Chapter 1

High Expectations: State-Dependence and Heterogeneous Pass-Through in an Augmented Euro Area Phillips Curve

Giovanni Borraccia (Geneva Graduate Institute)

Abstract

Firms' own selling-price expectations are the forward-looking force at the heart of the sectoral Phillips curve, yet their pass-through into producer prices is not necessarily uniform across sectors or across phases of the wider expectations cycle. This paper asks how that pass-through operates across euro-area sectors, and whether it intensifies when the international expectations climate turns hot. On harmonised quarterly data for seventeen euro-area economies and nineteen NACE sectoral aggregates over 2000–2023, I estimate a sectoral Phillips curve with the Dynamic Common Correlated Effects Mean Group (DCCE–MG) estimator, which accommodates heterogeneous slopes and pervasive cross-sectional dependence. Baseline pass-through is about 0.28 standard deviations (roughly 1 percentage point of PPI inflation per typical expectations swing), and rises to roughly 0.45 (approximately 1.7 percentage points), an amplification of about two-thirds, once a dynamic factor model of cross-country expectations crosses its 75th percentile. This result is robust to alternative thresholds and factor extractions. The amplification is broad-based across manufacturing and absent only in Electronics, where long-horizon procurement and globally administered pricing mute the channel. Across countries, transmission is strongest in the Core, intermediate in the Periphery, and muted in Central and Eastern Europe; an Oaxaca–Blinder decomposition attributes this gradient predominantly to within-sector differences in price-setting behaviour rather than to sectoral composition.

Keywords: inflation expectations, amplification, Phillips curve, DCCE-MG, cross-sectional dependence, state dependence.

JEL: E31; C23; E52; F41.

1.1 Introduction

This paper seeks to refine the Phillips Curve framework by offering an empirical sectoral approach that integrates international and domestic dimensions of firms' selling-price expectations, with the central focus placed on the heterogeneity of pass-through across sectors and country groups within the euro area. The question it addresses is how firms' selling-price expectations translate into producer prices, and whether this transmission strengthens when the international expectations environment is elevated. These are firms' own selling-price expectations from the European Commission's business surveys: a sectoral, forward-looking measure of expected producer-price changes, which I treat as the firm-side primitive underlying the expectation term of the New Keynesian Phillips curve rather than as a proxy for aggregate consumer-price or professional-forecaster inflation expectations. Using harmonised sectoral panel data for euro-area economies, I find that the pass-through from firms' own selling-price expectations to PPI inflation rises from roughly 0.28 to 0.45 standard deviations once a common factor of cross-country expectations crosses its 75th percentile. This is an amplification of about two-thirds, equivalent to a rise from roughly 1 to 1.7 percentage points of PPI inflation per typical expectations swing. This result is not a uniform euro-area phenomenon: amplification clusters in specific manufacturing sectors and varies sharply across Core, Periphery, and CEE economies, and an Oaxaca-Blinder decomposition shows that these cross-country differences are rooted in how firms respond to inflation signals rather than in sectoral composition: the behavioural component accounts for close to 88% of the cross-group gap in amplification, while the compositional component is statistically indistinguishable from zero. Heterogeneity is therefore the paper's central object. An aggregate baseline first documents state-dependent amplification at the euro-area level. A sectoral analysis then maps the heterogeneity of this amplification across manufacturing sectors and ties it to differences in pricing autonomy and exposure to globally traded inputs. A country-group analysis closes the sequence, with an Oaxaca-Blinder decomposition separating behavioural from compositional sources of the cross-group gradient.

The Phillips Curve has long served as a foundational framework in macroeconomics, linking inflation dynamics to fluctuations in economic activity. However, as global economic integration deepens, the traditional application of the Phillips Curve faces continuously evolving challenges, as firms' pricing behaviour is increasingly shaped by a complex interplay of both domestic and international factors. These dynamics have revealed significant limitations in the traditional model, especially in capturing inflation behaviour amidst periods of economic turbulence. In the US, in the aftermath of the Great Financial Crisis, inflation was expected to fall by more than it did, judging by the predictions of standard models such as reduced-form Phillips curves e.g. (Ball and Mazumder, 2011). This so-called 'missing disinflation' puzzle also characterises other advanced economies where inflation proved to be rather resilient given the weakness in economic activity (Friedrich, 2016; Del Negro et al., 2015). Since 2012, particularly in the euro area, inflation has surprised in the opposite direction, falling to very low levels and giving birth to the so-called 'missing inflation' puzzle (Constâncio, 2015). To face these challenges over time, the traditional Phillips Curve model has been in some cases revisited and enriched to incorporate open economy elements, such as trade openness and international price competition, recognising their influence on inflation dynamics. Additionally, this process has benefited from the rise of refined surveys capturing firms' and households' inflation expectations, offering an increasingly nuanced view of price-setting behaviour.

Despite these advancements, the level of generality of these measures often remains rather high, limiting the model's ability to fully account for sector-specific and forward-looking dynamics that are critical in an increasingly interconnected global economy. In an open economy, firms make pricing decisions not only based on immediate domestic costs but also on expectations about future price changes within their own sector, across related industries, and in the global marketplace. Expectations of sectoral price movements, anticipated consumer price inflation, and international competitive pressures all interact to shape firms' behaviour. Capturing these interactions requires a model that can unpack firms' selling-price expectations into distinct components that align with the specific economic factors influencing firms' decisions.

By decomposing expectations into multiple layers (sector-specific selling-price forecasts, input-cost linkages across industries, and broader cross-country demand and supply signals), this study lays the foundation for a richer understanding of how inflation expectations form and propagate across sectors. The empirical strategy estimates a dynamic common correlated effects mean group (DCCE-MG) Phillips-curve framework on harmonised sectoral panel data for euro-area economies, examining how the transmission of selling-price expectations into producer prices varies across sectors and across phases of the expectations cycle. Three main findings emerge.

First, the pass-through of expectations is heterogeneous, yet systematic: firm pricing responses to their own expectations vary between sectors but follow a clear hierarchy consistent with differences in exposure to international competition and cost structure (for instance, Metals and Chemicals, where input costs track globally traded commodities, are at the high end of the distribution, whereas Electronics, where pricing is dominated by multi-period contracts and global value-chain competition, is at the low end). This finding complements sectoral evidence on incomplete pass-through in European price-setting (e.g. Bartiloro et al., 2022; Fabiani et al., 2006).

Second, when the international expectations factor is elevated and explains a larger share of cross-sectoral variation in firms' beliefs, the domestic pass-through becomes significantly stronger, revealing a state-dependent amplification mechanism. This suggests that firms internalize not only their own expectations but also the broader expectations environment, consistent with recent work on expectation coordination and the global dimension of inflation persistence (Forbes, 2019; Ha et al., 2019a).

Third, pooling results across country groups using an Empirical-Bayes estimator highlights persistent asymmetries between Core, Periphery, and Central-Eastern European economies. Core countries exhibit stronger and more persistent transmission of inflation expectations, Periphery economies display weaker and more short-lived effects, while CEE sectors show higher volatility and sensitivity to external shocks. The subsequent Oaxaca-type decomposition indicates that these differences are primarily behavioural rather than compositional, owing to underlying heterogeneity in adjustment mechanisms and expectation formation.

These results carry a direct monetary-policy implication. When aggregate expectations enter a high-DFM regime, the inflationary pressure they generate is broadly distributed across manufacturing sectors, significant in every sector except Electronics, and concentrated in Core economies. Aggregate inflation forecasts that rest on economy-wide expectations measures may therefore understate sectoral price pressures during exactly the periods in which expectations matter most, posing a signal-extraction problem for central banks calibrating policy to euro-area aggregates.

Overall, this paper contributes to the literature on inflation dynamics by developing a refined, state-dependent Phillips Curve framework that integrates international and domestic forces shaping expectations. By modelling how these forces interact and are transmitted across sectors, the framework provides a theoretically motivated and empirically tractable tool to interpret the uneven propagation of inflation expectations in an open economy context.

The remainder of the paper is organised as follows. Section 1.2 reviews the literature on the expectations-augmented Phillips curve and on heterogeneity in price-setting behaviour. Section 1.3 describes the data and the construction of the key variables: the sectoral PPI panel, the firm-side selling-price expectations, the input–output-weighted other-sector expectations, and the international expectations factor. Section 1.4 details the DCCE–MG estimator and the augmented Phillips curve specification, and reports the aggregate baseline. Section 1.5 turns to heterogeneous transmission, examining the sectoral pattern, the Core/Periphery/CEE country-group gradient, and the Oaxaca–Blinder decomposition that separates behavioural from compositional sources of cross-group differences. Section 1.6 concludes.

1.2 Literature Review

The Phillips Curve originated with Phillips (1958), who documented an inverse relation between wage inflation and unemployment in the United Kingdom. Samuelson and Solow (1960) extended this relationship to prices, suggesting that policymakers could trade off unemployment for inflation. The subsequent monetarist critique, led by Friedman (1968) and Phelps (1967), replaced this static trade-off with an expectations-augmented version: when agents anticipate policy-induced inflation, the curve shifts, eliminating the trade-off in the long run. This shift placed inflation expectations at the centre of macroeconomic analysis and set the foundation for later rational-expectations models.

Following the Lucas (1976) critique, macroeconomic models incorporated forward-looking behaviour, leading to the Full-Information Rational Expectations (FIRE) framework (Sargent, 1987; Woodford, 2003). Within this approach, the New Keynesian Phillips Curve (NKPC) linked current inflation to expected future inflation and real marginal costs (Galí and Gertler, 1999; Clarida et al., 1999; Roberts, 1995). However, empirical evidence soon revealed that inflation adjusts more sluggishly than predicted, implying that agents update information or prices only intermittently. The Sticky-Information model of Mankiw and Reis (2002) assumes that only a fraction of agents update expectations each period, while Milani (2007) introduces adaptive learning, where expectations evolve gradually with experience. Coibion and Gorodnichenko (2015a,b) show using survey data that both firms and households exhibit systematic forecast errors inconsistent with FIRE. More recent behavioural models emphasise selective or state-dependent attention (Afrouzi, 2024; Bordalo et al., 2018), which indicates that expectation formation itself amplifies or dampens macroeconomic shocks depending on the information environment.

Together, these contributions demonstrate that expectations are heterogeneous, incomplete, and shaped by firm- and sector-specific information. Nonetheless, most empirical Phillips-Curve estimates continue to rely on aggregate measures of expectations, often masking substantial cross-sectional variation in how firms perceive and react to economic conditions.

Building on this, recent empirical work has focused on the observed flattening of the Phillips Curve since the mid-1990s (Blanchard, 2016; Del Negro et al., 2020; Stock and Watson, 2019; Jørgensen and Lansing, 2025). A branch of the literature stresses structural and global factors: intensified import competition, global value chains, and production relocation have constrained firms' pricing power and reduced domestic cost pressures (Auer et al., 2017; Forbes, 2019; Kohlsheen and Moessner, 2022). As Borio and Filardo (2007) and Ha et al. (2019b) argue, global slack and commodity prices increasingly influence domestic inflation, challenging closed-economy formulations. Another branch highlights policy credibility: stronger monetary frameworks have anchored inflation expectations and stabilised prices despite cyclical variation. Empirical studies show that credible policy responses flatten the inflation–activity relationship (Ball and Mazumder, 2019; Hooper et al., 2020; McLeay and Tenreyro, 2019), while euro-area evidence attributes the "missing inflation" episodes of the 2010s to policy-driven anchoring (Bobeica and Jarociński, 2019; Barnichon and Mesters, 2021).

Firm-level and sectoral analyses complement this evidence by uncovering the microeconomic origins of these aggregate patterns. Using Italian survey data, Bartiloro et al. (2022) show that firms' inflation expectations are dispersed and cyclical, arising from differing degrees of information rigidity and uncertainty. Sectoral studies further show that persistence and sensitivity to economic slack vary substantially across industries. Imbs et al. (2011) demonstrate that aggregation biases the estimated Phillips-Curve slope by concealing sector-specific responses. Luengo-Prado et al. (2018) identify a post-crisis decline in persistence in several sectors, while Seydl and Spittler (2016) find that international competition weakens the link between prices and domestic demand in tradable sectors. Evidence from the ECB's Inflation Persistence Network (Altissimo et al., 2006; Vermeulen et al., 2007) confirms large cross-sectoral differences in price-setting frequencies and cost pass-through, tied to input structures, contract rigidity, and exposure to external shocks. Cross-country studies further reveal persistent asymmetries between Core and Peripheral euro-area economies (Forbes et al., 2021; Borio et al., 2023), pointing towards unequal transmission of expectations into prices and adjustment speeds across institutional environments.

These findings resonate with recent research on nonlinear Phillips-Curve specifications, which emphasises state dependence and asymmetric transmission mechanisms. In particular, Barnichon and Mesters (2021) and Hazell et al. (2022) show that the slope of the Phillips Curve steepens in high-demand or high-inflation regimes, while flattening during periods of slack or anchored expectations. Similarly, Forbes et al. (2021) highlight amplification effects driven by interactions between expectations, global shocks, and local cost pressures. This evidence implies that heterogeneity across sectors and countries is not merely compositional but structural, reflecting nonlinear adjustment processes that conventional linear models fail to capture.

Taken together, this literature underscores a growing disjunction between the theoretical recognition of heterogeneity in expectations and the aggregate treatment these receive in empirical work. While models increasingly emphasise forward-looking and information-constrained agents, most empirical Phillips-Curve applications, particularly those using cross-country or aggregate data, still proxy inflation expectations with economy-wide measures such as professional forecasts (e.g., SPF, Consensus), consumer surveys, or lagged inflation. This aggregate treatment is grounded in both modelling convention and data limitations: long, harmonised time series of firm- or sector-level expectations have historically been scarce. Country-specific studies that exploit firm survey data, such as those based on Italy's Bank of Italy survey (Bartiloro et al., 2019), Germany's IFO survey, or the UK Decision Maker Panel (Yotzov, 2024), have provided valuable microeconomic insights but remain limited in scope and cross-country comparability.

By contrast, this paper leverages harmonised sectoral and subsectoral data on selling-price expectations from the European Commission's Business and Consumer Surveys (European Commission, 2024). This dataset offers a unique opportunity to analyse expectation formation and transmission mechanisms across a broad range of sectors and euro-area economies within a single empirical framework. Such granularity allows for the joint assessment of domestic (cross-sectoral) and external (cross-country) expectation channels in the context of a monetary union, where national inflation dynamics are increasingly interdependent but institutional and structural heterogeneities persist. This euro-area perspective provides a crucial setting to explore how firms' expectations interact across sectors and borders, shaping inflation propagation under a common monetary regime.

Ultimately, this paper contributes to this evolving field by condensing in a single specification several of the main sources of dynamism and heterogeneity introduced progressively into the standard Phillips curve model by different branches of the literature. It does so by estimating a dynamic common correlated-effects mean-group (DCCE-MG) Phillips-Curve model that incorporates both domestic (cross-sectoral) and external (cross-country) expectation channels at the sectoral level. This framework allows for the identification of non-linear, state-dependent transmission and systematic cross-group asymmetries in how firms translate expectations into price adjustments across the euro area.

1.3 Data and variable construction

1.3.1 Main variables

This analysis combines sectoral and macroeconomic data for euro-area economies, drawing on a range of official sources to construct a harmonised quarterly panel covering the period 2000–2023. The dataset integrates information on prices, output, and expectations, harmonised under the NACE Rev. 2 industry classification, and extends the analysis across 17 euro-area countries.[1]

Price indicators are sourced from *Eurostat*, specifically the Producer Price Index (PPI) for manufacturing and selected market services. Real activity data, used to construct the sectoral output gap, are obtained from *Eurostat's Short-Term Statistics (STS)* database, specifically the *Production in Industry (volume)* index. This series measures the chain-linked volume of industrial output (2015 = 100) at the NACE Rev. 2 division level and serves as a high-frequency proxy for real value added across sectors and countries.

The expectations variables, central to the construction of both domestic and external expectation components, are drawn from the *European Commission's Business and Consumer Surveys (BCS)*. The BCS collects monthly qualitative assessments of firms' selling price expectations at the NACE Rev. 2 *division level* (two-digit), providing a finer granularity than is typically available in macroeconomic datasets. These data are converted to quarterly frequency to align with the temporal resolution of the macro variables.

To preserve detail while ensuring consistency and cross-country comparability, all NACE Rev. 2 divisions are grouped into nineteen aggregated sectors (see Table 1 and Appendix 1.A.1). This level of aggregation, intermediate between the NACE section and division, balances granularity with full data coverage, allowing the analysis to capture cross-sectoral and cross-country interactions. Aggregation weights are based on sectoral *nominal* output shares derived from *Eurostat's Structural Business Statistics (SBS)*. Annual SBS turnover is first converted into sectoral shares and then smoothed over five quarters using centred moving averages. The resulting smoothed weights are renormalised within each country–sector–quarter group so that they sum to unity, providing a coherent and stable basis for sectoral aggregation across time.

The final dataset spans the period 2000Q1–2023Q4, covering 17 euro-area economies and nineteen harmonised macro-sectors constructed from NACE Rev. 2 division-level inputs. This structure enables a detailed examination of price and expectation dynamics across industries and countries within the monetary union, while supporting the identification of common components, state dependence, and cross-sectional interactions in inflation transmission.

[1] The euro-area countries included are Austria, Belgium, Croatia, Estonia, Finland, France, Germany, Greece, Ireland, Italy, Latvia, Lithuania, the Netherlands, Portugal, Slovakia, Slovenia, and Spain.

Table 1: Aggregated sector definitions (NACE Rev. 2 classification)

Aggregated sector	NACE Rev. 2 codes
Mining and Quarrying	B05–B09
Food, Beverages and Tobacco	C10–C12
Textiles, Apparel and Leather	C13–C15
Wood, Paper and Printing	C16–C18
Chemicals and Pharmaceuticals	C20–C22
Metals and Fabrication	C24–C25
Electronics and Electrical Equipment	C26–C27
Machinery and Vehicles	C28–C30
Other Manufacturing	C23, C31–C33
Energy and Utilities	D35, E36, C19
Transport and Storage	H49–H53
Wholesale and Retail Trade	G45–G47
Accommodation and Food Services	I55–I56
Information and Communication	J58–J63
Financial and Insurance Activities	K64–K66
Real Estate Activities	L68
Professional Services	M69–M75, N77
Other Services	N78–N82
Arts, Entertainment and Recreation	R90–R94

Note: Sectors correspond to harmonised groupings of NACE Rev. 2 divisions used for aggregation and weighting. Full mapping details are provided in Appendix 1.A.1.

1.3.2 Construction of the sectoral slack measure

Sectoral economic slack is measured as the cyclical component of real output at the country–industry level. Specifically, real output is proxied by the *Production in Industry (volume)* index from Eurostat's Short-Term Statistics (STS) database. This index provides chain-linked volume measures of industrial production, expressed as an index with base year 2015 = 100, and covers NACE Rev. 2 sections B to E (Mining, Manufacturing, and Energy-related activities). The series are seasonally and calendar adjusted and available at quarterly frequency, ensuring comparability across sectors and countries.

Let $y_{i,s,t}$ denote the log of the real production index for country i, NACE Rev. 2 division s, in quarter t, and decompose it into a trend $\tau_{i,s,t}$ and a cyclical component using the Hodrick–Prescott filter ($\lambda = 1600$) applied separately to each country–sector panel:

$$y_{i,s,t} = \tau_{i,s,t} + c_{i,s,t}, \qquad \text{gap}_{i,s,t} \equiv c_{i,s,t} = y_{i,s,t} - \tau_{i,s,t}.$$

A positive value indicates activity above trend (tight conditions), a negative value below-trend activity (slack).[2]

1.3.3 Sectoral inflation and inflation expectations

This work leverages the extensive European Commission Business and Consumer Survey (BCS), which provides sectoral-level insights into firms' short-term expectations. The dataset is comprehensive, containing monthly data on managers' selling price expectations in manufacturing, services, retail, and construction for all countries in the European Union. Specifically, the survey question on **selling price expectations over the next three months** is used as a proxy for sectoral inflation expectations.

Firms are presented with a qualitative, three-option question, such as: "How do you expect your selling prices to change over the next 3 months?" Responses are aggregated into a "balance" statistic, the standard EC quantification method. This index is computed as the difference between the weighted percentage of firms expecting an "increase" (P) and those expecting a "decrease" (M):

$$B = P - M$$

The resulting raw index is bounded by -100 and $+100$. To ensure representativeness, responses are weighted at the national level to reflect both firm size (e.g., turnover) and the sector's share of total value added. The final series are seasonally adjusted by the Commission.

Mapping to theoretical NKPC model

The empirical specification is anchored in the forward–looking New Keynesian Phillips Curve. The full Calvo derivation is reproduced in Appendix 1.B; here it is enough to recall the structural

[2]Panels with fewer than six valid quarterly observations were excluded to avoid instability in the estimated cycle. Robustness checks using the band-pass adjustment of Ravn and Uhlig (2002) and the regression-based approach of Hamilton (2018) yielded qualitatively similar results.

object the empirical model is meant to identify and the role of firm-side expectations within it.

In the standard Calvo setting, a fraction $1 - \theta$ of firms is drawn each period to re-optimise, and a re-optimising firm chooses a reset price p_t^* that solves a forward-looking optimality condition (Appendix 1.B, eq. (d)): the reset-price gap $p_t^* - p_{t-1}$ equals a discounted sum of expected future real marginal costs and future aggregate prices. Aggregating this reset-price rule across firms via the aggregate price-setting equation $\pi_t = (1 - \theta)(p_t^* - p_{t-1})$ (eq. (f)) and substituting an output-gap expression for marginal cost delivers the canonical NKPC,

$$\pi_t = \beta E_t \pi_{t+1} + \kappa \varphi (y_t - y_t^n), \tag{1}$$

with $\kappa = (1 - \theta)(1 - \beta\theta)\xi/\theta$. Equation (1) is the theoretical benchmark for the empirical model.

The expectations measure used in this paper is not a survey proxy for the aggregate inflation forecast $\mathbb{E}_t[\pi_{t+1}]$; it is the firm-side primitive from which that forecast is built. The optimal reset-price condition (d) shows that what disciplines p_t^* is the firm's own expectation of its future selling price, not its forecast of aggregate inflation. The aggregation of these own-price expectations under Calvo, via the aggregate price-setting equation (f), is the step that converts the micro-level price-setting expectation into the aggregate term $E_t[\pi_{t+1}]$ appearing in the NKPC (i). The European Commission Business Survey asks each firm exactly this question, how it expects its *own selling prices* to evolve over the next quarter, and is therefore a direct measurement of the price-setting expectation $\mathbb{E}_{t-1}[\Delta p_{i,t}]$ that the NKPC takes as the underlying object, rather than a derived forecast of aggregate consumer-price inflation extracted from professional forecasters or households. In a homogeneous equilibrium in which all firms face the same environment and the same law of motion for aggregate prices, this firm-side expectation aggregates one-for-one into the inflation term appearing in the NKPC: theoretically, the survey is closer to the structural object than aggregate-inflation surveys of forecasters or consumers, not further from it.

The timing of the business survey reinforces this theoretical mapping. Firms report their selling-price expectations at the beginning of quarter $t-1$, referring explicitly to planned price changes for quarter t. In many sectors, these plans are determined before the quarter begins through catalogues, production scheduling, and procurement contracts; the producer-price index for quarter t then tracks their realised outcome. Formally, the survey measures $\mathbb{E}_{t-1}[\pi_{i,t}]$, a one-period re-indexing of the NKPC expectation term $\mathbb{E}_t[\pi_{t+1}]$, with the survey reading collected before the quarter to which it refers begins.

In this context, the timing claim is one of *predetermination*, not *structural causation*: as in any forecast-then-realise sequence, the fact that expectations precede prices does not by itself rule out a common signal that drives both. This is why the empirical specification is built around the DCCE–MG estimator, which augments each country–sector regression with cross-sectional averages of all variables and absorbs unobserved common factors non-parametrically. Combined with the lagged regressor structure, this ensures that the slope coefficients on $E^{\text{own}}_{i,t-1}$ and on its high-DFM interaction are identified from *idiosyncratic, sector-specific* variation rather than from any signal common to expectations and prices. Predetermination plus common-factor purging is the identifying combination flagged by Coibion et al. (2018) as the operational requirement when working with firm-level survey expectations; it sustains a forward-looking interpretation of the pass-through and amplification coefficients without claiming structural identification beyond what timing and the estimator deliver.

The balance statistic is a qualitative, bounded index rather than a cardinal measure of firms' desired reset prices, but on theoretical grounds this is not a limitation under the structure of the NKPC. Aggregate inflation in the model evolves according to $\pi_t = (1-\theta)(p^*_t - p_{t-1})$, which depends only on the *average desired price change* among firms that re-optimise. What matters for the forward-looking term $E_t[\pi_{t+1}]$ is therefore the anticipated *sign* and *prevalence* of upcoming price adjustments across firms, not the cardinal magnitude of any individual reset price.

A natural concern is that sign alone is uninformative because a 1% and a 50% increase share the same sign. The diffusion index $B = P - M$ addresses this through the breadth dimension. Under Calvo with symmetric adjustment costs, the share of firms in a sector expecting to raise their selling prices is, to a first-order approximation, proportional to the sectoral average desired price change: a larger desired increase in p^*_t raises both the probability that any individual firm reports "up" and the depth of the average upward adjustment, in the same direction. The balance statistic therefore captures both who would adjust and how widespread the adjustment is, and consequently scales monotonically with the sectoral average desired price change that the NKPC requires. Axioglou et al. (2024) validate this empirically by showing that selling-price expectations from the EC Business and Consumer Surveys forecast realised producer-price inflation up to six to seven months ahead with stable cross-country precision; classical measurement error in the qualitative-to-balance transformation, where it remains, biases the estimated coefficients toward zero, so the reported $\hat{\beta}$ and $\hat{\delta}$ are best read as conservative lower bounds on the true pass-through and amplification elasticities.

Choice of dependent variable

Empirically, this expectations measure is highly informative. The European Commission's recent study (Axioglou et al., 2024) finds that selling-price expectations predict inflation up to 6–7 months ahead and remain strongly correlated with producer-price inflation even during volatile periods. In our data, the series closely co-moves with sectoral PPI inflation and exhibits explanatory power consistent with a supply-side Phillips Curve grounded in the firm-side NKPC mechanism.

The expectations variable used here is theoretically consistent with the firm-side forward-looking component of the NKPC, empirically validated, and timed in a way that preserves causal interpretation within a quarterly setting. The use of selling price expectations as an explanatory variable and PPI as the dependent one is supported by both theoretical coherence and empirical alignment. Selling price expectations reflect firms' assessments of cost conditions and market prospects at the production stage, directly corresponding to PPI, which measures price changes received by producers. Unlike HICP, which captures consumer-level dynamics, indirect taxes, and behavioural adjustments, PPI is more sensitive to sectoral cost pressures and thus better suited to analysing inflation transmission through production chains.

To verify that this theoretical mapping is borne out empirically, I conduct a simple calibration exercise that compares the survey-based selling-price expectations against realised inflation. Figures 1 and 2 plot the year-on-year diffusion index of selling-price expectations against PPI and HICP inflation in manufacturing and services for the euro area, and the full grid of lead–lag correlations is reported in Appendix 1.C. The exercise serves two purposes: it documents that the survey is informative about realised producer-price dynamics, and it identifies PPI rather than HICP as the appropriate dependent variable for a firm-side Phillips curve.

Figure 1: Year-on-Year Inflation and Selling Price Expectations
Manufacturing (Euro Area)

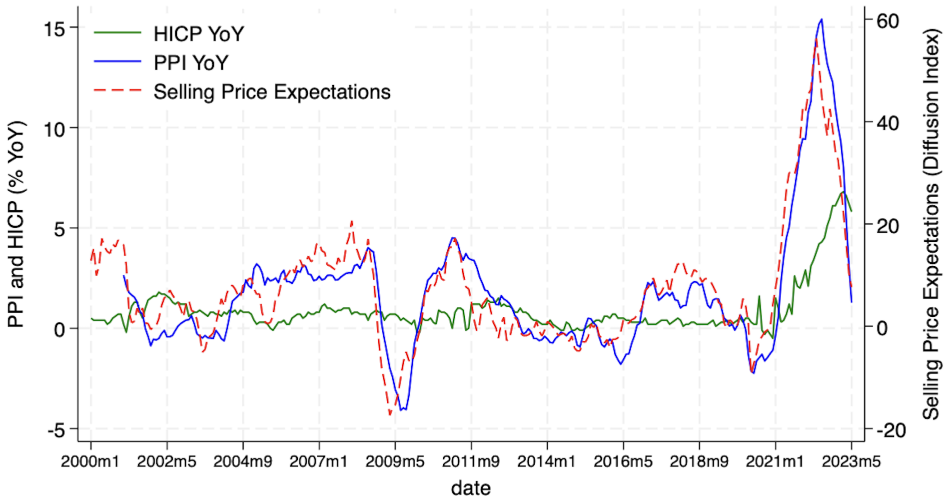

Figure 2: Year-on-Year Inflation and Selling Price Expectations
Services (Euro Area)

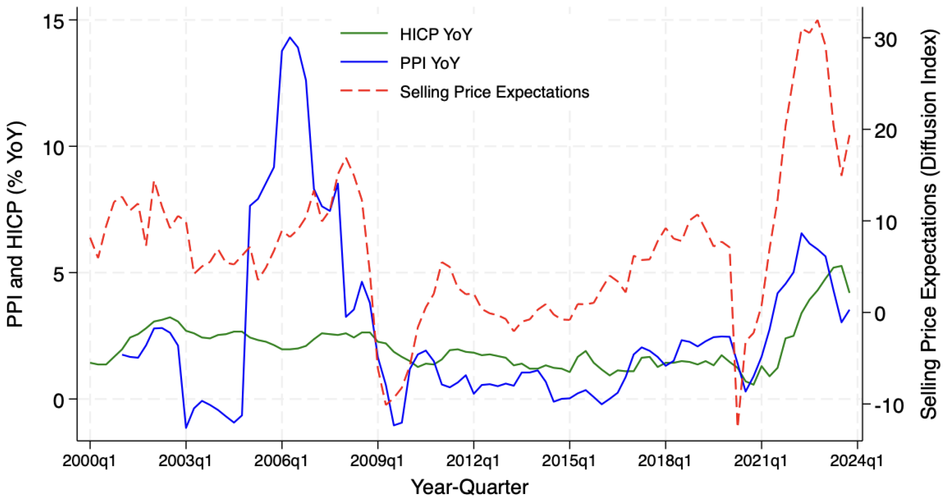

In manufacturing (Figure 1), the lag-correlation table reported in Appendix 1.C (Table 1.C.1) shows that selling-price expectations and PPI inflation reach a correlation of 0.94 when expectations lead PPI by one to two quarters, and remain above 0.85 over a ±five-quarter window. The correlation with HICP peaks lower, at 0.78, and only at much longer leads of seven to ten quarters, consistent with HICP incorporating consumption-side and fiscal components that respond on a slower timing.

In services (Figure 2), the visual link is milder, in line with thinner survey panels and lower price-setting frequencies. The contemporaneous correlation between expectations and PPI is 0.46, with a peak of 0.46 at $h = +1$ and values around 0.44 at $h = +2$ (Appendix 1.C,

Table 1.C.2). The HICP relationship is essentially absent at short horizons (correlations near zero or slightly negative through $h = +4$) and only becomes large at long lags, exceeding 0.80 from $h = +10$ onwards, again lining up with the slower adjustment of consumer-price aggregates relative to producer-side dynamics.

Together, these results justify PPI as the appropriate dependent variable in a Phillips Curve with selling price expectations. PPI aligns with the sectoral focus of the data, tracks firms' cost-based price-setting processes, and captures inflation dynamics most relevant to short-term expectations—ensuring theoretical consistency and empirical robustness.

Within-country other-sectors inflation expectations

To capture within-country spillovers from related industries, I construct a *leave-one-out* (LOO) measure of *other sectors'* selling-price expectations. This proxy represents the pricing climate faced by a sector through its input linkages, avoiding mechanical simultaneity with the sector's own expectation. Using input–output (IO) weights keeps the measure economically meaningful (cost-linked) rather than a naïve equal-weight average. Data are sourced from the OECD Inter-Country Input–Output (ICIO) Tables, 2023 Edition.

Bilateral sector weights come from the annual IO use table, harmonised to the aggregated sectors used in the analysis.[3] For each country–using-sector–year, I compute column-normalised shares (sum to one across source sectors). Because the main panel is quarterly, the IO year is aligned to each country–sector–quarter via an exact-year match when available; otherwise by boundary assignment (nearest available year in the panel) or nearest-neighbour year within the country–sector, with a final mode fallback. IO weights are merged *once* and never compounded with additional averaging.

Let i index countries, s sectors, and t quarters. Let $w_{i,s,r}$ be the (year-aligned) IO weight for inputs from source sector r used by sector s (column share), and $E_{i,r,t}$ the selling-price expectation for sector r. Within each (i, t), I first normalise weights over the set of sectors with

[3] See the mapping in Appendix 1.A.1.

non-missing $E_{i,r,t}$ so that they sum to one. The LOO "other sectors" expectation for sector s is:

$$E_{i,s,t}^{\text{other}} = \frac{\sum_r w_{i,s,r} E_{i,r,t} - w_{i,s,s} E_{i,s,t}}{1 - w_{i,s,s}}$$

defined when $1 - w_{i,s,s}$ is sufficiently large. In practice, I set $E_{i,s,t}^{\text{other}}$ to missing in corner cases where the own-share is too dominant (e.g., $w_{i,s,s} \geq 0.90$) or when the normalising denominator is ill-conditioned.

(i) Weights are column-normalised at the using-sector level and re-normalised within (i, t) over observed sectors to prevent composition bias from missing data; (ii) no double-weighting is applied (IO weights enter only once); (iii) the LOO transformation ensures $E_{i,s,t}^{\text{other}}$ is orthogonal to the sector's own $E_{i,s,t}$ by construction.

$E_{i,s,t}^{\text{other}}$ is an input-weighted signal of the pricing environment that sector s faces via its upstream and related industries within the same country at t. It is designed to proxy cost–competition pressures that are external to s yet macro-relevant for its price-setting decisions, while mitigating simultaneity and aggregation biases.

International expectations climate

To capture the cross-country component of firms' expectations, a measure of international inflation expectations is derived by extracting the common dynamic factor underlying the co-movement of sectoral selling price expectations across euro-area countries. The rationale for including this component is that, in a highly integrated monetary union, firms' expectations are not formed in isolation: common shocks, global supply conditions, and the alignment of monetary policy expectations tend to synchronise inflation outlooks across countries. Such interdependence implies that part of the variation in national or sectoral expectations maps onto a shared international signal, rather than purely domestic factors. In standard panel Phillips Curve models, this common variation would typically be absorbed by time fixed effects or left as residual cross-sectional dependence. By isolating it explicitly, the model disentangles the influence of the shared expectations environment from domestic cyclical dynamics, yielding a cleaner interpretation of the mechanisms driving price adjustment. Extracting the international expectations climate as an explicit variable also allows state-dependent pass-through to be introduced later in the main specification.

Formally, the international expectations factor is obtained by applying a one-factor Dynamic Factor Model (DFM) to the standardised sectoral selling-price expectations from the European Commission's Business and Consumer Surveys. Prior to estimation, each sectoral series is transformed into a panel-consistent z-score, ensuring comparability across sectors and preventing scale differences from mechanically driving the common component. The DFM extracts the latent dynamic factor governing the co-movement of expectations across sectors and countries, while allowing for idiosyncratic shocks and serial correlation in the factor and idiosyncratic components. The resulting factor is then aligned in sign with the first principal component (PC1) from a parallel PCA performed on the same standardised input, which serves as a robustness benchmark. Because both DFM and PCA are fed the same standardised inputs, the extracted factor is directly comparable in magnitude and dynamic structure. Both methods yield nearly identical dynamics, with a contemporaneous correlation above 0.9, confirming that the factor effectively summarises the shared international expectations cycle. This factor thus represents the international dimension of inflation expectations and enters the Phillips Curve specification as an external conditioning variable capturing the global expectations environment.

Figure 3 illustrates the strong empirical relationship between the extracted international expectations factor and realised euro-area PPI inflation (average across countries). The tight co-movement (contemporaneous correlation of 0.92) provides face validity that the DFM successfully captures the common inflation dynamics driving price pressures across the monetary union. The broader EU sample (rather than restricting to EA countries only) is used to maximise the precision of factor extraction and avoid mechanical correlation with EA-specific outcomes, while still capturing the relevant common European inflation expectations environment. For economies with substantial non-EU trade exposure, notably Ireland and the Baltic states, the EU-based factor may capture the relevant external climate only imperfectly. Adding the lagged nominal effective exchange rate (ECB EER-18) as a separate control leaves the main estimates essentially unchanged: $\hat{\beta} = 0.277$ and $\hat{\delta} = 0.160$ against the baseline values of 0.275 and 0.173, with the NEER coefficient itself insignificant (-0.002, SE 0.014), as one would expect once the cross-sectional averages already absorb the relevant external level. Further methodological details and insights about the factor extraction procedure, heterogeneity in sector exposure and empirical validation are laid out in Appendix 1.E.

Figure 3: International expectations factor and PPI inflation

(Euro Area)

Notes: Both series are standardised (mean zero, unit variance). The contemporaneous correlation of 0.92 shows that the extracted international expectations factor closely tracks euro-area producer-price inflation.

1.4 Empirical strategy

Before detailing the augmented Phillips Curve specification, it is essential to outline the key features of the estimator, since its treatment of cross-sectional dependence and heterogeneous dynamics directly shapes the structure of the empirical model.

1.4.1 Dynamic Common Correlated Effects Mean Group (DCCE–MG) Estimator

The empirical analysis relies on the Dynamic Common Correlated Effects–Mean Group (DCCE–MG) estimator developed by Chudik and Pesaran (Pesaran, 2006; Chudik and Pesaran, 2015b,a). This estimator is designed for panels characterised by strong cross-sectional dependence, heterogeneous behavioural relationships, and dynamic adjustment, features that are central to sectoral inflation data in a monetary union.

A first motivation concerns the presence of pervasive common shocks. Euro-area sectors co-move in response to union-wide monetary policy signals, global supply conditions, and shared disturbances such as energy or supply-chain disruptions. In standard fixed-effects or pooled specifications these latent factors contaminate the error structure and bias the slope coefficients. DCCE–MG deals with this problem by augmenting each sector–country regression with cross-sectional averages of the dependent variable and all regressors, along with their lags. Here, a practical issue concerns the choice of the number of lags of the cross-sectional averages to include. Chudik and Pesaran (2015a) show that consistency requires augmenting the regression with p_T lags of the cross-section averages, where $p_T = \lfloor T^{1/3} \rfloor$. This slow–growing sequence guarantees that the augmented regression absorbs the dynamic effects of the unobserved common factors while preserving enough degrees of freedom in finite samples. Under the joint limit $(N, T, p_T) \to \infty$ with $p_T^3/T \to 0$ and $N/T \to \delta > 0$, the unit-specific coefficients $\hat{\pi}_i$ and the Mean Group estimator $\hat{\pi}^{MG}$ are both consistent, provided the factor loadings have full rank. In practice, this implies that a small number of cross-sectional lags, such as those recommended in Ditzen (2018), is sufficient for empirical work with moderate–to–large T. These additional terms absorb the impact of unobserved global forces in a non-parametric manner, ensuring that the estimated coefficients capture truly idiosyncratic sectoral responses.

A second motivation is the clear heterogeneity in price-setting behaviour across sectors. Differences in openness, cost structures, intermediate-input intensity, and competitive pressures imply that the slope of the Phillips Curve varies substantially across industries. Imposing a common coefficient through pooled estimation would obscure these differences and introduce pooling bias. The DCCE–MG estimator avoids this by estimating a separate dynamic regression for each sector–country unit and then averaging the resulting coefficients. This mean-group structure preserves slope heterogeneity while maintaining consistent estimation of the common parameters.

Dynamic adjustment is another central feature of producer-price inflation. The inclusion of lagged regressors is important to control for persistence or delayed effects, yet doing so in traditional short-T estimators introduces well-known biases. In contrast, the panel used here has both a large number of time periods and a sizeable cross-section, placing it exactly in the asymptotic regime for which the dynamic CCE framework was designed. Under these conditions, lagged regressors are weakly exogenous and the autoregressive structure can be estimated without the distortions that characterise short-panel GMM or pooled fixed-effects methods.

Beyond consistency, the estimator also facilitates the post-estimation analysis carried out in later sections. Given that each sector–country coefficient is explicitly estimated, it becomes possible to study the distribution of pass-through elasticities, construct amplification measures, and perform counterfactual aggregation exercises. The CCE augmentation ensures that these coefficients are purged of the influence of union-wide expectation shocks, thereby separating the sector-specific expectation channel from the broader expectations climate.

For these reasons, the DCCE–MG estimator provides the most appropriate econometric environment for the augmented sectoral Phillips Curve developed in the next section. Its combination of dynamic adjustment, control for latent common factors, and heterogeneous-slope estimation matches the empirical challenges posed by sectoral inflation dynamics in the euro area and yields coefficients which are interpretable within the theoretical framework introduced above.

1.4.2 Baseline sectoral Phillips curve specification

The resulting baseline sectoral Phillips curve specification takes the following configuration:

$$\pi_{i,t} = \kappa_i\, \tilde{y}_{i,t-1} \;+\; \beta_i\, E^{\text{own}}_{i,t-1} \;+\; \delta_i \left(E^{\text{own}}_{i,t-1} \times \mathbf{1}\{\text{DFM}_{t-1} \geq p75\}\right) \;+\; \theta_i\, E^{\text{other}}_{i,t-1} \;+\; \text{CSA}_t \;+\; \varepsilon_{i,t}. \quad (2)$$

where i indexes a country–sector unit and $\pi_{i,t}$ denotes quarterly PPI inflation. Equation (2) attaches a unit-specific slope to each of the four behavioural channels of the augmented Phillips curve. The term $\kappa_i\, \tilde{y}_{i,t-1}$ captures the response of inflation to the lagged sectoral output gap and identifies the standard slack channel. The coefficient β_i measures the baseline pass-through from firms' one-quarter-ahead selling-price expectations $E^{\text{own}}_{i,t-1}$ to realised producer-price inflation; because the survey is collected at the beginning of quarter $t-1$ and refers to planned price changes in quarter t, this regressor is interpreted as $\mathbb{E}_{t-1}[\pi_{i,t}]$, so that β_i in equation (2) is the firm-level analogue of the forward-looking term in the canonical New Keynesian Phillips curve. The interaction term entering equation (2) with coefficient δ_i, $E^{\text{own}}_{i,t-1} \times \mathbf{1}\{\text{DFM}_{t-1} \geq p75\}$, isolates the additional pass-through that activates when the regime indicator switches on, so that the elasticity of inflation to expectations is β_i in normal periods and $\beta_i + \delta_i$ in high-DFM regimes. The threshold is defined as the 75th percentile of the international dynamic factor model (DFM) of inflation expectations; both the factor and the expectations regressor enter equation (2) lagged by one quarter, ensuring that the regime classification is predetermined with respect to innovations in $E^{\text{own}}_{i,t-1}$ and avoiding mechanical correlation between the interaction and contemporaneous expectations.

The fourth slope of equation (2), θ_i, attaches to $E^{\text{other}}_{i,t-1}$, an input–output weighted measure of spillover expectations from other domestic sectors: positive values would point to cost transmission along the supply chain, negative values to strategic-substitution effects in pricing. The vector CSA_t collects cross-sectional averages of the dependent and explanatory variables and enters equation (2) as the Common Correlated Effects controls. Note that the international DFM factor enters equation (2) only through the regime indicator, with no stand-alone level term: the cross-sectional averages already absorb the influence of common global drivers on the level of PPI inflation, so what δ_i identifies is a shift in the *elasticity* of inflation to firms' own expectations across regimes, separately from any direct effect of the global factor on the level. The unit-specific coefficients $(\kappa_i, \beta_i, \delta_i, \theta_i)$ are estimated by DCCE–MG and subsequently averaged across units to obtain the mean-group effects reported in Table 2.

The four specifications reported in Table 2 follow a progressive structure designed to provide robustness around the baseline specification. The signs are fairly disciplined ex ante: a positive β is the standard expectations-augmented prediction, a positive δ is what models of state-dependent or selective attention (Bordalo et al., 2018; Afrouzi, 2024) would imply once firms internalise the broader expectations environment more strongly under unsettled conditions, and the output-gap slope κ should retain some bite, since a vanishing κ would mean that expectations had absorbed the entire cyclical channel, which would sit awkwardly with the supply-side logic of the model. Magnitudes, not signs, are the empirical question. Column (1) presents a restricted version of the baseline that excludes the input–output spillover term and lagged inflation; it serves to benchmark the core forward-looking mechanism. Column (2) reports the baseline preferred specification and delivers three central results, indexed to the slopes of equation (2). First, the mean-group estimate of β_i — the baseline pass-through from firms' own-price expectations $E^{\text{own}}_{i,t-1}$ to realised inflation — is strong and precisely estimated. Second, the amplification coefficient δ_i on the interaction term $IE^{\text{high}}_{i,t-1}$ is positive and highly significant, so that the high-regime pass-through $\beta_i + \delta_i$ is substantially larger than the baseline. Third, the slack coefficient κ_i on the sectoral output gap remains statistically significant and economically relevant, indicating that domestic cyclical conditions continue to exert discipline on price-setting even when expectations play a dominant role. The estimate of θ_i in equation (2), which carries the input–output weighted spillover term $E^{\text{other}}_{i,t-1}$, is statistically weak, implying that the dominant forward-looking channel operates through firms' own sectoral outlook rather than through network-mediated amplification. Despite this, the inclusion of this term helps account for residual cross-sectional dependence, as evidenced by a pronounced improvement in the CD statistic relative to the restricted model in Specification (1), which barely passes the CD test.

Columns (3) and (4) augment equation (2) with a lagged-inflation term to nest a hybrid NKPC formulation in which inflation inertia is explicitly accounted for. The estimated coefficient on π_{t-1}

is large and tightly estimated, consistent with the empirical literature on price inertia. Importantly, both the baseline pass-through β_i and the amplification slope δ_i remain statistically significant and economically meaningful, even though the inclusion of lagged inflation naturally reduces their magnitude. This stability indicates that the observed state-dependent pass-through does not arise from omitted persistence or from misspecification of the inflation process. The spillover coefficient θ_i on $E^{\text{other}}_{i,t-1}$ remains statistically weak throughout, supporting the interpretation that firms primarily internalise their own sectoral environment when forming price expectations. The negative sign may stem from strategic price-setting in supply chains: when firms anticipate stronger upstream or downstream price increases, they may hold back on their own adjustments to preserve demand or market share. In this sense, E^{other}_{t-1} acts as a stabilising rather than amplifying force. Further assessment at the individual sector level is warranted to validate this hypothesis.

The robustness checks in Table 3 further reinforce these findings. Replacing the DFM-based high-regime indicator with the first principal component of sectoral expectations yields amplification effects that are again positive, precisely estimated, and comparable in magnitude. The amplification is also robust to which economies enter the factor: re-extracting the international factor on a sample that excludes the six largest EU economies (Germany, France, Italy, Spain, the Netherlands, and Poland) leaves the effect positive and significant ($\hat{\delta} = 0.140$, SE 0.038 under the DFM; $\hat{\delta} = 0.150$, SE 0.034 under the PCA), so the mechanism does not rest on the largest economies dominating the common factor. The consistency of these results across alternative constructions of the expectations climate provides additional support for the interpretation that state-dependent coordination among firms is a robust feature of sectoral inflation dynamics rather than an artefact of a particular factor extraction method.

Table 2: Inflation Expectations and the Phillips Curve

	(1) Restricted	**(2) Baseline**	(3) Hybrid	(4) Full
π_{t-1}			0.588^{***}	0.588^{***}
			(0.019)	(0.019)
E^{own}_{t-1}	0.246^{***}	0.275^{***}	0.070^{***}	0.065^{***}
	(0.036)	(0.036)	(0.018)	(0.020)
IE^{high}_{t-1}	0.263^{***}	0.173^{***}	0.118^{***}	0.124^{***}
	(0.046)	(0.036)	(0.033)	(0.031)
$\tilde{y}_{t-1}$	0.049^{**}	0.053^{**}	0.035^{**}	0.042^{**}
	(0.022)	(0.021)	(0.017)	(0.018)
E^{other}_{t-1}		-0.032		-0.039^{*}
		(0.046)		(0.023)
Observations	9,479	9,456	9,288	9,226
Groups	136	136	130	128
R-squared (MG)	0.69	0.69	0.88	0.88
RMSE	0.57	0.57	0.36	0.36
CD p-value	0.088	0.126	0.910	0.408

Notes: Dynamic Common Correlated Effects Mean Group (DCCE-MG) estimates. The dependent variable is sectoral PPI inflation (π_t). E^{own}_{t-1} denotes firms' own-price inflation expectations; IE^{high}_{t-1} is the interaction between own expectations and a high-inflation regime indicator (DFM $\geq$ 75th percentile); E^{other}_{t-1} captures other sectors' inflation expectations; $\tilde{y}_{t-1}$ is the output gap. All specifications include cross-sectional averages of regressors to control for unobserved common factors. Standard errors robust to heteroskedasticity and cross-sectional dependence in parentheses. CD p-value from the Pesaran (2015) test for residual cross-sectional dependence. Column (2) presents the preferred specification. Columns (3)–(4) augment the model with lagged inflation to nest a hybrid NKPC.

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 3: Robustness: Alternative High-Inflation Regime Indicators

	(1) DFM	(2) PCA
E^{own}_{t-1}	0.275^{***}	0.297^{***}
	(0.036)	(0.042)
IE^{high}_{t-1}	0.173^{***}	0.123^{***}
	(0.036)	(0.032)
$\tilde{y}_{t-1}$	0.053^{**}	0.058^{**}
	(0.021)	(0.023)
E^{other}_{t-1}	−0.032	−0.026
	(0.046)	(0.046)
Observations	9,456	9,456
Groups	136	136
R-squared (MG)	0.69	0.69
RMSE	0.57	0.57
CD p-value	0.126	0.099

Notes: Dynamic Common Correlated Effects Mean Group (DCCE-MG) estimates. The dependent variable is sectoral PPI inflation. Both columns use the preferred baseline specification; they differ only in the construction of the high-inflation regime indicator. Column (1) defines high-inflation regimes using the 75th percentile of a Dynamic Factor Model (DFM) extracted from cross-country inflation expectations. Column (2) uses the first principal component (PCA) of sectoral inflation expectations. IE^{high}_{t-1} is the interaction between own-price expectations and the respective regime indicator. The amplification effect is positive and highly significant under both approaches.

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Interpreting the Amplification Mechanism

The amplification term in equation (2) captures how strongly firms' own-price expectations translate into realised inflation when the global expectations climate is elevated. The total pass-through in high-expectation states is given by the sum of the baseline pass-through coefficient β in equation (2) and the state-dependent component δ on the interaction term. This sum therefore measures the expectations–inflation elasticity in periods when the international expectations factor lies above the 75th percentile.

Across all specifications in Table 2, the combined effect $\beta + \delta$ is economically large and statistically significant, indicating that firms place substantially more weight on their price expectations when the shared expectations environment is tense. In contrast, the baseline pass-through β remains moderate in calmer periods. The amplification mechanism therefore captures a shift in the behavioural rule used by firms: expectations are not only predictive of inflation but become *more* influential precisely when the macroeconomic climate is unsettled.

In quantitative terms, interpreting the baseline specification in Column (2): since all variables are standardised, the coefficient on E^{own}_{t-1} of 0.275 means that a one–standard deviation increase in firms' own selling-price expectations is associated with a 0.28 standard deviation increase in sectoral PPI inflation when the global expectations climate is in a "normal" state. In high-expectation regimes, the marginal effect of expectations is given by the sum of the baseline and interaction coefficients, 0.275 + 0.173 = 0.448, so that the pass-through is roughly two-thirds larger than in calmer periods. Put differently, the sensitivity of prices to firms' expectations rises from about 0.28 to 0.45 standard deviations across regimes, implying a strong amplification of expectations when the international expectations climate turns "hot". Translating into price-level terms via the median within-sector standard deviations of PPI inflation (3.8 percentage points) and selling-price expectations (13 balance points), a one-standard-deviation swing in expectations raises sectoral PPI inflation by roughly 1 percentage point in normal times and by approximately 1.7 percentage points when the DFM regime is elevated, with the state-dependent component alone ($\hat{\delta} = 0.173$) contributing about 0.7 percentage points. For perspective, the same one-standard-deviation movement in the output gap ($\hat{\kappa} = 0.053$) generates approximately 0.2 percentage points, placing the expectations channel roughly five times more potent than domestic slack at driving short-run sectoral price fluctuations. In the hybrid specifications (Columns (3)–(4)), the contemporaneous coefficients shrink as lagged inflation absorbs part of the adjustment, but the implied long-run effects remain highly state-dependent: using Column (3), the long-run pass-through of expectations rises from $\beta_i/(1-\phi_i) \approx 0.17$ in normal times to $(\beta_i+\delta_i)/(1-\phi_i) \approx 0.46$ in high-expectation regimes, again pointing to a markedly stronger amplification when the expectations climate is elevated.

Figure 4 provides robustness on the discretionary choice of percentile for the "high DFM state" interaction term. The chart shows how the amplification coefficient remains largely stable and statistically significant across a wide range of percentile thresholds for defining the high-DFM state. This indicates that the state-dependent pass-through is not driven by the specific 75th percentile cutoff used in the baseline. Only at the very highest thresholds do the confidence bands widen appreciably, since these extreme DFM realisations correspond to short, volatile episodes with far fewer observations; the point estimate itself shows no systematic decline and the mechanism does not break down.

Figure 4: Amplification coefficient across DFM thresholds
(Baseline specification)

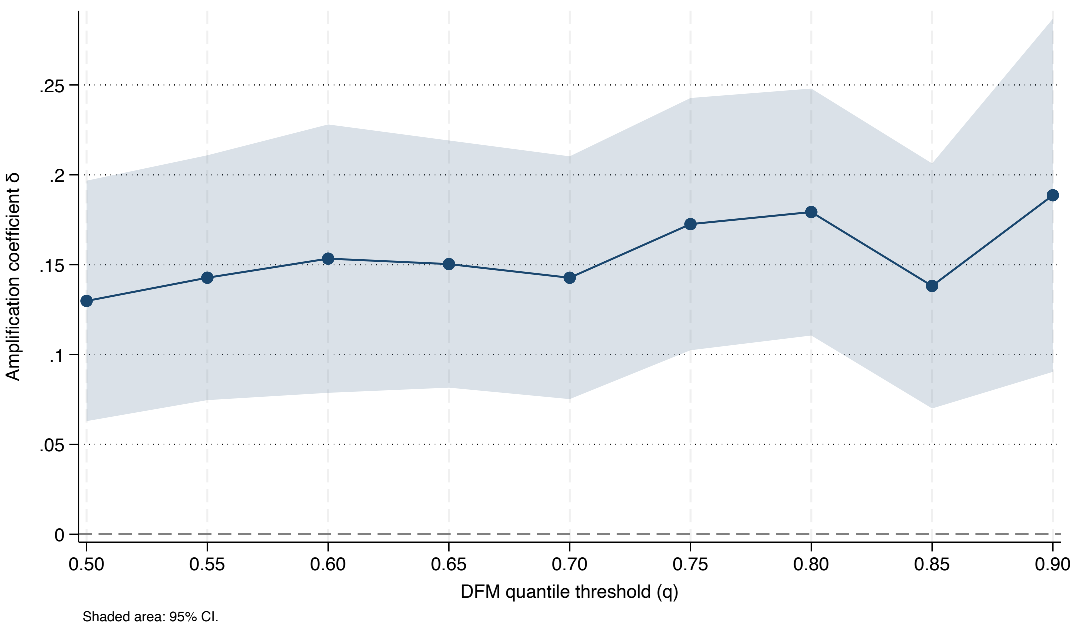

Notes: Each point reports the mean-group estimate of the interaction term between firms' selling-price expectations and a high-DFM regime indicator, where the regime is defined using different quantile cutoffs q of the international DFM distribution. Whiskers denote 95% confidence intervals.

A potential concern is that the estimated amplification might arise mechanically from stronger co-movement in expectations during high-DFM periods. Some pieces of evidence speak directly against this interpretation. (i) amplification is identified within sectoral regressions, not across sectors: the DCCE–MG estimator recovers heterogeneous slopes (β_i, δ_i) for each country–sector unit, and the dispersion of δ_i across sectors remains large even within the same global "hot" state, as will be shown in the heterogeneity analysis. If the phenomenon were mechanical, we would expect the sensitivity of inflation to expectations to rise uniformly across sectors, which is not what the estimates reveal.

(ii) the DCCE framework partials out common global shocks through cross-sectional averages and recursive demeaning, ensuring that the interaction effect is identified from idiosyncratic sector-level variation rather than shared movements in expectations or prices. If amplification merely reflected global co-movement, the coefficient δ_i would collapse once CSA terms are included; instead, it remains strongly significant.

(iii) a two-step diagnostic exercise shows that high-DFM periods are not characterised by mechanical synchronisation of expectations. The variance decomposition in Figure 5 reveals that the DFM factor explains a substantially larger share of sectoral expectations in high-DFM regimes, indicating that global signals become more informative. However, the off-diagonal correlation analysis (Table 4) shows that cross-sector correlations remain modest and do not rise in high states; if anything, they are slightly weaker. Thus, even when global signals strengthen, sectoral expectations retain substantial idiosyncratic variation rather than collapsing into a single common pattern.

(iv) the DFM factor is extracted from the broader EU sample rather than the EMU subsample used in estimation. This ensures that the regime indicator tracks a genuinely international signal rather than EMU-specific dynamics, providing a harder identification test: amplification is driven by exposure to a common European environment, not by mechanical correlation with the dependent variable's own sample.

(v) because both expectations and the DFM regime indicator enter the model lagged, the interaction term $E^{\text{own}}_{i,t-1} \times \mathbf{1}\{DFM_{t-1} \geq p75\}$ cannot reflect contemporaneous co-movement or reverse causality, further ruling out mechanical explanations.

Figure 5: Variance decomposition across DFM regimes

(Share of sectoral expectations explained by common factor)

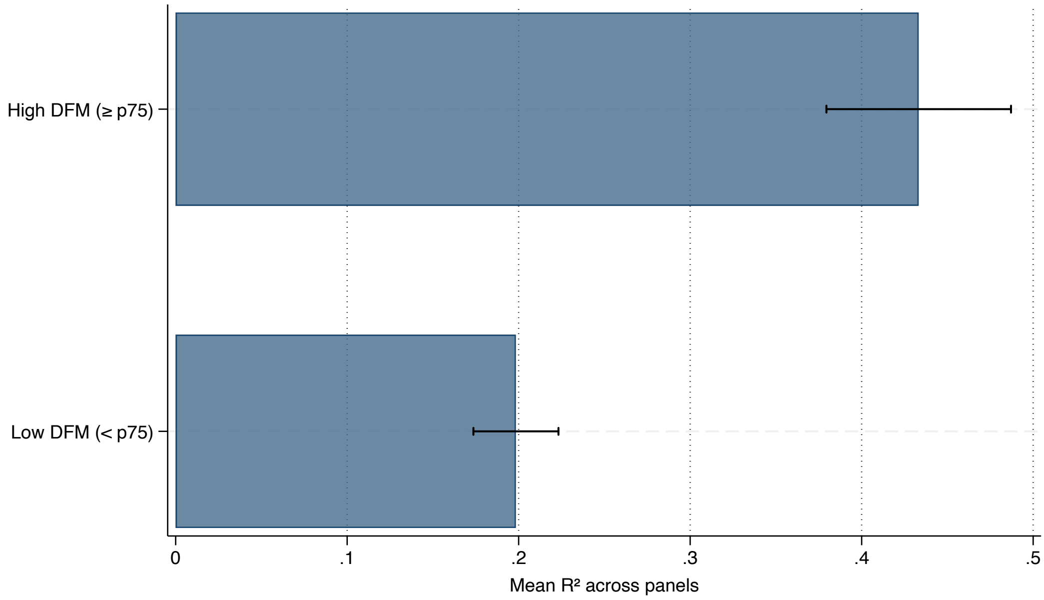

Notes: The figure reports the R^2 obtained from regressing sectoral expectations on the international DFM factor separately in low- and high-DFM states, illustrating how the explanatory power of the global component varies across regimes.

Finally (vi), as described earlier, the threshold grid results show a remarkably stable amplification plateau across percentiles, with a decline only at extreme thresholds where high-DFM observations become too sparse. A mechanical story would predict a monotonic increase in $\delta(q)$ as more firms report rising prices. Instead, the data reveal a state-dependent behavioural shift in how firms translate expectations into prices.

These diagnostics reinforce the interpretation that the amplification mechanism is structural rather than mechanical.

The aggregate results in Table 2 are deliberately a baseline rather than the paper's central finding. As a robustness check, replicating the same specification on a country-level panel, collapsing the 136 country–sector units to 17 country–quarter means and estimating with two-way fixed effects and clustered standard errors, yields $\hat{\beta} \approx 0.27$ and $\hat{\delta} \approx 0.42$, consistent with the sectoral estimates (Appendix 1.H). The aggregate pass-through and amplification mechanism is therefore not an artefact of the sectoral data structure.

The sectoral panel is not used to obtain these aggregate numbers. It is used for three reasons

Table 4: Cross-Sectoral Expectation Coordination by Inflation Regime

	Low Inflation (DFM < p75)	High Inflation (DFM ≥ p75)
Mean pairwise correlation	0.379	0.340
Valid pairs	16,770	9,684
Panels with ≥10 obs	130	99
Quarters	72	23

Notes: Average off-diagonal correlation of sectoral inflation expectations across country-sector pairs, computed separately for low- and high-inflation regimes. High inflation defined as quarters where the lagged DFM factor exceeds its 75th percentile. Correlations restricted to pairs with at least 8 overlapping observations.

that compound on one another. First, pooled fixed-effects estimation is inconsistent when slopes are heterogeneous and cross-sectional dependence is present, precisely the conditions that motivate DCCE–MG. These conditions hold at the country level too, but DCCE–MG's asymptotic guarantees require large N: with 17 countries the cross-sectional averages are imprecise and the mean-group averaging is noisy, which is why standard errors at country level are roughly fourfold larger than in the sectoral panel. Second, the sectoral panel provides $N = 136$ units, placing the estimator firmly in the asymptotic regime for which it was designed. Third, and most importantly, the heterogeneity in $\hat{\beta}$ and $\hat{\delta}$ across units is overwhelmingly a within-country phenomenon: 82% of the cross-unit variance in $\hat{\beta}$ and 92% of the variance in $\hat{\delta}$ reflects differences across sectors *within* the same country rather than between countries. This variation is structurally invisible to any country-level analysis. The remainder of the paper is devoted to unpacking it.

1.5 Heterogeneous transmission

Having estimated the baseline specification, the structural features of the DCCE-MG allow one to retrieve individual panel coefficients with their standard errors: these can then be re-aggregated at the desired level in order to gain insights grounded in the heterogeneous nature of sectors and, more broadly, countries. This enables a systematic exploration of how transmission differs across the structure of the economy. The next subsections examine these patterns at the sectoral and country-group level, building toward a single question: whether the cross-group gradient in transmission reflects how firms respond to inflation signals or merely which sectors each economy contains. The sectoral and country-group results supply the unit-level coefficients; the Oaxaca–Blinder decomposition that closes the section uses them to show the gradient is overwhelmingly behavioural—close to 88% of the cross-group difference in amplification—rather than compositional.

1.5.1 Sectoral analysis

Where the heterogeneity should sit is not an open question. The four slopes of equation (2) admit predictions tied to two structural features of European manufacturing: a sector's exposure to globally traded inputs, and its pricing autonomy, the speed at which firms re-price in response to cost and expectation shocks rather than carrying them through inventory adjustment, multi-period contracts, or globally administered prices. Sectors that combine high commodity exposure with rapid re-pricing (Metals, Chemicals, Food and Beverages, and Wood and Paper) should sit on the upper end of the β_i distribution, since their selling-price expectations track input-cost dynamics tightly and propagate downstream with little delay. Sectors organised around long-horizon procurement and globally negotiated pricing (Electronics being the canonical case) should look conspicuously muted on the same dimension. The amplification term δ_i is conjectured to track the same logic: where own expectations pass through quickly, the regime-conditional intensification triggered by a hot international expectations climate should bite as well; where they do not, δ_i should also be small. The slack slope κ_i should be sharpest in the capital-intensive sectors where capacity utilisation enters steeply convex marginal-cost regions (Machinery, Metals, Chemicals), and the cross-sector spillover θ_i is theoretically ambiguous in sign, positive under cost-pass-through dominance, negative under strategic substitution. Whether the data trace these contours cleanly or whether other forces cut across them is the empirical question.

Figure 6: Sector-level coefficient estimates — Manufacturing

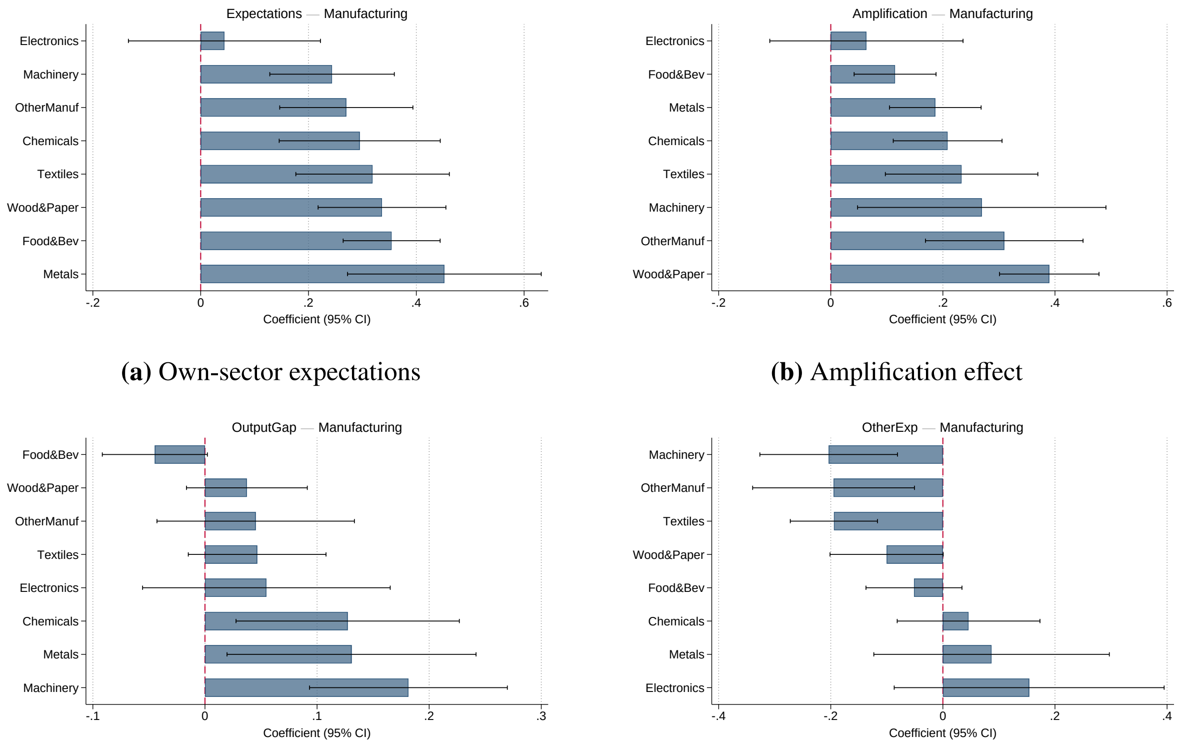

(a) Own-sector expectations

(b) Amplification effect

(c) Output gap

(d) Other sectors' expectations

Notes: Sector-level mean group coefficients from the baseline specification. Coefficients are inverse-variance weighted averages of the unit-level $\hat{\beta}_i$, $\hat{\delta}_i$, $\hat{\kappa}_i$, $\hat{\theta}_i$ across the country panels within each sector, with weights winsorised at the 95th percentile to limit the influence of any single very precisely estimated panel. Standard errors are clustered by country. Bars show point estimates; whiskers show 95% confidence intervals. All bars use a single neutral colour; significance is read off the 95% confidence intervals (whiskers).

Figure 6 presents the sector-level mean-group coefficients extracted from the DCCE–MG baseline specification. Because sectors differ sharply in their upstream/downstream position, external exposure, and pricing horizons, the heterogeneity in coefficients is a key object of interest. The four panels highlight how expectations, cross-country amplification, domestic slack, and cross-sector spillovers operate along these structural dimensions within manufacturing sectors. Estimates for services sectors are largely insignificant and reported in Appendix 1.F; this muted response is consistent with the thinner survey panels and lower price-setting frequencies in services, where the expectations–PPI correlation peaks near 0.46 against 0.94 in manufacturing, so that quarterly selling-price expectations carry correspondingly less information about realised producer prices. Country–sector coefficients are aggregated using inverse-variance weighted means with standard errors clustered by country. To avoid a handful of very precisely estimated panels dominating the sector average, inverse-variance weights are capped at the 95th percentile within each sector. As an additional robustness check, a leave-one-country-out (LOCO) diagnostic is implemented: for each sector and regressor, the pooled coefficient is recomputed excluding each country in turn, and the maximum absolute change in the sector mean is recorded. The main manufacturing sectors exhibit extremely small LOCO shifts, while larger shifts are confined to a few thin service sectors with wide confidence intervals, which are flagged in the sector summary table and not used for the core interpretation. Sector-level LOCO diagnostics and sample sizes are reported in Tables 1.F.1 and 1.F.2.

Own-sector expectations (Panel a). The prediction is borne out cleanly. The four commodity-exposed upstream sectors (Metals, Food & Beverages, Wood & Paper, and Chemicals) all post $\hat{\beta}_i$ between 0.30 and 0.45, each significant at the 1% level, with Metals at the top (0.45) and Chemicals at the bottom (0.30). At the other end, Electronics is the single manufacturing sector for which $\hat{\beta}_i$ is statistically indistinguishable from zero (0.04, $p = 0.64$), consistent with multi-period procurement, globally negotiated reference prices, and pricing horizons long enough that quarterly expectations carry little independent information beyond what is already absorbed by common factors. Textiles, Other Manufacturing, and Machinery occupy a middle band between 0.24 and 0.32, all individually significant: pricing in these sectors is forward-looking but operates over horizons longer than in the upstream commodity-exposed group. The cross-sector spread of $\hat{\beta}_i$ is therefore a graded representation of how tightly each sector's pricing horizon is bound to short-run cost dynamics.

Amplification effect (Panel b). Amplification turns out to be a broad-based feature of manufacturing rather than a property of any narrowly defined sub-group. $\hat{\delta}_i$ is positive and significant in seven of the eight manufacturing sectors, with magnitudes spanning 0.12 (Food & Beverages) to 0.39 (Wood & Paper). The lone exception is again Electronics, where $\hat{\delta}_i = 0.06$ ($p = 0.49$): the same sector that does not pass through own expectations in calm periods also does not amplify in hot ones, exactly as the unified prediction would have it. Among the sectors that do amplify, the cross-sector ranking does not mirror the ranking of $\hat{\beta}_i$. Wood & Paper and Other Manufacturing show the largest $\hat{\delta}_i$ (0.39 and 0.31 respectively) despite occupying intermediate positions in the $\hat{\beta}_i$ ordering: both are sectors whose inputs — timber and pulp; miscellaneous metals, plastics, and chemical components — are heavily traded internationally, and whose pricing therefore tightens disproportionately when the international expectations climate co-moves. That the levels and the regime sensitivities of pass-through do not rank identically is itself informative: baseline transmission and amplification capture related but distinct margins of price-setting behaviour, and the same sector can be moderate in calm periods and aggressive in hot ones.

Output gap (Panel c). The slack channel is, as predicted, narrow. $\hat{\kappa}_i$ is significant only in Machinery (0.18, $p < 0.01$), Metals (0.13, $p = 0.04$), and Chemicals (0.13, $p = 0.03$) — the three capital-intensive sectors where capacity utilisation translates most directly into marginal-cost pressure. The remaining sectors show small or insignificant gap sensitivities: in Electronics the link between domestic demand and price-setting is broken by global pricing; in Food & Beverages it is dampened by administered and policy-influenced pricing structures, notably the Common Agricultural Policy; and in Wood & Paper, Textiles, and Other Manufacturing capacity constraints simply do not bite tightly in the data. Domestic slack matters where domestic capacity is the binding constraint and not where pricing is set further away from the production margin.

Cross-sector expectations (Panel d). The cross-sector spillover term is mostly negative or insignificant once own expectations and common factors are absorbed, in line with strategic substitution dominating cost-pass-through at the inter-sectoral margin. $\hat{\theta}_i$ is significantly negative in Machinery (-0.20), Other Manufacturing (-0.20), and Textiles (-0.19): sectors that absorb upstream price signals as competitive pressure rather than as a cost mark-up to be passed forward. In the upstream commodity-exposed sectors, $\hat{\theta}_i$ is small and statistically indistinguishable from zero in Metals, Chemicals, and Food & Beverages, and weakly negative in Wood & Paper, indicating that cross-sector signals carry little marginal information once own and global components are accounted for. The takeaway is that, conditional on a sector's own expectations and on the common global factor, network-mediated propagation plays only a secondary role;

expectations formation and price-setting remain principally within-sector mechanisms.

These four panels point to a single organising margin: the speed at which a sector re-prices to its own expectations also determines whether it amplifies them when the international expectations climate turns hot, and whether domestic capacity matters at all. Electronics is the cleanest illustration — long horizons, globally administered prices, no pass-through, no amplification, no slack — while the upstream commodity-exposed sectors sit at the opposite end, with Wood & Paper and Other Manufacturing as additional cases in which heavy exposure to globally traded inputs pulls the regime sensitivity above what the baseline pass-through alone would point towards. The Phillips curve is therefore not a single relationship but a sector-specific mapping conditioned by pricing autonomy and exposure to globally traded inputs.

1.5.2 Country group analysis

A similar logic applies once the unit-level coefficients are pooled by country group. Pricing autonomy in the euro area is structurally weakest in CEE economies, which sit downstream of upstream Core producers as assembly hubs and largely take their pricing cues from abroad. Pass-through and state-dependent amplification therefore have an obvious place to be largest — the Core — and an obvious place to be muted, with the Periphery somewhere in between; whether the gradient turns out to be sharp or compressed is what the data are asked to tell us. Domestic slack carries some force where capacity-pressure channels remain meaningful (Core, and to a lesser extent Periphery) and is unlikely to do much work in CEE, where prices are largely determined by external cost structures.

Figure 7 reports country-group mean coefficients from the DCCE–MG baseline specification, where sector-level estimates are aggregated into three standard groups: Core (Austria, Belgium, Finland, France, Germany, Netherlands), Periphery (Greece, Ireland, Italy, Portugal, Spain), and CEE (Croatia, Estonia, Latvia, Lithuania, Slovakia, Slovenia). Group-level coefficients are constructed using an Empirical-Bayes inverse-variance weighting scheme, which accounts for the heterogeneous precision of unit-level estimates and delivers stable, interpretable summaries of cross-country patterns. Alternative pooling strategies, equal-weight averages, raw inverse-variance means, and balanced-panel restrictions, produce closely aligned results. These

Figure 7: Country-group coefficient estimates (empirical Bayes means)

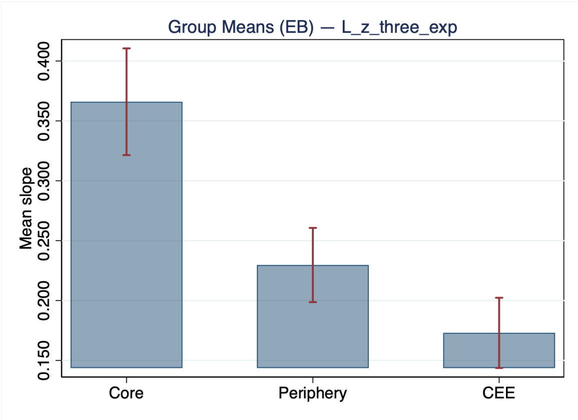

(a) Own-sector expectations

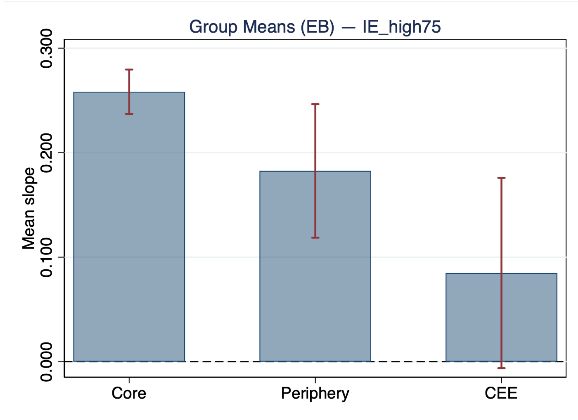

(b) Amplification effect

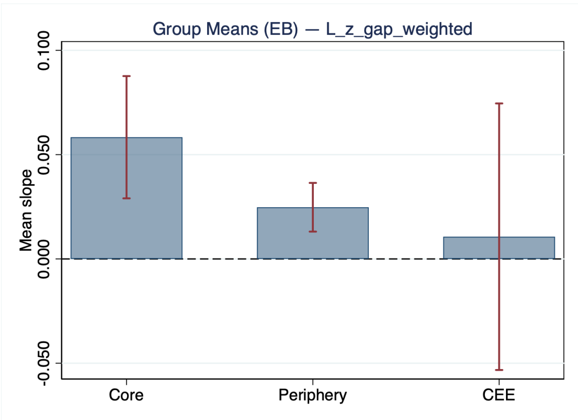

(c) Output gap

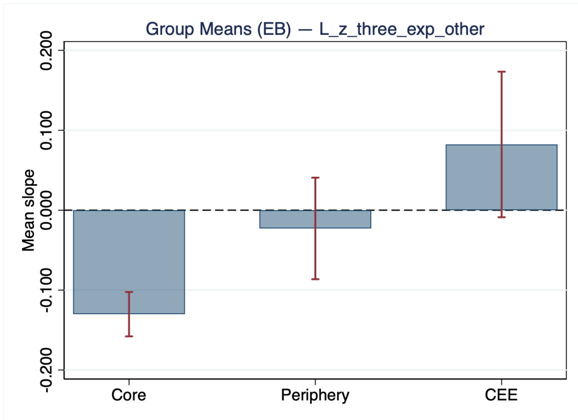

(d) Other sectors' expectations

Notes: Group-level empirical Bayes means of the panel coefficients from the baseline specification. Bars show point estimates; whiskers show 95% confidence intervals. Groups follow the Core/Periphery/CEE classification.

robustness checks are documented in Table 1.G.1. While sectoral heterogeneity captures vertical differentiation along production stages and external exposure, country-group heterogeneity highlights how the same underlying mechanisms operate across economies with distinct institutional architectures, financial structures, and positions in European production networks.

Own-sector expectations (Panel a). Pass-through from expectations to PPI inflation declines monotonically from Core to CEE economies. Core countries exhibit the strongest coefficients, consistent with deeper integration in euro-denominated value chains, greater pricing power in upstream production, and firms whose forward-looking behaviour is anchored by long established inflation-targeting credibility. Periphery economies show intermediate pass-through, rooted in their dual character: significant manufacturing capacity but also structural features (labour market rigidities, sovereign risk premia) that can lead to a decoupling of expectations from realised inflation. CEE economies display the weakest pass-through despite their manufacturing intensity, likely due to their assembly-hub role in European production: these economies host downstream activities where pricing is dictated by Western European clients and global competition, limiting firms' ability to translate expectations into autonomous price adjustments. The gradient confirms that expectations transmission is strongest where firms have pricing autonomy and credible nominal anchors.

Amplification effect (Panel b). The amplification pattern reinforces the baseline expectations hierarchy. Core economies show the strongest amplification when international expectations heat up, compatibly with their centrality in intra-European trade networks and exposure to globally-priced intermediate inputs. When the DFM regime shifts, Core firms, which operate at the nexus of European supply chains, respond most aggressively because cross-country shocks compress their pricing horizon and threaten margin stability across multiple downstream markets. Periphery economies exhibit moderate amplification, while CEE economies show the weakest response, reinforcing their role as price-takers: even when global expectations surge, assembly-oriented firms lack the market power to amplify cost pressures into proportionate price increases. The parallel gradients in baseline pass-through and amplification is most naturally read as pricing autonomy determining both the level and regime-sensitivity of expectations transmission: Core firms not only respond more to their own expectations but also intensify that response more sharply when global conditions tighten.

Output gap (Panel c). Slack effects are clearly present in both Core and Periphery economies, but vanish in CEE. Core coefficients are the largest and precisely estimated, consistent with capital-intensive, midstream manufacturing structures where capacity constraints bind quickly and marginal costs rise with demand. The Periphery also displays a positive and statistically significant slope, which squares with the fact that Southern European manufacturing sectors respond to slack even if service-sector rigidities dampen aggregate sensitivity. By contrast, the CEE estimate is close to zero and statistically indistinguishable from it. This ties well with the small open-economy logic: price setting is strongly anchored to external demand and upstream cost conditions, so domestic slack carries little information for pricing. This gradient highlights that slack-driven inflation dynamics require domestic markets and cost structures that are sufficiently insulated from global pricing.

Cross-sector expectations (Panel d). Cross-sector spillovers are statistically insignificant in the Periphery, significantly negative in the Core, and only marginally positive in CEE. The negative spillovers in Core economies point to strategic complementarity within tightly integrated production networks: when other sectors revise expectations downward, firms adjust pricing cautiously to maintain competitiveness and preserve long-term relationships in coordinated supply chains. The Periphery's near-zero and imprecisely estimated coefficients indicate weak or inconsistent cross-sector signalling, in line with more fragmented production structures and less synchronised cost pressures. In CEE economies, the small but borderline-significant positive coefficient allows for a downstream-dependence interpretation: assembly-focused industries adjust pricing in response to upstream signals originating abroad, yielding mild co-movement once own-sector expectations and global factors are controlled for. In general, cross-sector signalling channels remain secondary but informative, revealing that production network position shapes whether inter-sector expectation spillovers dampen (Core) or amplify (CEE) pricing dynamics.

These gradients are statistically, not just visually, present: the Core–CEE gap is significant for both baseline pass-through (0.20, $p < 0.001$) and amplification (0.18, $p < 0.01$), as is the corresponding Core–Periphery gap, whereas the narrower Periphery–CEE difference is at most marginal. The country-group analysis shows that expectations transmission is mediated by each economy's structural position within European production networks and by its degree of monetary integration. Only Core economies display the full Phillips-curve architecture: strong baseline pass-through, sizeable amplification in high-expectations states, and a significant slack effect. These patterns align with their role as pricing leaders with deep nominal anchors and substantial domestic cost-pressure channels. CEE economies, in contrast, exhibit weak pass-through, modest amplification, and statistically insignificant slack effects. This configuration links to their position as downstream, price-taking assembly hubs: a large share of their inflation dynamics is imported from upstream Core producers, which leaves firms with limited marginal pricing autonomy. As a result, even when the international expectations climate turns “hot”, CEE firms do not amplify their own expectations in the same way as pricing leaders do. Periphery economies fall in between: expectations matter and domestic slack retains some influence, but structural rigidities, external vulnerability, and a mixed sectoral composition temper the strength of these channels.

The broader implication is that the empirical content of the Phillips curve is not uniform within a currency union. In economies where pricing decisions are tightly embedded in international production networks, inflation dynamics are increasingly shaped by external cost structures and imported pricing norms rather than by domestic slack. As a result, the expectations-augmented Phillips curve retains structural relevance primarily in economies with sufficient pricing autonomy and credible nominal anchors, while becoming progressively weaker as production dependence rises.

This pattern does not sit in tension with the rationale for treating the international expectations factor as a state variable. The DFM is built on the observation that the *level* of inflation expectations co-moves strongly across euro-area economies, which is what allows a common indicator to classify regimes for all of them; the country-group results show that the *transmission* of those expectations into prices is much less synchronous, and that the gap between groups widens when the international signal is strongest. Co-movement of levels and divergence in pass-through are therefore complementary observations and they justify, respectively, the choice of state variable and the use of a heterogeneous-slope estimator.

It is important to note that only part of what Figure 7 shows is genuinely country-level: within each group, the unit-level $\hat{\beta}_i$ and $\hat{\delta}_i$ display sizeable dispersion across sectors, and it is this within-group sectoral variation, rather than the group means alone, that the Oaxaca decomposition in the next subsection is designed to interpret. Country size and trade openness also correlate with the Core/CEE classification and likely account for some of the spread: Core economies tend to be both larger and more deeply embedded in euro-denominated value chains, which plausibly raises pricing autonomy on its own, separately from the institutional channel emphasised here. Disentangling the two is a natural extension I leave to future work.

1.5.3 Oaxaca-Blinder decomposition

The sectoral and country-group analyses above document substantial heterogeneity in inflation expectations transmission, both within Europe's production network and across EMU country groups. A natural question is whether the cross-country differences reflect genuine variation in pricing behaviour (captured by differing slope coefficients) or instead arise from differences in sectoral composition across groups. Both readings are admissible in principle: a compositional story would be implied if Core economies happened to be over-represented in the upstream industries that already display the largest sectoral pass-through, whereas a behavioural story is more consistent with the fact that the Core/Periphery/CEE gradient lines up with pricing autonomy more closely than with industrial mix. The decomposition that follows puts numerical weight on these competing readings rather than choosing between them on intuition.

To disentangle these channels, we apply an Oaxaca-Blinder style decomposition to the group-level coefficient gaps. For any two groups A and B, the difference in weighted-average coefficients can be written as:

$$\bar{\beta}_A - \bar{\beta}_B = \underbrace{\sum_s \bar{W}_s \left(\beta_{As} - \beta_{Bs}\right)}_{\text{Within (coefficients)}} + \underbrace{\sum_s \bar{\beta}_s \left(W_{As} - W_{Bs}\right)}_{\text{Mix (composition)}} \tag{3}$$

where β_{gs} denotes the coefficient for sector s in group g, W_{gs} is the weight of sector s in group g, and $\bar{W}_s = (W_{As} + W_{Bs})/2$ and $\bar{\beta}_s = (\beta_{As} + \beta_{Bs})/2$ are the average weights and coefficients across groups.

The first term captures differences attributable to **within-sector** variation in coefficients, evaluated at common sectoral weights. The second term captures **compositional differences**, how the distribution of sectors differs across groups, evaluated at common coefficient values. This decomposition provides a descriptive accounting of the sources of cross-group heterogeneity.

The decomposition is implemented using the empirical Bayes shrinkage estimates from Specification 2 (Column (2) of Table 2), which controls for cross-sectional dependence. For each sector-country pair, inverse-variance weights are constructed from the posterior standard errors, with weights winsorised at the 95th percentile to limit the influence of imprecisely estimated coefficients. Sectoral coefficients are then aggregated within each EMU group using these weights, yielding group-sector level estimates β_{gs} and corresponding sectoral weight shares W_{gs}. Pairwise decompositions are computed for all three group comparisons (Core versus Periphery, Core versus CEE, and Periphery versus CEE) with results averaged across pairs.

To move beyond a point accounting, I attach inference to each component. Holding the precision weights fixed, both the Within and the Mix terms are linear combinations of the underlying unit-level coefficients, which the mean-group framework treats as independent across country–sectors; their sampling variances therefore follow in closed form from the estimated unit standard errors, and the absolute-value shares and their confidence intervals are obtained by simulating the implied joint normal of the two components. The resulting standard errors and intervals are reported in Table 5.

Figure 8: Oaxaca-Blinder Decomposition of Cross-Group Coefficient Differences

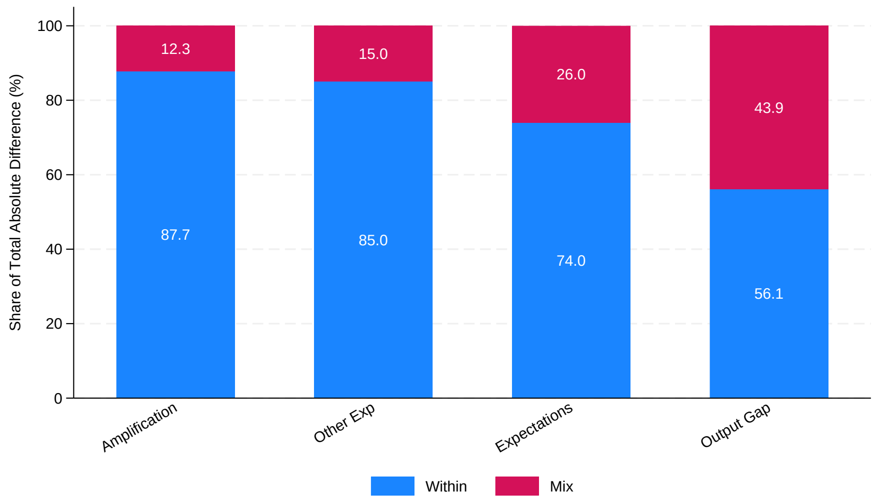

Notes: Decomposition of average pairwise coefficient differences across EMU groups into Within (coefficient heterogeneity at common weights) and Mix (compositional differences at common coefficients) components. Shares based on absolute values, averaged across all pairwise comparisons.

Figure 8 presents the point decomposition and Table 5 the associated inference. For the expectation channels the Within component dominates: it accounts for 88% of the cross-group difference in amplification, 85% for other-sector expectations, and 74% for own-sector expectations. These shares are not merely descriptive. The Within components are individually

Table 5: Inferential Oaxaca–Blinder decomposition of cross-group coefficient gaps

	Total gap (avg. pairwise)	Within (behavioural)	Mix (composition)	Within share [95% CI]
Own-sector expectations	0.132***	0.098***	0.034**	74.0%
	(0.021)	(0.025)	(0.015)	[47.9, 95.2]
Amplification	0.106***	0.093**	0.013	87.7%
	(0.037)	(0.041)	(0.025)	[28.0, 99.1]
Output gap	0.045***	0.026	0.020	56.1%
	(0.018)	(0.025)	(0.020)	[3.6, 97.7]
Other-sector expectations	−0.137***	−0.116***	−0.021	85.0%
	(0.024)	(0.029)	(0.019)	[55.4, 99.2]

Notes: Midpoint Oaxaca–Blinder decomposition of the average pairwise cross-group difference (Core/Periphery/CEE) in each empirical-Bayes coefficient, split into a Within (behavioural, common-weight) and a Mix (compositional, common-coefficient) component. Standard errors in parentheses: holding the precision weights fixed, each component is a linear combination of the independent unit-level coefficients, so its variance follows in closed form from their standard errors. The Within share is $|\text{Within}|/(|\text{Within}| + |\text{Mix}|)$, with a 95% interval from simulating the implied joint normal of (Within, Mix). $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

significant for all three expectation channels (0.093, SE 0.041 for amplification; 0.098, SE 0.025 for own expectations; −0.116, SE 0.029 for other-sector expectations), whereas the compositional components are statistically indistinguishable from zero for amplification and for other-sector expectations; only for own-sector expectations does composition make a smaller, separately significant contribution. The output gap behaves differently: its total cross-group gap is significant, but neither the Within nor the Mix component is individually well-determined (Within share 56%), so the decomposition cannot apportion it cleanly between behaviour and composition.

The decomposition remains an accounting identity rather than a causal model, but the inference shows that its behavioural reading is not an artefact of sampling noise: for amplification and other-sector expectations the cross-group gap is carried by a significant within-sector component and a compositional component that cannot be distinguished from zero. The point shares are themselves estimated imprecisely — the 95% interval for the amplification share runs from 28% to 99%, since a ratio inherits the variability of both its parts — so the component-level tests, rather than the headline percentages, are the firmer basis for interpretation. The shares are nonetheless robust to alternative weighting schemes (Appendix Table 1.G.2): the Within component dominates for amplification (85–95%), own-sector expectations (63–79%), and other-sector expectations (72–89%), while the output-gap share is method-sensitive, ranging from 16% under equal weights to 56% under precision-based weighting. Taken together, the evidence indicates that cross-group differences in expectations transmission stem primarily from how firms respond to inflation signals rather than from the sectoral composition of each economy.

1.6 Conclusion

This paper has revisited the Phillips Curve from a sectoral, forward-looking perspective, placing firms' selling-price expectations at the centre of price-setting behaviour in an open and highly integrated monetary union. By combining harmonised sectoral survey data with a dynamic common correlated effects framework, the analysis departs from aggregate representations of inflation dynamics and instead recovers how expectations are transmitted into prices across sectors, countries, and macroeconomic states.

The central result is that expectations matter, but unevenly. Across euro-area sectors, firms' own selling-price expectations pass through to producer prices even after global shocks, cross-sectional dependence, and inflation persistence are accounted for; the magnitude of that pass-through, however, varies by an order of magnitude across industries. Upstream, cost-sensitive manufacturing re-prices quickly on its expectations. Contract-intensive sectors barely move. A model that treats the euro area as a single price-setter misses this entirely, and the dispersion is large enough that it cannot be dismissed as estimation noise.

Second, the transmission of expectations appears to be state-dependent. When the international expectations climate is elevated, as captured by a common factor extracted from cross-country sectoral expectations, the pass-through from firms' beliefs to realised prices strengthens markedly. This amplification effect is stable across alternative definitions of the global expectations regime and persists in hybrid specifications that allow for intrinsic inflation inertia. The evidence points to a behavioural shift rather than a mechanical correlation: firms place greater weight on their own expectations precisely when the broader expectations environment becomes tense. In this sense, expectations are not only predictive but conditional, with their influence increasing endogenously in periods of heightened uncertainty or synchronised inflationary pressure.

Third, aggregating sector-level coefficients using empirical-Bayes methods reveals a clear gradient from Core to Periphery to Central and Eastern European economies. Core countries display strong baseline pass-through, pronounced state-dependent amplification, and a meaningful role for domestic slack, in line with their pricing leadership within European value chains. Peripheral economies occupy an intermediate position: expectations transmit into prices and amplify in high-expectations regimes, but both channels are weaker and less persistent than in the Core, while domestic slack retains only a modest influence. Central and Eastern European economies, by contrast, exhibit weaker pass-through, limited amplification, and negligible slack effects, which is grounded in their downstream, price-taking position in European production networks. An Oaxaca-type decomposition shows that these differences are driven primarily by within-sector coefficient heterogeneity rather than by sectoral composition: for amplification and other-sector expectations the within-sector component is statistically significant while the compositional component is indistinguishable from zero, indicating that institutional and structural features of price-setting dominate compositional explanations.

All in all, these results show that the Phillips Curve in an integrated monetary union is neither dead nor uniform. Instead, it is fragmented along sectoral and geographical lines and conditioned by the international expectations environment. These selling-price expectations do not propagate mechanically through the economy: they are filtered through production structures, pricing power, and firms' exposure to global conditions. A single aggregate Phillips Curve therefore obscures the mechanisms through which inflation pressures actually emerge and spread.

From a methodological perspective, the findings underscore the importance of modelling cross-sectional dependence and slope heterogeneity explicitly. Treating global forces as latent common factors rather than as residual noise is crucial for identifying the expectation channel cleanly. Likewise, exploiting sectoral expectations data allows the forward-looking component of the New Keynesian Phillips Curve to be anchored more closely to the firm-level decision problem from which it originates.

More broadly, the analysis highlights that inflation dynamics in the euro area are inherently state-contingent and uneven. Periods of synchronised expectations are the ones in which inflation becomes harder to contain, not because slack suddenly disappears, but because firms adjust their pricing rules. Understanding inflation therefore requires tracking not only levels of expectations, but also their coordination and the environments in which they become behaviourally salient.

Several extensions follow naturally. The most direct is to endogenise expectation formation rather than take the survey expectations as given, deriving the amplification mechanism from information frictions and the salience of the international climate instead of estimating it as a regime-conditional slope. A more policy-relevant path would link amplification to monetary-policy credibility, using anchoring measures or differential responses to ECB communication shocks to separate the credibility channel from the production-network one. Further work could embed sectoral wage-setting and labour-market institutions, assessing the expectations channel alongside the wage channel that disciplines it, and could formalise country size and trade openness as continuous moderators of the gradient flagged in Section 1.5 rather than the discrete Core/Periphery/CEE split.

What this paper establishes is that inflation expectations matter, but how much they matter depends on where firms operate, how they are embedded in global production networks, and when they choose to listen to them.

Appendix

1.A Sectoral Mapping

This appendix documents the aggregation used to harmonise sectoral data across countries. We map NACE Rev. 2 divisions into nineteen macro-sectors (`sector_agg`) to ensure cross-country comparability and stable coverage across sources (Eurostat STS/SBS, ECB Main Aggregates, and the EC Business and Consumer Surveys). Aggregation weights are quarterly output shares from Eurostat's Structural Business Statistics, smoothed with a centred five-quarter moving average and renormalised within each country–sector–quarter. These harmonised categories underpin all sectoral panels used in the empirical analysis.

Table 1.A.1: Mapping between aggregated sectors and NACE Rev. 2 divisions

Aggregated sector (`sector_agg`)	**Included activities (NACE Rev. 2 divisions)**	**NACE Rev. 2 codes**
Mining and Quarrying	Extraction of crude petroleum and natural gas; mining of metal ores; other mining and quarrying; mining support activities	B05–B09
Food, Beverages and Tobacco	Manufacture of food products; beverages; tobacco products	C10–C12
Textiles, Apparel and Leather	Manufacture of textiles; wearing apparel; leather and related products	C13–C15
Wood, Paper and Printing	Manufacture of wood, paper and paper products; printing and recorded media	C16–C18
Chemicals and Pharmaceuticals	Manufacture of chemicals; pharmaceuticals; rubber and plastic products	C20–C22
Metals and Fabrication	Manufacture of basic and fabricated metal products (excluding machinery)	C24–C25
Electronics and Electrical Equipment	Manufacture of computer, electronic, optical and electrical equipment	C26–C27
Machinery and Vehicles	Manufacture of machinery; motor vehicles; transport equipment	C28–C30
Other Manufacturing	Furniture; other non-metallic minerals; repair and installation of machinery and equipment	C23, C31–C33
Energy and Utilities	Electricity, gas, steam and air conditioning supply; water collection and treatment; coke and refined petroleum products	D35, E36, C19
Transport and Storage	Land, water and air transport; warehousing and postal activities	H49–H53
Wholesale and Retail Trade	Wholesale and retail trade; repair of motor vehicles and motorcycles	G45–G47
Accommodation and Food Services	Accommodation; food and beverage service activities	I55–I56
Information and Communication	Publishing, broadcasting, telecommunications, and ICT-related services	J58–J63
Financial and Insurance Activities	Financial services; insurance; reinsurance; auxiliary financial activities	K64–K66
Real Estate Activities	Real estate services	L68
Professional Services	Legal, accounting, consulting, architecture, advertising, R&D, and related professional services	M69–M75, N77
Other Services	Employment, security, cleaning, and business support activities	N78–N82
Arts, Entertainment and Recreation	Creative arts, cultural, gambling, and sports activities; membership organisations	R90–R94

Note: Aggregated sectors (`sector_agg`) are harmonised groupings of NACE Rev. 2 divisions. Weights are based on SBS output shares, smoothed over five quarters and renormalised within country–sector–quarter.

1.B Derivation of the New Keynesian Phillips Curve

This appendix reproduces the standard derivation of the forward-looking Phillips Curve following Galí (2015). All variables are in log deviations from steady state.

1.B.1 Firm's Optimal Price Setting

In period t, a firm that is allowed to re-optimise chooses a reset price P_t^* to maximise the present value of profits while that price remains in effect. Formally, it solves

$$\max_{P_t^*} \sum_{k=0}^{\infty} \theta^k \, \mathbb{E}_t \Big\{ Q_{t,t+k} \big(P_t^* Y_{t+k|t} - \Psi_{t+k}(Y_{t+k|t}) \big) \Big\}, \tag{a}$$

subject to the sequence of demand constraints

$$Y_{t+k|t} = \left(\frac{P_t^*}{P_{t+k}} \right)^{-\varepsilon} C_{t+k}, \qquad k = 0, 1, 2, \ldots, \tag{b}$$

where $Q_{t,t+k} \equiv \beta^k \left(\frac{C_{t+k}}{C_t}\right)^{-\sigma} \left(\frac{P_t}{P_{t+k}}\right)$ is the stochastic discount factor for nominal payoffs, $\Psi_{t+k}(\cdot)$ is the cost function, and $Y_{t+k|t}$ denotes output in period $t+k$ for a firm that last reset its price in period t.

The first order condition associated with the problem above takes the form:

$$\sum_{k=0}^{\infty} \theta^k \, \mathbb{E}_t \big\{ Q_{t,t+k} \, Y_{t+k|t} \, \big(P_t^* - \mathcal{M} \, \psi_{t+k|t} \big) \big\} = 0, \tag{c}$$

By log–linearizing the firm's optimal price–setting condition in equation (c) around the zero-inflation steady state, rewriting it in terms of real marginal costs, and performing a first–order Taylor expansion, we obtain the expression in equation (d) linking the reset–price gap to a discounted sum of expected future marginal costs and price movements.

$$p_t^* - p_{t-1} = (1 - \beta\theta) \sum_{k=0}^{\infty} (\beta\theta)^k E_t \big\{ \widehat{mc}_{t+k|t} + (p_{t+k} - p_{t-1}) \big\}, \tag{d}$$

$$p_t^* = \mu + (1 - \beta\theta) \sum_{k=0}^{\infty} (\beta\theta)^k E_t\left\{ mc_{t+k|t} + p_{t+k} \right\}, \tag{e}$$

Equation (d) is the log–linear optimal price–setting condition for a Calvo–resetting firm: the newly chosen price p_t^* differs from last period's aggregate price p_{t-1} by a weighted average of expected future real marginal costs and expected future price levels, where the weights $(\beta\theta)^k$ reflect the probability that the newly set price will remain effective for k periods. The rearranged expression (e) makes transparent that the optimal reset price equals the desired (log) markup μ over a discounted stream of expected nominal marginal costs and future aggregate prices.

1.B.2 Equilibrium Conditions

The following constitute the key equilibrium relationships needed for the derivation of the New Keynesian Phillips Curve.

Aggregate price dynamics. Under Calvo pricing, a fraction $1 - \theta$ of firms reset their price in any period. Log–linearizing the aggregate price index around the zero–inflation steady state gives:

$$\pi_t = (1 - \theta)\,(p_t^* - p_{t-1}), \tag{f}$$

where p_t^* is the log of the optimal reset price.

Real marginal cost of a resetting firm. The marginal cost faced by a firm resetting its price in t differs from the aggregate marginal cost because its relative output differs from aggregate output. Using the demand schedule and the production technology:

$$\widehat{mc}_{t+k|t} = \widehat{mc}_{t+k} - \frac{\lambda_p}{1 + \lambda_p}\,(p_t^* - p_{t+k}), \tag{g}$$

where $\lambda_p \equiv \dfrac{\varepsilon}{\varepsilon - 1}$ is the steady–state markup, and hats denote log–deviations from steady state.

Replacing $\widehat{mc}_{t+k|t}$ in (d) delivers the difference equation connecting the optimal reset price to aggregate inflation :

$$p_t^* - p_{t-1} = \beta\theta\, E_t(p_{t+1}^* - p_t) + (1-\beta\theta)\,\xi\,\widehat{mc}_t + \pi_t, \tag{h}$$

where $\xi \equiv \dfrac{1-\alpha}{1-\alpha+\alpha\varepsilon}$.

Substituting (h) into the aggregate price evolution (f) eliminates p_t^* and yields the log–linear inflation equation :

$$\pi_t = \beta E_t \pi_{t+1} + \kappa\,\widehat{mc}_t, \tag{i}$$

with slope $\kappa = \dfrac{(1-\theta)(1-\beta\theta)}{\theta}\,\xi$.

Aggregate real marginal cost can be expressed in terms of the **output gap** (defined as $y_t - y_t^n$, the deviation from the flexible–price efficient level). Using household optimality, production technology, and the definition of the natural output level, we obtain:

$$\widehat{mc}_t = \varphi\,(y_t - y_t^n), \tag{j}$$

where $\varphi > 0$ collects structural parameters.

New Keynesian Phillips Curve. Substituting (j) into the inflation equation (i) yields:

$$\pi_t = \beta E_t \pi_{t+1} + \kappa\,\varphi\,(y_t - y_t^n), \tag{k}$$

which is the canonical New Keynesian Phillips Curve.

1.C PPI and HICP lag-correlation

Table 1.C.1: Lagged correlations between inflation and selling price expectations (**Manufacturing**)

Lag	PPI Corr	PPI Abs Corr	HICP Corr	HICP Abs Corr
-12	0.4224342	0.4224342	0.7395272	0.7395272
-11	0.5034704	0.5034704	0.7616532	0.7616532
-10	0.5849395	0.5849395	0.7711651	0.7711651
-9	0.6616828	0.6616828	0.7744756	0.7744756
-8	0.7251012	0.7251012	0.7799655	0.7799655
-7	0.7830220	0.7830220	0.7797647	0.7797647
-6	0.8333257	0.8333257	0.7750828	0.7750828
-5	0.8737217	0.8737217	0.7630090	0.7630090
-4	0.9057574	0.9057574	0.7419083	0.7419083
-3	0.9287478	0.9287478	0.7094881	0.7094881
-2	0.9423525	0.9423525	0.6705821	0.6705821
-1	0.9426472	0.9426472	0.6262667	0.6262667
0	0.9266878	0.9266878	0.5744480	0.5744480
1	0.8917142	0.8917142	0.5396489	0.5396489
2	0.8415231	0.8415231	0.5061596	0.5061596
3	0.7848935	0.7848935	0.4746102	0.4746102
4	0.7263156	0.7263156	0.4408326	0.4408326
5	0.6661497	0.6661497	0.4067340	0.4067340
6	0.6026813	0.6026813	0.3666371	0.3666371
7	0.5361307	0.5361307	0.3289770	0.3289770
8	0.4696566	0.4696566	0.2946883	0.2946883
9	0.4021867	0.4021867	0.2610083	0.2610083
10	0.3365808	0.3365808	0.2203696	0.2203696
11	0.2720931	0.2720931	0.1663401	0.1663401
12	0.2088272	0.2088272	0.1129724	0.1129724

Table 1.C.2: Lagged correlations between inflation and selling price expectations (**Services**)

Lag	PPI Corr	PPI Abs Corr	HICP Corr	HICP Abs Corr
-12	0.1630179	0.1630179	0.0164565	0.0164565
-11	0.1025081	0.1025081	0.0297916	0.0297916
-10	0.0599905	0.0599905	-0.0343244	0.0343244
-9	0.0403249	0.0403249	-0.0356765	0.0356765
-8	0.0230747	0.0230747	-0.0662590	0.0662590
-7	0.0356542	0.0356542	0.0732447	0.0732447
-6	0.0712129	0.0712129	-0.1523122	0.1523122
-5	0.1202215	0.1202215	0.1678383	0.1678383
-4	0.1862224	0.1862224	-0.1764840	0.1764840
-3	0.2747793	0.2747793	-0.1888359	0.1888359
-2	0.3580036	0.3580036	-0.2199416	0.2199416
-1	0.4251046	0.4251046	0.2390882	0.2390882
0	0.4609951	0.4609951	-0.2681412	0.2681412
1	0.4618566	0.4618566	-0.2768661	0.2768661
2	0.4438031	0.4438031	-0.2835430	0.2835430
3	0.4057891	0.4057891	0.3472965	0.3472965
4	0.3436642	0.3436642	0.3860435	0.3860435
5	0.2690317	0.2690317	0.5009751	0.5009751
6	0.2017219	0.2017219	0.5325245	0.5325245
7	0.1320758	0.1320758	0.6610930	0.6610930
8	0.0595837	0.0595837	0.6643444	0.6643444
9	-0.0348967	0.0348967	0.7669227	0.7669227
10	-0.1275597	0.1275597	0.7666925	0.7666925
11	-0.2114548	0.2114548	0.8080058	0.8080058
12	-0.2869335	0.2869335	0.8269584	0.8269584

1.D Construction and Relevance of "Other-Sector" Expectations

This appendix documents the construction of the "other-sector" inflation expectations variables used in the empirical analysis. These variables measure, for each sector (s) in country (i) and quarter (t), the expected inflation prevailing in the remaining domestic sectors, weighted by input–output linkages. The procedure follows three key steps:

(i) **Column-normalised input–output weights:** weights are normalised at the using-sector level so that, for each country–year matrix, the shares of inputs across supplying sectors sum to one. They are then re-normalised within (i, t) over observed sectors to prevent composition bias from missing data.

(ii) **Single weighting:** no double-weighting is applied; input–output (IO) weights enter only once in constructing the cross-sectoral expectation component.

(iii) **Leave-one-out transformation:** the "other-sector" expectation $E^{\text{other}}_{i,s,t}$ is defined as the weighted mean of all other sectors' expectations within the same country and period, excluding the sector's own $E_{i,s,t}$. By construction, this ensures orthogonality between the two series.

The resulting indicators capture the contemporaneous exposure of each sector to the expectations environment of the domestic production network, weighted by its average input structure computed from the OECD Inter-Country Input–Output Tables.[4]

[4] Source: OECD Inter-Country Input–Output (ICIO) Tables, 2023 Edition.

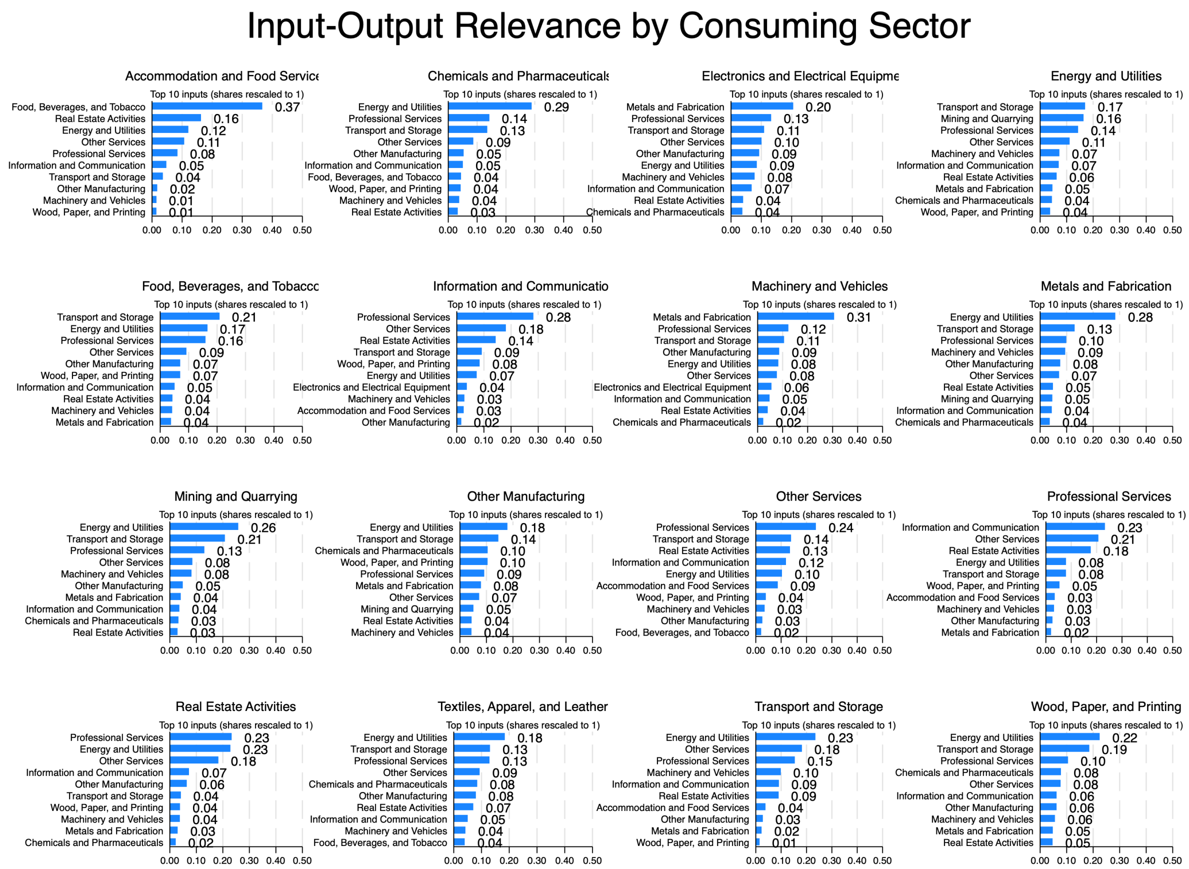

Figure 1.D.1: Average relevance of other sectors based on input–output linkages (Top 10 inputs per sector). Bars represent the average share of inputs from other sectors in each consuming sector's total intermediate consumption. Shares are column-normalised within each consuming sector and averaged across countries and years.

1.E International Expectations Factor: Methodology

1.E.1 Factor Extraction

The international expectations factor is derived from sectoral selling price expectations data covering 23 European countries and multiple sectors over the period 2000Q1–2023Q4, sourced from the European Commission's Business and Consumer Surveys. Two complementary methods are employed to extract the common component: Principal Component Analysis (PCA) and a one-factor Dynamic Factor Model (DFM) with AR(1) dynamics.

The DFM specification is:

$$X_{i,s,t} = \lambda_{i,s} F_t + \varepsilon_{i,s,t} \tag{1.E.1}$$

$$F_t = \rho F_{t-1} + \eta_t \tag{1.E.2}$$

where $X_{i,s,t}$ denotes the standardised expectation in country i, sector s, at time t; F_t is the unobserved common factor; $\lambda_{i,s}$ are the factor loadings; and $\varepsilon_{i,s,t}$ captures idiosyncratic variation. The factor follows an autoregressive process with persistence parameter ρ, allowing for temporal dynamics in the latent common component. Estimation is performed via maximum likelihood using the Kalman filter.

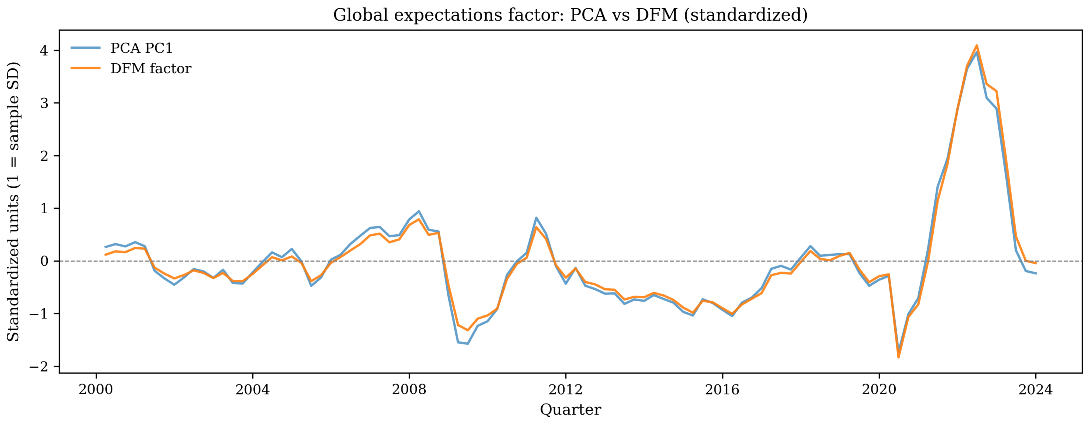

Figure 1.E.1: Comparison of first principal component (PCA) and dynamic factor model (DFM) estimates. Both methods yield nearly identical dynamics (correlation = 0.993), confirming robustness of the extracted common factor. Series are standardised to mean 0, unit variance.

As shown in Figure 1.E.1, both methods yield nearly identical estimates, with a contemporaneous correlation of 0.993. This confirms the robustness of the extracted common component to alternative identification strategies. The DFM is preferred for the main analysis as it explicitly models the temporal dynamics of the latent factor through an autoregressive structure, making it better suited to capture persistence in inflation expectations. The first principal component explains 52.6% of the total variance in standardised sectoral expectations, indicating substantial co-movement across countries and sectors.

1.E.2 Sectoral Heterogeneity in Factor Exposure

The factor loadings $\lambda_{i,s}$ from equation (1.E.1) measure each country-sector unit's sensitivity to the international expectations factor. These loadings exhibit substantial heterogeneity both across and within macro-sectors, arising from differences in exposure to global inflation dynamics.

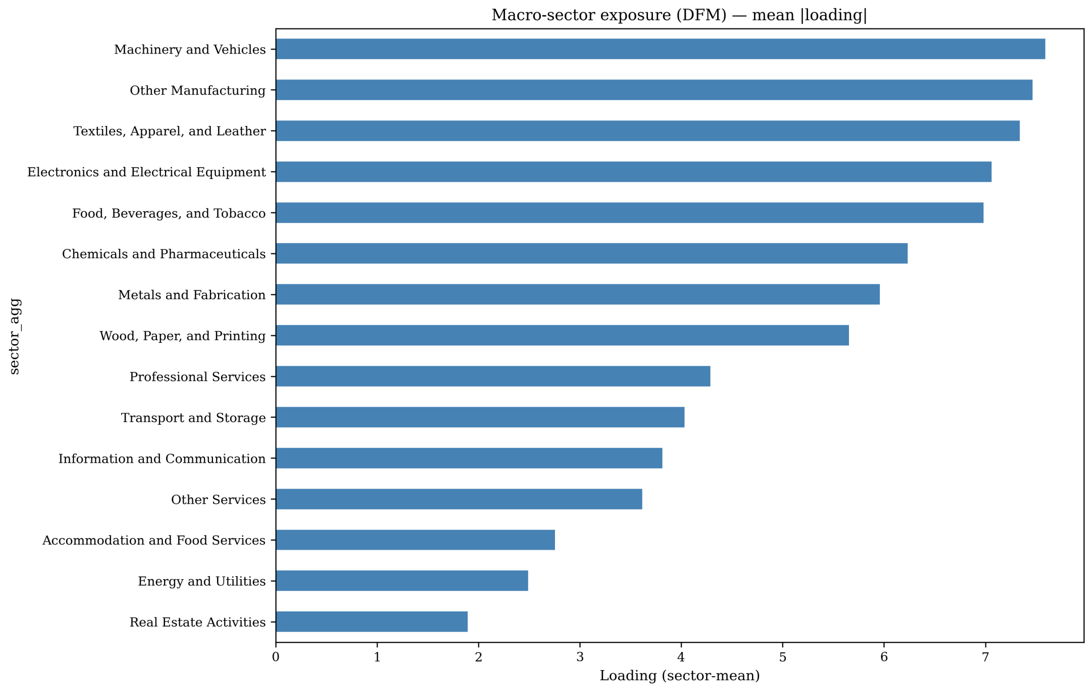

Figure 1.E.2: Mean absolute factor loadings by macro-sector. Bars show the average sensitivity of each sector to the international expectations factor, revealing systematic differences in exposure to global inflation dynamics. Higher values indicate greater responsiveness to the common international component.

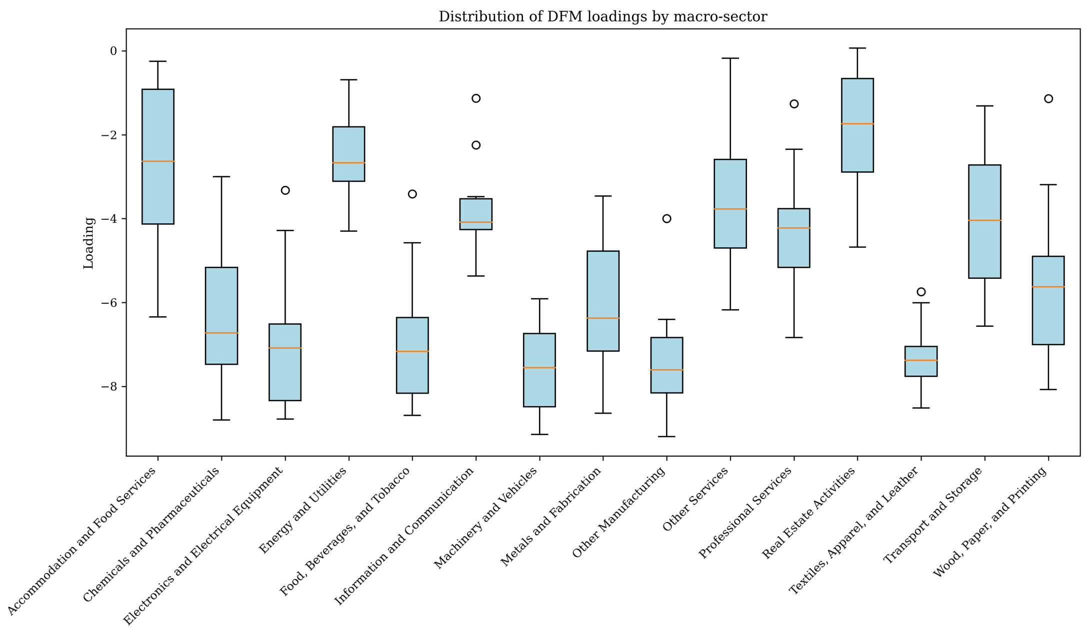

Figure 1.E.3: Distribution of factor loadings by macro-sector. Box plots show median, interquartile range, and outliers, illustrating both between-sector and within-sector heterogeneity in exposure to the international factor. The dispersion within sectors mirrors country-specific characteristics and differences in production structure.

Figures 1.E.2 and 1.E.3 document this heterogeneity. Manufacturing sectors exhibit the highest loadings, echoing their deep integration into international supply chains and exposure to global demand conditions. Service sectors, especially real estate, energy and utilities, and accommodation services, display substantially lower loadings, consistent with their more domestically-oriented nature. The within-sector variation (evident in the box plots) captures country-specific differences in production structure and trade openness.

1.E.3 Empirical Validation

The extracted factor's relevance for inflation dynamics is validated through its strong correlation with realised euro-area PPI inflation and its demonstrated predictive content.

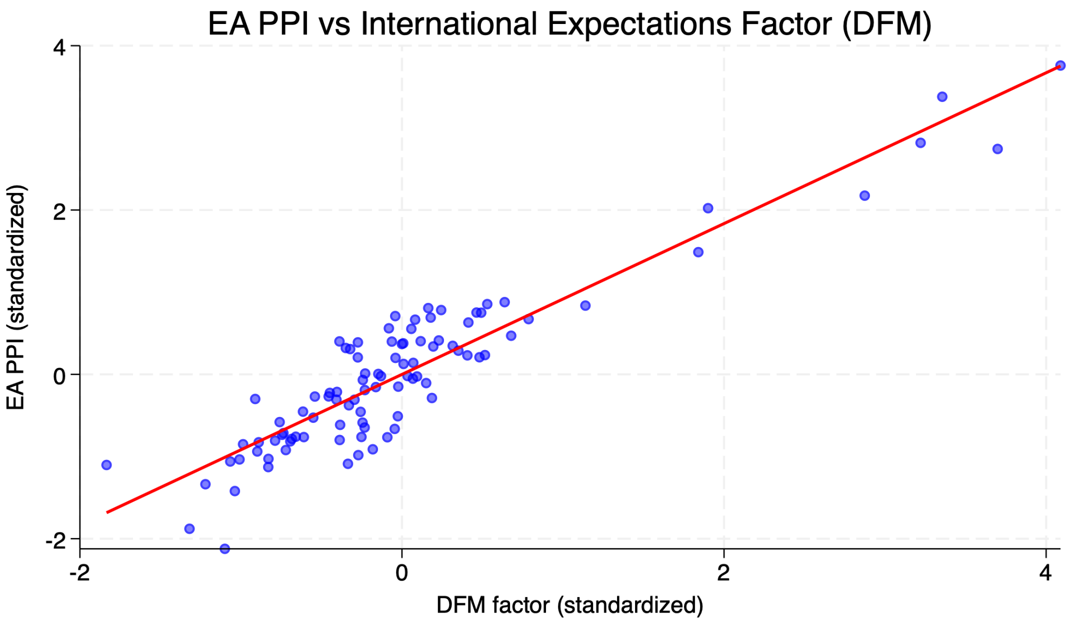

Figure 1.E.4: Scatter plot of euro-area PPI inflation and international expectations factor (both standardised). The contemporaneous correlation of 0.92 demonstrates that the extracted factor effectively tracks aggregate inflation dynamics. Each point represents one quarter (2000Q1–2023Q4).

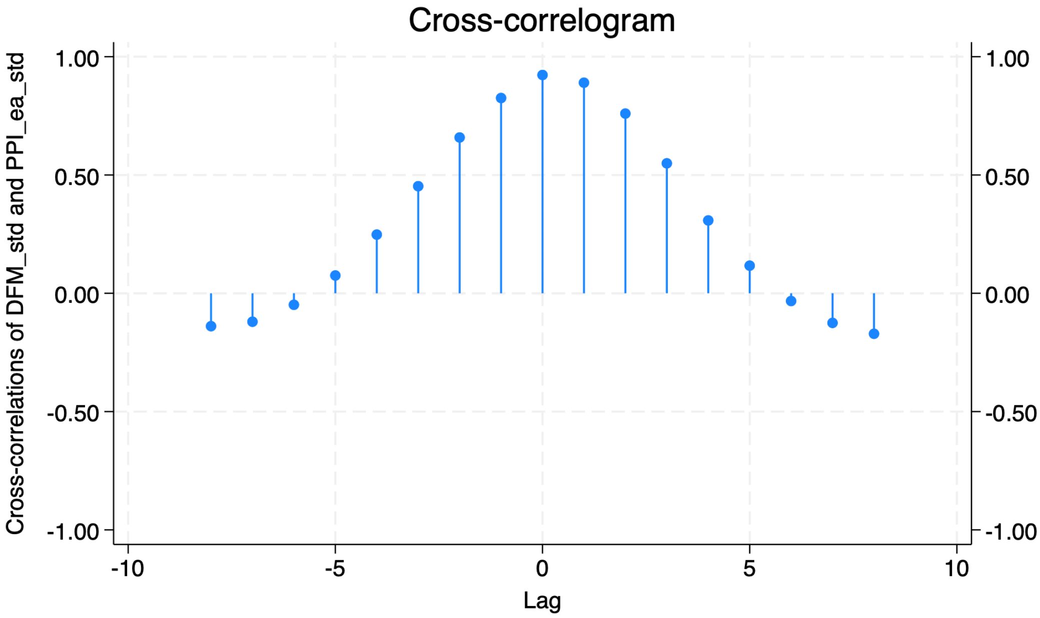

Figure 1.E.5: Cross-correlation between international expectations factor and EA PPI inflation at various leads and lags. Positive correlation at negative lags indicates the expectations factor contains forward-looking information about subsequent inflation. The peak at lag 0 confirms strong contemporaneous co-movement.

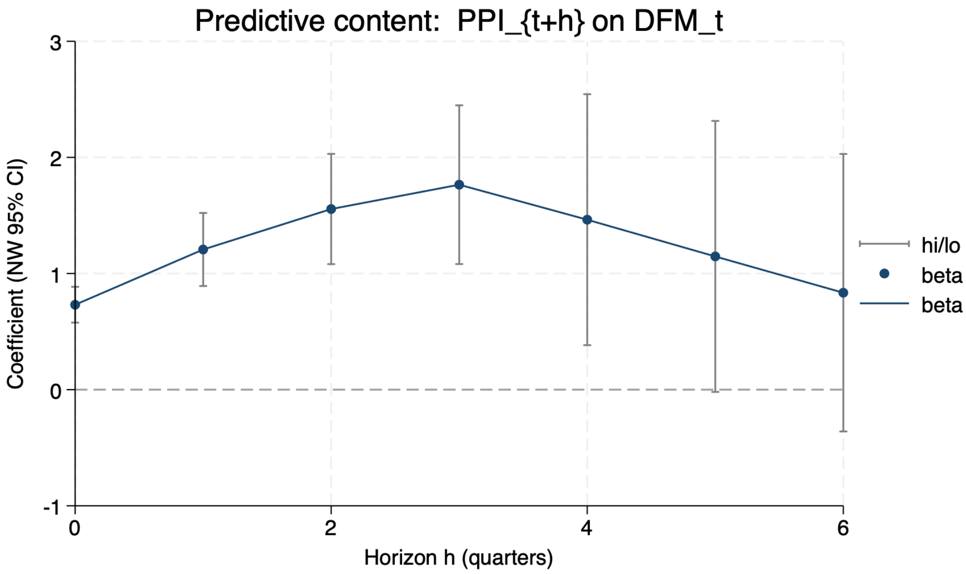

Figure 1.E.6: Local projection estimates: effect of a one-standard-deviation shock to the international expectations factor on EA PPI inflation at horizons $h = 0$ to 6 quarters ahead. Estimated via Newey-West regression with 4 lags, controlling for lagged PPI and lagged DFM. The 95% confidence intervals (gray bars) show statistically significant predictive content extending several quarters ahead.

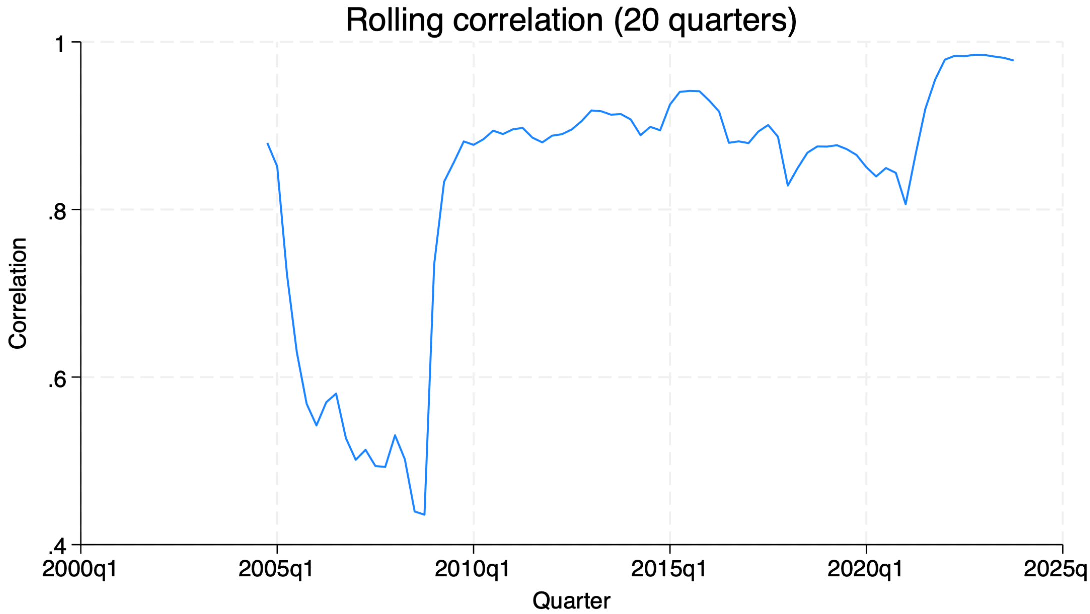

Figure 1.E.7: Rolling 20-quarter correlation between international expectations factor and EA PPI inflation. The consistently high correlation demonstrates temporal stability of the relationship across different subperiods, including the 2008 financial crisis and the 2021–2022 inflation surge.

Figure 1.E.4 shows the tight relationship between the factor and realised inflation, with a correlation of 0.92. The cross-correlation function (Figure 1.E.5) reveals that the expectations factor moves broadly contemporaneously with PPI inflation, with the strongest correlation at lag 0, consistent with forward-looking expectations incorporating information about current and near-term price pressures. Positive correlations at small negative lags suggest the factor contains modest leading information for subsequent inflation movements. Local projection estimates (Figure 1.E.6) confirm that shocks to the expectations factor have significant and persistent effects on subsequent PPI inflation, with effects remaining statistically significant through horizon $h = 4$; the contemporaneous impact is around 0.7 standard deviations and the response peaks near 1.8 at a horizon of three quarters. Finally, the rolling correlation analysis (Figure 1.E.7) indicates a strong and persistent association between the international expectations factor and EA PPI inflation. While the relationship weakens temporarily around the 2008–2009 financial crisis, it re-establishes quickly and remains high thereafter, including during the 2021–2022 inflation surge.

1.F Sectoral aggregation

Figure 1.F.1: Sector-level coefficient estimates — Services

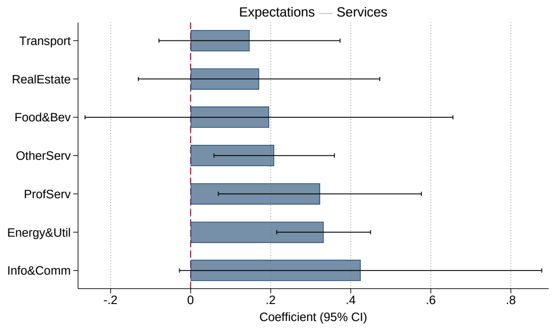

(a) Own-sector expectations

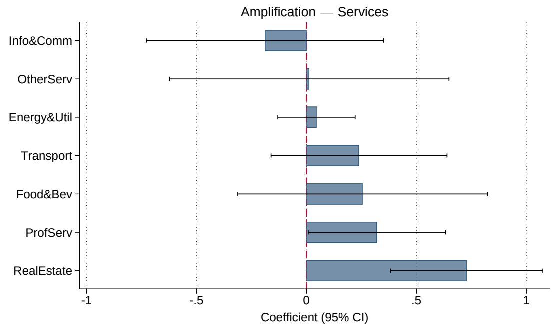

(b) Amplification effect

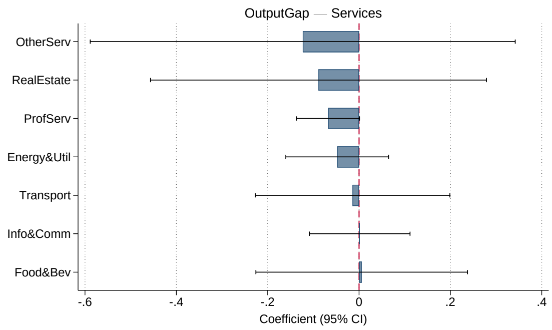

(c) Output gap

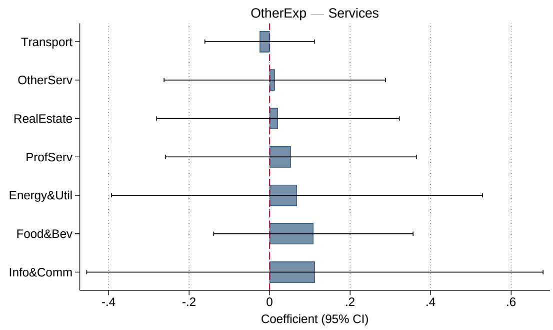

(d) Other sectors' expectations

Notes: Sector-level mean group coefficients from the baseline specification. Coefficients are inverse-variance weighted averages of the unit-level $\hat{\beta}_i$, $\hat{\delta}_i$, $\hat{\kappa}_i$, $\hat{\theta}_i$ across the country panels within each sector, with weights winsorised at the 95th percentile to limit the influence of any single very precisely estimated panel. Standard errors are clustered by country. Bars show point estimates; whiskers show 95% confidence intervals. All bars use a single neutral colour; significance is read off the 95% confidence intervals (whiskers).

Table 1.F.1: Sector-Level Coefficients with LOCO Diagnostics

Sector	Coef.	s.e.	p-value	max $\|\Delta\|$	Influential
Output gap (lagged)					
Wood, Paper, and Printing	0.037	0.027	0.201	0.023	
Transport and Storage	−0.014	0.109	0.900	0.097	
Textiles, Apparel, and Leather	0.047	0.031	0.168	0.017	
Real Estate Activities	−0.089	0.188	0.683	0.294	France
Professional Services	−0.068	0.035	0.095	0.021	
Other Services	−0.123	0.237	0.622	0.240	
Other Manufacturing	0.045	0.045	0.336	0.028	
Metals and Fabrication	0.131	0.057	0.044	0.039	
Machinery and Vehicles	0.182	0.045	0.002	0.033	
Information and Communication	0.002	0.056	0.979	0.038	
Food, Beverages, and Tobacco	−0.045	0.024	0.088	0.018	
Energy and Utilities	−0.048	0.057	0.443	0.042	France
Electronics and Electrical Equip.	0.055	0.056	0.354	0.064	
Chemicals and Pharmaceuticals	0.127	0.051	0.031	0.035	
Accommodation and Food Services	0.006	0.118	0.963	0.098	
Cross-sector expectations (lagged)					
Wood, Paper, and Printing	−0.100	0.052	0.077	0.027	
Transport and Storage	−0.025	0.069	0.732	0.110	
Textiles, Apparel, and Leather	−0.194	0.040	0.001	0.025	
Real Estate Activities	0.021	0.154	0.905	0.180	
Professional Services	0.053	0.159	0.748	0.231	Slovenia
Other Services	0.013	0.140	0.930	0.124	
Other Manufacturing	−0.195	0.074	0.023	0.042	
Metals and Fabrication	0.087	0.107	0.436	0.137	
Machinery and Vehicles	−0.204	0.063	0.009	0.036	
Information and Communication	0.113	0.289	0.709	0.357	Slovakia
Food, Beverages, and Tobacco	−0.052	0.044	0.263	0.031	
Energy and Utilities	0.068	0.235	0.784	0.338	Greece
Electronics and Electrical Equip.	0.154	0.123	0.239	0.113	
Chemicals and Pharmaceuticals	0.046	0.065	0.497	0.039	
Accommodation and Food Services	0.109	0.126	0.438	0.190	Slovenia
Own-sector expectations (lagged)					
Wood, Paper, and Printing	0.336	0.060	<0.001	0.058	
Transport and Storage	0.147	0.115	0.243	0.103	
Textiles, Apparel, and Leather	0.319	0.073	0.001	0.049	
Real Estate Activities	0.171	0.154	0.382	0.435	France
Professional Services	0.323	0.129	0.041	0.112	
Other Services	0.209	0.077	0.035	0.058	
Other Manufacturing	0.270	0.063	0.001	0.046	
Metals and Fabrication	0.452	0.092	0.001	0.095	
Machinery and Vehicles	0.244	0.059	0.002	0.045	Italy
Information and Communication	0.425	0.231	0.109	0.293	France
Food, Beverages, and Tobacco	0.354	0.046	<0.001	0.027	
Energy and Utilities	0.332	0.060	0.003	0.062	
Electronics and Electrical Equip.	0.044	0.091	0.638	0.066	
Chemicals and Pharmaceuticals	0.295	0.076	0.003	0.055	
Accommodation and Food Services	0.196	0.235	0.451	0.222	

Table 1.F.2: Sector-Level Coefficients with LOCO Diagnostics (continued)

Sector	Coef.	s.e.	p-value	max $\|\Delta\|$	Influential
High inflation regime (75th percentile)					
Wood, Paper, and Printing	0.390	0.045	<0.001	0.027	
Transport and Storage	0.239	0.204	0.280	0.250	
Textiles, Apparel, and Leather	0.233	0.069	0.007	0.045	
Real Estate Activities	0.729	0.177	0.054	0.126	Slovenia
Professional Services	0.321	0.159	0.084	0.139	
Other Services	0.013	0.324	0.970	0.304	Slovakia
Other Manufacturing	0.310	0.072	0.001	0.041	
Metals and Fabrication	0.187	0.042	0.001	0.030	
Machinery and Vehicles	0.269	0.113	0.038	0.136	
Information and Communication	−0.189	0.275	0.514	0.269	France
Food, Beverages, and Tobacco	0.115	0.037	0.011	0.019	
Energy and Utilities	0.046	0.090	0.632	0.088	
Electronics and Electrical Equip.	0.064	0.088	0.486	0.044	
Chemicals and Pharmaceuticals	0.209	0.050	0.002	0.029	
Accommodation and Food Services	0.255	0.291	0.430	0.207	Slovenia

1.G Country aggregation

Table 1.G.1: Robustness to Alternative Aggregation Methods

	Country Group		
Specification	Core	Periphery	CEE
Amplification effect (75th percentile)			
(1) EB with IVW	0.266***	0.185***	0.085*
	(0.008)	(0.033)	(0.046)
(2) EB equal weights	0.246***	0.154***	0.091***
	(0.017)	(0.029)	(0.019)
(3) Raw IVW	0.283***	0.210***	0.105
	(0.011)	(0.053)	(0.081)
(4) EB IVW balanced	0.266***	0.192***	0.071
	(0.008)	(0.034)	(0.042)
Own-sector expectations (lagged)			
(1) EB with IVW	0.369***	0.229***	0.173***
	(0.024)	(0.016)	(0.015)
(2) EB equal weights	0.348***	0.232***	0.201***
	(0.025)	(0.022)	(0.025)
(3) Raw IVW	0.381***	0.223***	0.149***
	(0.033)	(0.024)	(0.029)
(4) EB IVW balanced	0.372***	0.238***	0.174***
	(0.025)	(0.013)	(0.017)
Cross-sector expectations (lagged)			
(1) EB with IVW	−0.130***	−0.018	0.078*
	(0.014)	(0.035)	(0.044)
(2) EB equal weights	−0.130***	−0.003	0.080
	(0.022)	(0.026)	(0.053)
(3) Raw IVW	−0.131***	−0.023	0.071
	(0.020)	(0.048)	(0.062)
(4) EB IVW balanced	−0.137***	−0.016	0.080*
	(0.016)	(0.038)	(0.042)
Output gap (lagged)			
(1) EB with IVW	0.058***	0.023***	0.011
	(0.015)	(0.006)	(0.033)
(2) EB equal weights	0.071***	0.027***	0.021
	(0.015)	(0.006)	(0.038)
(3) Raw IVW	0.054**	0.016	−0.002
	(0.021)	(0.015)	(0.034)
(4) EB IVW balanced	0.059***	0.023***	0.001
	(0.015)	(0.007)	(0.038)

Notes: Standard errors in parentheses. Specifications: (1) Empirical Bayes shrinkage with inverse-variance weighting; (2) Empirical Bayes with equal country weights; (3) Raw country-level estimates with inverse-variance weighting (no shrinkage); (4) Empirical Bayes IVW restricted to balanced panels. Country groups: **Core** = AT, BE, DE, FI, FR, NL; **Periphery** = EL, ES, IE, IT, PT; **CEE** = EE, HR, LT, LV, SI, SK. (LU, CY and MT belong to the Core/Periphery classification scheme but are absent from the estimation panel for data-coverage reasons.) $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Figure 1.G.2: Country-group mean coefficients - robustness across pooling specifications

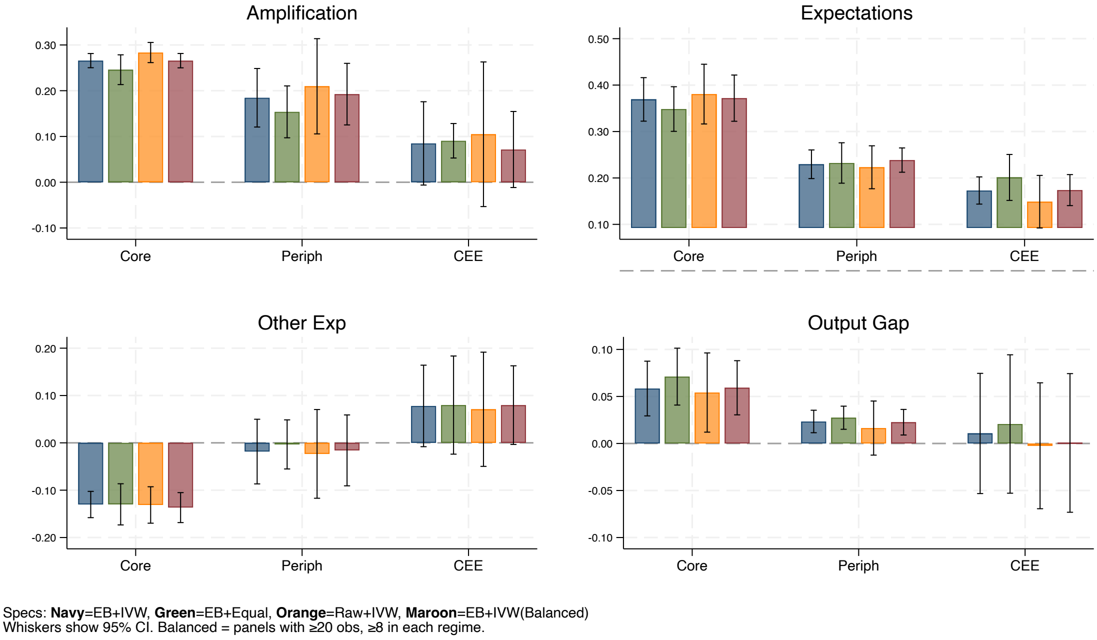

Notes: The figure compares country-group mean coefficients across alternative pooling strategies: Empirical Bayes with inverse-variance weighting (EB+IVW), Empirical Bayes with equal weights (EB+Equal), raw inverse-variance weighting (Raw+IVW), and EB+IVW restricted to balanced panels. Whiskers denote 95% confidence intervals. Balanced panels require at least 20 observations and at least 8 observations in each regime.

Table 1.G.2: Oaxaca-Blinder Decomposition - Robustness to Weighting

Variable	Component	EB+IVW	Equal	Raw IVW
Amplification	Within	87.7	94.5	84.7
	Mix	12.3	5.5	15.3
Expectations	Within	74.0	79.2	62.7
	Mix	26.0	20.8	37.3
Other Exp	Within	85.0	89.1	72.3
	Mix	15.0	10.9	27.7
Output Gap	Within	56.1	15.9	56.1
	Mix	43.9	84.1	43.9

Notes: Entries show the share (%) of total absolute group differences attributable to each component, averaged across all pairwise comparisons. Within tracks coefficient heterogeneity; Mix maps onto sectoral composition differences. EB+IVW uses empirical Bayes shrinkage estimates with inverse-variance weights; Equal uses equal weights; Raw IVW uses unshrunk estimates with inverse-variance weights.

1.H Country-level comparison

This appendix reports the country-level robustness check referenced in the discussion that follows Table 2. The exercise asks whether the aggregate pass-through and amplification documented at the sectoral tier require sectoral disaggregation or could in principle be obtained from country-level data alone. The answer has two parts. First, the aggregate state-dependent pass-through reproduces at the country tier, confirming that the mechanism is not an artefact of the sectoral panel structure. Second, the heterogeneity that constitutes the paper's main contribution is overwhelmingly within-country and is therefore invisible to any estimation that operates on country-level means.

Country-tier panel and estimation

The country-level panel is constructed by collapsing the sectoral panel used in the baseline analysis to country–quarter cells, taking equal-weight means across the nineteen aggregated sectors within each country–quarter. The variables of interest are then re-standardised at the country level so that the resulting series carry the same z-score interpretation as the sectoral counterparts. Two estimators are then applied to the resulting panel of euro-area countries over the same horizon. The first is a two-way fixed-effects regression with country and quarter dummies and standard errors clustered by country. The second is the DCCE–MG estimator with the same configuration as the sectoral baseline (`cr_lags(0) recursive`), with the regime indicator, own and other-sector expectations, and the lagged sectoral output gap entering as in Specification 2 of Table 2.

Table 1.H.1 reports the three sets of estimates side by side. The country two-way fixed-effects model returns an own-expectation coefficient of 0.27 and an amplification coefficient of 0.42, the latter significant at the one-percent level. The country-tier DCCE–MG estimator returns 0.30 and 0.18 respectively. Both columns are economically aligned with the sectoral baseline ($\hat{\beta} = 0.275$, $\hat{\delta} = 0.173$): the state-dependent amplification mechanism survives a country-level reduction. Standard errors at the country tier are roughly fourfold larger than in the sectoral panel, which is the expected consequence of moving from $N = 136$ to $N \approx 17$ when the asymptotic

Table 1.H.1: Sectoral vs. country-level estimation of the augmented Phillips curve

	(1) Sectoral DCCE–MG	(2) Country DCCE–MG	(3) Country two-way FE
E^{own}_{t-1}	0.275^{***}	0.303^{**}	0.272^{*}
	(0.036)	(0.131)	(0.156)
IE^{high}_{t-1}	0.173^{***}	0.181^{*}	0.419^{***}
	(0.036)	(0.101)	(0.141)
$\tilde{y}_{t-1}$	0.053^{**}	0.076	0.043
	(0.021)	(0.049)	(0.035)
Observations	9,456	1,092	1,092
Cross-sectional units	136 (country×sector)	16 (countries)	17 (countries)
CD p-value	0.126	0.731	—

Notes: Column (1) reproduces the baseline sectoral DCCE–MG estimates from Table 2, column (2). Columns (2) and (3) collapse the sectoral panel to country–quarter equal-weight means and re-standardise at country level. Column (2) re-estimates the same DCCE–MG specification at the country tier (`cr_lags(0) recursive`); one country is dropped by the estimator for insufficient within-panel coverage, leaving $N = 16$. Column (3) is a two-way fixed-effects regression with standard errors clustered by country. All three columns include E^{other}_{t-1} as a regressor (coefficient suppressed for compactness). CD p-value from the Pesaran (2015) test for residual cross-sectional dependence; not applicable to two-way FE. Standard errors in parentheses.

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

justification of mean-group estimation rests on a large cross-section. The DCCE–MG estimator is consistent in both panels but markedly more efficient at the sectoral tier.

Variance decomposition

A more decisive question is whether the heterogeneity exploited later in the paper has any country-level analogue. Using the unit-level coefficients $(\hat{\beta}_i, \hat{\delta}_i)$ recovered from the sectoral DCCE–MG baseline, Table 1.H.2 decomposes the cross-unit variance into a between-country component (the variance of country-level means) and a within-country component (the residual, capturing dispersion across sectors of the same country). The within-country share is 82.2% for the own-expectations coefficient and 91.9% for the amplification coefficient. The bulk of the cross-unit variation in pass-through and amplification therefore lives across sectors within the same country, not between countries.

Table 1.H.2: Within-country vs. between-country variance of unit-level coefficients

	Total variance	Between-country	Within-country	Within share
Own expectations ($\hat{\beta}_i$)	0.1768	0.0315	0.1453	82.2%
Amplification ($\hat{\delta}_i$)	0.1743	0.0141	0.1603	91.9%

Notes: Variance decomposition of the country–sector–level coefficients $\hat{\beta}_i$ and $\hat{\delta}_i$ recovered from the sectoral DCCE–MG baseline (Specification 2 in Table 2), pooled across the $N = 136$ country–sector panels. Total variance is the cross-unit variance of each coefficient; between-country variance is the variance of country-level means computed as $\bar{x}_c = \text{mean}_{s \in c}\, \hat{x}_{c,s}$; within-country variance is the residual. The within-country share is the percentage of total cross-unit variation that lives across sectors of the same country and is therefore unrecoverable from country-level estimation.

This decomposition makes the cost of aggregation explicit. An estimator that operates only on country-level means averages out roughly four-fifths of the cross-unit variation in $\hat{\beta}_i$ and over nine-tenths of the variation in $\hat{\delta}_i$ before estimation begins. The sectoral, country-group, and Oaxaca–Blinder analyses developed in Appendix 1.G and the corresponding part of the main text are designed to recover precisely this within-country variation, which a country-level reduction cannot.

References

Afrouzi, H. 2024. Strategic inattention, inflation dynamics, and the nonneutrality of money. Journal of Political Economy. 132 (10), 3378–3420.

Altissimo, F., Ehrmann, M., Smets, F. 2006. Inflation persistence and price-setting behaviour in the euro area: A summary of the inflation persistence network. ECB Occasional Paper Series. (46).

Auer, R., Borio, C., Filardo, A. 2017. The globalisation of inflation: The growing importance of global value chains. BIS Working Papers. (602).

Axioglou, C., Friz, R., Morice, F., Wozniak, P. 2024. Selling price expectations and headline inflation: Developing predictive models for hicp inflation using eu business survey data on selling price expectations. technical report, European Commission, Draft presented at the CIRET Conference.

Ball, L. M., Mazumder, S. 2011. Inflation dynamics and the great recession. Brookings Papers on Economic Activity. 2011 (1), 337–405.

Ball, L., Mazumder, S. 2019. A Phillips curve with anchored expectations and short-term unemployment. Journal of Money, Credit and Banking. 51 (1), 111–137.

Barnichon, R., Mesters, G. 2021. The Phillips multiplier. Journal of Monetary Economics. 117, 689–705.

Bartiloro, L., Bottone, M., Rosolia, A. 2019. Firms' inflation expectations and pricing strategies. International Journal of Central Banking. 15 (5), 233–283.

Bartiloro, L., Bottone, M., Rosolia, A. 2022. The heterogeneity of the inflation expectations of italian firms along the business cycle. Occasional Papers699, Bank of Italy.

Blanchard, O. 2016. The Phillips curve: Back to the '60s? American Economic Review. 106 (5), 31–34.

Bobeica, E., Jarociński, M. 2019. Missing disinflation and missing inflation: A VAR perspective. International Journal of Central Banking. 15 (1), 199–232.

Bordalo, P., Gennaioli, N., Shleifer, A. 2018. Diagnostic expectations and credit cycles. Journal of Finance. 73 (1), 199–227.

Borio, C., Filardo, A. 2007. Globalisation and inflation: New cross-country evidence on the global determinants of domestic inflation. BIS Working Papers. (227).

Borio, C., Filardo, A., Lombardi, M. 2023. Globalisation and inflation: Re-examining the global Phillips curve. BIS Working Papers. (1097).

Chudik, A., Pesaran, M. H. 2015a. Common correlated effects estimation of heterogeneous dynamic panel data models with weakly exogenous regressors. Journal of Econometrics. 188 (2), 393–420.

Chudik, A., Pesaran, M. H. 2015b. Large panel data models with cross-sectional dependence: A survey. In B. H. Baltagi ed. The Oxford Handbook of Panel Data, 2–45, Oxford, Oxford University Press.

Clarida, R., Galí, J., Gertler, M. 1999. The science of monetary policy: A new keynesian perspective. Journal of Economic Literature. 37 (4), 1661–1707.

Coibion, O., Gorodnichenko, Y. 2015a. Information rigidity and the expectations formation process: A simple framework and new facts. American Economic Review. 105 (8), 2644–2678.

Coibion, O., Gorodnichenko, Y. 2015b. Is the phillips curve alive and well after all? inflation expectations and the missing disinflation. American Economic Journal: Macroeconomics. 7 (1), 197–232.

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

Constâncio, V. 2015. Understanding inflation in the euro area. Speech at the 2015 US Monetary Policy Forum, New York.

Del Negro, M., Lenza, M., Primiceri, G. E., Tambalotti, A. 2015. The missing disinflation and the 'great recession'. In NBER International Seminar on Macroeconomics 2015, 231–277, University of Chicago Press.

Del Negro, M., Lenza, M., Primiceri, G., Tambalotti, A. 2020. What's up with the Phillips curve? Brookings Papers on Economic Activity. 301–373.

Ditzen, J. 2018. Estimating dynamic common-correlated effects in Stata. The Stata Journal. 18 (3), 585–617.

European Commission. 2024. Working with business and consumer survey (bcs) data. technical report, Directorate-General for Economic and Financial Affairs, Explains construction of sectoral “selling price expectations” balance indicators.

Fabiani, S., Druant, M., Hernando, I. et al. 2006. What firms’ surveys tell us about price-setting behavior in the euro area. In S. Fabiani, C. Loupias, F. Room eds. Monetary Policy and the Inflation Process, 38–89, London, Palgrave Macmillan.

Forbes, K., Gagnon, E., Collins, C. 2021. Global disinflation and the flattening Phillips curve. Brookings Papers on Economic Activity.

Forbes, K. J. 2019. Has globalization changed the inflation process? BIS Working Papers. 791.

Friedman, M. 1968. The role of monetary policy. American Economic Review. 58 (1), 1–17.

Friedrich, C. 2016. The ’missing’ disinflation in the euro area: a quantitative investigation. technical report, Bank of Canada Staff Working Paper.

Galí, J. 2015. Monetary Policy, Inflation, and the Business Cycle. 2nd edition, Princeton University Press.

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

Ha, J., Kose, M. A., Ohnsorge, F. 2019a. Global inflation synchronization. World Bank Policy Research Working Paper. (8768).

Ha, J., Kose, M. A., Ohnsorge, F. 2019b. Inflation in Emerging and Developing Economies: Evolution, Drivers, and Policies. World Bank.

Hamilton, J. D. 2018. Why you should never use the Hodrick–Prescott filter. The Review of Economics and Statistics. 100 (5), 831–843.

Hazell, J., Herreño, J., Nakamura, E., Steinsson, J. 2022. The slope of the phillips curve: Evidence from u.s. states. Quarterly Journal of Economics. 137 (3), 1299–1344.

Hooper, P., Mishkin, F. S., Sufi, A. 2020. Prospects for inflation in a high pressure economy: Is the Phillips curve dead or is it just hibernating? Research in Economics. 74 (1), 26–62.

Imbs, J., Jondeau, E., Pelgrin, F. 2011. Sectoral Phillips curves and the aggregate Phillips curve. Journal of Monetary Economics. 58 (4), 328–344.

Jørgensen, P. L., Lansing, K. J. 2025. Anchored inflation expectations and the flattening Phillips curve. Journal of Monetary Economics. 178.

Kohlsheen, E., Moessner, R. 2022. Globalisation and the slope of the phillips curve. Economics Letters. 216.

Lucas, R. E. 1976. Econometric policy evaluation: A critique. Carnegie-Rochester Conference Series on Public Policy. 1, 19–46.

Luengo-Prado, M. J., Rao, N., Sheremirov, V. 2018. Sectoral inflation and the phillips curve: What has changed since the great recession? Economics Letters. 172, 63–68.

Mankiw, N. G., Reis, R. 2002. Sticky information versus sticky prices: A proposal to replace the new keynesian phillips curve. Quarterly Journal of Economics. 117 (4), 1295–1328.

McLeay, M., Tenreyro, S. 2019. Optimal inflation and the identification of the Phillips curve. NBER Macroeconomics Annual. 33, 199–255.

Milani, F. 2007. Expectations, learning and macroeconomic persistence. Journal of Monetary Economics. 54 (7), 2065–2082.

Pesaran, M. H. 2006. Estimation and inference in large heterogeneous panels with a multifactor error structure. Econometrica. 74 (4), 967–1012.

Pesaran, M. H. 2015. Testing weak cross-sectional dependence in large panels. Econometric Reviews. 34 (6–8), 1089–1117.

Phelps, E. S. 1967. Phillips curves, expectations of inflation, and optimal unemployment over time. Economica. 34 (135), 254–281.

Phillips, A. W. 1958. The relation between unemployment and the rate of change of money wage rates in the united kingdom, 1861–1957. Economica. 25 (100), 283–299.

Ravn, M. O., Uhlig, H. 2002. On adjusting the Hodrick–Prescott filter for the frequency of observations. The Review of Economics and Statistics. 84 (2), 371–376.

Roberts, J. M. 1995. New keynesian economics and the phillips curve. Journal of Money, Credit and Banking. 27 (4), 975–984.

Samuelson, P. A., Solow, R. M. 1960. Analytical aspects of anti-inflation policy. American Economic Review. 50 (2), 177–194.

Sargent, T. J. 1987. Dynamic Macroeconomic Theory. Harvard University Press.

Seydl, J., Spittler, M. 2016. Did globalization flatten the Phillips curve? U.S. consumer price inflation at the sectoral level. Journal of Post Keynesian Economics. 39 (3), 387–410.

Stock, J. H., Watson, M. W. 2019. Slack and cyclically sensitive inflation. NBER Working Papers. 25987, 223–282.

Vermeulen, P. et al. 2007. Price setting in the euro area: Some stylised facts from individual producer prices data. ECB Working Papers Series. 727.

Woodford, M. 2003. Interest and Prices: Foundations of a Theory of Monetary Policy. Princeton University Press.

Yotzov, I. 2024. The speed of firm response to inflation. Presentation based on UK Decision Maker Panel (DMP) data, European Central Bank Workshop on Expectation Surveys.