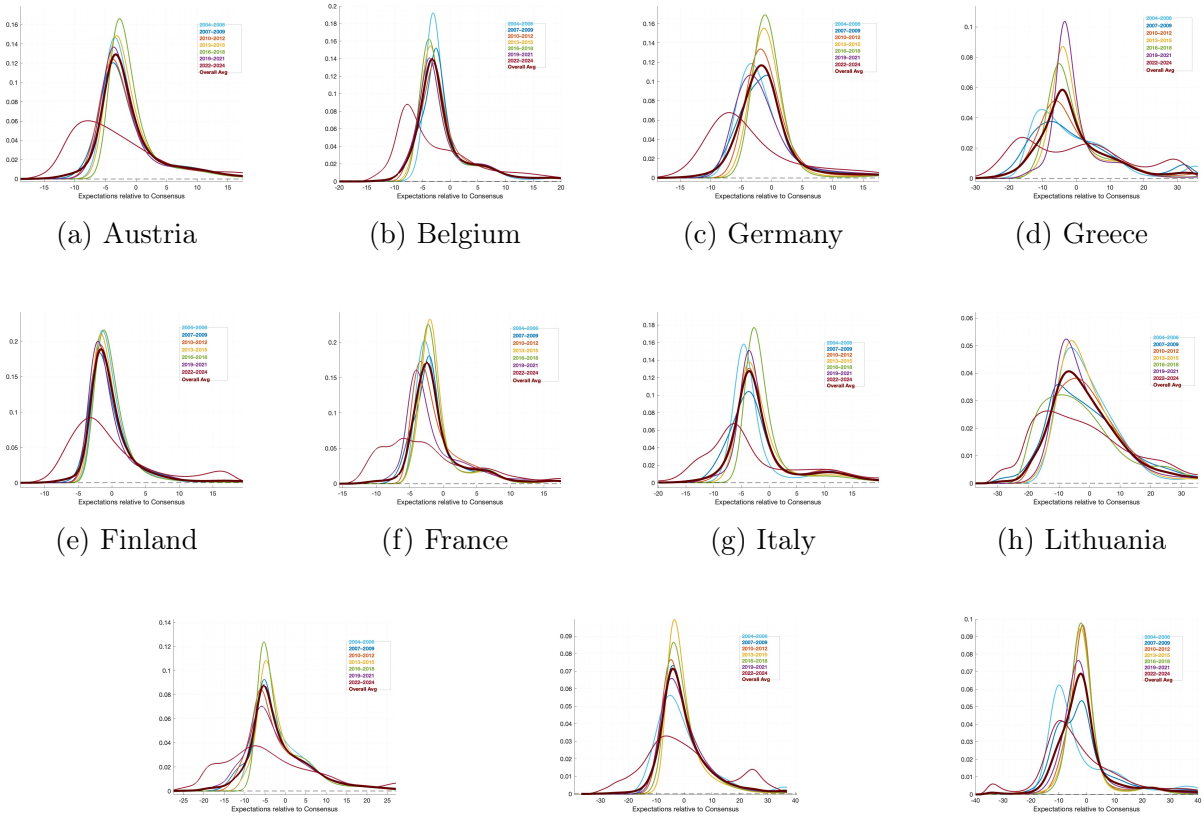
In response to these limitations, Meeks and Monti (2023) propose a functional linear regression approach to address the aggregation problem of incorporating individual survey forecasts into a macroeconomic model of price-setting behaviour. Rather than relying on a few numerical moments, their method maps the entire distribution of expectations at each point in time to realized inflation, capturing how changes in the distribution of beliefs translate into inflation dynamics when studying the effects of various shocks on inflation expectations.[1]

The above paper focuses on the US economy, leaving open two important questions: first, whether the empirical finding that the entire distribution of expectations, beyond its mean, matters for actual inflation represents a form of American exceptionalism; and second, whether, beyond within-country heterogeneity across households, there also exists substantial cross-country heterogeneity in belief dynamics that could influence inflation differently across economies. To the best of our knowledge, no work has systematically answered these questions. In Figure 1, we present a visualization of the degree of cross-country heterogeneity in the distributions of inflation expectations across households.

We pursue our analysis using the quantitative questions on price expectations from the Business and Consumer Survey (BCS) of the European Commission, which provide household-level data at a monthly frequency for the period 2004:1–2023:11 across euro-area economies. We build on Meeks and Monti (2023) by extending their functional linear regression framework to allow for an unknown structural break, capturing potential regime shifts in the relationship between inflation and expectations. The estimation proceeds in two steps: first, we extract the leading functional principal components (FPCs) summarizing the cross-sectional distributions of expectations; and second, we estimate a functional linear regression of realized inflation on these FPC scores, the consensus expectation, and standard New Keynesian

[1] Close in spirit, Chang et al. (2022) highlight the importance of accounting for the entire distribution.

Figure 1: Distributions of Inflation Expectations



(a) Austria (b) Belgium (c) Germany (d) Greece

(e) Finland (f) France (g) Italy (h) Lithuania

(i) Slovenia (j) Slovakia (k) Spain

This figure displays the three-year average densities of household inflation expectations for each country, together with the overall sample-average density shown in dark red. All densities are estimated using the penalized maximum-likelihood procedure described in Section 3.

Phillips Curve controls, allowing all coefficients to differ between the low- and high-inflation regimes using an exogenously imposed break that reflects the onset of the post-2020 high-inflation period.

We find that the cross-sectional distribution of inflation expectations varies over time within each country and across the euro area, and that this variation can be effectively summarized by a small number of functional factors. We then examine the relevance of these factors for inflation realizations and show that they exert a significant influence across euro-area economies.

These findings indicate that the U.S. experience is not unique: similar distributional mechanisms operate across other economies, though with notable cross-country differences in strength and timing. This points to a role for the distribution of inflation expectations in the Phillips curve relationship, beyond the role of consensus expectations implied by standard macroeconomic models that abstract from such heterogeneity. Moreover, the inflationary impact of these functional components differs between high- and low-inflation environments.

Importantly, empirical models that account for the distribution of past expectations of current inflation, alongside the distribution of current expectations of future inflation, perform substantially better during the period under study than models including only the forward-looking component, especially in turbulent times. This suggests that rational inattention (Sims (2003), Mackowiak and Wiederholt (2009)) with heterogeneity could provide a useful foundation for building models that can generate empirically relevant Phillips curve dynamics.

The next section presents our data and preliminary analysis. We then introduce the functional regression framework applied in our paper and present evidence of cross-country heterogeneity in belief dynamics. Subsequently, we conduct the functional regression analysis to investigate the role of the cross-sectional distribution of householdsâ inflation expectations in determining inflation across euro-area economies. Finally, we examine the drivers of cross-country heterogeneity in belief dynamics before concluding.

2 Data

2.1 Data description

We use household-level data from the Business and Consumer Survey (BCS) of the European Commission, available monthly from January 2004 to December 2023 for eleven euroâarea economies: Austria (AT), Belgium (BE), Germany (DE), Greece (EL), Finland (FI), France

(FR), Italy (IT), Lithuania (LT), Slovenia (SI), Slovakia (SK), and Spain (ES). The surveys are nationally representative and capture demographic and socioeconomic heterogeneity, providing a rich basis for analysing the link between inflation expectations and inflation. For most countries, data are collected through computer-assisted telephone interviews.[2]

The BCS includes both qualitative and quantitative questions on households perceptions and expectations of prices. Our analysis focuses on the quantitative items reporting expected and perceived price changes over the next and past 12 months. Each monthly cross-section contains between 700 and 2,000 observations per country. Further details on variable construction, winsorization, and treatment of missing values can be found in Section A.2 of the Online Appendix.

Aggregate variables are measured at monthly frequency. Inflation is computed as the annualised month-on-month percentage change in the seasonally adjusted Harmonised Index of Consumer Prices (HICP) from the ECB. Unemployment rates are obtained from Eurostat; the unemployment gap equals the deviation of unemployment from its HP-filtered trend ($\lambda = 14{,}400$). Further information on seasonal adjustment procedures and data sources is reported in Online Appendix B.

To account for global price pressures, we include the monthly percentage change in Brent crude oil prices, its first and second lags, and the Global Supply Chain Pressure Index (GSCPI) of Benigno et al. (2022). These capture both direct cost effects and broader supply-side shocks that may influence inflation expectations and realised inflation.

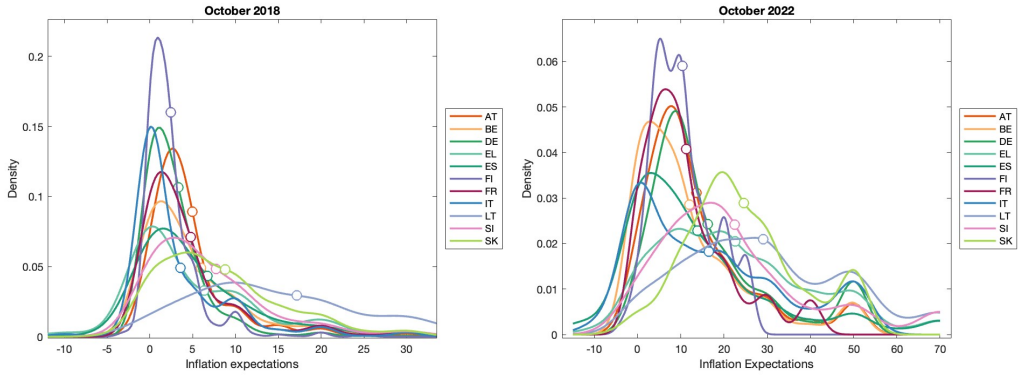
In our robustness analysis, we consider three additional sources of information. First, we compute higher moments of the cross-sectional distribution of inflation expectations (its standard deviation, skewness, and kurtosis) to capture changes in dispersion and asymmetry that may convey information beyond the mean. Second, we incorporate households' perceived inflation, obtained from the quantitative survey question asking respondents to report their assessment of current inflation; these responses are transformed into smooth densities using the same procedure applied to inflation expectations. Third, we construct lagged expectations by shifting the one-year-ahead expectations series twelve months backward, allowing us to investigate whether past beliefs contain predictive content for current inflation. Summary statistics and additional descriptive figures are presented in Online Appendix G and Appendix H.

[2]See Duca-Radu et al. (2021) and Geiger et al. (2025) for other applications of these data. Full question wording and additional details on the survey design are provided in Online Appendix A.

2.2 The high-inflation period

As mentioned earlier, the available sample spans twenty years across euro-area countries. Following the expansionary fiscal and monetary policies in response to the COVID-19 pandemic and the Russian invasion of Ukraine, inflation in the euro area surged, marking a transition from a period of low and stable inflation to one of high price growth. This shift raised concerns about a potential de-anchoring of inflation expectations.

Figure 2: Cross-Country Heterogeneity in Slack and Expectations Channels



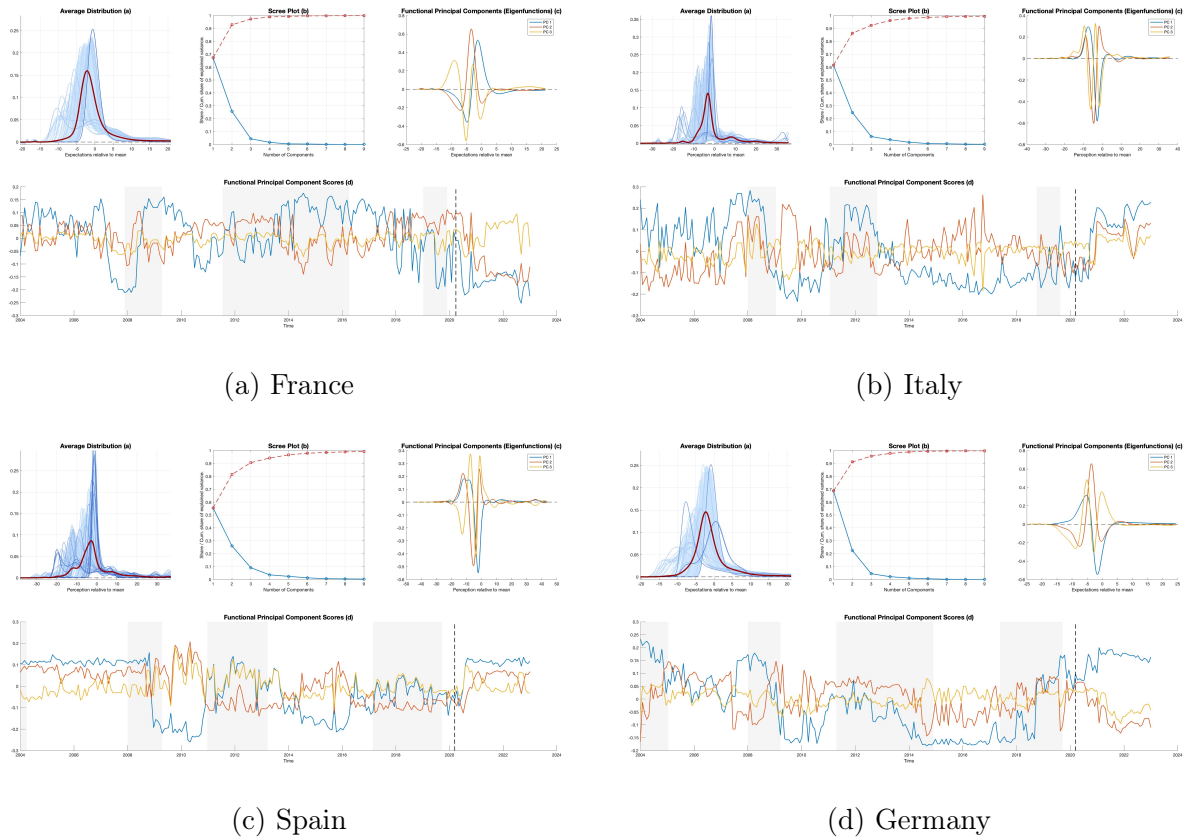
Distributions of inflation expectations using kernel density estimation with Silverman's "rule of thumb" bandwidth. The plot on the left illustrates cross-country differences in inflation expectations for October 2018, while the plot on the right shows the corresponding distributions for October 2022.

Figure 2 shows the estimated distributions of inflation expectations across euro-area economies before and after the inflation surge. The densities, estimated using kernel methods, reveal a pronounced widening and rightward shift of the distributions after 2020. During the low-inflation environment (e.g., as of October 2018), most households expected stable or mildly rising prices and cross-country differences were limited. By contrast, four years later, at the peak of the recent high-inflation episode in October 2022, the distributions were substantially more dispersed and heterogeneous across countries.

Summary statistics and additional descriptive evidence are reported in Online Appendix A (Table A1 and Figure A1). After 2020, the distributions became more dispersed, less symmetric, and markedly leptokurtic, indicating heavier tails and greater disagreement among households. The largest increases in dispersion and skewness occurred in Greece, Slovenia, Slovakia, and Lithuania.

Overall, these descriptive patterns indicate that the recent period of high inflation was accompanied by a sharp rise in cross-country heterogeneity of expectations, consistent with the notion that macroeconomic uncertainty amplifies dispersion in beliefs.

Figure 3: Functional Principal Component Analysis



(a) France (b) Italy

(c) Spain (d) Germany

This figure presents the FPCA analysis for four euro-area economies: France (panel (a)), Italy (panel (b)), Spain (panel (c)), and Germany (panel (d)). For each country, the top-left subplot displays the mean-centred distributions of inflation expectations over time; the dark red curve denotes the time-average distribution, while darker blue curves indicate larger deviations from this average shape. The top-middle subplot reports the proportion of total functional variation explained by each principal component, together with the cumulative variance share. The top-right subplot shows the first three eigenfunctions. The bottom subplot presents the corresponding time-varying principal component scores, along with recession periods (shaded) and the exogenous inflation-break date (vertical line).

Focusing on the four major euro area economies in Figures 3a-3d, we observe substantial time heterogeneity in the cross-sectional distribution of inflation expectations relative to its sample mean.[3] In the first panel, the mean-centered distributions over time (blue) vary markedly from their time average (dark red curve). In particular, the dark blue curves correspond to periods with low band depth in the sense of Lopez-Pintado and Romo (2009), indicating that these distributions lie far from the central region of the functional sample. Their position outside the central bands highlights pronounced deviations from typical distributional behavior and underscores the strong time variation in cross-sectional belief heterogeneity. Moreover, the noticeable cross-country differences in the time averages of these distributions reinforce the importance of cross-country heterogeneity in belief dynam-

[3]Comprehensive FPCA results for all countries are provided in Online Appendix-Section G

ics. The second panel shows that FPCA provides an effective way to summarize the cross-sectional heterogeneity of these distributions using a finite number of functional principal components. A larger number of components naturally yields a closer approximation to the underlying curves. For the selected countries, the first (dominant) component explains roughly 60 percent of the total variation in the data, while the first three components jointly account for approximately 90 percent. The shapes of these three leading components are presented in panel (c), and their time-varying scores in panel (d), shown alongside recession periods and the exogenously identified break associated with the high-inflation episode.

3 Functional Phillips Curve with Regime Dependence

For each month $t = 1, \ldots, T$, let $\{\pi^e_{i,t}\}_{i=1}^{N_t}$ denote the cross-section of one-year-ahead household point forecasts of inflation. We interpret these realizations as draws from an underlying cross-sectional probability density of beliefs $f_t(x)$ on a compact domain $\mathcal{X} \subset \mathbb{R}$. As documented in Section 2, the distributions widen and shift rightward after 2020, consistent with a transition from a low- to a high-inflation regime.

Motivated by this structural shift, we impose an exogenous break in 2020 allowing regime dependence in three components: (i) the intercept, (ii) the consensus (mean) expectation, and (iii) the functional distributional term.[4] Define the post-2020 indicator $d_t \equiv \mathbb{I}\{t \geq 2020{:}12\}$. Let π_t denote realized inflation, $\bar{\pi}^e_t = N_t^{-1} \sum_{i=1}^{N_t} \pi^e_{i,t}$ the consensus expectation, and $Z_t \in \mathbb{R}^q$ a vector of controls. The latter includes the unemployment gap, the current monthly percentage change in Brent oil prices and its first two lags, the contemporaneous GSCPI, and seven lags of inflation. Each specification additionally incorporates country-specific outlier dummies and up to two level-shift dummies capturing structural breaks in the inflation series, as documented in Online Appendix Table A2.[5]

The model is specified as

$$\begin{aligned} \pi_t = \alpha + \beta^e \, \bar{\pi}^e_t + \int_{\mathcal{U}} \gamma^e(u) \, f_t\big(u - \bar{\pi}^e_t\big) \, du + Z_t' \kappa \\ + \left(\alpha_\Delta + \beta^e_\Delta \, \bar{\pi}^e_t + \int_{\mathcal{U}} \gamma^e_\Delta(u) \, f_t\big(u - \bar{\pi}^e_t\big) \, du \right) d_t + e_t, \end{aligned} \tag{1}$$

where $\gamma^e, \gamma^e_\Delta \in L^2(\mathcal{U})$ are the pre- and post-2020 slope functions, respectively. The coefficients κ on the control vector Z_t are restricted to be common across regimes to maintain a

[4] Our choice of break dates is consistent with the structuralâbreak evidence reported in Online Appendix C. There, we estimate break points using the dataâdriven algorithm of Bai and Perron (1998), which confirms the presence of a major break around late 2020 for most countries.

[5] Seven lags are imposed uniformly across countries and eliminate serial autocorrelation in all specifications. This control set is used throughout the paper.

parsimonious specification focused on shifts in expectations and their cross-sectional shape.

The disturbance sequence $\{e_t\}$ is assumed to form a martingale difference sequence (MDS) with respect to the information set $\mathcal{F}_{t-1}$, $\mathbb{E}[e_t \mid \mathcal{F}_{t-1}] = 0$ and $\sup_t \mathbb{E}(|e_t|^{2+\delta}) < \infty$, allowing for weak serial dependence and conditional heteroskedasticity.[6] In addition, the joint process $\{(\pi_t, Z_t, f_t)\}$ is assumed to be piecewise strictly stationary and ergodic within each regime, with all roots of the inflation lag polynomial lying outside the unit circle. We estimate this model in steps.

Step 1: From discrete point forecasts to functional objects. We transform the discrete survey responses into smooth functional objects that represent each month's cross-sectional distribution of inflation expectations. Following Meeks and Monti (2023), we estimate a smooth density $\widehat{f}_t$ for the mean-centred responses using a penalised maximum-likelihood (pML) estimator.[7] The estimator imposes smoothness by penalising the integrated squared third derivative of the density:

$$\widehat{f}_t = \arg\max_{f \in \mathcal{F}} \left\{ \sum_{i=1}^{N_t} \log f(\pi_{i,t}^e - \bar{\pi}_t^e) - \lambda\, J(f) \right\}, \qquad J(f) = \int_{\mathcal{U}} \left[f^{(3)}(x) \right]^2 dx, \tag{2}$$

where $\lambda > 0$ is the smoothing parameter and $\mathcal{F}$ denotes the class of sufficiently smooth densities. Penalising the third derivative stabilises the estimator and prevents overfitting to sampling noise while allowing for flexible shapes, including multimodality when supported by the data.[8] The resulting sequence $\{\widehat{f}_t\}_{t=1}^T$ provides a smooth functional time-series representation of the evolving cross-sectional distribution of inflation expectations.

Step 2: Functional principal component representation. To obtain a finite-dimensional representation of the functional objects $\{\widehat{f}_t\}_{t=1}^T$, we apply functional principal component analysis (FPCA). Let the sample mean function be $\bar{f}(u) = T^{-1} \sum_{t=1}^T \widehat{f}_t(u)$, and define the associated covariance operator $\mathcal{C}$ by $\mathcal{C}(v) = T^{-1} \sum_{t=1}^T \langle \widehat{f}_t - \bar{f},\, v \rangle (\widehat{f}_t - \bar{f})$ for any $v \in L^2(\mathcal{U})$, where $\langle g, h \rangle = \int_{\mathcal{U}} g(u)\, h(u)\, du$ denotes the usual inner product on $L^2(\mathcal{U})$.

The operator $\mathcal{C}$ is compact, self-adjoint, and positive semi-definite, so FPCA is obtained from its spectral decomposition. Specifically, the eigenfunctions $\{\phi_k(u)\}_{k \geq 1}$ and eigenvalues

[6] Throughout all specifications in equations (1)-(6), the regression disturbance u_t is assumed to be a martingale difference sequence with respect to the information set at time t, so that all predictable components of inflation are captured by the included covariates.

[7] For each survey month, the domain bounds are set to the sample-wide minimum and maximum of observed responses, expanded slightly to avoid truncation where probability mass may be non-negligible.

[8] As in Meeks and Monti (2023), the smoothing parameter is set according to $\lambda = 10^{(M/2 - N_t)/10}$, where M is the number of basis functions and N_t is the cross-sectional sample size in month t.The number of basis functions was set proportional to the range of the data.

$\{\lambda_k\}_{k\geq 1}$ solve $\mathcal{C}(\phi_k) = \lambda_k \phi_k$ with orthonormality $\langle \phi_k, \phi_{k'} \rangle = \delta_{kk'}$. The eigenfunctions provide the dominant modes of variation in the cross-sectional distributions, and the corresponding scores are $\xi_{k,t} = \langle \widehat{f}_t - \bar{f},\, \phi_k \rangle$.

Truncating the expansion to the first K principal components yields the scalar-on-function approximation used in estimation:

$$\pi_t = \alpha + \beta^e\, \bar{\pi}_t^{\,e} + \sum_{k=1}^{K} \gamma_k^e\, \xi_{k,t} + Z_t'\kappa + \alpha_\Delta d_t + \beta_\Delta^e\, \bar{\pi}_t^{\,e} d_t + \sum_{k=1}^{K} \gamma_{\Delta k}^e\, \xi_{k,t} d_t + u_t, \tag{3}$$

where the composite error is $u_t = e_t + r_t$ and $r_t = \sum_{k>K} \gamma_k \xi_{k,t}$ denotes the truncation remainder.

When the retained components explain a high proportion of total functional variation (typically exceeding 95%), the truncation error r_t is negligible, and (3) provides a consistent and tractable representation of the underlying scalar-on-function model with regime-dependent coefficients. Estimation of (3) proceeds by ordinary least squares, with inference based on Newey–West heteroskedasticity- and autocorrelation-consistent standard errors.

4 Empirical Implementation

4.1 Baseline analysis

We now estimate the heterogeneous-beliefs Phillips curve using the functional regression framework described above. Our baseline empirical specification is the break model in Equation (1), which allows the intercept, the coefficient on average expectations, and the functional slope $\gamma(u)$ to differ between the low- and high-inflation regimes. Before turning to the main results, we also estimate a restricted model with constant coefficients, obtained by imposing $\alpha_\Delta = \beta_\Delta^e = \gamma_\Delta^e(u) = 0$ in Equation (1). This simpler specification serves as a benchmark for evaluating the empirical importance of allowing for regime dependence and provides a direct comparison to the functional Phillips curve of Meeks and Monti (2023), which assumes stability in the expectationsâinflation relationship. Instead of presenting the coefficients of the functional principal components scores we provide a summary measure in the form of the studentized range statistic that measures how much the effect varies across the distribution relative to its overall variability.[9] For all models and countries, we include

[9]The studentized range of an estimated slope function $\gamma(u)$ is defined as $\mathrm{SR}(\gamma) = \left[\max_{u\in\mathcal{U}} \widehat{\gamma}(u) - \min_{u\in\mathcal{U}} \widehat{\gamma}(u)\right] \Big/ \sqrt{|\mathcal{U}|^{-1} \int_{\mathcal{U}} \left(\widehat{\gamma}(u) - \bar{\gamma}\right)^2 du}, \qquad \bar{\gamma} = |\mathcal{U}|^{-1} \int_{\mathcal{U}} \widehat{\gamma}(u)\, du$. The denominator is the L^2 standard deviation of the functional coefficient over its support.

the full set of controls and lagged inflation terms, chosen to ensure that the residuals are not serially correlated.[10]

Table 1, Panel A reports the results from the constant-coefficient specification. Although this model restricts all parameters to be stable over the full sample, several notable patterns emerge. The coefficient on average expectations is positive and statistically significant in all economies, consistent with the notion that higher consensus inflation expectations tend to be associated with higher realized inflation. The unemployment-gap coefficient is generally negative, as standard slack-based Phillips curve mechanisms would predict, but its magnitude varies widely across countries, reflecting substantial heterogeneity in the role of labor-market tightness. The estimated sum of lag coefficients, which measures inflation persistence in this specification, also differs materially across the euro area: some countries (such as Austria, Lithuania, and Slovakia) exhibit significant persistence, whereas others display near-zero persistence.

Importantly, the joint Wald tests for the relevance of the functional principal components reject the null $H_0 : \gamma^e(u) = 0$ in all euro-area economies except Belgium and Finland. In all other countries the rejection occurs at the 1% significance level. This result aligns closely with Meeks and Monti (2023), who show that the entire distribution of inflation expectations, not merely its mean, contains independent information for forecasting inflation. Distributional features such as skewness, dispersion, and tail mass capture elements of belief heterogeneity that standard Phillips curve models abstract from. Consistent with this interpretation, the studentized-range statistics reveal sizeable cross-country differences in the magnitude of $\gamma^e(u)$, indicating that heterogeneity in household beliefs plays an uneven and economically meaningful role across the euro area.

Table 1, Panel B presents the central empirical results of the paper. The break specification in Equation (1) allows the intercept, the coefficient on average expectations, and the functional slope $\gamma^e(u)$ to vary between the low- and high-inflation regimes, providing a direct test of whether the expectations-inflation relationship underwent a structural change after 2020. Interestingly, the joint Wald tests strongly reject the null $H_0 : \alpha_\Delta = \beta^e_\Delta = \gamma^e_\Delta(u) = 0$ for nearly all countries (weakly so for Austria and France, with p-values of 0.03 and 0.071 respectively), offering clear evidence that the link between household expectations and realized inflation shifts materially in the post-2020 period. In most of these economies, the coefficient on average expectations becomes larger in the high-inflation regime, indicating that inflation responds more strongly to household beliefs when inflation is elevated. This heightened sensitivity is consistent with the interpretation in Meeks and Monti (2023) that

[10]For each country, we control for a set of outlier dummies to account for extreme observations in inflation, as well as level-shift dummies to capture structural changes in the inflation intercept. Full details on the construction and treatment of these dummies are provided in Appendix C.

greater disagreement and a larger mass of high-expectation households, amplify the transmission of expectations into inflation.

Table 1: Heterogeneous-Beliefs Phillips Curve

	AT	BE	DE	EL	FI	FR	IT	LT	SI	SK	ES
Panel A: Without Breaks											
Unemployment gap	-1.282***	-0.225	-1.050	-0.636	-0.947**	-0.694	-0.497*	0.244	-1.068*	-0.680**	-0.291
	(0.333)	(0.876)	(1.273)	(0.435)	(0.452)	(0.595)	(0.284)	(0.345)	(0.545)	(0.295)	(0.193)
Average expectations	2.296***	3.208**	0.979**	0.280*	1.303*	1.322***	0.425**	0.743***	0.897***	0.930***	1.079*
	(0.695)	(1.412)	(0.448)	(0.142)	(0.689)	(0.178)	(0.213)	(0.196)	(0.260)	(0.230)	(0.582)
Sum of lag coefficients	-0.795**	-0.386	-0.021	0.302*	0.026	0.034	0.156	0.362***	0.125	0.287*	0.115
	[0.027]	[0.283]	[0.94]	[0.062]	[0.918]	[0.849]	[0.191]	[0.009]	[0.597]	[0.07]	[0.438]
Testing for distributional effects											
$H_0 : \gamma^e(u) = 0$	52.839	4.75	25.71	22.757	3.884	30.81	16.512	22.477	28.218	33.098	29.873
	[0.000]	[0.191]	[0.000]	[0.004]	[0.422]	[0.000]	[0.002]	[0.013]	[0.000]	[0.000]	[0.000]
Studentized range of $\gamma(u)$	6.246	8.163	7.787	7.274	5.038	6.961	7.541	6.953	7.479	6.248	6.749
Controls	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
R^2_{Adj}	0.536	0.501	0.603	0.515	0.537	0.548	0.710	0.722	0.514	0.617	0.468
Ljung–Box	[0.892]	[0.117]	[0.580]	[0.957]	[0.794]	[0.736]	[0.457]	[0.221]	[0.728]	[0.784]	[0.540]
N Obs	239	239	239	232	239	239	239	239	239	238	239
Panel B: With Breaks											
Unemployment gap	-1.007***	0.001	-1.504	-0.179	-0.967*	-0.718	-0.812**	0.165	-1.276**	-0.696**	-0.416**
	(0.325)	(0.762)	(1.223)	(0.495)	(0.506)	(0.641)	(0.340)	(0.397)	(0.552)	(0.269)	(0.164)
Average expectations	1.984***	0.909	0.301	0.161	0.849*	1.578**	-0.122	0.957***	1.190***	0.506*	0.547**
	(0.540)	(1.053)	(0.295)	(0.119)	(0.453)	(0.689)	(0.525)	(0.178)	(0.400)	(0.278)	(0.270)
Average expectations × high-inflation dummy	3.552***	2.883*	2.895**	1.916*	3.705***	0.125	-0.440	1.377***	1.534	1.151***	1.192
	(1.230)	(1.638)	(1.138)	(0.972)	(1.111)	(0.686)	(0.811)	(0.431)	(0.992)	(0.395)	(0.772)
High-inflation dummy	-19.877***	2.367	-25.753***	-16.499*	-13.115***	-2.237	2.962	-19.449***	-10.545	-8.168**	-10.668**
	(6.922)	(11.511)	(4.814)	(8.955)	(4.393)	(3.120)	(5.164)	(7.137)	(14.595)	(3.484)	(4.591)
Sum of lag coefficients	-0.715**	-0.362	-0.318	0.225	-0.158	0.010	0.127	-0.092	0.192	0.090	0.116
	[0.016]	[0.221]	[0.570]	[0.330]	[0.379]	[0.646]	[0.409]	[0.835]	[0.927]	[0.413]	[0.959]
Testing for breaks											
$H_0 : \alpha_\Delta = \beta^e_\Delta = \gamma^e_\Delta = 0$	[0.030]	[0.000]	[0.000]	[0.000]	[0.000]	[0.071]	[0.000]	[0.000]	[0.001]	[0.000]	[0.000]
$H_0 : \alpha_\Delta = \beta_\Delta = 0$	[0.014]	[0.000]	[0.000]	[0.142]	[0.004]	[0.364]	[0.846]	[0.001]	[0.085]	[0.014]	[0.066]
$H_0 : \gamma^e_\Delta = 0$	[0.014]	[0.000]	[0.001]	[0.032]	[0.000]	[0.096]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
Testing for distributional effects											
$H_0 : \gamma^e(u) = \gamma^e_\Delta(u) = 0$	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.002]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
$H_0 : \gamma^e(u) = 0$	[0.001]	[0.600]	[0.000]	[0.000]	[0.181]	[0.204]	[0.426]	[0.000]	[0.000]	[0.056]	[0.000]
Studentized range of $\gamma^e(u)$	5.806	8.822	6.536	6.691	5.738	6.968	7.453	5.904	5.865	5.67	7.234
Studentized range of $\gamma^e_\Delta(u)$	6.319	8.599	7.844	6.67	5.243	6.437	6.953	5.886	9.727	6.373	6.85
Controls	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
R^2_{Adj}	0.571	0.554	0.635	0.540	0.577	0.549	0.718	0.764	0.530	0.675	0.579
Ljung–Box	[0.534]	[0.234]	[0.462]	[0.766]	[0.776]	[0.764]	[0.544]	[0.225]	[0.753]	[0.902]	[0.193]
N	239	239	239	232	239	239	239	239	239	238	239

This table reports estimates of the heterogeneous-beliefs Phillips curve in equation (1). Panel A presents results for the constrained specification with coefficients held constant across regimes, while Panel B reports estimates allowing for an exogenous post-2020 break in the coefficients of equation (1). Newey–West (13-lag) standard errors appear in round parentheses, and p-values for the corresponding tests appear in square brackets. We also report p-values for break tests and for tests of distributional effects. Significance levels: *, **, and *** denote the 10%, 5%, and 1% levels, respectively. Studentized ranges summarize the variation of $\gamma(u)$ and $\gamma_\Delta(u)$ across their support.

The individual tests for $H_0 : \gamma^e_\Delta(u) = 0$ reject the null in all countries (in Austria, Greece and France, weakly so, with p-values at 0.014, 0.032 and 0.96 respectively), providing com-

pelling evidence that the shape of the cross-sectional distribution of expectations plays a different role in the post-2020 environment. The studentized-range statistics reveal pronounced changes in $\gamma^e_\Delta(u)$ relative to the pre-2020 slope functions, suggesting that the regime shift is accompanied by distributional transformations, such as increased dispersion, right-skewness, or heavier upper tails, that are empirically important for inflation dynamics. These features are precisely those emphasized by Meeks and Monti (2023) as amplifying the influence of high-expectation households. Again, the test for the relevance of the functional principal components rejects the null $H_0 : \gamma^e(u) = 0$ in most of these euro-area economies, suggesting that the entire distribution of inflation expectations, not merely its mean, contains independent information for forecasting inflation.

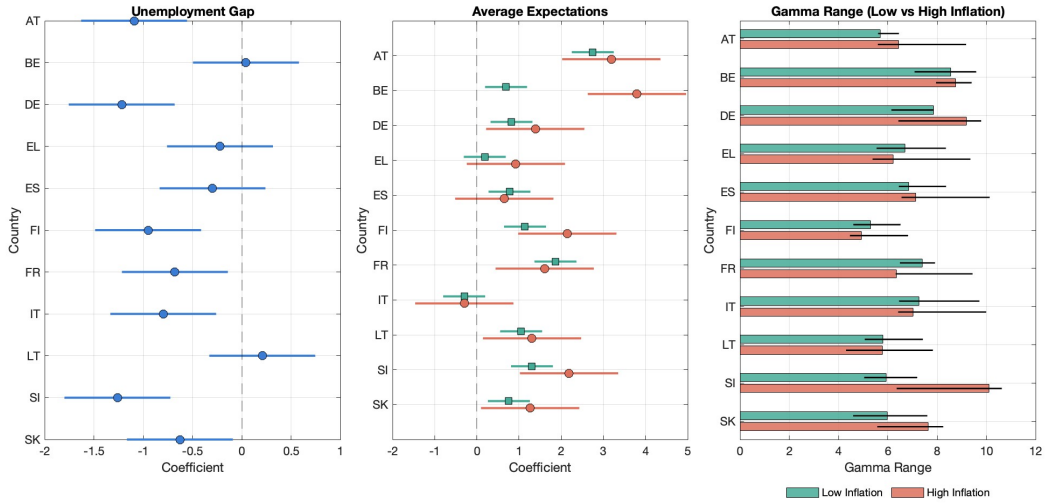
Finally, the estimated intercept adjustments α_Δ are also large and statistically significant in several countries, indicating a discrete upward shift in inflation levels after 2020 that cannot be fully attributed to changes in expectations or their distributional properties. This suggests that the post-2020 inflation regime is characterized not only by altered expectation dynamics but also by a structural shift in the underlying inflation process itself.

Furthermore, the heterogeneity in these effects becomes even clearer in Figure 6, which summarizes cross-country variation in three components: the unemployment-gap coefficient, the sensitivity to average expectations, and the studentized range of the functional slope. Three patterns emerge.

First, slack effects vary widely across countries, strong in some economies but weak or insignificant in others. Second, the response to average expectations increases almost uniformly in the high-inflation regime, reflected in the right-shift of the red confidence intervals and consistent with the Wald-test evidence that $\beta^e_\Delta > 0$ in most cases. Third, the studentized-range statistic widens sharply after 2020: red bars exceed green bars for most countries indicating that distributional heterogeneity becomes more influential in high-inflation periods. This pattern aligns with Meeks and Monti (2023), who emphasize the role of the upper tail of inflation expectations. Overall, the graphical evidence reinforces the formal tests: heterogeneity in expectations, both in mean and shape, gains importance in the high-inflation environment, with substantial variation across euro-area economies.

Taken together, these results underscore the importance of modelling regime dependence: the break specification captures substantial structural changes in the inflation process that the constant-coefficient model in Panel A necessarily masks.

Figure 4: Phillips Curve Coefficient Plots Across Euro-Area Economies



This figure reports the estimated coefficients from the regime-dependent functional Phillips curve in equation (1). The panels show point estimates and 95% wild-bootstrap confidence intervals (20,000 replications) for the unemployment-gap coefficient, the coefficient on mean expectations, and the studentized range of the functional slope. Green bars correspond to the low-inflation regime and red bars to the high-inflation regime. Black lines denote the associated confidence intervals.

4.2 Robustness and Extensions

4.2.1 Controlling for Higher Moments of Inflation Expectations

A natural question is whether the predictive power of the functional principal components reflects information beyond what is already summarized by conventional measures of dispersion and asymmetry. To examine this, we augment the break specification in Equation (1) by including the standard deviation, skewness, and kurtosis of inflation expectations as controls, allowing their effects to differ across regimes through interactions with the post-2020 dummy.[11]

Let $M_t = (m_{2t},\, m_{3t},\, m_{4t})' \in \mathbb{R}^3$ denote these three higher moments, and let $\lambda, \lambda_\Delta \in \mathbb{R}^3$ be their corresponding coefficients. The resulting extended model is

$$\begin{aligned}\pi_t = \alpha + \beta^e\, \bar{\pi}_t^e + M_t'\lambda + \int_{\mathcal{U}} \gamma^e(u)\, f_t\big(u - \bar{\pi}_t^e\big)\, du + Z_t'\kappa \\ + \Big(\alpha_\Delta + \beta_\Delta^e\, \bar{\pi}_t^e + M_t'\lambda_\Delta + \int_{\mathcal{U}} \gamma_\Delta^e(u)\, f_t\big(u - \bar{\pi}_t^e\big)\, du\Big)\, d_t + u_t\end{aligned} \tag{4}$$

where u_t is a martingale difference sequence (MDS) with respect to the information set at time t, ensuring that all predictable components of inflation are captured by the included covariates.

[11] With slight abuse of notation, parameter and error symbols are kept the same as in Equation (4) despite the augmented specification.

Table 2 reports a set of p-values for Wald statistics that summarize the implications of this enriched specification. The joint break tests continue to reject the stability of the expectations–inflation relationship in all countries, indicating that the structural shift documented earlier persists even after accounting for movements in higher moments. Thus, the post-2020 regime change is not driven solely by changes in the variance or skewness of beliefs.

The functional distributional tests show that the FPCA-based coefficients remain highly significant across most economies. This holds both for the joint test of functional effects across regimes and for the separate pre-2020 and post-2020 components. In particular, the post-2020 functional coefficients remain strongly rejected, reinforcing the conclusion that the shape of the expectations distribution matters differently during the high-inflation period.

Table 2: Break and Functional Coefficient Tests with Higher Moments Added

	AT	BE	DE	EL	FI	FR	IT	LT	SI	SK	ES
Testing for breaks											
$H_0:\ \alpha_\Delta = \beta_\Delta = \lambda_\Delta = \gamma_\Delta(u) = 0$	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
$H_0:\ \alpha_\Delta = \beta_\Delta = \gamma_\Delta(u) = 0$	[0.000]	[0.000]	[0.002]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.090]	[0.000]	[0.000]
$H_0:\ \alpha_\Delta = \beta_\Delta = 0$	[0.000]	[0.543]	[0.006]	[0.003]	[0.002]	[0.001]	[0.021]	[0.017]	[0.101]	[0.115]	[0.006]
$H_0:\ \gamma_\Delta(u) = 0$	[0.014]	[0.000]	[0.014]	[0.019]	[0.000]	[0.000]	[0.002]	[0.000]	[0.013]	[0.000]	[0.000]
$H_0:\ \alpha_\Delta = 0$	[0.012]	[0.275]	[0.063]	[0.454]	[0.017]	[0.000]	[0.043]	[0.115]	[0.699]	[0.412]	[0.657]
Testing for distributional effects											
$H_0: \gamma^e(u) = \gamma^e_\Delta(u) = 0$	[0.046]	[0.004]	[0.001]	[0.000]	[0.000]	[0.000]	[0.001]	[0.000]	[0.000]	[0.000]	[0.000]
$H_0: \gamma(u) = 0$	[0.609]	[0.231]	[0.005]	[0.000]	[0.137]	[0.032]	[0.936]	[0.000]	[0.000]	[0.003]	[0.438]

This table reports break and functional coefficient tests for the higher-moments specification in equation (4). All entries are p-values reported in square brackets.

Overall, adding higher moments does little to attenuate the predictive content of the functional principal components. This confirms that the FPCA representation captures richer, nonparametric features of belief heterogeneity, such as local tail mass, multimodality, and nonlinear shape changes, that cannot be summarized by dispersion, skewness, or kurtosis. The influence of heterogeneity on inflation dynamics thus reflects features of the expectations distribution that go beyond its first four moments, which the model includes only as additive controls.

4.2.2 Controlling for Perceptions of Inflation

A further robustness question is whether the heterogeneous-beliefs Phillips curve is driven solely by forward-looking expectations or whether householdsâ perceptions of *past* inflation also contain predictive information. Perceived inflation may matter for two reasons. First, if

households anchor their expectations partly on recently experienced price changes, then heterogeneity in perceptions can translate into heterogeneity in forecasting behaviour. Second, perceptions of past inflation may directly influence current price-setting beliefs, independently of their effect on expectations of future inflation. To examine these channels, we augment our benchmark break model by including both the average perceived inflation rate and the functional principal components constructed from the cross-sectional distribution of perceived inflation. As before, all perception-related coefficients are allowed to shift after 2020 through interactions with the high-inflation dummy.

Using the same procedure employed for expectations to transform the discrete survey responses on perceived inflation into smooth density functions, the extended model is given by

$$\begin{aligned}\pi_t = \alpha + \beta^e\,\bar{\pi}_t^{\,e} + \beta^p\,\bar{\pi}_t^{\,p} + \int_{\mathcal{U}} \gamma^e(u)\, f_t(u)\, du + \int_{\mathcal{U}} \gamma^p(u)\, g_t(u)\, du + Z_t'\kappa \\ + \Big(\alpha_\Delta + \beta_\Delta^e\,\bar{\pi}_t^{\,e} + \beta_\Delta^p\,\bar{\pi}_t^{\,p} + \int_{\mathcal{U}} \gamma_\Delta^e(u)\, f_t(u)\, du + \int_{\mathcal{U}} \gamma_\Delta^p(u)\, g_t(u)\, du\Big)\, d_t + u_t,\end{aligned} \tag{5}$$

where $f_t(u)$ and $g_t(u)$ denote the mean-centered densities of expected and perceived inflation, respectively; $\gamma^e(u)$ and $\gamma_\Delta^e(u)$ are the corresponding distributional slope functions for expectations; and $\gamma^p(u)$ and $\gamma_\Delta^p(u)$ are the analogous slope functions for perceived inflation.[12]

The results in Table 3 show that adding households' perceptions of past inflation does not alter the main conclusions from the baseline specification. The break tests reject the joint restriction $\alpha_\Delta = \beta_\Delta^e = \beta_\Delta^p = \gamma_\Delta^e(u) = \gamma_\Delta^p(u) = 0$ in every country at the 1% significance level, indicating that the relationship between inflation, current expectations, and perceived past inflation shifts markedly in the post-2020 period. The sub-tests reported in the table are broadly consistent with this overall rejection, although a few countries exhibit weaker evidence for certain components of the break.

A natural next step is to assess whether distributional heterogeneity, taken jointly across both expectations and perceptions, continues to matter once perceptions are introduced. The joint functional-distribution test $H_0 : \gamma^e(u) = \gamma_\Delta^e(u) = \gamma^p(u) = \gamma_\Delta^p(u) = 0$ is decisively rejected in every country at the 1% significance level. The results therefore reinforce the central conclusion from the baseline analysis that heterogeneity matters not only through the average of beliefs, but also through the full distribution of expectations.

[12] The comprehensive results of the FPCA for all countries are provided in the Online Appendix (Section H).

Table 3: Break and Functional Coefficient Tests with Perceived Inflation

	AT	BE	DE	EL	FI	FR	IT	LT	SI	SK	ES
Testing for breaks											
$H_0:\ \alpha_\Delta = \beta^e_\Delta = \beta^p_\Delta = \gamma^e_\Delta(u) = \gamma^p_\Delta(u) = 0$	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
$H_0:\ \alpha_\Delta = \beta^e_\Delta = \gamma^e_\Delta(u) = 0$	[0.018]	[0.014]	[0.014]	[0.000]	[0.000]	[0.002]	[0.000]	[0.000]	[0.008]	[0.000]	[0.000]
$H_0:\ \alpha_\Delta = \beta^e_\Delta = 0$	[0.127]	[0.011]	[0.021]	[0.000]	[0.008]	[0.423]	[0.000]	[0.000]	[0.963]	[0.745]	[0.036]
$H_0:\ \alpha_\Delta = 0$	[0.052]	[0.849]	[0.018]	[0.000]	[0.002]	[0.663]	[0.000]	[0.000]	[0.810]	[0.680]	[0.488]
Testing for distributional effects (jointly)											
$H_0: \gamma^e(u) = \gamma^e_\Delta(u) = \gamma^p(u) = \gamma^p_\Delta(u) = 0$	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
Testing for inflation expectations											
–jointly											
$H_0:\ \beta^e = \beta^e_\Delta = \gamma^e(u) = \gamma^e_\Delta(u) = 0$	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
–distributional effects only											
$H_0: \gamma^e(u) = \gamma^e_\Delta(u) = 0$	[0.073]	[0.000]	[0.135]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
$H_0: \gamma^e(u) = 0$	[0.143]	[0.511]	[0.042]	[0.002]	[0.017]	[0.039]	[0.582]	[0.018]	[0.000]	[0.199]	[0.059]
Testing for perceived inflation											
–jointly											
$H_0:\ \beta^p = \beta^p_\Delta = \gamma^p(u) = \gamma^p_\Delta(u) = 0$	[0.027]	[0.275]	[0.099]	[0.000]	[0.000]	[0.018]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
–distributional effects only											
$H_0: \gamma^p(u) = \gamma^p_\Delta(u) = 0$	[0.032]	[0.303]	[0.044]	[0.000]	[0.000]	[0.032]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
$H_0: \gamma^p(u) = 0$	[0.570]	[0.286]	[0.048]	[0.199]	[0.029]	[0.264]	[0.984]	[0.018]	[0.236]	[0.560]	[0.611]

This table reports Wald tests for the heterogeneous-beliefs Phillips curve with perceived inflation, corresponding to the specification in equation (5) augmented with perceived inflation terms. Superscripts e and p denote current expectations and perceived past inflation, respectively. Functional coefficients are $\gamma^e(u)$ and $\gamma^p(u)$, with break-adjusted counterparts $\gamma^e_\Delta(u)$ and $\gamma^p_\Delta(u)$. All p-values appear in square brackets.

Turning to current expectations of inflation, the joint tests $H_0 : \beta^e = \beta^e_\Delta = \gamma^e(u) = \gamma^e_\Delta(u) = 0$ are rejected in all eleven countries at the 1% significance level, reaffirming that forward-looking beliefs remain a dominant driver of inflation even after controlling for perceptions of past inflation. The functional-distribution restrictions also show broad rejection, with a few notable exceptions.

The joint test for the role of perceived inflation, $H_0 : \beta^p = \beta^p_\Delta = \gamma^p(u) = \gamma^p_\Delta(u) = 0$, is rejected at the 5% significance level in all countries except Belgium. Thus, in most euro-area economies, perceptions of past inflation have independent explanatory power for current inflation, over and above the contribution of current expectations of future inflation. The functionalâdistribution tests indicate that the restriction $H_0 : \gamma^p(u) = \gamma^p_\Delta(u) = 0$ is rejected at the at the 5% level in all countries except Belgium, whereas the pre-break restriction $H_0 : \gamma^p(u) = 0$ is rejected at the 5% significance level only for Germany, Finland and Lithuania. Taken together, these results suggest that heterogeneity in perceived inflation becomes relevant for inflation realizations mainly in the post-2020 high-inflation regime for most economies.

Overall, these findings indicate that the influence of perceived inflation on current inflation is

strongly regime-dependent, emerging primarily during the post-2020 high-inflation period. Their influence is weak or absent during the low-inflation period, becomes economically relevant primarily through distributional adjustments in the post-2020 regime, and remains quantitatively smaller than the contribution of forward-looking expectations. This pattern reinforces the broader conclusion that, even when perceptions of recent inflation are taken into account, heterogeneity in inflation expectations remains the central channel through which household beliefs shape inflation dynamics in the euro area.

4.2.3 Extending the baseline to include past inflation expectations

Some macroeconomic models, including rational inattention frameworks, imply that current inflation may depend not only on householdsâ current expectations of future inflation but also on their earlier expectations of the inflation rate observed today.[13] Motivated by this insight, we extend the baseline specification to incorporate past expectations of current inflation alongside current expectations of future inflation. Beyond adding mean lagged expectations, we also allow the entire cross-sectional distribution of lagged expectations to enter the model, enabling us to test whether dispersion and asymmetry in past beliefs contain predictive information for inflation and whether such effects differ across inflation regimes.

As before, the discrete survey responses are converted into smooth densities. Let $\bar{\pi}_t^e$ and $f_t(u)$ denote the mean and mean-centered density of current one-year-ahead expectations at time t, and let $\bar{\pi}_t^\ell$ and $f_t^\ell(u)$ denote the corresponding quantities for expectations formed twelve months earlier.

The resulting specification with lagged expectations and a regime break is

$$\begin{aligned}\pi_t = \alpha + \beta^e\,\bar{\pi}_t^e + \int_{\mathcal{U}} \gamma^e(u)\, f_t(u)\, du + \beta^\ell\,\bar{\pi}_t^\ell + \int_{\mathcal{U}} \gamma^\ell(u)\, f_t^\ell(u)\, du + Z_t'\kappa \\ + \Big(\alpha_\Delta + \beta_\Delta^e\,\bar{\pi}_t^e + \beta_\Delta^\ell\,\bar{\pi}_t^\ell + \int_{\mathcal{U}} \gamma_\Delta^e(u)\, f_t(u)\, du + \int_{\mathcal{U}} \gamma_\Delta^\ell(u)\, f_t^\ell(u)\, du\Big)\, d_t + e_t.\end{aligned} \tag{6}$$

In this formulation, $\gamma^e(u)$ and $\gamma_\Delta^e(u)$ capture the distributional effects of current expectations in the low- and high-inflation regimes, while $\gamma^\ell(u)$ and $\gamma_\Delta^\ell(u)$ represent the analogous effects for lagged expectations. As in the previous specifications, the covariates and control terms are chosen so that predictable components of inflation are absorbed within the model. The estimation results for this empirical model are reported in Table 4.

Before turning to the contribution of lagged expectations, it is useful to verify that the inclusion of past inflation expectations leaves the main conclusions from the baseline specification

[13] See Coibion et al. (2018), Table 5.

intact. First, the break tests reject the joint restriction $\alpha_\Delta = \beta^e_\Delta = \beta^\ell_\Delta = \gamma^e_\Delta(u) = \gamma^\ell_\Delta(u) = 0$ at the 1% significance level for all countries, pointing to a clear shift in the relationship between inflation and both current and lagged expectations after 2020. The remaining sub-tests display some country-specific variation but largely support this overall rejection. Second, the test for $H_0 : \gamma^e(u) = \gamma^e_\Delta(u) = \gamma^l(u) = \gamma^l_\Delta(u) = 0$ is rejected in all countries at the 1% significance level, confirming the importance of the full distribution of expectations for capturing the heterogeneity in beliefs beyond average expectations. Third, the joint tests for current expectations, $H_0 : \beta^e = \beta^e_\Delta = \gamma^e(u) = \gamma^e_\Delta(u) = 0$, are rejected for the majority of the countries at the 5% level. This null is not rejected for Belgium, Finland, Italy, Slovakia and Spain. By contrast, the test for the distributional effect $H_0 : \gamma^e(u) = \gamma^e_\Delta(u) = 0$ is rejected for all countries. This implies that, even after controlling for lagged expectations, current one-year-ahead expectations continue to play a central role in inflation dynamics, particularly through the shape of their cross-sectional distribution.

Table 4: Break and Functional Coefficient Tests for Past Inflation Expectations

	AT	BE	DE	EL	FI	FR	IT	LT	SI	SK	ES
Break tests											
$H_0:\ \alpha_\Delta = \beta^e_\Delta = \beta^\ell_\Delta = \gamma^e_\Delta(u) = \gamma^\ell_\Delta(u) = 0$	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
$H_0:\ \alpha_\Delta = \beta^e_\Delta = \gamma^e_\Delta(u) = 0$	[0.000]	[0.000]	[0.000]	[0.000]	[0.002]	[0.001]	[0.000]	[0.000]	[0.027]	[0.000]	[0.000]
$H_0:\ \beta^e_\Delta = \beta^\ell_\Delta = 0$	[0.426]	[0.056]	[0.749]	[0.141]	[0.003]	[0.026]	[0.494]	[0.002]	[0.369]	[0.106]	[0.066]
$H_0:\ \alpha_\Delta = 0$	[0.707]	[0.679]	[0.244]	[0.127]	[0.204]	[0.087]	[0.147]	[0.016]	[0.773]	[0.447]	[0.125]
Testing for distributional effects (jointly)											
$H_0 : \gamma^e(u) = \gamma^e_\Delta(u) = \gamma^l(u) = \gamma^l_\Delta(u) = 0$	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
Testing for current inflation expectations											
–jointly											
$H_0:\ \beta^e = \beta^e_\Delta = \gamma^e(u) = \gamma^e_\Delta(u) = 0$	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
–distributional effects only											
$H_0 : \gamma^e(u) = \gamma^e_\Delta(u) = 0$	[0.000]	[0.000]	[0.000]	[0.000]	[0.033]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
$H_0 : \gamma^e(u) = 0$	[0.001]	[0.381]	[0.001]	[0.000]	[0.125]	[0.088]	[0.161]	[0.000]	[0.000]	[0.422]	[0.003]
Tests for lagged expectations											
–jointly											
$H_0:\ \beta^\ell = \beta^\ell_\Delta = \gamma^\ell(u) = \gamma^\ell_\Delta(u) = 0$	[0.000]	[0.000]	[0.000]	[0.000]	[0.003]	[0.000]	[0.000]	[0.000]	[0.013]	[0.000]	[0.000]
–distributional effects only											
$H_0 : \gamma^\ell(u) = \gamma^\ell_\Delta(u) = 0$	[0.001]	[0.146]	[0.001]	[0.000]	[0.001]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
$H_0 : \gamma^\ell(u) = 0$	[0.000]	[0.808]	[0.008]	[0.061]	[0.054]	[0.510]	[0.167]	[0.000]	[0.497]	[0.715]	[0.152]

This table reports Wald tests for the heterogeneous-beliefs Phillips curve with current and lagged expectations, corresponding to equation (6). Superscripts e and ℓ denote current and lagged expectations. Functional coefficients are $\gamma^e(u)$ and $\gamma^\ell(u)$, with break-adjusted counterparts $\gamma^e_\Delta(u)$ and $\gamma^\ell_\Delta(u)$. P-values are shown in square brackets. Break tests jointly consider shifts in intercepts, mean-expectation coefficients, and functional components.

The central question in this extension is whether past expectations provide additional explanatory power once current expectations and macroeconomic controls are accounted for.

The joint tests $H_0 : \beta^\ell = \beta^\ell_\Delta = \gamma^\ell(u) = \gamma^\ell_\Delta(u) = 0$ are rejected at the 1% size of the test for every country, showing that lagged expectations retain systematic predictive content for current inflation. The functional-distribution restriction $H_0 : \gamma^\ell(u) = \gamma^\ell_\Delta(u) = 0$ is also rejected in all economies except Belgium, confirming that the role of lagged expectations is not driven solely by their average level but also by the distributional heterogeneity embodied in $f^\ell_t(u)$. Interestingly, the pre-break restriction $H_0 : \gamma^\ell(u) = 0$ is strongly rejected for Austria, Germany and Lithuania at the 1% level, and for Greece and Finland at the 5% level. By contrast, it is not rejected for the remaining economies, implying that for most countries the distributional effects of lagged expectations become significant only in the high-inflation regime.

Overall, these results indicate that heterogeneity in past expectations plays a salient but regime-dependent role in shaping inflation dynamics across the countries in our sample, providing explanatory power over and above that of current expectations, particularly during the high-inflation period.

4.2.4 Comparison across models: Evidence from Rolling regressions

Our last exercise pertains to robustness and re-estimates all models using a rolling-window approach. This alternative estimation strategy dispenses with the regime-break structure and allows the coefficients of the various inflation models to evolve flexibly over time, instead. The aim is to assess the models' relative performance based on their in-sample predictive performance.

Table 5: In-Sample Relative RMSE Comparison Across Models

Model	AT	BE	DE	EL	FI	FR	IT	SI	SK	ES	LT
Model 1: FPCs	0.8936	0.9298	0.9518	0.8394	0.9284	0.9425	0.9576	0.8762	0.8342	0.7931	0.7670
Model 2: Moments + FPCs	0.8360	0.8917	0.8920	0.7984	0.9009	0.8975	0.9085	0.8230	0.7747	0.7229	0.7177
Model 3: Perceptions	0.7001	0.7911	0.7430	0.6280	0.7944	0.7881	0.7178	0.6446	0.5619	0.5269	0.4852
Model 4: Current & Past Exp.	0.8137	0.8485	0.8746	0.6248	0.8378	0.8378	0.8169	0.7268	0.6482	0.6153	0.5058
Model 5: Past Exp. Only	0.9759	0.9673	0.9598	0.8323	0.9478	0.9237	0.9530	0.9052	0.8447	0.8776	0.7880
Model 6: Moments Only	0.9078	0.9365	0.9371	0.9324	0.9566	0.9591	0.9554	0.9553	0.9125	0.8755	0.9099

This table reports the ratio of each model's average rolling-window RMSE to that of the benchmark Model 0. Values below one indicate an improvement in fit relative to Model 0, whereas values above one indicate a deterioration. Model 1 (equation (1)) includes the mean and functional principal components (FPCs) of current inflation expectations. Models 2–4 (equations (4), (5), and (6)) augment Model 1 by adding, respectively, higher moments of the expectations distribution, perceived past inflation (mean and FPCs), and both current and lagged expectations (mean and FPCs). Model 5 includes only lagged expectations (mean and FPCs), while Model 6 includes only the higher moments of current expectations without their FPCs. Detailed estimation results for Models 5 and 6 are provided in the Online Appendix.

In each window of 48 monthly observations, we compute the root mean square error (RMSE)

for every specification.[14] We consider six models. Model 1 is the baseline functional Phillips curve with current expectations only (eq. (1)). Model 2 adds the higher moments of the expectations distribution (eq. (4)). Model 3 augments the baseline with perceived past inflation (eq. (5)). Model 4 jointly includes current and past expectations (eq. (6)).

For completeness, we also evaluate three additional specifications that are not nested within the above models. Model 5 includes only lagged expectations, both in their mean and functional components, but excludes current expectations entirely. This specification is suggested by the stickyâinformation models in the tradition of Mankiw and Reis (2002), in which agents update infrequently and current inflation is driven primarily by delayed expectations of today's inflation rate. Model 6 includes only the higher moments of the expectations distribution without their functional components. Finally, Model 0 serves as the benchmark hybrid Phillips curve, which includes lagged realized inflation and mean expectations but omits any distributional features of expectations, as in the sticky price framework of Gali and Gertler (1999).

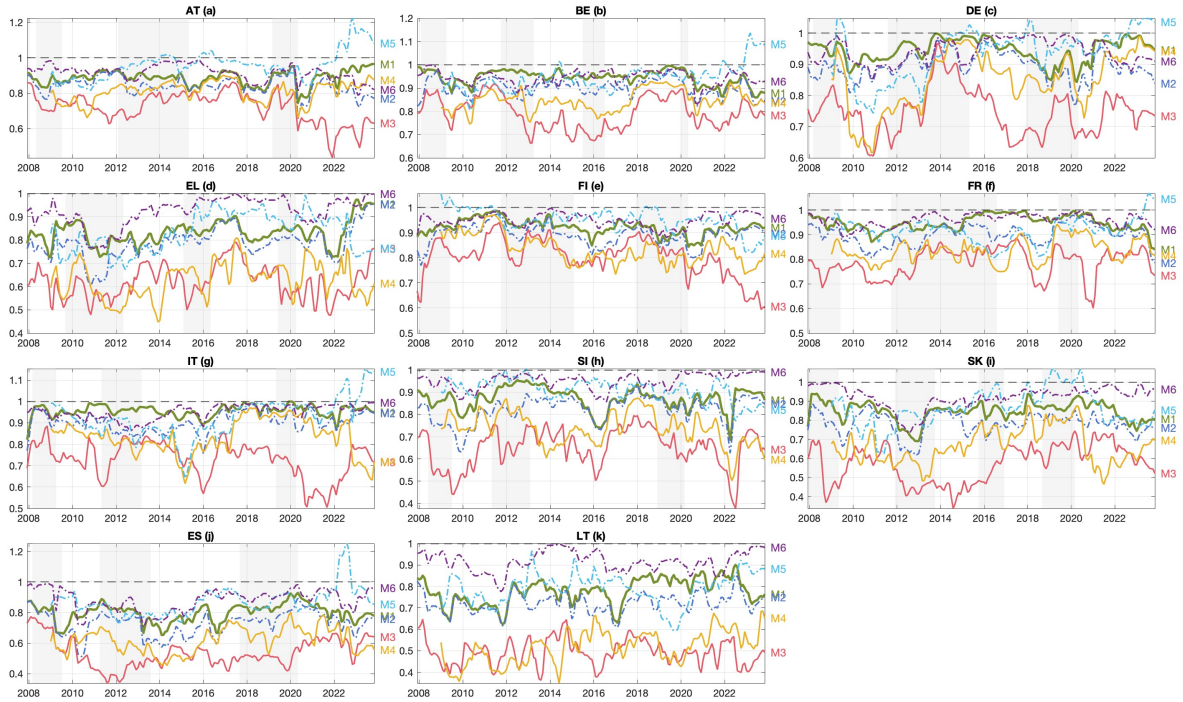
To ensure a transparent comparison across specifications, Table 5 reports relative RMSEs, defined as the ratio of each modelâs average RMSE to that of the benchmark Model 0. Values below one indicate an improvement in fit. Several clear findings emerge. First, all specifications that incorporate distributional information, whether through functional principal components (Model 1), higher moments combined with functional components (Model 2), or perceptions of past inflation (Model 3), achieve relative RMSEs below one across all countries, confirming that accounting for belief heterogeneity substantially enhances model fit. Second, the model that combines current and past expectations (Model 4) performs very well: it is the best-performing specification in Greece and ranks second in remaining countries. Third, the weakest performance consistently arises in models that omit functional information from current expectations. Model 6 (moments only) yields the highest relative RMSEs in nearly all economies, while Model 5 (past expectations only) is among the worst performers in several countries. These patterns indicate that models relying solely on delayed expectations, as in sticky-information frameworks, do not adequately capture the inflation-expectations relationship in our data.

[14] For window t,

$$\mathrm{RMSE}_t^{(m)} = \sqrt{\frac{1}{W}\sum_{j=0}^{W-1}(\pi_{t-j} - \widehat{\pi}_{t-j}^{(m)})^2},$$

where W is the window length, π_{t-j} is realized inflation, and $\widehat{\pi}_{t-j}^{(m)}$ is the fitted value of model m.

Figure 5: Rolling Relative RMSEs Across Model Specifications



This figure shows the relative RMSEs, computed as two-month moving averages of monthly data, from rolling regressions for six alternative model specifications, all estimated without accounting for regime breaks. Model 1 (dark green) is the baseline functional Phillips curve with current expectations only (equation (1)). Model 2 (dark blue) augments Model 1 with the higher moments of the expectations distribution (equation (4)). Model 3 (red) incorporates perceptions of past inflation (equation (5)). Model 4 (yellow) includes both current and past expectations (equation (6)). Model 5 (light blue) uses only lagged expectations, excluding current expectations entirely. Model 6 (purple) includes only the higher moments of the expectations distribution without their functional components. All RMSEs are reported relative to the benchmark Model 0. Recession periods for Lithuania(LT) are not displayed due to data limitations.

Furthermore, to evaluate the economic relevance of the full expectations distribution, Figure 6 plots the RMSE of Models 1–6 relative to Model 0. Model 1 (dark green), which embeds Phillips curves with heterogeneous beliefs, consistently delivers lower RMSEs around recession periods (indicated by the grey shaded regions), showing that belief heterogeneity helps explain realized inflation during economic disruptions. The improvement is especially visible for Spain, Greece, Germany, and Slovakia. These recession episodes highlight the importance of going beyond simple higher moments when analysing heterogeneity in beliefs. Model 1 also outperforms Model 6 (purple) for the majority of countries, indicating that the FPC-based representation captures heterogeneity more comprehensively; relying only on higher moments can potentially therefore understate the contribution of the full distribution of beliefs, particularly in turbulent periods.

Another notable pattern in Figure 6 is the systematic improvement in fit when incorporating perceived inflation (Model 3, red) or past expectations (Model 4, yellow), in addition to current expectations. This enhancement appears in every country. Much like the role of

heterogeneous beliefs about future inflation, these variables become especially informative around recession periods, where perceived inflation or lagged expectations help explain movements in realized inflation. Focusing on the post-COVID-19 inflation surge, the sizeable decline in RMSE for Models 3 and 4 points to an intuitive mechanism: agents relied more heavily on recent inflation experience when forming expectations. For several countries (most notably Austria, Belgium, Greece, Lithuania, and Slovenia) the RMSE patterns from Models 3 and 4 also reveal pronounced comovement between perceived inflation and lagged expectations, suggesting that in these economies perceived inflation conveys information closely aligned with historical expectations.

Last but not least, Model 5 (light blue) performs worse than all other specifications for several countries, particularly during the recent post-COVID-19 inflation surge. This underscores that, although incorporating past expectations or perceived inflation can enhance in-sample fit, omitting the forward-looking component leads to substantial informational loss and markedly poorer model performance. This pattern aligns with rational-inattention models in which agents optimally allocate attention to both current and past information, and is inconsistent with frameworks that rely exclusively on delayed expectations.

In sum, these results show that incorporating the full distributional structure of inflation expectations yields substantial improvements in in-sample fit over standard Phillips curves that rely solely on mean expectations. Models that omit the functional components of current expectations, whether relying only on higher moments or only on lagged expectations, perform markedly worse in most economies. This highlights the central role of belief heterogeneity, and specifically its functional representation, in explaining inflation dynamics.

4.3 Macroeconomic Correlates of Distributional Heterogeneity

Understanding the economic content of the functional principal components requires assessing the extent to which they are driven by three observable macroeconomic fundamentals, namely lagged inflation, the unemployment gap, and supply-chain disruptions.[15]

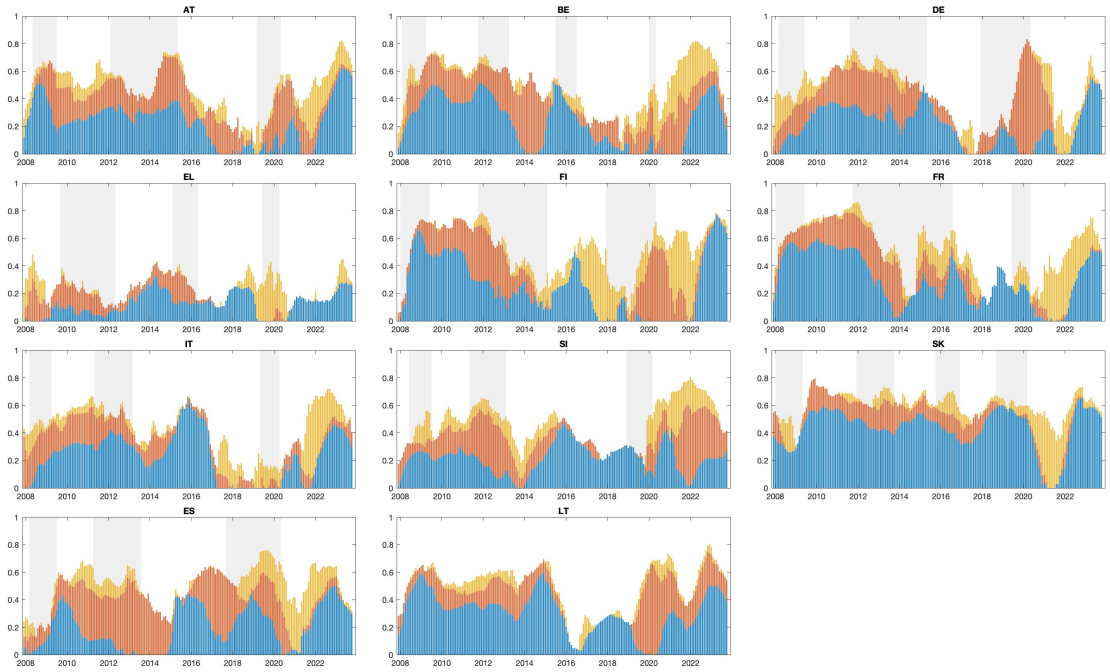
To this end, we adopt a function-on-scalar regression framework in which the entire cross-sectional distribution of inflation expectations, denoted by $f_t(u)$, is treated as a functional outcome and linked to a vector of macroeconomic aggregates that enter as scalar regressors. This allows us to quantify how much of the observed heterogeneity in household beliefs can be attributed to variation in macroeconomic conditions. The function-on-scalar model takes the form

[15]Since survey participants report their expectations before the realized values of the macroeconomic aggregates are known, we use lags of all independent variables.

$$f_t(u) = \beta_0(u) + X_t' \beta(u) + \varepsilon_t(u), \qquad u \in \mathcal{U}, \tag{7}$$

where X_t is a vector of macroeconomic predictors (taken in lagged form), $\beta_0(u)$ is an intercept function, and $\beta(u)$ is a vector of unknown coefficient functions describing how the scalar regressors shift the entire distribution of expectations.

Figure 6: Rolling Functional R^2 Decomposition Across Nested Specifications



Rolling functional R^2 for the three nested specifications described in Section 3. The blue area corresponds to $R^{2,(1)}_\tau$, the share of functional variation explained by lagged inflation alone. The orange area shows the incremental contribution $R^{2,(2)}_\tau - R^{2,(1)}_\tau$ obtained when the unemployment gap is added. The yellow area depicts the increase $R^{2,(3)}_\tau - R^{2,(2)}_\tau$ generated by including the supplyâchain index. The stacked areas sum to $R^{2,(3)}_\tau$, the coefficient of determination for the full model in window $\mathcal{T}_\tau$. Recession periods for Lithuania(LT) are not displayed due to data limitations.

To obtain an estimable representation, we project the above functional regression onto the empirical basis introduced in Section 3. As detailed in Section 4, we obtain an estimable representation by projecting $f_t(u)$ onto the empirical FPCA basis. This yields a system of scalar score regressions of the form $\xi_{k,t} = X_t' \beta_k + u_{k,t}, \quad k = 1, \ldots, K$, where β_k contains the coefficients describing how the k-th principal mode of variation in the expectations distribution responds to the macroeconomic predictors, and $u_{k,t}$ is the error term.

Building on the evidence from the main analysis that the explanatory role of fundamentals varies over time, we track how these relationships evolve by computing adjusted functional coefficients of determination for three nested rolling specifications. For each window $\mathcal{T}_\tau$ of length W, the total functional variation is defined as $\text{TSS}_\tau = \sum_{k=1}^{K} \sum_{t \in \mathcal{T}_\tau} \xi^2_{k,t}$. For

each model $m \in \{1, 2, 3\}$, where $m = 1$ includes lagged inflation only, $m = 2$ adds the unemployment gap, and $m = 3$ further adds the supplyâchain index, we obtain fitted scores $\widehat{\xi}_{k,t,\tau}^{(m)}$ and compute the residual sum of squares $\mathrm{RSS}_\tau^{(m)} = \sum_{k=1}^{K} \sum_{t \in \mathcal{T}_\tau} (\xi_{k,t} - \widehat{\xi}_{k,t,\tau}^{(m)})^2$.

The adjusted functional coefficient of determination for specification m in window τ is $\widetilde{R}_\tau^{2,(m)} = 1 - \frac{\mathrm{RSS}_\tau^{(m)}/(WK - p_m)}{\mathrm{TSS}_\tau/(WK-1)}$, where p_m is the number of regressors in model m.[16] Because the models are nested, these adjusted measures satisfy $\widetilde{R}_\tau^{2,(1)} \leq \widetilde{R}_\tau^{2,(2)} \leq \widetilde{R}_\tau^{2,(3)}$, so differences between successive values quantify the incremental explanatory power of adding the unemployment gap and supply-chain pressures. The resulting sequences of adjusted rolling $\widetilde{R}_\tau^{2,(m)}$ values, plotted in Figure 6, show how the contribution of each macroeconomic driver to the functional variation in expectations evolves over time.

Lagged inflation explains a substantial portion of the functional variation during periods of stable prices. The unemployment gap adds further explanatory power in several economies, particularly during labour-market stress, such as the COVID-19 period. Supply-chain disruptions become a key determinant from 2020 onwards, with their contribution rising markedly during global bottlenecks and the subsequent high-inflation episode. These patterns suggest that the drivers of belief heterogeneity shift over time: inflation persistence dominates in normal periods, whereas labour-market conditions and supply shocks gain prominence when inflation is elevated

5 Conclusion

We study the relationship between inflation and inflation expectations across euro-area economies using a functional data approach that incorporates the entire cross-sectional distribution of household beliefs. Three main findings emerge. First, the functional components of the expectations distribution play a central role in explaining inflation dynamics. Their contribution remains substantial even after controlling for higher moments or perceptions of past inflation, indicating that important information is embedded in distributional features that standard Phillips curve models overlook. Second, forward-looking expectations are generally more influential than lagged expectations, but empirical models that incorporate both perform better during periods of macroeconomic turbulence. This pattern is consistent with rational-inattention frameworks that allow joint dependence on forward-

[16] Here, we retain $K = 20$ functional principal components, which together account for essentially all of the variation in the estimated belief distributions and stack the K score regressions over the W observations in window $\mathcal{T}_\tau$ yields the system $\Xi_\tau = (I_K \otimes X_\tau)B + U_\tau$, where Ξ_τ is the $WK \times 1$ vector collecting all scores, X_τ is the $W \times p_m$ design matrix, B is the $p_m K \times 1$ vector of coefficients, and U_τ is the stacked error vector. The adjusted R^2 reported above is the standard adjusted coefficient of determination for this stacked system.

and backward-looking information, and inconsistent with sticky-information or sticky-price models. Third, there is pronounced cross-country heterogeneity: the magnitude and composition of the expectations channel differ across economies, with distributional heterogeneity becoming more informative in high-inflation periods when labour-market conditions and supply disturbances gain prominence. While our analysis does not make causal claims, the findings highlight the importance of modelling both within-country and cross-country heterogeneity in expectations, and of allowing the full distribution of beliefs, rather than only consensus expectations, to shape inflation dynamics.

References

Bai, J. and Perron, P. (1998). Estimating and testing linear models with multiple structural changes. *Econometrica*, pages 47–78.

Benigno, G., di Giovanni, J., Groen, J. J., and Noble, A. I. (2022). A new barometer of global supply chain pressures. Technical report, Federal Reserve Bank of New York.

Chang, Y., Gómez-Rodríguez, F., and Hong, M. G. H. (2022). *The Effects of Economic Shocks on Heterogeneous Inflation Expectations*. International Monetary Fund.

Coibion, O. and Gorodnichenko, Y. (2012). What can survey forecasts tell us about informational rigidities? *Journal of Political Economy*, 120:116–159.

Coibion, O., Gorodnichenko, Y., and Kamdar, R. (2018). The formation of expectations, inflation, and the phillips curve. *Journal of Economic Literature*, 56(4):1447–1491.

Duca-Radu, I., Kenny, G., and Reuter, A. (2021). Inflation expectations, consumption and the lower bound: Micro evidence from a large multi-country survey. *Journal of Monetary Economics*, 118:120–134.

Fofana, S., Patzelt, P., and Reis, R. (2024). Household disagreement about expected inflation. In Ascari, G. and Trezzi, R., editors, *Research Handbook of Inflation*. Edward Elgar Publishing.

Gali, J. and Gertler, M. (1999). Inflation dynamics: A structural econometric analysis. *Journal of Monetary Economics*, 44(2):195–222.

Geiger, M., Sterghides, I., and Zachariadis, M. (2025). When bad news breeds bias: Cross-country evidence on inflation-as-a-bad and overreaction in inflation expectations. *Working Paper 04-2025, Department of Economics, University of Cyprus*.

Lopez-Pintado, S. and Romo, J. (2009). On the concept of depth for functional data. *Journal of the American Statistical Association*, 104:718–734.

Mackowiak, B. and Wiederholt, M. (2009). Optimal sticky prices under rational inattention. *American Economic Review*, 99(3):769–803.

Mankiw, N. G. and Reis, R. (2002). Sticky information versus sticky prices: a proposal to replace the New Keynesian Phillips curve. *The Quarterly Journal of Economics*, 117(4):1295–1328.

Mavroeidis, S., Plagborg-Moller, M., and Stock, J. H. (2014). Empirical evidence on inflation expectations in the new keynesian phillips curve. *Journal of Economic Literature*, 52(1):124–188.

Meeks, R. and Monti, F. (2023). Heterogeneous beliefs and the phillips curve. *Journal of Monetary Economics*, 139(2):41â54.

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690.