Reading Between the Lines: Uncovering Inflation Expectations from Multilingual Media Coverage*

Alexandra Piller

Study Center Gerzensee & University of Bern

First version: November 20, 2025
This version: August 29, 2026

Latest version

Abstract

This paper examines how households' inflation beliefs respond to country-specific media coverage of European Central Bank (ECB) communication. Using 79,000 newspaper articles published between 2013 and 2025, I construct country-specific measures of frequency and direction of ECB-related coverage and link them to inflation perceptions and expectations. I uncover substantial cross-country heterogeneity in media transmission: German media frequently tie ECB communication to inflation, while Spanish and Italian media focus more on economic developments. Estimates from a BVAR suggest that frequency raises inflation perceptions by up to 1pp, in addition, direction shifts expectations by up to 0.4pp. A signal-extraction model shows that households rely more on media signals during the post-2021 inflation surge. Embedding these differences in a New Keynesian model opens a role for country-specific communication: Pairing targeted messaging with the common interest rate raises union welfare. The findings highlight the scope of targeted monetary policy communication in a multilingual currency union.

JEL Codes: E31, E52, D83, C32, C45

*I thank Isaac Baley, Luca Benati, Pierpaolo Benigno, Carola Binder, Martin Brown, Olivier Coibion, Francesco D'Acunto, Thomas Drechsel, Stefano Eusepi, Jordi Galí, Julian Ludwig, Michael McMahon (discussant), Galo Nuño, Oliver Pfäuti, Ricardo Reis, and Michael Weber for their helpful comments. I greatly benefited from comments from seminar participants at the EEA-ESEM 2026, Workshop on New Frameworks for Monetary Policy Analysis, ECONDAT 2026 Spring meeting, 2nd Kiel-CEPR Conference on Monetary Policy, Brown Bag Seminars at the Study Center Gerzensee and the University of Bern, the Macro Writing Workshop at the University of Texas at Austin, the Macro PhD Seminar at the University of Bern and the Research Seminar at the Riksbank. I thank the Data Science Lab of the University of Bern for their translation services. Email: alexandra.piller@szgerzensee.ch

1 Introduction

We are beginning to understand how households are exposed to and respond to central bank communication. It took sixteen years, from when Blinder et al. (2008) argued that central bank communication should extend beyond Wall Street to Main Street until Blinder et al. (2024) note that central banks have begun to act on this insight, aiming to shape public expectations directly. Yet households are only partially informed about monetary policy and central bank objectives (Van Der Cruijsen et al., 2015; Carvalho and Nechio, 2014; Binder, 2017), and few receive central bank communication directly (Afrouzi et al., 2015). When they do, the information is often covered by media outlets or shaped by pre-existing beliefs.

In the European context, central bank communication is particularly challenging, since some central banks must convey their messages across various media outlets and languages. For the European Central Bank (ECB), this means communicating across 21 countries in 24 languages, each with its own media landscape. As Draghi (2014) noted, this makes coherent and consistent messaging complex, but how the ECB communicates across this landscape is itself a policy choice. The ECB has begun to act on that choice, improving transparency and outreach, but most households still encounter information indirectly. And yet, the implications of monetary policy decisions are far from abstract: they affect the price of households' consumption baskets, real wages, and mortgage rates.

This paper studies how households' inflation beliefs change when ECB communication is covered by media separated by country and language. Although the ECB communicates in a uniform way, national media translate, frame, and distribute its messages unevenly. Using text analysis of newspaper coverage for the four largest euro-area economies Germany, Spain, France, and Italy, I document differences in how frequently outlets write about a set of macroeconomic topics and in the direction in which those topics are framed. This variation shapes household beliefs: the frequency of inflation coverage moves qualitative and quantitative inflation perceptions, while direction moves qualitative and quantitative inflation expectations, with effects that are pronounced in some countries and muted in others.[1] Moreover, households lean more heavily on media signals during the post-2021 inflation surge. Integrating these differences in media

[1] Inflation perceptions reflect how households believe prices have evolved over the past 12 months. In contrast, inflation expectations reflect how households believe prices will evolve over the next 12 months.

transmission into a New Keynesian model, I study how the central bank should optimally communicate across a currency union in which its message is mediated by national media to raise union welfare.

I collect print and online media coverage of the ECB from major daily business newspapers in Germany, Spain, France, and Italy, yielding over 79,000 articles from 2013 to 2025. Through their linguistic and editorial practices, national media shape how the ECB's message reaches households, so that topic and tone differ substantially across countries. I link this media data to national quantitative and qualitative inflation perceptions and expectations from surveys by the European Commission, together with other macroeconomic indicators. This allows me to assess how shifts in media coverage affect households' beliefs at the national level while accounting for macroeconomic environments.

I apply text analysis methods to the newspaper articles to extract two key measures: frequency and direction. Frequency measures how prominently macroeconomic topics (inflation, labor, housing, financial, and economy) appear in coverage of the ECB. Direction measures whether the coverage frames a topic as moving up, moving down, or holding steady. By separating frequency from direction, I can distinguish how prominent a topic is covered from how it is covered. I show that ECB communication is often associated with inflation in Germany but economic developments in Spain and Italy. France, by contrast, shows the largest directional changes when talking about labor market developments.

To understand the macroeconomic implications of cross-country heterogeneity in media coverage, I estimate a Bayesian Vector Autoregression (VAR) that includes the media-based measures alongside inflation perceptions, expectations, and other macroeconomic variables. Using inertial restrictions, I control for country-specific shocks and isolate the role of national media environments within the common monetary union. When inflation receives more frequent media coverage, median estimates of quantitative inflation perceptions rise by 0.2 to 1pp, depending on the country, while the qualitative balance statistic rises by 1 to 5 points. When coverage directs towards rising inflation, the effect extends to expectations: median estimates of quantitative inflation expectations rise by 0.2 to 0.4pp and qualitative measures by 0.7 to 2.1 points. Due to the inertial restrictions, these responses are orthogonal to inflation and the policy rate, so they point to a forward-looking media channel that operates more strongly in some countries. Differences in frequency and direction across countries thus explain part of the heterogeneity among the four major euro-area economies.

To trace how households update their inflation expectations, I develop a model with three information signals: prior beliefs, media reports, and private observations of current inflation. The prior signal is the previous period's expectation, the media signal is given by the measures of frequency and direction, and food and consumption prices proxy for private observed inflation. I estimate the model on observed inflation expectations by maximum likelihood, which recovers the precision of each signal and hence the weight households place on it. The estimated precisions reveal how strongly households weight each signal and how those weights evolve across countries and over time. I find that prior beliefs and media signals outweigh personal price observations. In the post-2021 inflation surge, in the four countries, households shift weight away from their priors and private observations toward media signals, although the magnitude of this shift differs substantially across countries.

Finally, I embed these estimated media differences in a New Keynesian model of a currency union, in which a single central bank sets one interest rate but can also communicate in a country-specific way. Because national media transmit the central bank's message unequally, the same policy reaches households differently, and the model asks how the central bank should respond optimally. I show that communication complements the policy rate: since the rate is common to the union while communication can be targeted, the central bank optimally directs stronger and more persistent communication to the countries whose media transmission is weakest, precisely where the interest rate alone cannot reach. The model also reproduces the post-2021 shift toward media-based beliefs. Pairing this country-specific communication with the union-wide interest rate yields a consumption-equivalent welfare gain of 0.068% for the currency union, which is largest where media transmission is weakest and communication compensates most for the limited reach of local coverage.

These findings underline the limitation of a uniform communication in a linguistically fragmented monetary union. Despite the ECB's substantial efforts, including dedicated communication departments in each country (Draghi, 2014), significant national differences in responsiveness persist, underscoring the challenges of conveying its message effectively across diverse media outlets and languages.

Related Literature. The paper relates to three strands of the literature. First, it contributes to the vast line of research on central bank communication and inflation expectations. Second, it adds to the literature of media as a channel for macroeconomic fluctuations. Third, the paper contributes to the rapidly expanding research on text analysis in the context of monetary policy.

A large literature has documented that modern monetary policy operates increasingly through the management of expectations, making communication a central policy instrument. Blinder et al. (2008) provide an early overview of how central bank communication can meaningfully shape market expectations but document that its effects on the general public remain underexplored. They emphasize that this gap matters for substantial reasons: the general public is the source of central banks' democratic legitimacy and, therefore, of their institutional independence. Since then, several strands of the literature have examined the more direct effects of central bank communication on the expectations of households and firms. Binder (2017) summarizes the evidence on the rationales for, and the effectiveness of, central bank communication with households. The author finds that such communication can clarify policy intentions and reduce uncertainty, although its effects are distributed unevenly.

Another line of research emphasizes the role of intermediaries. Lamla and Lein (2014) provide early evidence that the tone and intensity of inflation-related news influence households' inflation expectations. Bolliger (2024) demonstrates that newspaper inflation coverage affects households' inflation expectations and perceptions within a multi-language country. Gorodnichenko et al. (2024) show that the media are among the most active interactors with the U.S. central bank on social media. Blinder et al. (2024) review the central bank communication literature and note that television and newspapers remain the most effective channels for reaching the general public. They also stress that, although central banks increasingly recognize the value of engaging directly with the public, this effort is still at an early stage and comes with substantial challenges.

Against this background, recent contributions highlight that the effectiveness of communication varies across subgroups and contexts. Coibion et al. (2022) demonstrate that households' expectations differ when they are exposed to direct central bank communication compared with news articles covering the same topic. Coibion and Gorodnichenko (2025) extend this argument by showing that communication strategies must account for heterogeneous information structures if central banks aim to anchor household expectations effectively. Wabitsch (2025) shows that messenger identity matters for the effectiveness of central bank communication.

Together, this literature demonstrates that central bank communication influences inflation expectations, but its effectiveness hinges on how messages are filtered through intermediaries and interpreted across heterogeneous audiences. Yet despite these advances, we still know little about how intermediaries shape these heterogeneous interpretations, and how different media environments condition the transmission of central bank communication to households. This paper fills that gap by examining how inflation expectations respond to central bank communication once it is mediated through different media system within a monetary union.

Beyond their role as intermediaries for central bank communication, a parallel strand of research views media as independent sources of information frictions that can drive macroeconomic fluctuations. Theoretical work highlights how the media can generate macroeconomic fluctuations. In their model, Nimark and Pitschner (2019) show that agents' reliance on media reports can propagate shocks and produce comovement beyond what fundamentals imply. Chahrour et al. (2021) further argue that shared media sources can create information traps in which misperceptions become persistent and self-reinforcing. Empirical studies support these mechanisms. Chahrour et al. (2025) show that media's systematic selection of reporting creates asymmetric changes in expectations. In financial markets, Bybee et al. (2024) demonstrate that economic news coverage systematically influences asset prices.

Together, this literature establishes media as a key information channel affecting economic behavior. However, existing work largely treats media as homogeneous, leaving open how cross-country differences in coverage influence the transmission of economic information. This is a key issue addressed in this paper.

The paper also connects to the growing literature employing text analysis methods to analyze monetary policy communication. A foundational contribution in this area, Hansen and McMahon (2016) use Latent Dirichlet Allocation and dictionary methods to extract meaningful signals from Federal Open Market Committee (FOMC) statements that affects asset prices and macroeconomic outcomes. More recently, Handlan (2022) uses a neural network on FOMC and alternative FOMC statements to show effects of central bank communication on macroeconomic variables. Aruoba and Drechsel (2024) create sentiment measures from staff economist document prepared for FOMC meetings using the dictionary by Loughran and McDonald (2011). Larsen et al. (2025) also employ Latent Dirichlet Allocation in combination with the dictionary by Loughran and McDonald (2011) to extract inflation sentiments from media reporting about the ECB in different European countries.

The current paper is closely related to Larsen et al. (2025) but differs by focusing on broader cross-country differences in reporting on the ECB, not only inflation, employing a neural network to obtain sentiments and controlling for country-specific shocks. Furthermore, it expands the analysis to include a theoretical foundation to quantify welfare implications arising from heterogeneous transmission across countries.

The remainder of this paper is structured as follows. Section 2 describes the data. Section 3 explains the text analysis methodology in use to extract frequency and tone from the data and identifies country-specific differences in media coverage. Section 4 estimates the macroeconomic consequences of these differences in reporting. Section 5 traces how households update their inflation expectations and Section 6 presents the theoretical model. Finally, Section 7 concludes.

2 Data

To uncover the impact of central bank communication on inflation expectations of households, I collect print and internet media by newspapers. The analysis focuses on the four largest euro area economies: Germany, Spain, France and Italy. These regions are selected due to their linguistic distinctions, which facilitate a clearer separation of households in the different regions.

2.1 Newspaper Data

The study focuses on major daily economic newspapers from the four largest euro area economies: *Handelsblatt* (Germany), *Expansión* (Spain), *Les Echos* (France), and *Il Sole 24 Ore* (Italy). These outlets are selected because they share similar profiles in content, political orientation, readership, and market position, ensuring cross-country comparability.[2] All articles mentioning the ECB were retrieved from Factiva in April 2025. The sample begins in January 2013, corresponding to the earliest date for which all four newspapers provide consistent digital archives.

In Germany, the analysis includes 23,711 articles from Handelsblatt and Handelsblatt Online referring to the ECB. Handelsblatt is Germany's leading financial daily, focusing on business, finance, and economic policy. The outlet is generally regarded as center-right and pro-market in orientation. As of 2023, its circulation reached approximately 150,000 daily copies, most of which were paid subscriptions but the newspaper also holds a strong online presence.

The Spanish dataset comprises 16,414 articles from Expansión that mention the ECB. Expansión is Spain's foremost economic and financial newspaper, with a strong emphasis on business, market, and policy developments. The newspaper holds a center-right and pro-market editorial stance, similar to Handelsblatt. According to Newman et al. (2025), it maintains around 110,000 paid subscriptions.

For France, the study includes 10,042 articles from Les Echos and LesEchos.fr referencing the ECB. The newspaper provides broad coverage of business, finance, and market news targeted at professional and household audiences. It is widely regarded as center-right and pro-business. In 2024, Les Echos reported approximately 140,000 paid subscriptions.

In Italy, the sample consists of 29,415 articles from Il Sole 24 Ore, including both the Digital Replica and Online editions, that mention the ECB. The newspaper specializes in economic, financial, and market reporting and holds a center-right, pro-market editorial position. Il Sole 24 Ore also maintains a strong digital presence, reaching roughly 7% of weekly internet users (Newman et al., 2025). This broad online reach enhances its representativeness among Italian economic media.

Together, these four newspapers provide a consistent and comparable representation of economic media coverage across the largest euro area economies. Their similar editorial focus and political orientation help isolate linguistic and cultural differences in

[2]Research on intermedia agenda-setting shows that such elite media outlets often shape the agendas of other news organizations, including regional and general-interest newspapers (Zhang, 2018). Other work on the interaction between traditional and online media suggests that traditional media continue to play a dominant role in setting the news agenda that is later amplified on digital platforms and social media (Harder et al., 2017; Van Den et al., 2019).

how monetary policy communication is transmitted to the public. The resulting dataset captures the coverage of ECB-related reporting, forming a solid foundation for analyzing how media outlets shape expectations across countries.

2.2 Inflation Perceptions and Expectations

To connect multilingual media coverage with qualitative and quantitative inflation perceptions and expectations, I use survey data for Germany, Spain, France, and Italy. Monthly series are used to ensure comparability across countries. The sample spans from January 2013 to March 2025, a period covering both conventional and unconventional monetary policy phases.[3] This time frame captures diverse macroeconomic conditions, providing a robust setting for estimating the relationship between media coverage and belief formation.

Qualitative and quantitative inflation perceptions and expectation measures come from the European Commission's consumer surveys, which provide harmonized indicators across countries.[4] For qualitative inflation perception, survey participants were asked: *How do you think that consumer prices have developed over the last 12 months? They have....* The possible answer choices are given by: *risen a lot (++)* (PP), *risen moderately (+)* (P), *risen slightly (=)* (E), *stayed about the same (-)* (M), *fallen (–)* (MM) and *don't know* (N). These responses are summarized through a balance statistic $B = PP + P/2 - M/2 - MM$, which translates qualitative judgments into a continuous measure of perceived inflation. If survey participants answer with a rise or a fall in prices, they are asked a follow-up question: *By how many per cent do you think that consumer prices have gone up/down over the past 12 months? (Please give a single figure estimate).* The participants then provide a number, that presents the quantitative inflation perceptions.

Qualitative and quantitative inflation expectations are collected in a similar manner, asking survey participants the following: *By comparison with the past 12 months, how do you expect that consumer prices will develop in the next 12 months? They will....*

[3]For Spain, the sample is cut in September 2024 as inflation perceptions and expectations are unavailable from then on.

[4]Although the ECB provides more detailed series of inflation perceptions and expectations beginning in 2020, I rely on the data by the European Commission to capture the longer-run dynamics and to ensure comparability across an extended period in which substantial variations occur.

Possible answer choices are similar: *increase more rapidly (++)* (PP), *increase at the same rate (+)* (P), *increase at a slower rate (=)* (E), *stay about the same (-)* (M), *fall (–)* (MM) and *don't know* (N). The same balance formula is applied to this forward-looking question, yielding a consistent indicator of expected inflation across all countries. Also this question has a follow-up, if the participant answered with an increase or fall, namely, *By how many per cent do you expect consumer prices to go up/down in the next 12 months? (Please give a single figure estimate).* This number provides the quantitative inflation expectations.

Figure 1 shows the time series of qualitative and quantitative inflation perceptions and expectations for the four countries. In all countries, inflation perceptions declined from 2013 to 2016 and rose again after 2021. Qualitative and quantitative measures mainly follow the same patterns. A notable difference can be observed after 2016, when qualitative inflation perceptions and expectations rose in Germany, Spain, and France but the same pattern can only be observed in the quantitative measure in Spain. The subsequent analysis will investigate whether some of these movements can be attributed to changes in media coverage.

2.3 Macroeconomic Data

The study uses macroeconomic data that allow for isolating how country-specific reporting channels shape household expectations while accounting for national economic environments. To control for macroeconomic conditions that might influence inflation perceptions and expectations, I add six additional time series for each country. The first is the Harmonized Index of Consumer Prices (HICP) inflation rate from Eurostat, capturing actual price developments faced by households. The second variable is real Gross Domestic Product (GDP), taken from the ECB and interpolated with industrial production following the method of Bernanke et al. (1997), to provide monthly estimates of economic activity. The third control variable is the short-term policy rate, representing the main refinancing operations rate of the ECB. Because this rate was constrained by the zero lower bound during much of the sample, ECB's total assets and liabilities are included as a the fourth variable as a proxy for unconventional monetary policy. Fifth, stock prices are added as a forward-looking variable. Finally, the seasonally adjusted unemployment rate from the ECB is used to capture labor market conditions that can

Figure 1: Inflation Perceptions and Expectations per Country

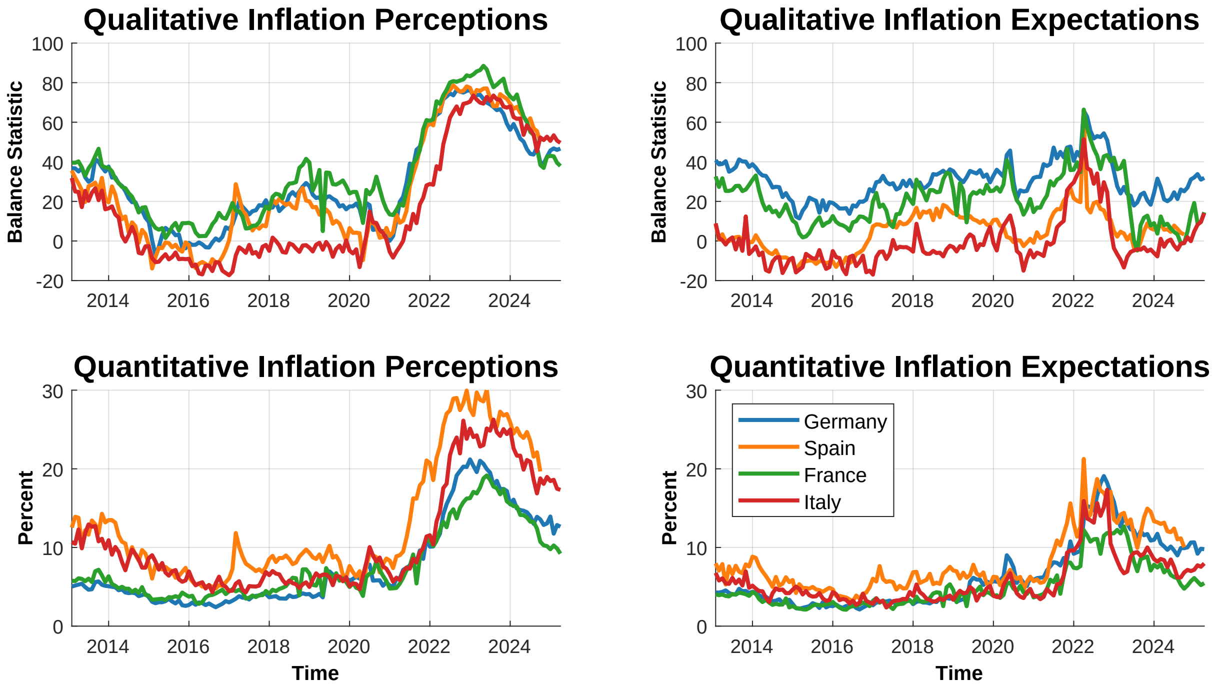

Note: The figure shows how qualitative (top) and quantitative (bottom) inflation perceptions (left) and expectations (right) evolved over time in the countries Germany (blue line), Spain (yellow line), France (green line) and Italy (red line). These survey responses are displayed using the balance statistic and percent provided by the European Commission described in the text. The horizontal axis shows the time horizon.

shape households' inflation perceptions and expectations.

Figure 2 displays inflation rates for the four countries along with the ECB's short-term policy rate. The countries exhibit broadly similar patterns: inflation declined until 2015, increased through 2017, and remained elevated until 2019, after which it began to fall and reached a low point in 2021. Following this trough, inflation rose to varying levels and evolved differently across countries. In Germany, Spain, and Italy, inflation increased rapidly. Spain's inflation began to decline in mid-2022, shortly after the ECB raised the short-term policy rate, with Germany and Italy following a similar pattern soon after. France experienced a more moderate increase, reaching levels below those of the other countries, and subsequently followed Germany and Italy in the later decline. The short-term policy rate continued to rise until late 2023 and began to fall again in mid-2024. Compared with Figure 1, the rise in inflation between 2017 and 2021 did not translate into higher inflation perceptions in Italy, nor into higher inflation expectations in Germany, France, and Italy. Moreover, inflation perceptions remain elevated in all countries even though actual inflation has fallen back to levels close to those observed in

2017.

Figure 2: Country-Specific Inflation Rates and ECB's Short-Term Policy Rate

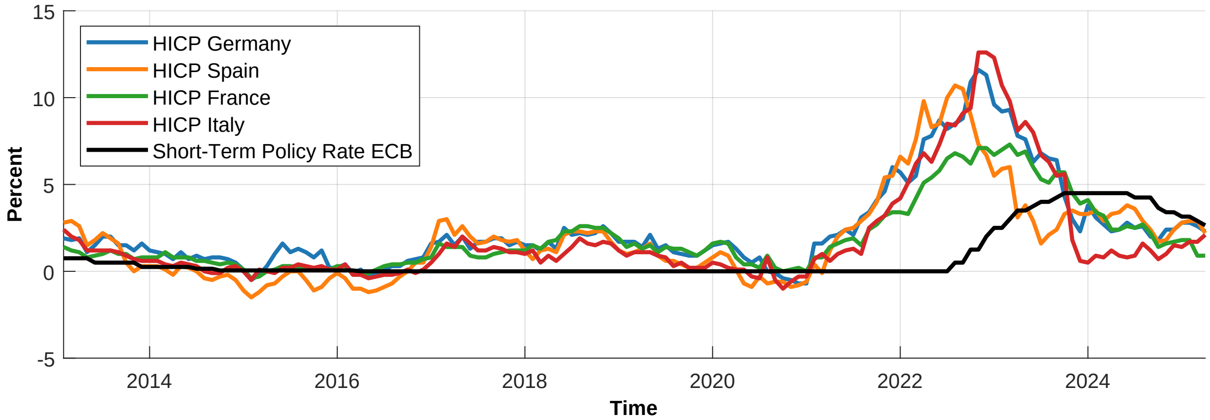

Note: The figure shows how the Harmonized Index of Consumer Prices (HICP) inflation rate per month in the countries Germany (blue line), Spain (yellow line), France (green line) and Italy (red line). In addition, the short-term policy rate by the ECB is displayed in black. The horizontal axis shows the time horizon.

3 Quantifying Central Bank Communication Across Media and Languages

To better understand how households are exposed and respond to central bank communication, I first observe how national media cover the message by the central bank. In a second step, using natural language processing (NLP) techniques applied to newspaper data, I construct two measures to quantify key features of media coverage: frequency and direction. From these two measures, I obtain substantial differences across the countries: media vary both in the frequency with which they report on core macroeconomic topics and in the direction they use to frame them.

3.1 ECB Communication and National Media Filtering

ECB communication is mostly issued in English and subsequently translated to official EU languages by the central bank itself. However, the dissemination of interpretation of these messages occurs largely through national media outlets, but also the structure of national media market influences how the communication is transmitted.

Understanding how the ECB's centralized communication is distributed across diverse media outlets is crucial for grasping its impact on public perception, especially since the effectiveness of this distribution hinges on how these messages are subsequently picked up and relayed by national media outlets. The ECB primarily communicates through centralized channels, such as press conferences, official statements, and speeches, which are predominantly issued in English and then translated into the official EU languages by the ECB itself to ensure consistency and accuracy (Draghi, 2014). The translated messages are distributed through various platforms, including the ECB's website and social media. This centralized approach aims to provide a uniform message across all member states, minimizing the risk of misinterpretation at the source and maintaining a coherent monetary policy message.

The role of national media in translating and distributing ECB messages introduces a layer of complexity to the communication process. National media outlets act as intermediaries, deciding whether to cover ECB communications, which aspects to highlight, and how prominently to feature them. This selective process means that households often receive mediated versions of the ECB's messages, shaped by the editorial priorities and biases of national media systems. The extent and manner of coverage can vary significantly between countries, depending on the national context. Consequently, the same ECB message can be framed differently across countries, leading to varied public perceptions and responses.

As an example, in the press release of June 5, 2014 the

> *ECB introduces a negative deposit facility interest rate. [...] When deciding to lower the key ECB interest rates at its meeting today, the Governing Council of the ECB took the decision to cut the interest rate on the deposit facility to -0.10%* (European Central Bank, 2014).

The Handelsblatt in Germany released the following news coverage:

> *The minus revolution. Banks will be penalised for hoarding money in future, the base rate is heading towards zero – ECB President Draghi is fighting deflation with every means at his disposal.*

The Expansión in Spain wrote:

> *Banks and construction companies, sectors favored by Draghi.[...] immediate positive effect on virtually all Ibex stocks, but especially on banking stocks.*[5]

[5] Text translated by the Data Science Lab at the University of Bern.

Les Echos in France states:

> *The ECB pulls out all the stops to revive the eurozone. Unique measures to avoid deflation. European leaders and financial markets welcome Mario Draghi's actions.*

Although the core message of the ECB was transmitted, the way the information is embedded differs across the three journals.[6]

The structure of national media markets determines how widely and consistently ECB communication is distributed. According to the Centre for Media Pluralism and Media Freedom (2022), the risk of news media concentration is medium in Germany and high in Spain, France, and Italy. In more concentrated media systems, a few dominant outlets may set the framing direction and amplify central bank messaging.[7] In contrast, fragmented or polarized media landscapes can lead to uneven or conflicting interpretations of the same message. Differences in public broadcasting roles, journalistic standards, and economic reporting cultures also affect how ECB content is selected and presented. These institutional features can significantly mediate the public's exposure to and understanding of monetary policy. In Appendix A, other national factors are discussed that might influence how households are exposed and respond to central bank communication.

3.2 Frequency

To measure the extent of potential exposure of households to central bank communication, I extract the frequency of coverage of macroeconomic topics from newspaper data using text processing. The key idea is that greater exposure to ECB news framed around specific topics, such as inflation, increases the likelihood that households associate central bank actions with these topics. Frequency thus captures the informational environment through which economic expectations may form differently across countries.

The dictionary forms the foundation of the exposure measure by identifying the linguistic markers of five major economic topics: inflation, housing, labor markets, financial

[6]These excerpts are illustrative and do not capture the full context of the original articles. They should therefore be interpreted with caution.

[7]Beckers et al. (2019) show that higher concentration of media ownership leads less diverse reporting over time.

markets, and the overall economy. The idea of organizing dictionary keywords into distinct macroeconomic topics follows Gardner et al. (2022), although the specific set of five topics is defined for this study. The dictionary itself is constructed by first extracting the most frequent n-grams from the article corpus and then selecting, from among these, the keywords most closely associated with each topic.[8] For instance, the inflation topic includes keywords such as cost, deflation, devaluation, inflation, prices, and purchasing power. The full dictionary is reported in Appendix B, which lists the complete set of keywords for each topic. Because media outlets publish in different languages, each newspaper article is analyzed in its original language (German, Spanish, French, and Italian) to preserve its linguistic and contextual nuances.

Using a predefined dictionary is suitable in this context as the texts are general-interest articles that follow no fixed narrative structure. Unlike the ECB's own speeches or press releases, articles can make unsupervised topic models (see, e.g., Hansen and McMahon (2016); Hansen et al. (2018)) less reliable for recovering consistent themes across countries. A dictionary approach provides a stable first filter that captures core economic concepts even when journalists use broader or more varied language. The cost of fixing categories in advance is that a dictionary cannot capture relevant terms or emerging themes it was not designed to include, although building the dictionary from the corpus's most frequent n-grams limits this risk. Relying on the same set of predefined topics facilitates cross-country comparability, which is central to studying how different media outlets report on a common underlying signal. This approach therefore allows the analysis to focus directly on meaningful variation in how national media emphasize economic topics rather than on variation generated by the topic-identification method itself.

The computation of frequency standardizes keyword occurrences to ensure that topic frequencies are comparable across articles and languages. Each article is searched for dictionary keywords, and each mention of a keyword increases the count of its associated topic. These counts are then aggregated at the article level and divided by the total number of words excluding stopwords in the article, thereby controlling for length differences.[9] This normalization ensures that longer articles do not automatically receive

[8]An n-gram is a sequence of n words appearing together in the text, ranging from single words to short multi-word phrases, so ranking n-grams by frequency reveals the terms media most commonly use when writing about each topic.

[9]Stopwords are common function words, e.g., *the*, *and*, *of*, and *to*, that carry little topical meaning.

higher topic counts simply because they contain more words. As a result, a short article that mentions inflation once contributes proportionally more to the inflation topic frequency than a long article mentioning it once.

Figure 3 shows the frequencies of the topics inflation and housing covered in the media outlets across the four countries over the sample period. In all countries, reporting about inflation increased during the post-2021 inflation surge, but magnitude differs across the countries. Germany features more attention to inflation and housing, especially after 2020, suggesting a stronger focus on these two topics in this country. Spanish, French, Italian media outlets devote comparatively less attention to inflation and housing over the whole timeseries.

Figure 3: Monthly Mean Frequencies of Topics Inflation and Housing in Media Coverage per Country

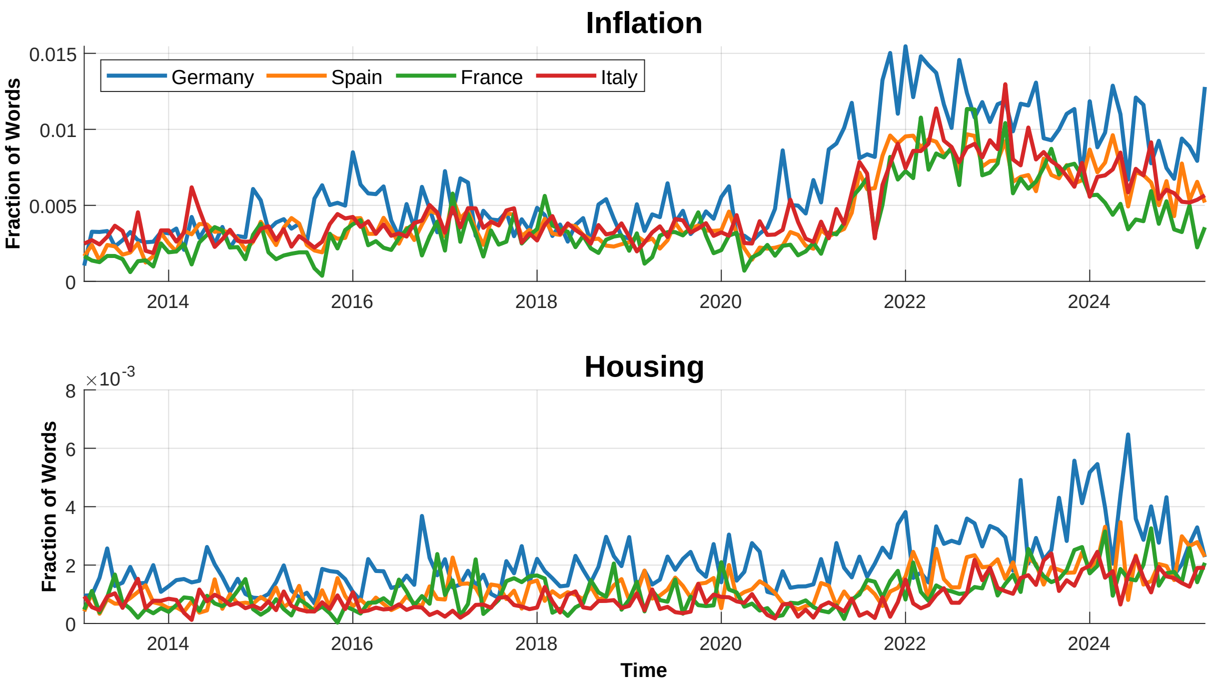

Note: The figure shows how frequently the topic inflation (top) and housing (bottom) is mentioned on average per month in the countries Germany (blue line), Spain (yellow line), France (green line) and Italy (red line). Both topics are normalized by the number of words per article that they are mentioned in. The horizontal axis shows the time horizon.

Table 1 displays how strongly the frequencies of economic topics in media coverage differ in the four countries. The topic financial receives the most coverage overall, ranking as the most frequently covered topic in Germany, Spain, and France, only in Italy coverage

They are excluded from the word count so that the normalization reflects the article's content length rather than its use of generic language.

of the topic economy is slightly higher. German media outlets devote comparatively more attention to inflation, housing and financial developments, pointing to an environment in which price stability and financial conditions receive heightened emphasis. Spain features the highest coverage of the topic labor.

Table 1: Mean Frequencies of Economic Topics in Media Coverage per Country

Topic	**Germany**	**Spain**	**France**	**Italy**
Inflation	0.0064 (0.0036)	0.0044 (0.0023)	0.0038 (0.0023)	0.0047 (0.0023)
Labor	0.0008 (0.0004)	0.0018 (0.0007)	0.0014 (0.0007)	0.0011 (0.0005)
Housing	0.0021 (0.0011)	0.0012 (0.0006)	0.0011 (0.0006)	0.0009 (0.0005)
Financial	0.0195 (0.0031)	0.0124 (0.0023)	0.0131 (0.0030)	0.0091 (0.0020)
Economy	0.0074 (0.0020)	0.0096 (0.0016)	0.0071 (0.0016)	0.0104 (0.0018)

Note: Mean values with standard deviations (in parentheses) of how frequent the topics inflation, labor, housing, financial and economy are covered in the media per country as a fraction of total words per article.

These cross-country patterns indicate that central bank communication reaches households through media environments that differ markedly in the coverage of core macroeconomic themes. Because reporting on inflation, labor, housing, financial developments and economic conditions varies both across countries and over time, the informational context surrounding central bank actions is not uniform. As a result, households interpret monetary policy signals within national media landscapes that emphasize some topics more than others, shaping how they process central bank communication. These differences in media emphasis ultimately contribute to heterogeneity in how households form and adjust their expectations.

3.3 Direction

Understanding how media frames central bank communication requires a systematic measure of direction.[10] To capture this framing, I develop a measure that assesses whether

[10] I use *direction* rather than *sentiment* or *tone* deliberately. Conventional sentiment measures capture valence—whether coverage is favorable or unfavorable—which is not monotone in inflation: both high inflation and deflation register as unfavorable, so a given sentiment score cannot be mapped to an expected inflation level. Direction instead captures whether coverage points toward higher or lower future inflation, a signed and monotone measure that maps directly into inflation expectations.

news coverage conveys a downward, neutral, or upward direction for each topic. The underlying assumption is that the direction associated with specific topics shapes how households form and adjust their expectations. For instance, *inflation is climbing further above target* and *prices are sliding toward deflation* would both read as unfavorable to a sentiment measure, yet they point in opposite directions for inflation, which is the distinction the direction measure is designed to capture. To construct this measure, I train a multilingual NLP model to interpret media coverage in a four-step procedure: selecting a pre-trained language model, creating a multilingual fine-tuning dataset, training the model, and applying it to obtain the direction.

A multilingual NLP model is particularly useful here because it allows the articles to be analyzed in their original language, preserving the wording that households actually encounter. While dictionary-based approaches are valuable, most widely used direction dictionaries, such as Loughran and McDonald (2011), are developed from English financial texts and therefore do not naturally reflect the broader vocabulary found in general-interest media outlets. Moreover, translating such dictionaries may miss linguistic nuances or country-specific expressions that shape how similar economic developments are described across languages. In contrast, a multilingual NLP model can learn direction directly from the phrasing and context used in German, Spanish, French, and Italian media. This approach ensures that direction is captured as it appears in everyday language, rather than through a financial-specific or translation-adjusted lens.

3.3.1 Select Pre-Trained Neural Network for Text Processing

A multilingual language model is essential for capturing the content of central bank communication as it appears across languages and contexts. For this reason, I use Multilingual RoBERTa, a natural language understanding model developed by Conneau et al. (2020).[11] The model was trained on 100 languages and is capable of detecting the input language without the need for separate language embeddings. Its training objective, masked language modeling, requires predicting randomly hidden words within a sentence, allowing the model to learn both word meanings and sentence structures simultaneously. This approach enables Multilingual RoBERTa to represent linguistic meaning

[11] The base model consists of 12 layers, a hidden size of 768, 12 attention heads, resulting in approximately 270 million parameters.

in a bidirectional way, making it particularly suitable for multilingual text analysis.

Because Multilingual RoBERTa already encodes general linguistic knowledge across 100 languages, it provides an ideal foundation for direction classification in multilingual data. The model can be fine-tuned to map text onto discrete categories such as downward, neutral, and upward direction. In this study, I use the base version's pre-trained word representations as a starting point, building upon its broad understanding of language to specialize it for economic contexts.[12] Such pre-trained architectures are widely used because training them from scratch would require extensive computational resources and massive datasets. This transfer learning approach allows the model to adapt efficiently to the task of identifying direction in media articles reporting about the ECB.

3.3.2 Create a Dataset for Fine-Tuning

Developing an accurate direction measure requires a dataset that reflects how media report on central bank communication in multiple languages. The model must learn to classify phrases containing specific keywords as downward, neutral, or upward, but no such financial or multilingual dataset currently exists. Although Multilingual RoBERTa understands many languages, it must be fine-tuned on text that reflects the stylistic and contextual features of media coverage about monetary policy. To address this gap, I construct a novel dataset designed specifically for this purpose.

The dataset generation process aims to approximate real newspaper language while avoiding direct overlap with the model's prediction data. I use articles from the Swiss Broadcasting Corporation, which publishes economic news in German, French, Italian, and Spanish through its multilingual platform *SWI swissinfo.ch*. Using these texts as templates, I generate additional examples with two large language models: GPT-4.5 by OpenAI and mistral-large-2411 by MistralAI. One example of setting up such a prompt is written in Appendix C. For each topic and language, I obtain 50 examples per direction category (downward, neutral and upward), resulting in 3,000 training phrases. Then I translate these phrases into all languages so that the model is trained on the same data for all languages. The generated labels were verified by two human annotators who reviewed and resolved any disagreements, ensuring a reliable ground truth for fine-tuning

[12]The model was pretrained on approximately 2.5 trillion tokens across 100 languages, derived from Common Crawl data.

the model.

3.3.3 Fine-Tune the Neural Network

Fine-tuning adapts the general multilingual model to the specific task of classifying direction in central bank–related news. Each phrase in the training dataset is first converted into tokens consistent with Multilingual RoBERTa's vocabulary, which represents words in an N-dimensional space where semantically similar terms lie close together. This tokenization process enables the model by Conneau et al. (2020) to interpret linguistic nuances such as contextual differences between words like *slight increase* and *surge*. The fine-tuning stage then trains the model to associate these tokenized representations with categories.

Moreover, we want our model to learn the mapping from the provided phrases to downward, neutral and upward categories. The objective is for the model to predict category based on an unknown phrase containing one of the keywords. The model provided by Conneau et al. (2020) can not yet execute this task. Thus, in this project, I train it further with the created database without changing its original structure. Appendix D explains the training algorithm in further detail.

To train the model, I apply five-fold cross-validation, splitting the dataset into five parts so that each subset serves once as a validation set while the remaining four are used for training. Notably, I always have a training set that the model adapts its internal parameters on and a test set that it runs the model on but does not adapt its internal parameters to. Such train and test splits are important in machine learning because the models are prone to overfitting, i.e., to learn too much from the training data, thereby being unable to predict data it has not yet seen. This procedure also implies that for the same hyperparameter described in Appendix D, I have five different parameter values, depending on the data the model was trained on.

To find the optimal hyperparameters for the model, I experiment with different combinations of the number of epochs[13] and the learning rate.[14] During these experiments, I monitor its learning ability on the training data by observing the loss to be able to correctly predict the labels of the test data (accuracy). The average loss on the training

[13]The number of epochs defines how often the model sees the same data set to adapt its parameters.

[14]The learning rate defines by how much the neural network adapts its internal parameters after each iteration.

data including the accuracy on the test data over the epochs with a fixed learning rate of 5e-07 is displayed in Table 2. Based on these results, the number of epochs for the training was set to 5, as the training loss keeps decreasing but test accuracy plateaus. The model has learned most of what it can learn without overfitting

Table 2: Training Loss and Test Accuracy Across Epochs

Epoch	Average Training Loss	Test Accuracy
0	0.5400	0.8177
1	0.3340	0.8469
2	0.2147	0.9003
3	0.1154	0.8983
4	0.0707	0.9134
5	0.0654	0.9154
6	0.0269	0.9144
7	0.0373	0.9265
8	0.0179	0.9134
9	0.0079	0.9204
10	0.0153	0.9144

Note: Average training loss reflects how well the model fits the training data during each epoch, while test accuracy measures how well it generalizes to unseen data. Monitoring both indicators helps determine when further training no longer improves performance: training loss may continue to fall even when test accuracy plateaus or declines, signaling overfitting.

3.3.4 Obtain Direction

After fine-tuning, the model is ready to predict the direction of media content related to central bank communication. Each article is scanned for the presence of keywords, and for every instance, the surrounding phrase is extracted and passed through the fine-tuned model, which classifies it as downward, neutral, or upward. The resulting direction categories are then converted into numeric values (-1 for downward, 0 for neutral and 1 for upward) and aggregated at the article level, with normalization again by total word count to account for length differences across texts. This procedure ensures that the direction measure is comparable both across languages and across time.

The resulting direction index provides a consistent qualitative indicator of how central bank communication is represented in different country-specific media environments. By aggregating this measure across time and languages, I capture the evolution of media framing that households are exposed to in their respective national contexts. This allows for systematic cross-country comparisons of qualitative exposure to central bank narratives, shedding light on how the direction of communication shapes expectations across countries.

Figure 4 shows the direction index for inflation and housing for the four countries over time. Overall, the direction used by media to talk about inflation became more intense during the post-2021 inflation surge. Media in Germany seem to use more intense upwards direction when talking about inflation and housing, showing higher scores. In Italy, media outlets used more downwards pointing language when talking about the housing market in the later part of the sample.

Figure 4: Monthly Mean Direction of Topics Inflation and Housing in Media Coverage per Country

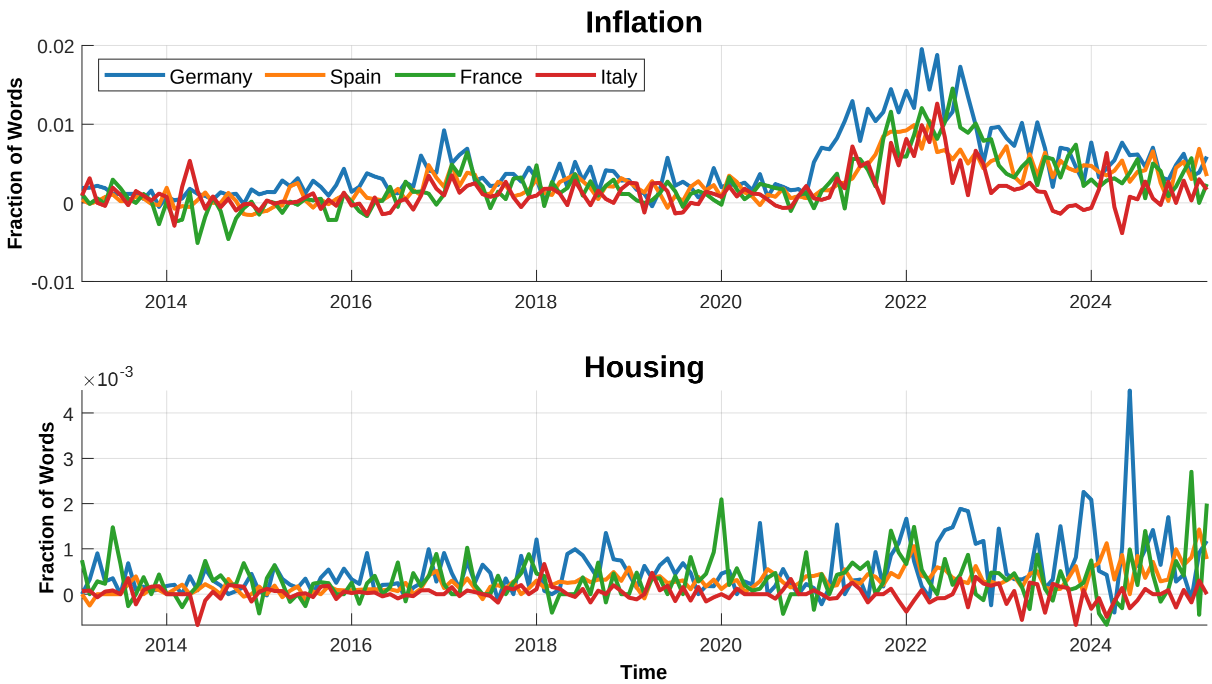

Note: The figure shows the direction used by media to report on the topic inflation (top) and housing (bottom) on average per month in the countries Germany (blue line), Spain (yellow line), France (green line) and Italy (red line). Both topics are normalized by the number of words per article that they are mentioned in. The horizontal axis shows the time horizon.

Table 3 highlights substantial cross-country differences in the direction with which key economic topics are discussed in the media. Directions are, in almost all cases, mildly upward rather than downward over the sample period. Germany leans slightly more upward on inflation and housing, while France records the most upward direction for labor and the general economy, along with the most variable labor direction of the four countries. Italy is consistently the least upward: it records the lowest direction for inflation, housing, and financial developments. Overall, the cross-country gaps in direction are narrower than the differences in coverage frequency.

Table 3: Mean Direction of Economic Topics in Media Coverage per Country

Topic	**Germany**	**Spain**	**France**	**Italy**
Inflation	0.0045 (0.0040)	0.0025 (0.0025)	0.0023 (0.0033)	0.0014 (0.0024)
Labor	0.0003 (0.0005)	0.0003 (0.0007)	0.0014 (0.0015)	0.0007 (0.0005)
Housing	0.0005 (0.0006)	0.0003 (0.0003)	0.0003 (0.0005)	0.0000 (0.0002)
Financial	0.0032 (0.0030)	0.0036 (0.0021)	0.0019 (0.0033)	0.0009 (0.0032)
Economy	0.0003 (0.0015)	0.0021 (0.0018)	0.0029 (0.0018)	-0.0001 (0.0011)

Note: Mean values with standard deviations (in parentheses) of direction used by the media to report on the topics inflation, labor, housing, financial and economy per country as a fraction of total words per article.

Taken together, frequency and direction results indicate that households encounter central bank communication within media landscapes that differ along two dimensions: how much media attention is devoted to key macroeconomic topics and if these topics are framed downward, neutral, or upward. National media not only emphasize different economic themes, such as Germany's stronger focus on inflation and Spain and Italy's heightened attention to the economy, but also attach distinct directional patterns to their reporting. These combined differences mean that identical central bank actions are embedded in informational environments that vary in both frequency and direction. Consequently, households interpret and update their inflation beliefs within national media contexts that shape not only what economic issues they notice but in what direction they point.

4 Macroeconomic Implications

Understanding how media frequency and direction of central bank communication influence households' belief is crucial for linking the media coverage to macroeconomic outcomes. To examine these effects, I estimate a Bayesian VAR model that includes both media measures alongside inflation perceptions, inflation expectations, and other key macroeconomic variables. In this section, I first explain the details of the Bayesian VAR model. Then, before presenting the results, I explain the inertial restrictions that I use for identification of the media measures, which are crucial to isolate media coverage from macroeconomic conditions.

4.1 Bayesian VAR Methodology

The Bayesian VAR framework provides a structured way to estimate how media direction and frequency affect inflation perceptions, inflation expectations, and broader macroeconomic dynamics. Following Giannone et al. (2015) and Lenza and Primiceri (2022), this approach incorporates prior information into the estimation and accommodates uncertainty in both parameters and hyperparameters.[15] As this framework forms the empirical foundation of the analysis, I summarize its key components below. Consider the following Bayesian VAR:

$$\begin{aligned} Y_t &= B_0 + B_1 Y_{t-1} + \cdots + B_p Y_{t-p} + u_t \qquad (1) \\ u_t &\sim N(0, \Sigma), \end{aligned}$$

where Y_t is an $n \times 1$ vector of macroeconomic variables as described in Section 2, u_t is an $n \times 1$ vector of the reduced-form innovations, and $B_0, B_1, \ldots, B_p$ and Σ are matrices of suitable dimensions containing the model's unknown parameters. N denotes a normal density. Consistent with the literature, I set the lag order, p, to 6. I assume that the reduced-form innovations, u_t, are a linear combination of n orthogonal structural disturbances, ε_t, which I write as $u_t = A_0 \varepsilon_t$. Here, A_0 is the impact matrix of the structural disturbances at $t = 0$.

I estimate the Bayesian VAR as described in Lenza and Primiceri (2022). As suggested by the authors, I define $\beta \equiv vec([B_0, B_1, \ldots, B_p]')$ and $X_t \equiv I_n \otimes x_t' \equiv$

[15] The key difference between Giannone et al. (2015) and Lenza and Primiceri (2022) is that the latter accounts for the increased volatility following the COVID-19 pandemic.

$I_n \otimes [1, Y'_{t-1}, \ldots, Y'_{t-p}]$, where I_n denotes a $n \times n$ identity matrix. Substituting these definitions, Equation 1 becomes:

$$Y_t = X_t \beta + u_t \tag{2}$$

This representation enables efficient sampling from the posterior distribution of β and Σ, given suitable priors.

The hierarchical Bayesian structure allows the data to inform the strength and tightness of prior beliefs. Following Giannone et al. (2015), I treat the hyperparameters as additional parameters to be estimated, with the prior distribution of the coefficients belonging to the Normal–Inverse-Wishart family:

$$\Sigma \sim IW(\Psi; d) \tag{3}$$

$$\beta|\Sigma \sim N(b, \Sigma \otimes \Omega), \tag{4}$$

where IW denotes the Inverse-Wishart density and the elements Ψ, d, b and Ω are functions of a lower dimensional vector of hyperparameters. Through this hierarchical approach, the prior hyperparameters are assigned hyperpriors, and the marginal likelihood is used to explore the full posterior hyperparameter space.

Setting the hyperparameters ensures both statistical coherence and empirical plausibility. I set the degrees of freedom of the Inverse-Wishart distribution to $d = n + 2$, the minimum value that guarantees the existence of the prior mean of Σ, equal to $\Psi/(d-n-1)$. The matrix Ψ is diagonal, with hyperparameters ψ on the main diagonal. For β, I use a Gaussian prior conditional on Σ, a version of the Minnesota prior introduced by Litterman (1979). This prior assumes that each variable follows a random walk process, possibly with drift, and defines the first and second moments as follows:

$$E\left[(B_s)_{ij}|\Sigma\right] = \begin{cases} \mu & \text{if } i = j \text{ and } s = 1 \\ 0 & \text{otherwise,} \end{cases}$$

$$Cov\left[(B_s)_{ij}(B_r)_{hm}|\Sigma\right] = \begin{cases} \lambda^2 \frac{1}{s^2} \frac{\Sigma_{ih}}{\psi_j/(d-n-1)} & \text{if } m = j \text{ and } r = s \\ 0 & \text{otherwise,} \end{cases}$$

The hyperparameters μ and λ govern persistence and tightness, respectively. I set μ to 0.5 for real GDP growth and inflation, reflecting their relatively low serial correlation during the sample period, and μ to 0.9 for other variables such as interest and unemployment rates, inflation perceptions and expectations, consistent with their high persistence. The hyperpriors for λ and ψ follow the data-driven approach proposed by Giannone et al. (2015). This specification ensures a balance between flexibility and empirical realism in capturing the effects of media communication.

I implement the estimation using the routines from Lenza and Primiceri (2022). My implementation diverges from the source code in one respect: I impose stationarity on the VAR throughout the Markov Chain Monte Carlo sampling. Specifically, I accept a draw for the VAR parameters only if it satisfies the stationarity condition; otherwise, I redraw. This ensures that the posterior inference reflects economically meaningful dynamics.

4.2 Identification

Identifying structural disturbances within the Bayesian VAR is essential for linking media-driven shocks to macroeconomic outcomes. I adopt a recursive identification strategy, implemented through a Cholesky decomposition of the covariance matrix of reduced-form innovations from Equation 1.[16] The decomposition yields mutually orthogonal shocks that can be interpreted as structural disturbances. The contemporaneous ordering of the variables encoding which shocks are allowed to affect which variables on impact. I order the media measures of frequency and direction after inflation, real GDP, and the monetary policy rate, so that the macroeconomic disturbances are identified before the media measures are introduced. This ordering embeds the inertial restrictions: within a period, the slower-moving macroeconomic aggregates, inflation, output, and the policy rate, are treated as inertial and do not respond on impact to media coverage, whereas media coverage may respond contemporaneously to them. At monthly frequency, such inertia is plausible, since coverage adjusts to macroeconomic news faster than these aggregates respond to coverage. The media measures are therefore orthogonal to inflation, real GDP, and the policy rate on impact, so the residual variation attributed to frequency and direction reflects information dissemination rather than contemporaneous macroeco-

[16]A recursive scheme imposes an explicit ordering assumption rather than the sign restrictions discussed by Baumeister and Hamilton (2015), who show that under sign restrictions the posterior distribution of the impulse responses can reflect the priors implicit in the imposed signs. A recursive ordering does not rely on such sign priors, though it substitutes its own identifying assumption about the contemporaneous timing of shocks.

nomic or country-specific structural effects. The resulting framework allows me to trace the macroeconomic consequences of media coverage filtered through national contexts.

4.3 Media Coverage, Inflation Perceptions and Expectations

The Bayesian VAR with inertial identification allows me to quantify how media-based measures of frequency and direction shape households' inflation perceptions and expectations across the four countries. These innovations are normalized to reflect a standard deviation increase on impact in each measure, which ensures a meaningful comparison across countries and across media. The setup isolates the effect of media coverage by constructing shocks that are orthogonal to contemporaneous changes in inflation, real GDP and monetary policy. As a result, the impulse responses reflect changes attributable solely to media coverage of central-bank communication and abstract from country-specific shocks.

The frequency shock captures how quantitative inflation beliefs react when inflation receives more media coverage. As shown in Figure 5, quantitative inflation perceptions rise by 0.2 to 1pp, with the strongest response in Spain, followed by Italy, Germany, and France. Quantitative inflation expectations respond most strongly in Spain (0.6pp), followed by Germany and Italy (0.3pp each) and France (0.1pp), where the response is significant only at the 16–84 percentile bands. Because the media measures are orthogonal to actual inflation, real GDP and the monetary policy rate on impact, these responses indicate that households react to more prominent inflation coverage over and above contemporaneous macroeconomic conditions. More frequent inflation reporting therefore raises both how households perceive current inflation and what they expect for future inflation. The full set of macroeconomic reactions are displayed in Appendix E.

The direction shock captures how households react when media coverage of inflation turns upward. As shown in Figure 6, quantitative inflation perceptions increase by 0.3 to 0.6pp and expectations by 0.2 to 0.4pp in Germany, France, and Italy, with the strongest effects in Germany, followed by Italy and France. In Spain, quantitative inflation expectations rise by 0.3pp, significant only at the 16–84 percentile bands. The shock can be read as follows: conditional on inflation already being covered more frequently, additionally reading that inflation is rising moves expectations by a further 0.2 to 0.4pp. Because inflation, real GDP, the policy rate, and reporting frequency have not moved on impact, this additional response reflects the direction of coverage rather than its volume or the

Figure 5: Impulse Response Functions from the Innovation to Media Frequency

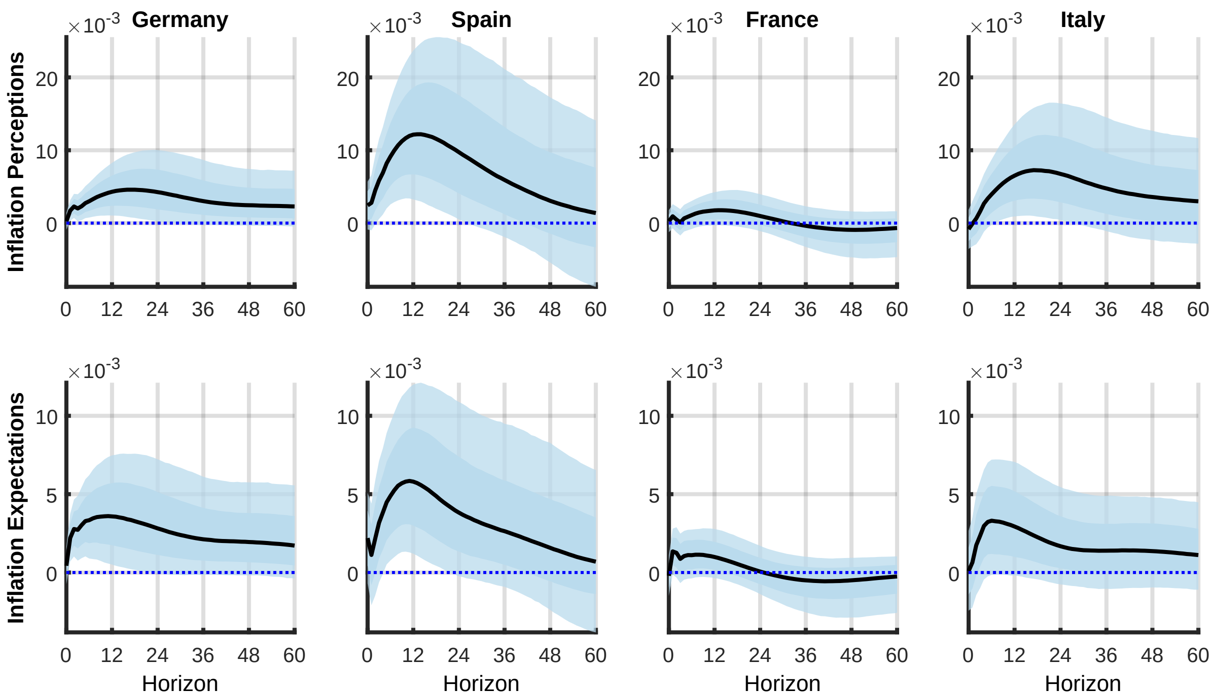

Note: The (dark) blue shaded areas represent the 5–95 (16-84) percentiles of the impulse responses identified with the innovation to media frequency about inflation. The impulse responses are normalized a standard deviation increase in the average number of mentions of the topic inflation per article per month. The figures display the responses of inflation perceptions and inflation expectations in the four countries: Germany, Spain, France, and Italy. The horizontal axis shows the time horizon in months.

underlying macroeconomic conditions.

In addition, displayed in Appendix F, I repeat the analysis using qualitative measures of inflation beliefs. The qualitative results confirm the quantitative pattern, moving beliefs in the same direction and preserving the same country ranking: more frequent inflation reporting raises the balance statistic of qualitative inflation perceptions by 1 to 5 points, while an upward shift in the direction of reporting about inflation raises qualitative inflation expectations by 0.7 to 2.1 points. The frequency shock, however, does not significantly move qualitative inflation expectations.

Taken together, the results show that both frequency and direction of media reporting about inflation shape households' inflation beliefs, and that they do so differently. More frequent reporting raises inflation perceptions across both the quantitative and qualitative measures and also moves quantitative expectations, whereas an upward shift in the direction of reporting raises both inflation perceptions and expectations. This difference between how much and how inflation is reported, together with its uneven pass-through

Figure 6: Impulse Response Functions from the Innovation to Media Direction

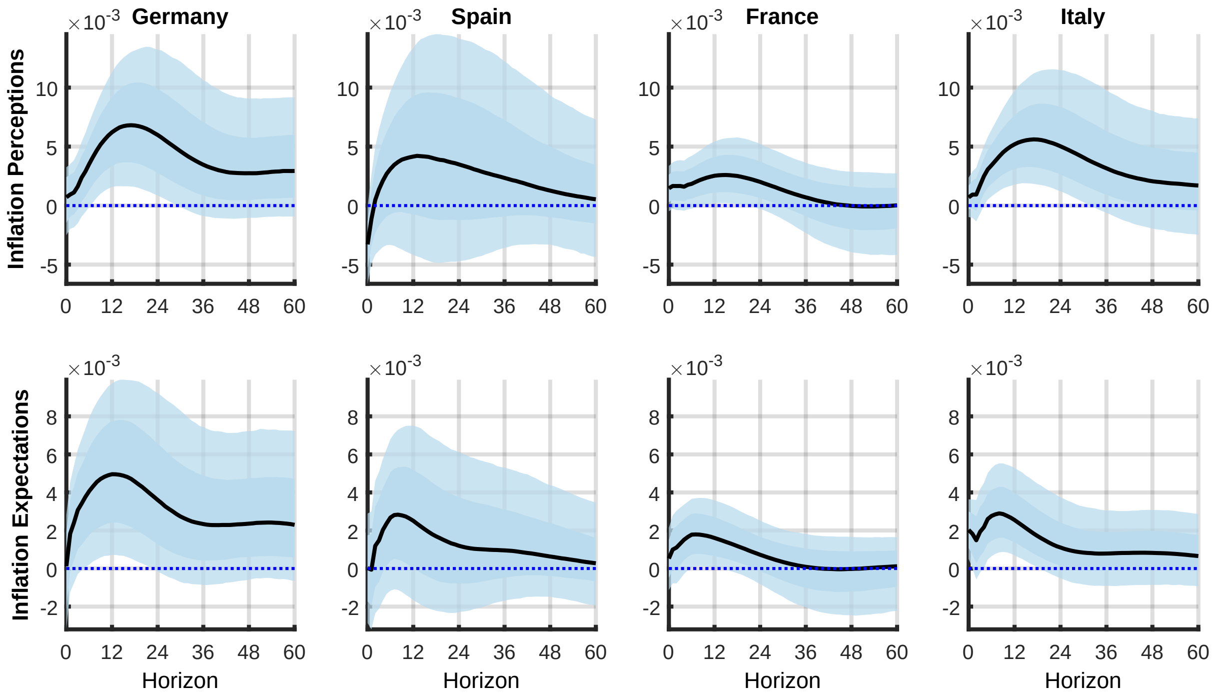

Note: The (dark) blue shaded areas represent the 5–95 (16-84) percentiles of the impulse responses identified with the innovation to media direction about inflation. The impulse responses are normalized a standard deviation increase in the average direction surrounding the topic inflation per article per month. The figures display the responses of inflation perceptions and inflation expectations in the four countries: Germany, Spain, France, and Italy. The horizontal axis shows the time horizon in months.

across the four economies, indicates that national media environments condition how a common euro-area signal reaches the public. The evidence therefore points to media structure as a distinct channel in the transmission of information to household inflation beliefs, operating alongside actual macroeconomic conditions. Thus, understanding these media environments is central to evaluating how monetary policy communication ultimately reaches households in a currency union.

5 Information Signals and Inflation Expectations

To uncover why households' inflation expectations react differently in the countries and over time, I develop a model that accounts for some of the changes. Similar to Lamla and Lein (2014), I assume consumer form expectations about future inflation π_{t+1} based on a prior belief and media reports. In contrast to the authors, she has her own observation of current inflation. The representative consumer in country i holds normally distributed prior beliefs,

$$\Pi_t^i \sim \mathcal{N}\left(\pi_t^{\text{prior},i}, \frac{1}{\tau_\pi^i}\right),$$

about future inflation π_{t+1}. Each period, the consumer also receives V_t^i independent normally distributed media reports about the state of the economy

$$\psi_{v,t}^i \sim \mathcal{N}\left(\pi_{t+1}, \frac{1}{\tau_\psi^i}\right), \qquad v_t^i = 1, \dots, V_t^i,$$

and M_t^i independent normally distributed private observations of inflation

$$\phi_{m,t}^i \sim \mathcal{N}\left(\pi_{t+1}, \frac{1}{\tau_\phi^i}\right), \qquad m_t^i = 1, \dots, M_t^i.$$

These private observations can be seen as personal shopping experience or personal price observations.

Let $\bar{\psi}_t^i = \frac{1}{V_t^i}\sum_{v_t^i=1}^{V_t^i} \psi_{v,t}^i$ and $\bar{\phi}_t^i = \frac{1}{M_t^i}\sum_{m_t^i=1}^{M_t^i} \phi_{m,t}^i$ denote the average signals from the media and own observations, respectively. Given conditional independence of all signals, the posterior distribution of π_{t+1} is normal with mean

$$\mathrm{E}_t^i(\pi_{t+1} \mid \bar{\psi}_t^i, \bar{\phi}_t^i) = \omega_{\pi,t}^i \pi_t^{\text{prior,i}} + \omega_{\psi,t}^i \bar{\psi}_t^i + \omega_{\phi,t}^i \bar{\phi}_t^i, \tag{5}$$

and precision

$$\tau_{\text{post},t}^i = \tau_\pi^i + V_t^i\,\tau_\psi^i + M_t^i\,\tau_\phi^i. \tag{6}$$

The weights on the prior, media signals, and private observations are given by

$$\omega_{\pi,t}^i = \frac{\tau_\pi^i}{\Lambda_t^i}, \qquad \omega_{\psi,t}^i = \frac{V_t^i\,\tau_\psi^i}{\Lambda_t^i}, \qquad \omega_{\phi,t}^i = \frac{M_t^i\,\tau_\phi^i}{\Lambda_t^i},$$

where

$$\Lambda_t^i \equiv \tau_\pi^i + V_t^i\,\tau_\psi^i + M_t^i\,\tau_\phi^i, \qquad \omega_{\pi,t}^i + \omega_{\psi,t}^i + \omega_{\phi,t}^i = 1.$$

Thus, the consumer's inflation expectation is a weighted average of her prior belief, the average media report, and her private observation of inflation. An increase in the number or precision of media reports (higher V_t^i or higher τ_ψ^i) raises the precision of the average media signal $\bar{\psi}_t^i$ and therefore increases the weight $\omega_{\psi,t}^i$, shifting weight away from the prior and from private information. Likewise, an increase in the number or precision of private observations (higher M_t^i or higher τ_ϕ^i) raises the precision of $\bar{\phi}_t^i$ and increases $\omega_{\phi,t}^i$.

To estimate the model parameter in Equation 6, I link each component of household information to a suitable empirical series and estimate the signal precision with maximum likelihood. The prior signal of inflation expectations is the expectation from the previous period. For the estimation, I use the quantitative inflation expectations measures from the European Commission's consumer surveys. The average media reports that the households observe are captured through the direction measure developed earlier and paired with the frequency measure that quantifies the volume of coverage. These two indicators jointly reflect the information set that reaches households through news outlets and therefore serve as the basis for the media signal in the model. The household's personal observation of inflation are represented by the Eurostat monthly food price monitor and the HICP, which closely tracks the price changes individuals encounter in everyday consumption. Together, these series allow the representative consumer in the model to update her beliefs about future inflation using one media signal and one private signal per period.

The estimated parameters in Table 4 indicate how strongly households respond to each of the three information components. The results show that the estimated variance of the prior is low across all countries, implying that households place relatively high weight on their prior beliefs when updating expectations. The estimated signal precision of the media is highest in France, followed by Italy, Germany and Spain, suggesting cross-country differences in the informativeness of media coverage. Private observations are less precise than media in all countries indicating that the salience of private observation is lower than media in the model. These comparisons reveal how the model allocates belief updating between media and personal observation, with notable heterogeneity across countries.

Table 4: Estimated Parameter Values by Country

Parameter	DEU	ESP	FRA	ITA
$\frac{1}{\tau^i_\pi}$	0.00034 (0.0000)	0.00062 (0.0000)	0.00033 (0.0000)	0.00040 (0.0000)
$\frac{1}{\tau^i_\psi}$	0.00037 (0.0000)	0.00055 (0.0000)	0.00025 (0.0000)	0.00033 (0.0000)
$\frac{1}{\tau^i_\phi}$	0.00280 (0.0000)	0.00240 (0.0000)	0.00160 (0.0000)	0.00230 (0.0000)

Note: The table shows the model estimates for the parameter values of $1/\tau^i_\pi$, $1/\tau^i_\psi$ and $1/\tau^i_\phi$ for the four countries including p-values in parentheses.

The model-implied inflation expectations displayed in Figure 7 reproduce the broad dynamics of quantitative inflation expectations but exhibit slight deviations. The estimated series track the rise in expectations from 2017 to 2021 in all countries, mirroring patterns in the survey data. The model slightly underestimates some peaks in 2022, e.g., mean inflation expectations in Germany were 15.3%, whereas the model can only replicate 13.9%, in Spain observations were 15.6% versus the model fits 13.0%, in France 12.2% versus 10.8% and in Italy 15.9% versus 14.0%. However, overall, the estimated expectations capture the main direction of movements but underestimate some observed in the actual data.

Figure 7: Estimated Inflation Expectations

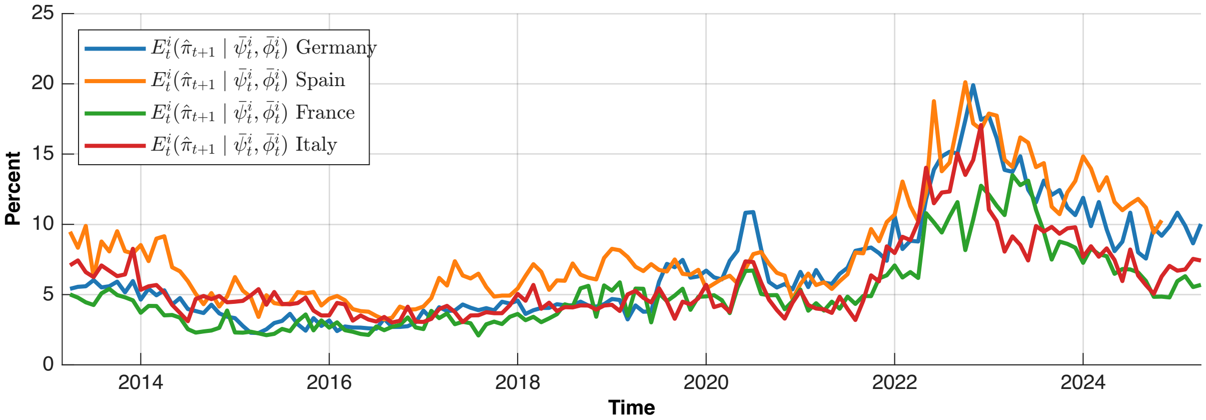

Note: The figure shows the model estimates of inflation expectations, $E_t^i(\hat{\pi}_{t+1} \mid \bar{\psi}_t^i, \bar{\phi}_t^i)$, per month in the countries Germany (blue line), Spain (yellow line), France (green line) and Italy (red line).

The evolution of the model-implied weights in Figure 8 shows how households shift attention across prior beliefs, media reports, and private observations over time. The estimated weights on prior beliefs remain consistently high for all countries, reinforcing the earlier evidence that prior beliefs contribute considerably to updating even once new information becomes available. The onset of the inflation surge in 2021 is associated with a pronounced rise in the weight on media reports, which simultaneously reduces the weight placed on prior beliefs and private observations. This is consistent with paper that focus on attention in times of higher inflation, see e.g. Pfäuti (2023) and Weber et al. (2025). In all countries, media observations contribute slightly less than prior beliefs before 2021, after which media dominate the updating process. Appendix G presents the contributions of model-implied weights.

Taken together, the results show that households update their inflation expectations

Figure 8: Estimated Parameter Weights for Consumer's Inflation Expectations

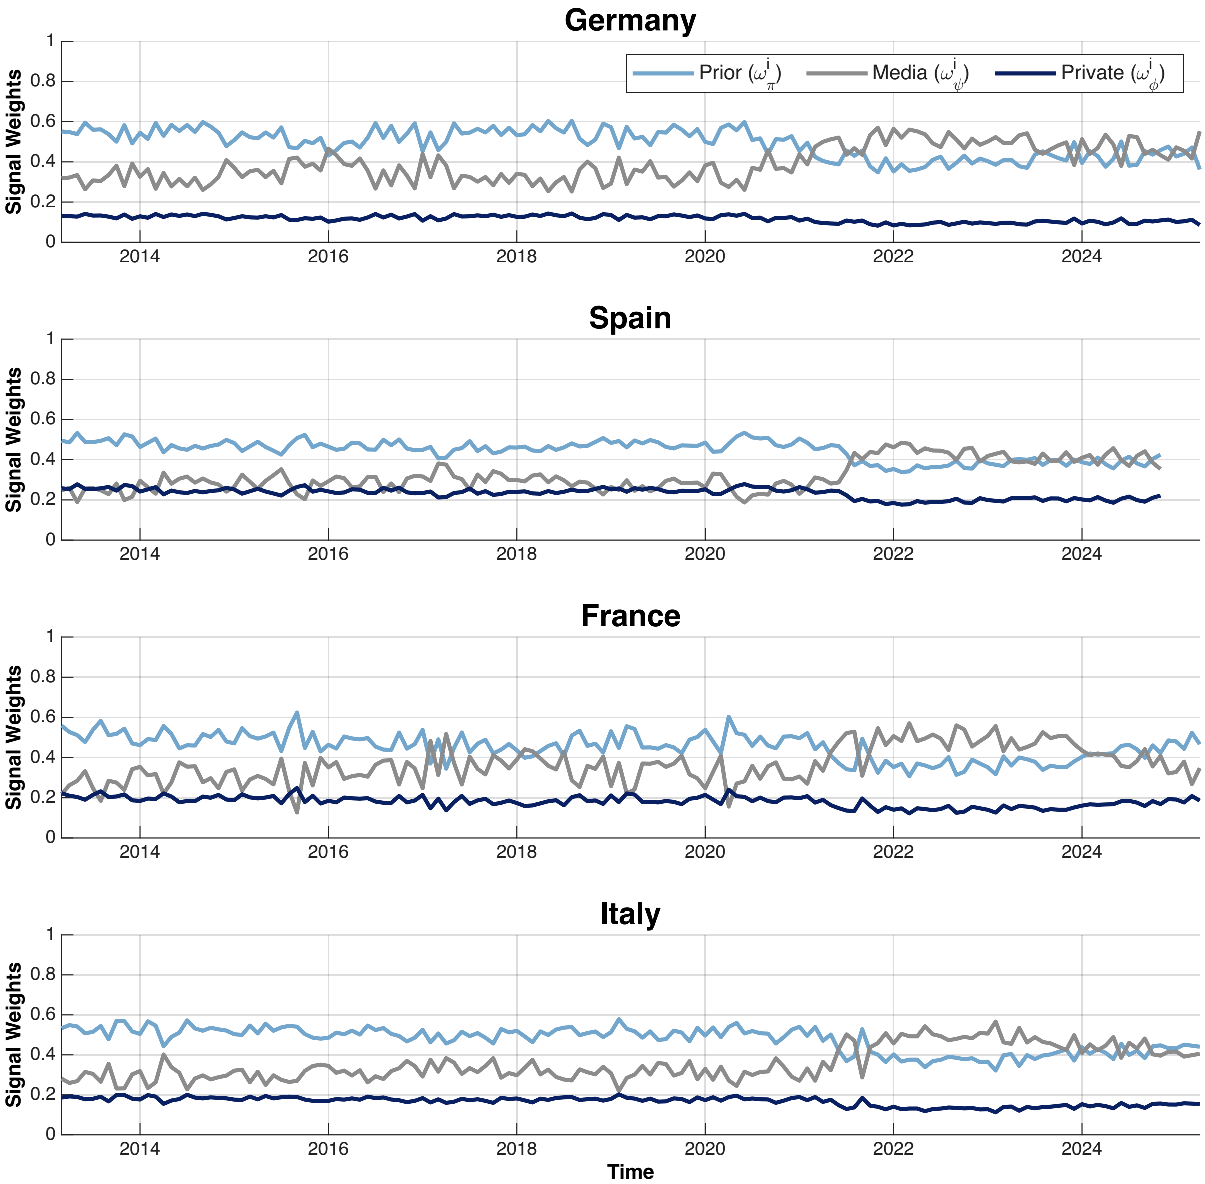

Note: The figure shows the model estimates for the parameter weights for the consumer's inflation expectations of the prior belief ($\omega^i_{\pi,t}$; light blue line), the media report ($\omega^i_{\psi,t}$; grey line) and her private observation of inflation ($\omega^i_{\phi,t}$; dark blue line) over time.

primarily through prior beliefs and media observations rather than through private shopping experience. Media become particularly influential during periods of rising inflation, especially after 2021, when the volume and direction of coverage intensify. Although the model tracks broad movements in survey expectations, it differs slightly in magnitude, highlighting the complexity of expectation formation. Overall, the evidence suggests that country-specific media environments meaningfully shape how households interpret inflation developments.

6 A New Keynesian Model with Heterogeneous News Coverage

Given the documented differences in the media signals across countries, I ask how the central bank can respond to them optimally. To do so, I develop a New Keynesian model along the lines of Galí and Monacelli (2008) and Bertasiute et al. (2020). I first present the currency union model, then define expectation formation as in Section 5, and finally formalize the media sector and monetary policy before turning to calibration and results. Throughout, I present the aggregated equations in the main text and relegate the microfoundations to Appendix H.

6.1 Currency Union Model

I consider a currency union of I countries and assume that the currency union does not interact with the rest of the world. The economic dynamics in the currency union can be described by the following equations:

$$\tilde{y}_t^i = \mathrm{E}_t^i \tilde{y}_{t+1}^i - \frac{1}{\sigma}\left(r_t - \mathrm{E}_t^i \pi_{t+1}^i - \rho_\beta\right) + \gamma\, \mathrm{E}_t^i(\Delta s_{t+1}^i) + \nu_t^i \tag{7}$$

$$\pi_t^i = \beta\, \mathrm{E}_t^i \pi_{t+1}^i + \kappa \tilde{y}_t^i + \xi_t^i \tag{8}$$

The superscript $i = 1, \ldots, I$ indexes countries. In country i, $\tilde{y}_t^i$ is the output gap, π_t^i is inflation, and r_t is the union-wide nominal interest rate, common to all members. The operator E_t^i denotes the (behavioral) expectation of agents in country i, formed conditional on information at the start of period t. The parameters σ (the inverse intertemporal elasticity of substitution), β (the discount factor, with $\rho_\beta := -\log\beta$), and κ (the Phillips-curve slope) are positive. For simplicity, I hold these structural parameters equal across countries, though in principle they may differ. The disturbances ν_t^i and ξ_t^i are random demand and supply shocks, respectively. Finally, γ is the parameter of economic integration, so that $\gamma\, \mathrm{E}_t^i(\Delta s_{t+1}^i)$ captures the expected change in country i's effective terms of trade relative to the rest of the union.

Equation 7 is the dynamic IS curve. As in the closed-economy benchmark, the output gap depends on the expected future output gap $\mathrm{E}_t^i \tilde{y}_{t+1}^i$, the real interest rate $r_t - \mathrm{E}_t^i \pi_{t+1}^i$, and a demand disturbance ν_t^i. It departs from the closed-economy case through the open-economy term $\gamma\, \mathrm{E}_t^i(\Delta s_{t+1}^i)$: when inflation is expected to be higher abroad than at home, country i's competitiveness improves, so that, for $\gamma > 0$, its output gap is higher than under equal expected inflation.[17]

Equation 8 is the New Keynesian Phillips curve, which takes the same form as in a closed economy: inflation depends on expected future inflation $\mathrm{E}_t^i \pi_{t+1}^i$, on the current output gap $\tilde{y}_t^i$ through the slope κ, and on a supply shock ξ_t^i.

The currency union consists of I countries that share a single nominal interest rate but experience country-specific inflation and output-gap dynamics, linked through their terms of trade. Each country is described by a behavioral IS curve and Phillips curve that coincide with the closed-economy equations except for an open-economy term in expected relative inflation. Because the policy rate is common while expectations and shocks are national, identical ECB actions can produce heterogeneous outcomes across countries. This is the environment in which the media-driven differences in expectation formation, introduced next, become the cross-country dispersion that optimal policy must weigh.

6.2 Expectation Formation

Full rationality requires that all agents understand the structure of the economy and can perform the calculations needed to form model-consistent expectations. This is an assumption that a large literature in economics finds empirically implausible; see, e.g. Carroll (2003); Pfajfar and Santoro (2010); Malmendier and Nagel (2016). Households instead rely on simpler heuristics to forecast economic variables. I therefore retain the expectation-formation rule of Section 5, now embedded in general equilibrium following Drenik and Perez (2020) and derived in Appendix H.2.

As in Section 5, the representative household in country i holds a prior about future inflation, $\pi_t^{\mathrm{prior},i} = \rho_\pi \, \mathrm{E}_{t-1}^i \pi_t^i$, its previous forecast projected one step forward under the AR(1)[18] law of motion; receives V_t^i independent media signals with noise variance $1/\tau_\psi^i$, and makes one private observation $\bar{\phi}_t^i$ of future inflation with noise variance $1/\tau_\phi^i$. The private signal is centered on π_{t+1}^i but observed with error, so the household does not observe future inflation directly. Applying Bayes' rule as in Equation 5, the posterior

[17] For further discussion on the value of γ and implications for the stability of the New Keynesian model the reader may refer to Bertasiute et al. (2020).

[18] Section 5 is the special case $\rho_\pi = 1$, in which each prior reduces to the previous period's expectation. The maximum-likelihood estimates place these persistences below one for every country, so the general-equilibrium formulation carries the projection explicitly.

mean is a precision-weighted average of the prior, the average media signal $\bar{\psi}_t^i$, and the private signal, with weights $\omega_{\pi,t}^i, \omega_{\psi,t}^i, \omega_{\phi,t}^i$ defined as in Section 5 (specialized to a single private signal):

$$\mathrm{E}_t^i \pi_{t+1}^i = \omega_{\pi,t}^i \pi_t^{\mathrm{prior},i} + \omega_{\psi,t}^i \bar{\psi}_t^i + \omega_{\phi,t}^i \bar{\phi}_t^i. \tag{9}$$

Output-gap expectations are formed analogously. The household combines a prior $\tilde{y}_t^{\mathrm{prior},i} = \rho_y \, \mathrm{E}_{t-1}^i \tilde{y}_t^i$, V_t^i media signals about economic activity with noise variance τ_φ^i, and one private signal ς_t^i, which it infers from its participation in the labor market, with noise variance τ_ς^i. The weights $\omega_{\tilde{y},t}^i, \omega_{\varphi,t}^i, \omega_{\varsigma,t}^i$ are the corresponding precision weights, defined exactly as in the inflation case and summing to one:

$$\mathrm{E}_t^i \tilde{y}_{t+1}^i = \omega_{\tilde{y},t}^i \tilde{y}_t^{\mathrm{prior},i} + \omega_{\varphi,t}^i \bar{\varphi}_t^i + \omega_{\varsigma,t}^i \bar{\varsigma}_t^i. \tag{10}$$

Households in each country form expectations not under full rationality but by combining three noisy signals about future inflation and activity: a prior, the reports they receive from national media, and their own observation of realized prices. Bayes' rule makes the posterior a precision-weighted average of these inputs, so the weight a household places on media rises with the number and precision of the reports it receives. Because that volume and precision differ across countries, identical central-bank information enters national expectations with different force. This signal-extraction structure is what translates the media differences documented earlier into the cross-country dispersion in beliefs that monetary policy must later confront.

6.3 Media

Newspapers receive information about the general economic condition of country i and the stance of the central bank, and I assume one news outlet per country informs its population. Reporting intensity responds to how newsworthy conditions are:

$$\chi_t^i = \rho_\chi \chi_{t-1}^i + \varrho_\pi (\pi_{t-1}^i - \bar{\pi})^2 + \varrho_{\tilde{y}} (\tilde{y}_{t-1}^i - \bar{y})^2 + \varrho_z (z_t^i)^2, \tag{11}$$

where ϱ_π, $\varrho_{\tilde{y}}$, and ϱ_z are positive weights on inflation, the output gap, and central-bank communication. Although these could differ across countries, I hold them equal for outlets of a similar type. The persistence term ρ_χ captures follow-up reporting, as outlets produce further articles explaining an event. The squared deviations $(\pi_{t-1}^i - \bar{\pi})^2$ and

$(\tilde{y}^i_{t-1} - \bar{y})^2$ make both high inflation and deflation newsworthy, and communication enters in squared form because any communication is newsworthy regardless of its direction; its sign affects only the content households receive. Newsworthiness thus rises with the squared deviation of inflation or the output gap and with the strength of the central bank's signal, and reporting frequency intensifies accordingly:

$$V^i_t = \bar{V}^i \exp\left(\eta^i_V \chi^i_t\right),$$

where $\bar{V}^i$ is steady-state media frequency in country i and η^i_V is the semi-elasticity of V^i_t with respect to χ^i_t.

This link between policy signals and coverage connects to a literature on public signaling by the monetary authority. The idea that the central bank sends public signals to an economy in which agents hold dispersed information goes back to Morris and Shin (2002). More recent work stresses the signaling effects of policy when the central bank is better informed than households and firms, yielding a less sluggish adjustment of inflation expectations (Melosi, 2016) or a more accurate reading of policy actions (Zhang, 2022). In the other direction, Gorodnichenko et al. (2024) show that media are among the most active listeners of the central bank, which is the mechanism the frequency equation above formalizes: coverage intensifies precisely when the central bank's signal is strong.

Media content is given by their expectations of future variables but it also takes into account the communication by the central bank. The media signal is valuable to households precisely because it carries information they do not already hold. Households observe their own prices and their prior, but they do not costlessly observe economic conditions of other countries or the central bank's communication. Acquiring this information directly would be too costly for the household. The media outlets, by contrast, monitor the economic conditions and the central bank closely (Gorodnichenko et al., 2024) and mediate information to the households. Media content therefore adds direction to household information, the orientation of future inflation implied by the outlet's reading of conditions and, crucially, by the central bank's signal, rather than solely echoing the inflation households already observe. For inflation, this mean content over a period is:

$$\bar{\psi}^i_t = \mathrm{E}^{N,i}_t(\pi^i_{t+1} | \mathcal{I}^{N,i}_t) + b_z z^i_t + \rho_\psi \varepsilon^{\psi,i}_{t-1} + \varepsilon^{\psi,i}_t,$$

whereas for the output-gap the mean content is:

$$\bar{\varphi}_t^i = \mathrm{E}_t^{N,i}(\tilde{y}_{t+1}^i | \mathcal{I}_t^{N,i}) + b_z z_t^i + \varepsilon_t^{\varphi,i}.$$

The term $\mathrm{E}_t^{N,i}$ is the outlet's own reading of where inflation is heading, formed on its information set $\mathcal{I}_t^{N,i}$. $b_z z_t^i$ is the central-bank signal it relays to the households (with $b_z \in (0,1)$ capturing incomplete pass-through), and $\varepsilon_t^{\psi,i}$ and $\varepsilon_t^{\varphi,i}$ are reporting noise. ρ_ψ is the persistence of the reporting noise. The disturbance to mean inflation reporting, $\varepsilon_t^{\psi,i}$, is the theoretical counterpart of the direction disturbance discussed in Section 4. The set-up underlines that information is mediated by media, thereby creating content.

National media translate economic conditions and central-bank signals into the coverage households see, along two dimensions that mirror the empirical measures: the frequency of reporting, which intensifies with the newsworthiness of inflation and output-gap deviations, and the direction of content, which orients beliefs toward higher or lower future inflation. Because households cannot cheaply observe the economic conditions of other countries or the central bank directly, media coverage is the channel through which its signal reaches the public, and both the intensity of that coverage and its direction differ across countries. This is the step that turns country-specific media behavior into the differential signals households receive, and thus into the belief dispersion that optimal policy must weigh.

6.4 Monetary Policy

I consider three different monetary policy set-ups: a Taylor rule for calibration and two welfare-based objectives. First, I set a Taylor Rule to calibrate some model parameters to the impulse response function obtained in Section 4. Second, I will compare two different welfare functions to calculate the welfare loss for each of them. These welfare functions are presented here and derived in Appendix H.3.

6.4.1 Taylor Rule

The central bank sets the union-wide nominal interest rate according to

$$r_t = \bar{\pi} + \rho_\beta + \rho_r r_{t-1} + \Phi_\pi(\pi_t^{CU}) + \Phi_y(\tilde{y}_t^{CU}), \tag{12}$$

for all i and t. ρ_r is the interest rate smoothing parameter. The positive parameters Φ_π and Φ_y determine the responsiveness of the central bank to currency-union deviations of inflation and the output-gap.

6.4.2 Central Bank Loss Functions

Whereas the Taylor rule serves only to calibrate the model to the empirical impulse responses, the loss functions define the central bank's optimal policy. Under the baseline objective the bank chooses the policy rate r_t to minimize $\mathcal{L}^B$, under the alternative objective it chooses the rate and country-specific communication $\{r_t, z_t^i\}$ jointly to minimize $\mathcal{L}^A$. The model equations are presented here and derived in Appendix H.3.

The central bank minimizes the welfare-theoretic baseline loss function, which includes cross-country dispersion terms

$$\begin{aligned}\mathcal{L}^B &= \lambda_\pi(\pi_t^{CU})^2 + \lambda_{\tilde{y}}(\tilde{y}_t^{CU})^2 \\ &\quad + \lambda_{\pi^i}\sum_{i=1}^{I}(\pi_t^i - \pi_t^{CU})^2 + \lambda_{\tilde{y}^i}\sum_{i=1}^{I}(\tilde{y}_t^i - \tilde{y}_t^{CU})^2, \end{aligned} \tag{13}$$

where the weights $\lambda_\pi = \epsilon/\lambda$ and $\lambda_{\tilde{y}} = \sigma+\zeta$ are standard to the welfare function. However, the central bank also internalizes within-union heterogeneity, which gives rise to welfare costs from cross-country dispersions λ_{π^i} and $\lambda_{\tilde{y}^i}$.

As mentioned before, the central bank obtains an additional tool, communication z_t^i, to combat union-level and within-union inflation and output-gap. Crucially, the policy rate, r_t, is common to the union while communication is country-specific, so communication can be directed at the cross-country dispersion that a single rate cannot reach. However, the additional tool is also costly so the central bank is penalized for the level of communication but also the change of communication over time:

$$\begin{aligned}\mathcal{L}^A &= \lambda_\pi(\pi_t^{CU})^2 + \lambda_{\tilde{y}}(\tilde{y}_t^{CU})^2 \\ &\quad + \lambda_{\pi^i}\sum_{i=1}^{I}(\pi_t^i - \pi_t^{CU})^2 + \lambda_{\tilde{y}^i}\sum_{i=1}^{I}(\tilde{y}_t^i - \tilde{y}_t^{CU})^2 \\ &\quad + \lambda_z\sum_{i=1}^{I}(z_t^i)^2 + \lambda_{\Delta z}\sum_{i=1}^{I}(z_t^i - z_{t-1}^i)^2, \end{aligned} \tag{14}$$

where λ_z and $\lambda_{\Delta z}$ represent the costs that the central bank pays for using the additional tool. These communication costs can be motivated as a reduced-form representation of reputational or operational costs of country-specific communication analogous to the instrument smoothing term in standard Taylor rules (Woodford, 2003).

Because communication is costly in both its level and its adjustment, the central bank uses it only up to the point where its marginal stabilization benefit equals its marginal cost, concentrating it where media transmission is weakest. Comparing the welfare loss under the two regimes lets the paper quantify the gains from using communication as a complement to the policy rate.

6.5 Calibration

As mentioned above, I use the Taylor rule to calibrate the model to the empirical results obtained in Sections 4 and 5, and use the values obtained to evaluate the two monetary policy set-ups. I will first discuss the parameters that are identical for all countries and then discuss the parameters that are different for the countries. For the country-specific calibration, I use the results obtained in Section 3 and 5. In addition, I calibrate some parameters to match the IRFs from Section 4. The model is calibrated around a zero inflation steady-state.

Table 5 summarizes the parameters and their values used in the numerical analysis that are common for all countries in the currency union. Most parameters are taken from Galí and Monacelli (2008). An exception is $\lambda_{\tilde{y}}$, which is not part of the numerical analysis of Galí and Monacelli (2008) and γ, the terms-of-trade channel. Strictly speaking, given the parameter value of $\sigma = 1$, the parameter for terms-of-trade should be $\gamma = 0$. However, I keep it slightly positive to let the countries adjust for trade, taking the value from Bertasiute et al. (2020). The small parameter value only affects the numerical results negligibly. The pass-through parameter b_z as well as the weights associated with communication in the welfare function $(\lambda_z, \lambda_{\Delta z})$ are not separately identified in the data. I set $b_z = 0.30$, reflecting the evidence that media transmit central-bank communication only partially and $\lambda_z = \lambda_{\Delta z} = 10$, reflecting some cost of communication. Because these parameters define how much the communication channel is used, I verify different values across the two regimes to check the welfare ranking.

In Table 6, I summarize the parameters that require country-specific values. The country weights, w^i, are each country's share of private final consumption (Eurostat). The noise variances are the results obtained from the maximum likelihood estimation in

Table 5: Core Parameters

Parameter	Description	Value	Target/Reference
Core New Keynesian Parameters			
β	Discount factor	0.99	Galí and Monacelli (2008)
σ	Curvature of utility of consumption	1.00	Galí and Monacelli (2008)
θ	Calvo parameter	0.75	Galí and Monacelli (2008)
ζ	Curvature of disutility of labor	3.00	Galí and Monacelli (2008)
α	Foreign-good / CPI weight	0.40	Galí and Monacelli (2008)
κ	Slope of Phillips curve	0.34	Implied by $\beta, \sigma, \theta, \zeta$
γ	Open-economy / terms-of-trade channel	0.01	Bertasiute et al. (2020)
ρ_r	Persistence of interest rate	0.50	Galí and Monacelli (2008)
Φ_π	Taylor-rule response to inflation	1.50	Galí and Monacelli (2008)
Φ_y	Taylor-rule response to output gap	0.50	Bertasiute et al. (2020)
ϵ	Elasticity between goods	6.00	Galí and Monacelli (2008)
$\lambda_\pi = \lambda_{\pi^i}$	Inflation weights in welfare function	70.0	Implied by ϵ, β, θ
$\lambda_{\tilde{y}} = \lambda_{\tilde{y}^i}$	Output-gap weights in welfare function	4.00	Implied by σ, ζ
$\lambda_z = \lambda_{\Delta z}$	Communication weights in welfare function	10.0	Assigned
Media Transmission			
ϱ_π	Weight on inflation newsworthiness	100.0	Calibrated to Section 4
$\varrho_{\tilde{y}}$	Weight on output-gap newsworthiness	100.0	Calibrated to Section 4
ϱ_z	Weight on central bank communication	100.0	Calibrated to Section 4
ρ_χ	Persistence of media reporting	0.30	Calibrated to Section 4
ρ_ψ	Persistence of reporting noise	0.80	Calibrated to Section 4
b_z	Pass-through of central bank communication	0.30	Assigned

Notes: The table reports parameters that are common across countries in the four-country currency-union model. Variables are expressed as deviations from steady state.

Section 5. For simplicity, the results obtained for inflation are also used for the output-gap. For steady-state media frequency, the mean reporting frequency of Section 3 is taken. The elasticity of newsworthiness η_V^i is set, together with the newsworthiness weights, to match the IRFs of the media direction disturbance obtained in Section 4.

Table 6: Country-Specific Parameters

Parameter	Description	Germany	Spain	France	Italy	Target/Reference
Currency Union Weights						
w^i	Country weight	0.370	0.150	0.270	0.210	Private final consumption
Expectation Formation and Media						
$\frac{1}{\tau_\pi^i} = \frac{1}{\tau_{\tilde{y}}^i}$	Prior noise variance	0.00034	0.00062	0.00033	0.00040	Estimates from Section 5
$\frac{1}{\tau_\psi^i} = \frac{1}{\tau_\varphi^i}$	Media noise variance	0.00037	0.00055	0.00025	0.00033	Estimates from Section 5
$\frac{1}{\tau_\phi^i} = \frac{1}{\tau_\varsigma^i}$	Private noise variance	0.0014	0.0012	0.0008	0.0012	Estimates from Section 5
$\bar{V}_i$	Steady-state media frequency	0.6977	0.5424	0.5396	0.4908	Estimates from Section 3
η_V^i	Elasticity of newsworthiness	70	60	50	70	Calibrated to Section 4

Notes: The table reports country-specific parameters. Countries are indexed as follows: $i = 1$ Germany, $i = 2$ Spain, $i = 3$ France, and $i = 4$ Italy.

Figure 9 presents the calibration results under the Taylor rule. I compare the reaction of inflation expectations from the estimation results on the innovation to media direction presented in Section 4 with the results on the disturbance to media reporting, $\varepsilon_t^{\psi,i}$. The results from the theoretical model (black dotted line) do not perfectly match the median results from the VAR (black line). In Spain and Italy, the peak response from the model is slightly lower than the empirical results. For Germany, Spain, and Italy, the empirical results are more persistent than the model can replicate. However, all theoretical results lie within the 5-95 percentiles (blue shaded areas) of the empirical counterpart. Therefore, I conclude that the values of ϱ_π, $\varrho_{\tilde{y}}$ and η_V^i are plausible, generating the reactions empirically observed.

6.6 Results

This subsection presents the model's results in three steps. I first examine optimal central-bank communication and household belief formation under a country-specific supply disturbance, then repeat the exercise for a union-wide disturbance, where the model reproduces the shift toward media-based beliefs that the data display after 2021. Fi-

Figure 9: Impulse Response Functions from the Innovation to Media Direction (Section 4) and Inflation Reporting (Section 6)

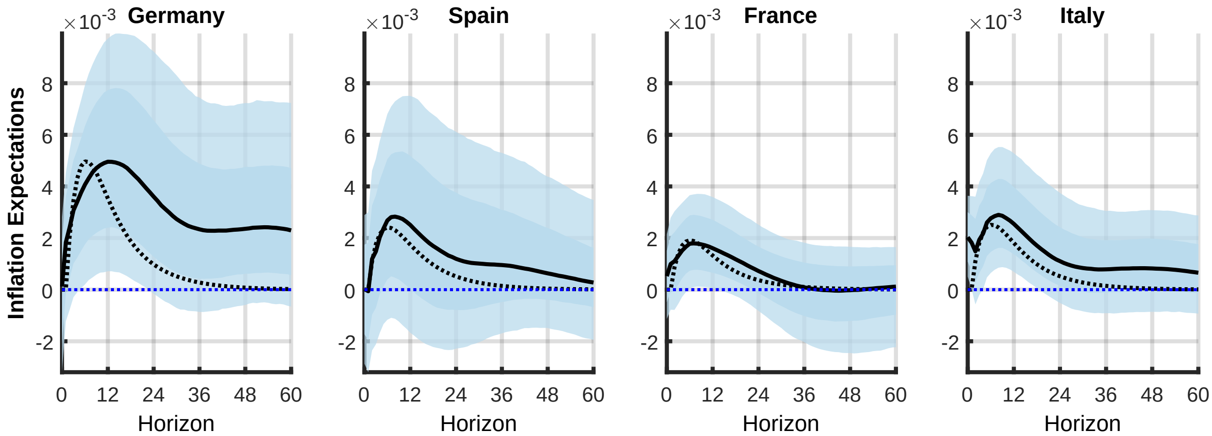

Note: The black solid lines are the median response and blue shaded areas represent the 5–95 percentiles of the impulse responses identified with the innovation to media direction about inflation. The black dotted lines mark the impulse response from the media reporting about inflation from the New Keynesian Model. The impulse responses are normalized to a standard deviation increase in the average direction surrounding the topic inflation per article per month. The figures display the responses of inflation expectations in the four countries: Germany, Spain, France, and Italy. The horizontal axis shows the time horizon in months.

nally, I compare welfare under the two policy regimes, quantifying the gain from adding country-specific communication to the union-wide policy rate.

Figure 10 shows the optimal communication response to a country-specific supply disturbance (one percent standard deviation) for Germany, Spain, France, and Italy. Communication is negative in every country: facing an inflationary supply shock, the central bank uses communication to lower household inflation expectations everywhere. Its magnitude differs across countries, and for two distinct reasons. It is strongest for Germany (-0.0020) and France (-0.0018), weaker for Italy (both -0.0015) and Spain (-0.0014). For France, the strong response reflects weak media transmission (low reporting frequency and low precision) so the central bank must supply a stronger signal to reach households. For Germany, it reflects instead the country's large consumption weight in the union, which the central bank internalizes. Because the baseline response mixes these two forces, transmission strength and country size, I isolate the transmission channel under equal weights in Appendix H.4.1: there, Spain and France (the weakest-transmission countries) receive the strongest communication (-0.0018 each), while Germany and Italy receive less (-0.0016 and -0.0017). Communication is also more persistent where reporting frequency is low.

Figure 10: Impulse Response Functions from Country-Specific Supply Disturbance

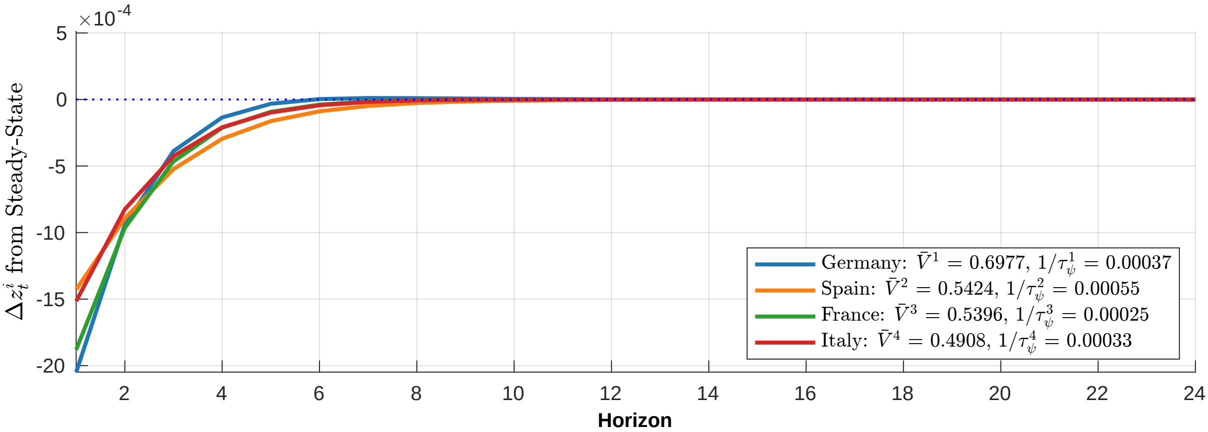

Note: The figure shows the impulse response functions of central bank communication, z_t^i, in response to a country-specific supply disturbance, ξ_t^i, for the countries Germany (blue line), Spain (yellow line), France (green line) and Italy (red line).

Figure 11 shows the response of the household's signal weights—on the prior, media, and private observation—to the country-level supply disturbance. In every country the weight on media rises on impact, because more communication raises reporting frequency, which raises the precision of the media signal and shifts belief formation toward it. The increase is largest in France (+0.033), followed by Spain (+0.02), Germany (+0.016), and Italy (+0.011). Even so, the media weight does not overtake the prior in every country.

The optimal communication response to a union-wide supply disturbance, in Figure 12, differs in an instructive way. Communication is now strongest for Germany (-0.0005), reflecting both its large union share and its strong media frequency and precision. France and Italy receive almost equally strong signals, despite the latter having a smaller union weight. In Italy, media transmission is precise, raising communication raises reporting frequency and delivers a relatively precise signal to households. France and Spain instead receive communication over a more extended horizon, again underscoring that persistent communication is optimal where media frequency and precision are low.

The signal-weight response to the union-wide disturbance, in Figure 13, is the most telling result. The weight on media rises strongly in all countries. Crucially, the media weight rises above the prior, the same reordering that the estimates of Section 5 identify in the data after 2021. The model thus reproduces part of the belief dynamics observed once the currency union was hit by a succession of supply disturbances after 2021, linking

Figure 11: Bayesian Weights from Country-Specific Supply Disturbance

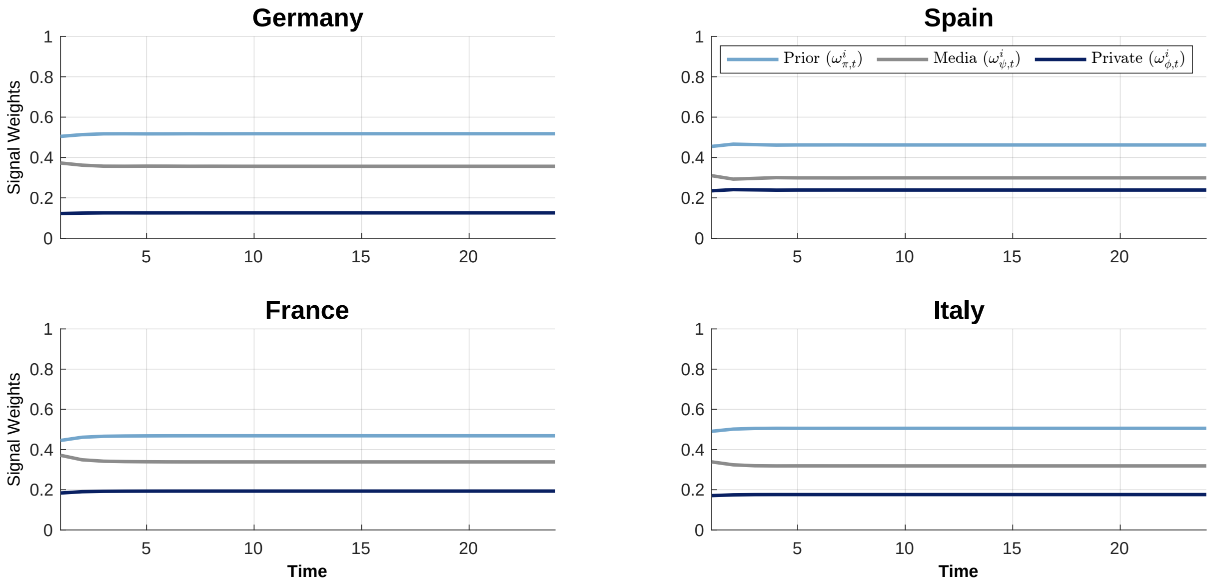

Note: The figure shows the Bayesian parameter weights for the household's inflation expectations of the prior belief (ω^i_π; light blue line), the media report (ω^i_ψ; grey line) and her private observation of inflation (ω^i_ϕ; dark blue line) in response to a country-specific supply disturbance.

Figure 12: Impulse Response Functions from Currency-Union Supply Disturbance

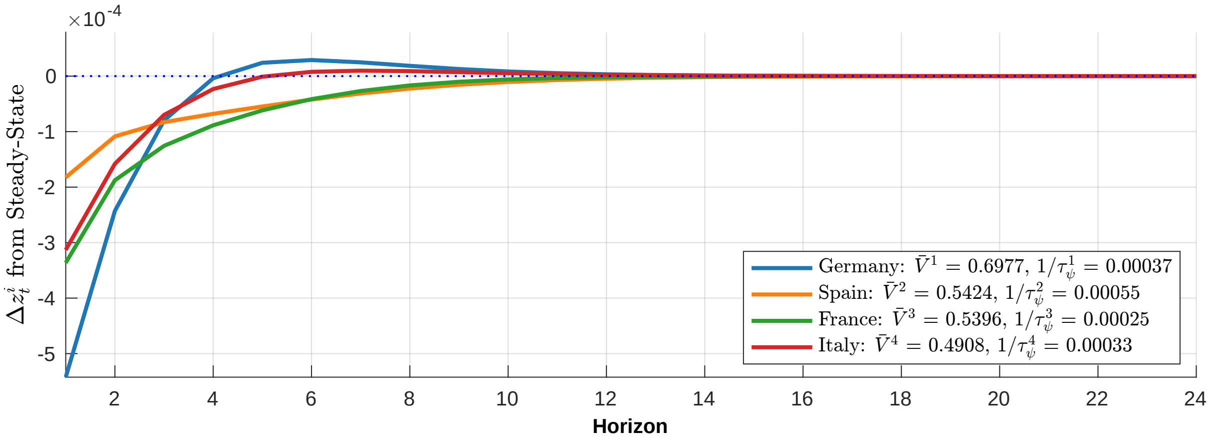

Note: The figure shows the impulse response functions of central bank communication, z^i_t, in response to a currency-union supply disturbance, ξ^i_t, for the countries Germany (blue line), Spain (yellow line), France (green line) and Italy (red line).

the model's mechanism to the empirical shift documented earlier.

Taken together, the central bank uses country-specific communication in addition to the union-wide policy rate when transmission is heterogeneous across media. Moreover, across both the country-specific and union-wide disturbances, the communication regime generates a permanent consumption-equivalent gain of 0.068% for the union relative to the policy rate alone. The baseline calibration uses intermediate values for the two

Figure 13: Bayesian Weights from Currency-Union Supply Disturbance

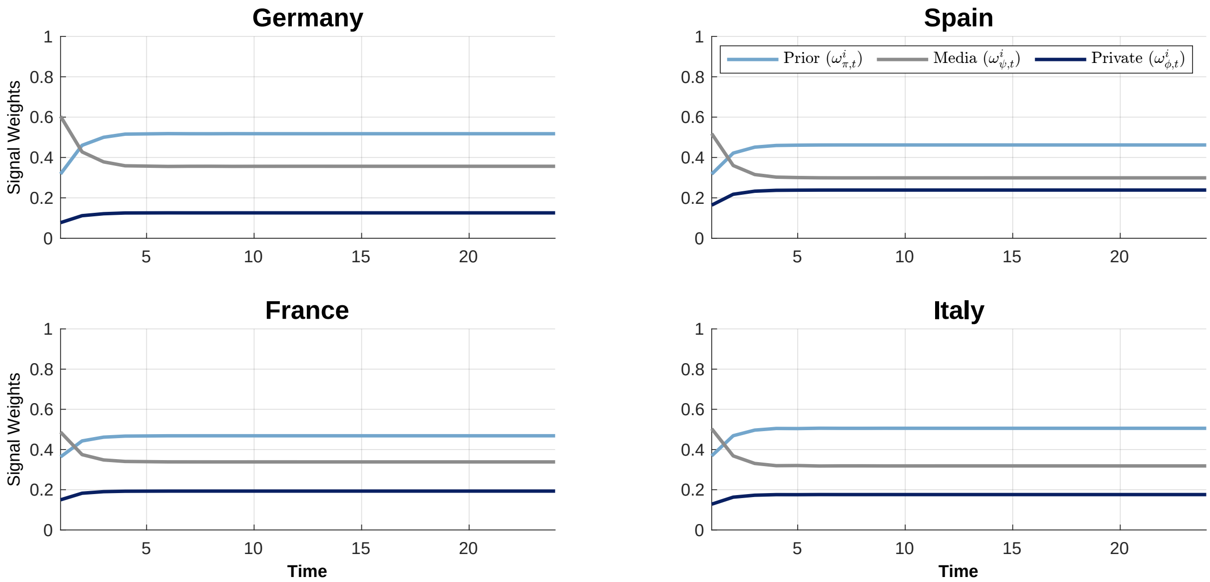

Note: The figure shows the Bayesian parameter weights for the household's inflation expectations of the prior belief (ω^i_π; light blue line), the media report (ω^i_ψ; grey line) and her private observation of inflation (ω^i_ϕ; dark blue line) in response to a currency-union supply disturbance.

assigned parameters, a communication pass-through of $b_z = 0.3$, so that media relay only part of the central bank's signal, and cost weights $\lambda_z = \lambda_{\Delta z} = 10$, high enough that communication is used with discipline rather than freely. The gain is therefore not an artifact of favorable parameters: Appendix H.4.2 shows that the welfare ranking favors communication across the entire grid of pass-through and cost values, falling as communication becomes costlier or its pass-through weaker and rising to as much as 0.78% when pass-through is strong and costs are low. Consistent with the mechanism, communication delivers a larger welfare gain in France than in Germany, the weakest-transmission country benefits more than the strongest, a contrast that is robust to using equal country weights. Under equal weights the ordering extends across all four countries, with welfare gains rising monotonically as media transmission weakens.

7 Conclusion

The analysis demonstrates that the way households are exposed and respond to central bank communication depends critically on the media environments through which they receive it. Even when the ECB delivers a centralized and consistent message, national media mediate this information in ways that introduce substantial variation across countries. By separating frequency from direction, the analysis reveals that frequency primarily affects inflation perceptions, while direction can shape both inflation perceptions and expectations in certain countries. These differences are not driven by underlying macroeconomic developments but by the way national media outlets frame and prioritize economic topics.

A behavioral model of household belief formation shows that these media environments feed directly into expectations: during inflationary periods, households shift weight away from their own price observations and prior beliefs toward media signals, though the strength and timing of this shift differ markedly across countries. Embedding this mechanism in a New Keynesian model of the currency union shows that these differences also matter for policy. Because a single interest rate reaches households unevenly while communication can be targeted, the central bank optimally directs stronger and more persistent communication to the countries whose media transmit its signal least effectively, and pairing such country-specific communication with the common policy rate raises union welfare.

As a result, despite the ECB's substantial efforts to strengthen its outreach, significant national differences in households' responsiveness still persist, underscoring the challenges of conveying a centralized message effectively across diverse media outlets, languages, and editorial environments.

8 References

Afrouzi, Hassan, Olivier Coibion, Yuriy Gorodnichenko, and Saten Kumar, "Inflation Targeting Does Not Anchor Inflation Expectations: Evidence from Firms in New Zealand," *Brookings Papers on Economic Activity*, 2015.

Aruoba, S Boragan and Thomas Drechsel, "Identifying Monetary Policy Shocks: A Natural Language Approach," *NBER Working Paper 32417*, 2024.

Baumeister, Christiane and James D. Hamilton, "Sign Restrictions, Structural Vector Autoregressions, and Useful Prior Information," *Econometrica*, 2015, *83* (5), 1963–1999.

Beckers, Kathleen, Andrea Masini, Julie Sevenans, Miriam van der Burg, Julie De Smedt, Hilde Van den Bulck, and Stefaan Walgrave, "Are newspapers' news stories becoming more alike? Media content diversity in Belgium, 1983–2013," *Journalism*, 12 2019, *20* (12), 1665–1683.

Bernanke, Ben S, Mark Gertler, and Mark Watson, "Systematic Monetary Policy and the Effiects of Oil Price Shocks," *Brookings Papers on Economic Activity*, 1 1997, pp. 91–157.

Bertasiute, Akvile, Domenico Massaro, and Matthias Weber, "The behavioral economics of currency unions: Economic integration and monetary policy," *Journal of Economic Dynamics and Control*, 3 2020, *112*, 103850.

Binder, Carola, "Fed speak on main street: Central bank communication and household expectations," *Journal of Macroeconomics*, 6 2017, *52*, 238–251.

Blinder, Alan S., Michael Ehrmann, Jakob de Haan, and David Jan Jansen, "Central Bank Communication with the General Public: Promise or False Hope?," *Journal of Economic Literature*, 6 2024, *62* (2), 425–457.

Blinder, Alan S, Michael Ehrmann, Marcel Fratzscher, Jakob De Haan, and David-Jan Jansen, "Central Bank Communication and Monetary Policy: A Survey of Theory and Evidence," *Journal of Economic Literature*, 11 2008, *46* (4), 910–945.

Bolliger, Elio, "Inflation Expectations, Perceptions and News Media: Regional Differences in Switzerland," Technical Report 2024.

Brouwer, Nils and Jakob de Haan, "Trust in the ECB: Drivers and consequences," *European Journal of Political Economy*, 9 2022, *74*, 102262.

Bybee, Leland, Bryan Kelly, Asaf Manela, and Dacheng Xiu, "Business News and Business Cycles," *The Journal of Finance*, 10 2024, *79* (5), 3105–3147.

Calvo, Guillermo A., "Staggered prices in a utility-maximizing framework," *Journal of Monetary Economics*, 9 1983, *12* (3), 383–398.

Carroll, C. D., "Macroeconomic Expectations of Households and Professional Forecasters," *The Quarterly Journal of Economics*, 2 2003, *118* (1), 269–298.

Carvalho, Carlos and Fernanda Nechio, "Do people understand monetary policy?," *Journal of Monetary Economics*, 9 2014, *66*, 108–123.

Centre for Media Pluralism and Media Freedom, "Monitoring media pluralism in the digital era: application of the Media Pluralism Monitor in the European Union, Albania, Montenegro, the Republic of North Macedonia, Serbia and Turkey in the year 2021," Technical Report, Centre for Media Pluralism and Media Freedom, Media Pluralism Monitor, 2022 2022.

Chahrour, Ryan, Adam Hale Shapiro, and Daniel Wilson, "News Selection and Household Inflation Expectations," *NBER WP*, 5 2025.

_ , Kristoffer Nimark, and Stefan Pitschner, "Sectoral Media Focus and Aggregate Fluctuations," *American Economic Review*, 12 2021, *111* (12), 3872–3922.

Christelis, Dimitris, Dimitris Georgarakos, Tullio Jappelli, and Maarten van Rooij, "Trust in the Central Bank and Inflation Expectations," *International Journal of Central Banking*, 12 2020, *16* (6), 1–37.

Coibion, Olivier and Yuriy Gorodnichenko, "Inflation, Expectations and Monetary Policy: What Have We Learned and to What End?," Technical Report, National Bureau of Economic Research, Cambridge, MA 5 2025.

_ , _ , and Michael Weber, "Monetary Policy Communications and Their Effects on Household Inflation Expectations," *Journal of Political Economy*, 6 2022, *130* (6), 1537–1584.

Conneau, Alexis, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzman, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov, "Unsupervised Cross-lingual Representation Learning at Scale," 4 2020.

Cruijsen, Carin Van Der, David-Jan Jansen, and Jakob De Haan, "How Much Does the Public Know about the ECB's Monetary Policy? Evidence from a Survey of Dutch Households *," *International Journal of Central Banking*, 12 2015, *42.*

Den, Linda Van, Heijkant Martine, Van Selm, Iina Hellsten, and Rens Vliegenthart, "Intermedia Agenda Setting in a Policy Reform Debate," *International Journal of Communication*, 2019, *13*, 1890–1912.

Draghi, Mario, "Central bank communication," 8 2014.

Drenik, Andrés and Diego J. Perez, "Price setting under uncertainty about inflation," *Journal of Monetary Economics*, 12 2020, *116*, 23–38.

European Central Bank, "ECB introduces a negative deposit facility interest rate," 6 2014.

European Commission, "Standard Eurobarometer 95: Standard Eurobarometer 95 Spring 2021," Technical Report, European Commission. Directorate-General for Communication., Brussels 9 2021.

Galí, Jordi and Tommaso Monacelli, "Monetary Policy and Exchange Rate Volatility in a Small Open Economy," *Review of Economic Studies*, 7 2005, *72* (3), 707–734.

_ and _ , "Optimal monetary and fiscal policy in a currency union," *Journal of International Economics*, 9 2008, *76* (1), 116–132.

Gardner, Ben, Chiara Scotti, and Clara Vega, "Words Speak as Loudly as Actions: Central Bank Communication and the Response of Equity Prices to Macroeconomic Announcements," *Journal of Econometrics*, 7 2022, *231* (2), 387–409.

Giannone, Domenico, Michele Lenza, and Giorgio E. Primiceri, "Prior Selection for Vector Autoregressions," *Review of Economics and Statistics*, 5 2015, *97* (2), 436–451.

Gorodnichenko, Yuriy, Tho Pham, and Oleksandr Talavera, "Central bank communication on social media: What, to whom, and how?," *Journal of Econometrics*, 10 2024, p. 105869.

Handlan, Amy, "Text Shocks and Monetary Surprises: Text Analysis of FOMC Statements with Machine Learning," Technical Report 2022.

Hansen, Stephen and Michael McMahon, "Shocking language: Understanding the macroeconomic effects of central bank communication," *Journal of International Economics*, 3 2016, *99*, S114–S133.

_ , _ , and Andrea Prat, "Transparency and Deliberation Within the FOMC: A Computational Linguistics Approach*," *The Quarterly Journal of Economics*, 5 2018, *133* (2), 801–870.

Harder, Raymond A., Julie Sevenans, and Peter Van Aelst, "Intermedia Agenda Setting in the Social Media Age: How Traditional Players Dominate the News Agenda in Election Times," *The International Journal of Press/Politics*, 7 2017, *22* (3), 275–293.

Klapper, Leora and Annamaria Lusardi, "Financial literacy and financial resilience: Evidence from around the world," *Financial Management*, 9 2020, *49* (3), 589–614.

Lamla, Michael J. and Sarah M. Lein, "The role of media for consumers' inflation expectation formation," *Journal of Economic Behavior & Organization*, 10 2014, *106*, 62–77.

Larsen, Vegard Høghaug, Nicolò Maffei-Faccioli, and Laura Pagenhardt, "Where do they care? The ECB in the media and inflation expectations," *Applied Economics Letters*, 4 2025, *32* (7), 945–950.

Lenza, Michele and Giorgio E. Primiceri, "How to estimate a vector autoregression after March 2020," *Journal of Applied Econometrics*, 6 2022, *37* (4), 688–699.

Litterman, Robert B, "Techniques of forecasting using vector autoregressions," Technical Report 115, Federal Reserve Bank of Minneapolis 1979.

Loughran, Tim and Bill McDonald, "When Is a Liability Not a Liability? Textual Analysis, Dictionaries, and 10-Ks," *Journal of Finance*, 7 2011, *66* (1), 35–65.

Malmendier, Ulrike and Stefan Nagel, "Learning from Inflation Experience," *The Quarterly Journal of Economics*, 2 2016, *131* (1), 53–88.

Melosi, Leonardo, "Signalling Effects of Monetary Policy," *The Review of Economic Studies*, 9 2016, p. rdw050.

Morris, Stephen and Hyun Song Shin, "Social Value of Public Information," *American Economic Review*, 11 2002, *92* (5), 1521–1534.

Newman, Nic, Amy Ross Arguedas, Craig T. Robertson, Rasmus Kleis Nielsen, and Richard Fletcher, "Reuters Institute Digital News Report 2025," Technical Report, Reuters Institute for the Study of Journalism 2025.

Nimark, Kristoffer P. and Stefan Pitschner, "News media and delegated information choice," *Journal of Economic Theory*, 5 2019, *181*, 160–196.

Pfajfar, Damjan and Emiliano Santoro, "Heterogeneity, learning and information stickiness in inflation expectations," *Journal of Economic Behavior & Organization*, 9 2010, *75* (3), 426–444.

Pfäuti, Oliver, "The Inflation Attention Threshold and Inflation Surges," *American Economic Review (Forthcoming)*, 2023.

Wabitsch, Alena, "The Messenger Matters," Technical Report 2025.

Walder, Maxime, Daniel Schwarz, Nathalie Giger, and Jan Fivaz, "Trust in political institutions: The central role of the traditional media," *DeFacto*, 2 2024.

Weber, Michael, Bernardo Candia, Hassan Afrouzi, Tiziano Ropele, Rodrigo Lluberas, Serafin Frache, Brent Meyer, Saten Kumar, Yuriy Gorodnichenko, Dimitris Georgarakos, Olivier Coibion, Geoff Kenny, and Jorge Ponce, "Tell Me Something I Don't Already Know: Learning in Low- and High-Inflation Settings," *Econometrica*, 2025, *93* (1), 229–264.

Woodford, Michael, *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton University Press, 9 2003.

Zhang, Donghai, "Term Structure, Forecast Revision, and the Signaling Channel of Monetary Policy," *Journal of the European Economic Association*, 8 2022, *20* (4), 1522–1553.

Zhang, Xiaoqun, "Intermedia agenda-setting effect in corporate news: Examining the influence of The New York Times and The Wall Street Journal on local newspapers," *Journal of Applied Journalism & Media Studies*, 7 2018, *7* (2), 245–263.

A Understanding National Context

The difference observed between countries may reflect not just linguistic filtering but broader structural, cultural and institutional factors. Such factors can shape how households and firms receive, understand and respond to central banks' communication. Understanding these national contexts is crucial for interpreting the effects of ECB communication.

Trust in media is central to whether households engage with information about monetary policy presented by media outlets. When audiences distrust traditional media outlets, they are less likely to absorb or act upon coverage of central bank decisions, even when exposure is high. Walder et al. (2024) show that trust in political institutions is strongly tied to traditional media use, underlining the media's role in shaping public understanding. In countries where media outlets are perceived as reliable, messages from the ECB are more likely to influence inflation expectations effectively. Conversely, where trust is low, the link between central bank communication and household behavior weakens. As shown in Figure 14, trust in media varies notably across countries, with Germany (46%) and Italy (39%) reporting higher levels than Spain (31%) and France (27%), illustrating an uneven foundation for communication effectiveness.

Trust in the central bank itself shapes how households interpret and react to monetary policy messages. Christelis et al. (2020) show that trust in the ECB reduces inflation expectations on average and significantly reduces uncertainty about future inflation. This indicates that even clear and consistent communication cannot succeed if the public doubts the institution's credibility. Similarly, Brouwer and de Haan (2022) show that those who trust the ECB align their expectations more closely with its inflation target, reinforcing the role of confidence in policy transmission. Data from European Commission (2021) reveal that general trust in the ECB differs only slightly across countries. However, 22% of the participants in France didn't know whether they should trust the ECB compared to 14% in Spain, 12% in Germany and 10% in Italy. In the same study, participants were also asked if they have heard of the ECB. In Germany 94% had heart of the ECB, compared to 90% in Spain, 84% in France and 80% in Italy.

Financial literacy provides the foundation for understanding and acting on monetary policy communication. According to Klapper and Lusardi (2020) financial literacy rates

Figure 14: Trust in the Institutions

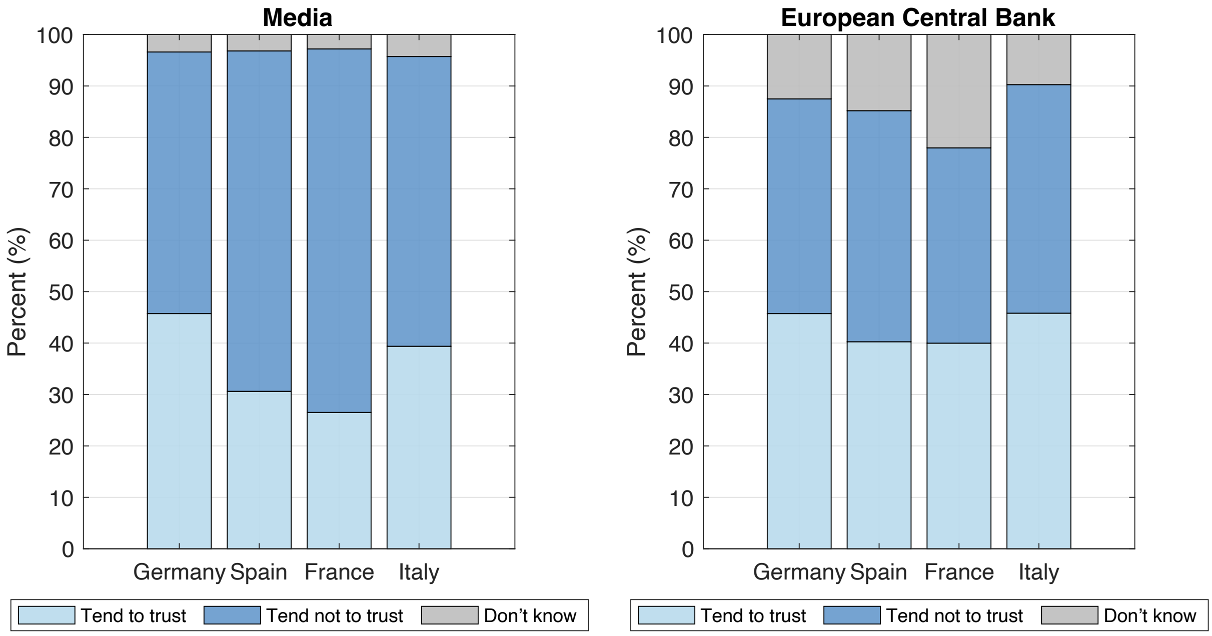

Note: Distribution of trust from European Commission (2021). For the left graph, participants where asked *How much trust do you have in certain institutions? For each of the following institutions, do you tend to trust it or tend not to trust it? The media.* Possible answer choices were *Tend to trust* (light blue), *Tend not to trust* (darker blue) or *Don't know* (grey). For the right graph, participants where asked *And do you tend to trust or tend not to trust these European institutions? The European Central Bank.* Answer choices were the same.

[19] differ widely across Europe, with Germany (66%) ahead of France (52%), Spain (49%), and Italy (37%), reflecting unequal capacity to interpret economic messages. These variations suggest that communication may resonate more effectively in countries where basic economic concepts are more familiar. Moreover, Binder (2017) argues that when information is too complex or abstract, individuals often disengage entirely, leaving the message unprocessed. Low financial literacy therefore weakens the transmission of central bank communication by limiting comprehension of key terms such as inflation or interest rates. This variation in understanding constrains how well households can adjust expectations in response to policy signals.

Together, the national context shapes how households interpret and internalize monetary policy communication. Even when information is available and well-formulated,

[19] Klapper and Lusardi (2020) calculate the share of respondents who answer at least three out of the five questions on financial literacy correctly. The questions assess basic knowledge of four fundamental concepts in financial decision making: knowledge of interest rates, interest compounding, inflation, and risk diversification

its impact depends on whether audiences find the source credible and the content comprehensible. Differences across countries in these factors help explain why central bank messages may shift inflation expectations unevenly. Understanding and addressing these soft foundations is thus essential for effective policy communication in a multilingual and diverse information environment.

B Frequency Dictionary for Topics

The table below presents the set of keywords employed for the text-based analysis explained in Section 3.2. For each topic category, the table lists word stems or frequent linguistic variants in English (ENG), German (DEU), Spanish (ESP), French (FRA), and Italian (ITA). During text processing, these stems are used to detect topic-related terms in the media corpus, allowing for consistent classification across languages.

Table 7: Keywords by language and topic.

ENG	**DEU**	**FRA**	**ITA**	**ESP**	**Topic**
Cost	kost	cost	coût	cost	Inflation
Deflation	deflation	deflac	déflat	deflaz	Inflation
Devaluation	entwert	devalu	dévalu	svalut	Inflation
Inflation	inflat	inflac	inflat	inflaz	Inflation
Inflation	teuerung				Inflation
Price	preis	preci	prix	prezz	Inflation
Purchasing power	kaufkraft	poder adquisit	pouvoir d'achat	potere d'acquisto	Inflation
Employers	arbeitgeb	empleador	employeur	datori	Labor
Employment	beschäftig	empleo	emploi	occupaz	Labor
Income	einkomm	ingreso	revenu	reddit	Labor
Job gains	beschäftig	empleo	emploi	occupaz	Labor
Job gains	stelle	contrat	embauch	assunz	Labor
Job losses	verlust	pérdida	perte	perd	Labor
Jobs	arbeitsplatz	empleo	emploi	lavoro	Labor
Labor market	arbeitsmarkt	mercado laboral	marché de travail	mercato del lavoro	Labor
Layoffs	entlass	licenci	licenci	licenz	Labor
Remuneration	vergüt	remuner	rémunér	retribuz	Labor
Unemployment	arbeitslos	desemple	chôm	disoccup	Labor
Wages	lohn	salari	salair	salar	Labor
Workers	arbeitnehm	trabajad	travaill	lavorat	Labor
Construction	bau	constru	construct	costruzion	Housing
Housing	wohn	viviend	logement	allogg / abit	Housing
Mortgages	hypothek	hipotec	hypoth	mutu	Housing
Real estate	immobil	inmobil	immobil	immobil	Housing

Continued on next page

ENG	DEU	ESP	FRA	ITA	Topic
Rent	miet	alquil	loyer	affitt	Housing
Bonds	anlei	bon	oblig	obblig	Financial
Capital	kapital	capital	capital	capital	Financial
Credit	kredit	crédit	crédit	credit	Financial
Credit		credit			Financial
Defaults	ausfall	incumpl	défaut	insolv	Financial
Debt	schuld	deud	dett	debit	Financial
Exchange rate	wechsel	cambi	change	cambi	Financial
Financial	finanz	financi	financi	finanz	Financial
Interest rates	zins	tipo	taux	tasso	Financial
Saving	spar	rispar	éparg	ahorro	Financial
Securities	wertpapier	valor	titre	titol	Financial
Stock	akti	accion	action	azion	Financial
Swaps	swap	swap	swap	swap	Financial
Swaps				scamb	Financial
Yields	rendit	rendim	rendem	rendiment	Financial
Crisis	kris	cris	cris	crisi	Economy
Confidence	vertrau	confianz	confianc	fiduc	Economy
Consumption	konsum	consum	consomm	consum	Economy
Economy	wirtsch	económ	économ	econom	Economy
Economy	konjunkt		eonom		Economy
Export	export	export	export	esport	Economy
GDP	bip	pib	pib	pib	Economy
Import	import	import	import	import	Economy
Indicators	indikator	indic	indicateur / indic	indic	Economy
Investment	invest	invest	invest	invest	Economy
Investment		invert			Economy
Production	produkt	produc	product	produz	Economy
Recession	rezess	reces	récess	recess	Economy
Uncertainty	unsicher	incertid	incertitud	incertezz	Economy

Note: Keywords used per topic in English (ENG), German (DEU), Spanish (ESP), French (FRA) and Italian (ITA). The keywords are constructed by using the most frequent n-grams per topic from the article corpus.

C Example Prompt to Generate Dataset for Fine-Tuning

Figure 15 displays an example prompt used to generate phrases that can be used for fine-tuning the language model on predicting the correct sentiment.

Figure 15: French prompt used to generate dataset for fine-tuning

```
J'entraine un modele linguistique avec des actualites financieres.
Pourrais-tu me fournir 25 phrases [label] sur le theme [topic] en
francais ? Inspire-toi de phrases d'exemples ci-dessous et
concentre-toi sur les actualites concernant la banque centrale.
Merci de fournir 25 phrases exemplaires differentes.

Example 0: [FRA_topic_example0]
Example 1: [FRA_topic_example1]
Example 2: [FRA_topic_example2]
Example 3: [FRA_topic_example3]
Example 4: [FRA_topic_example4]
Example 5: [FRA_topic_example5]
Example 6: [FRA_topic_example6]
Example 7: [FRA_topic_example7]
Example 8: [FRA_topic_example8]
Example 9: [FRA_topic_example9]
```

Note: French prompt used to generate 25 negative, neutral and positive phrases per topic. Words in squared brackets are dynamic, i.e., [label] was replaced by negative, neutral and positive; the [topic] adjusted to the five topics used in this study as well as the examples given to the model.

D Overview of the Training Algorithm

1. **Label encoding:** The categorical labels *downward*, *neutral*, and *upward* are converted into numerical values 0, 1, and 2, respectively. This transformation allows the model to process the categories as discrete classes during training.

2. **Definition of hyperparameters:** The main hyperparameters, which determine the learning behavior of the model, are defined as follows:

   (a) **Learning rate:** Specifies how much the model's internal parameters are adjusted after each iteration. A relatively small learning rate is chosen to ensure gradual and stable learning.

   (b) **Number of epochs:** Indicates how many times the model iterates over the entire training dataset to refine its parameters.

(c) **Loss function:** The loss function measures the difference between the model's predictions and the true labels. For this classification task, a cross-entropy loss function is used, as it is particularly suited for categorical outcomes. In this setting, the model predicts the probability of each class, and cross-entropy quantifies how close these predicted probabilities are to the correct class.

(d) **Batch size:** Determines how many input samples are processed together before updating the model parameters.

3. **Cross-validation:** The dataset is divided into five equally sized subsets (five-fold cross-validation). In each iteration, one subset serves as the validation set, and the remaining four are used for training. This procedure ensures that every observation is used once for validation and provides a more reliable estimate of the model's ability to generalize.

4. **Training phase:** The model is set to training mode, allowing it to adjust its internal parameters. During training, the model iteratively minimizes the loss function using an optimization algorithm (e.g., stochastic gradient descent). The process of adjusting the parameters based on the loss gradient is known as backpropagation.

5. **Validation phase:** After each training iteration (epoch), the model's performance is evaluated on the validation subset. This allows monitoring of potential overfitting, which occurs when the model learns the training data too precisely and performs poorly on unseen data.

6. **Early stopping:** If the validation performance no longer improves or begins to deteriorate while training performance continues to improve, the training process is stopped early to prevent overfitting.

7. **Repetition cross folds:** Steps 4 to 6 are repeated until all five folds have been used for validation. The overall performance of the model is then computed as the average across all folds.

8. **Label Decoding:** Finally, the numerical predictions (0, 1, 2) are converted back into their corresponding categorical labels: downward, neutral, and upward.

E Full Set of Macroeconomic Reactions

Figures 16 to 19 display the impulse response functions to the five shocks identified through inertial restrictions for Germany, Spain, France and Italy, respectively. The set includes the quantitative inflation perceptions and expectations, the results using the qualitative measures can be requested from the author. The five shocks are a shock to inflation, real GDP, monetary policy, media frequency and direction. The reactions from all macroeconomic variables are displayed: inflation, log real GDP, the policy rate, media frequency and direction, inflation perception and expectations, unemployment rate, asset growth as well as log stock prices. The shocks to inflation and real GDP are normalized to an 100-basis-point increase of the target series on impact. The monetary policy rate is normalized with 250-basis-points. The media measures are normalized to have a one standard deviation impact.

F Results using Qualitative Inflation Beliefs

Figures 20 to 21 display the impulse response functions of qualitative inflation perceptions and expectations to the frequency and direction shocks. The results mirror those in Section 5. The qualitative results confirm the pattern documented for the quantitative measures in Section 5: Both shocks move inflation beliefs and the ordering of the country ranking is preserved, although the frequency shock produces a more muted and less uniformly significant response. Following an increase in the frequency of inflation reporting, the balance statistic of qualitative inflation perceptions rises by 1 to 5 points, with the strongest response in Spain, followed by Italy, Germany, and France. The responses in Spain and Italy are significant at the 5-95 percentile band, whereas in Germany and France they are significant at the 16–84 percentile bands. When the direction of inflation reporting turns upward, the median balance statistic of qualitative inflation expectations rises by 0.7 to 2.1 points. Here, the effects are strongest in Germany and Italy, significant at the wider 5–95 percentile bands, while the weaker effects in Spain and France remain significant at the 16–84 bands. Qualitative inflation perceptions respond to the direction shock in the same way.

Figure 16: Impulse Responses to all Shocks in Germany

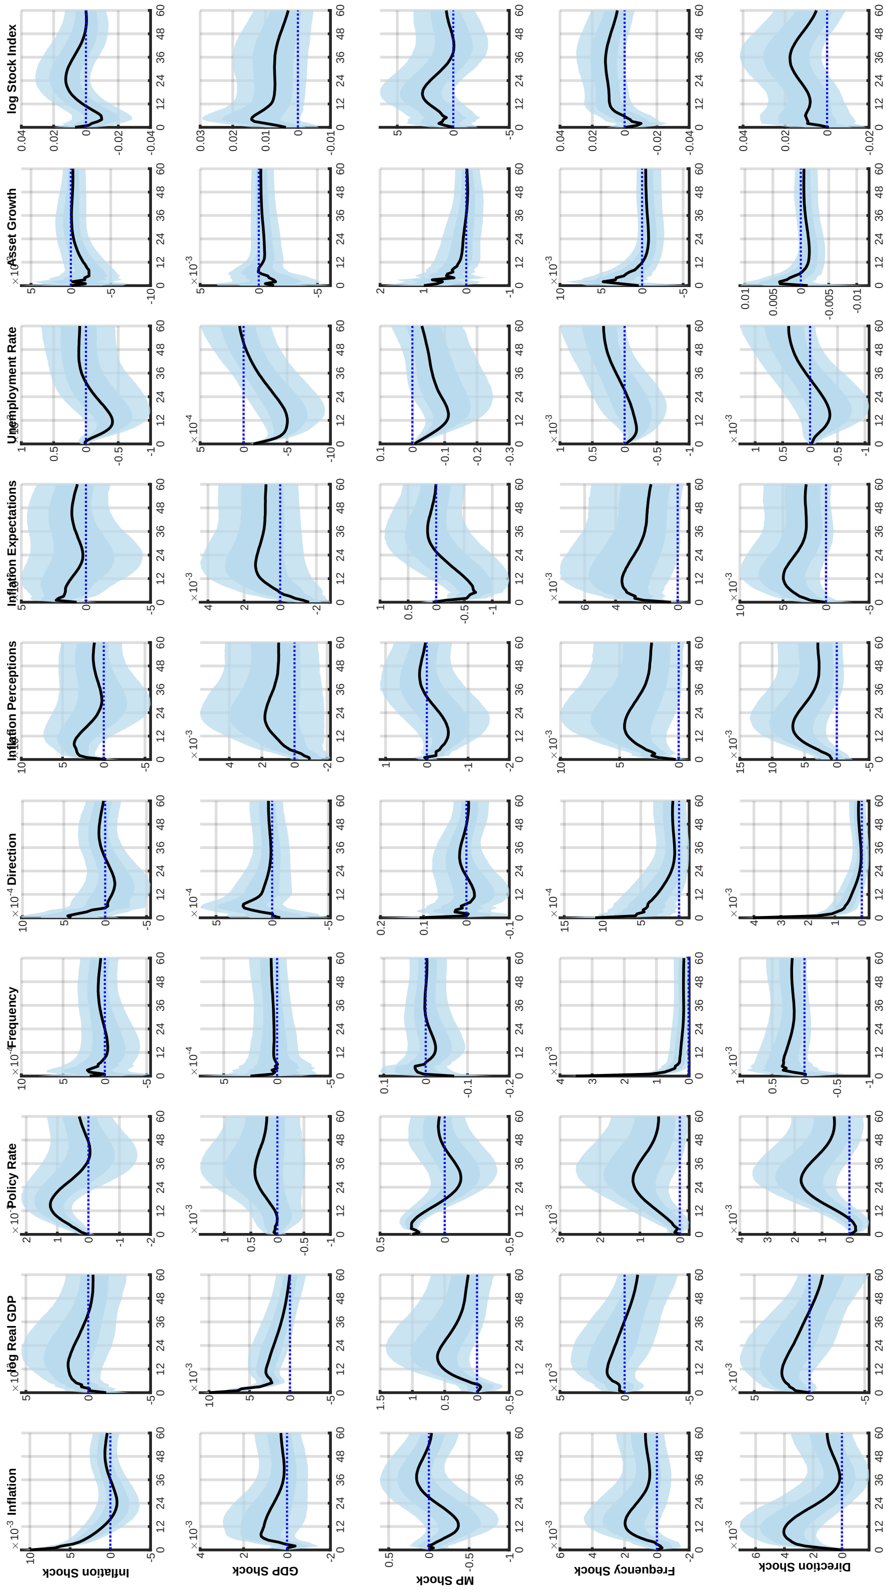

Note: The (dark) blue shaded areas represent the 5–95 (16-84) percentiles of the impulse responses identified using inertial restrictions. The horizontal axis shows the time horizon in months.

Figure 17: Impulse Responses to all Shocks in Spain

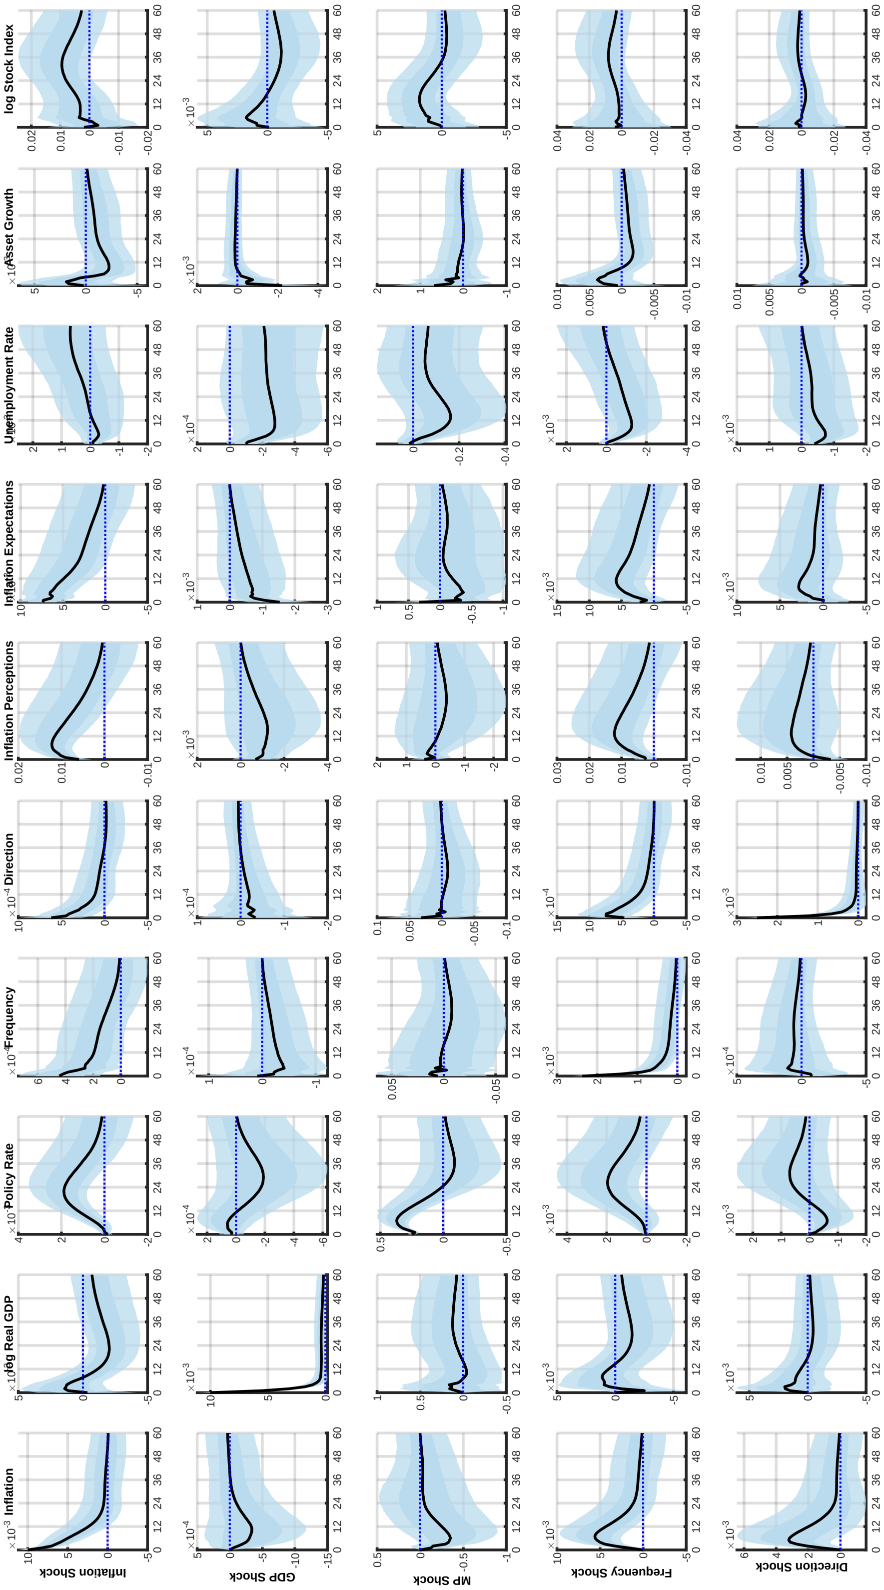

Note: The (dark) blue shaded areas represent the 5–95 (16-84) percentiles of the impulse responses identified using inertial restrictions. The horizontal axis shows the time horizon in months.

Figure 18: Impulse Responses to all Shocks in France

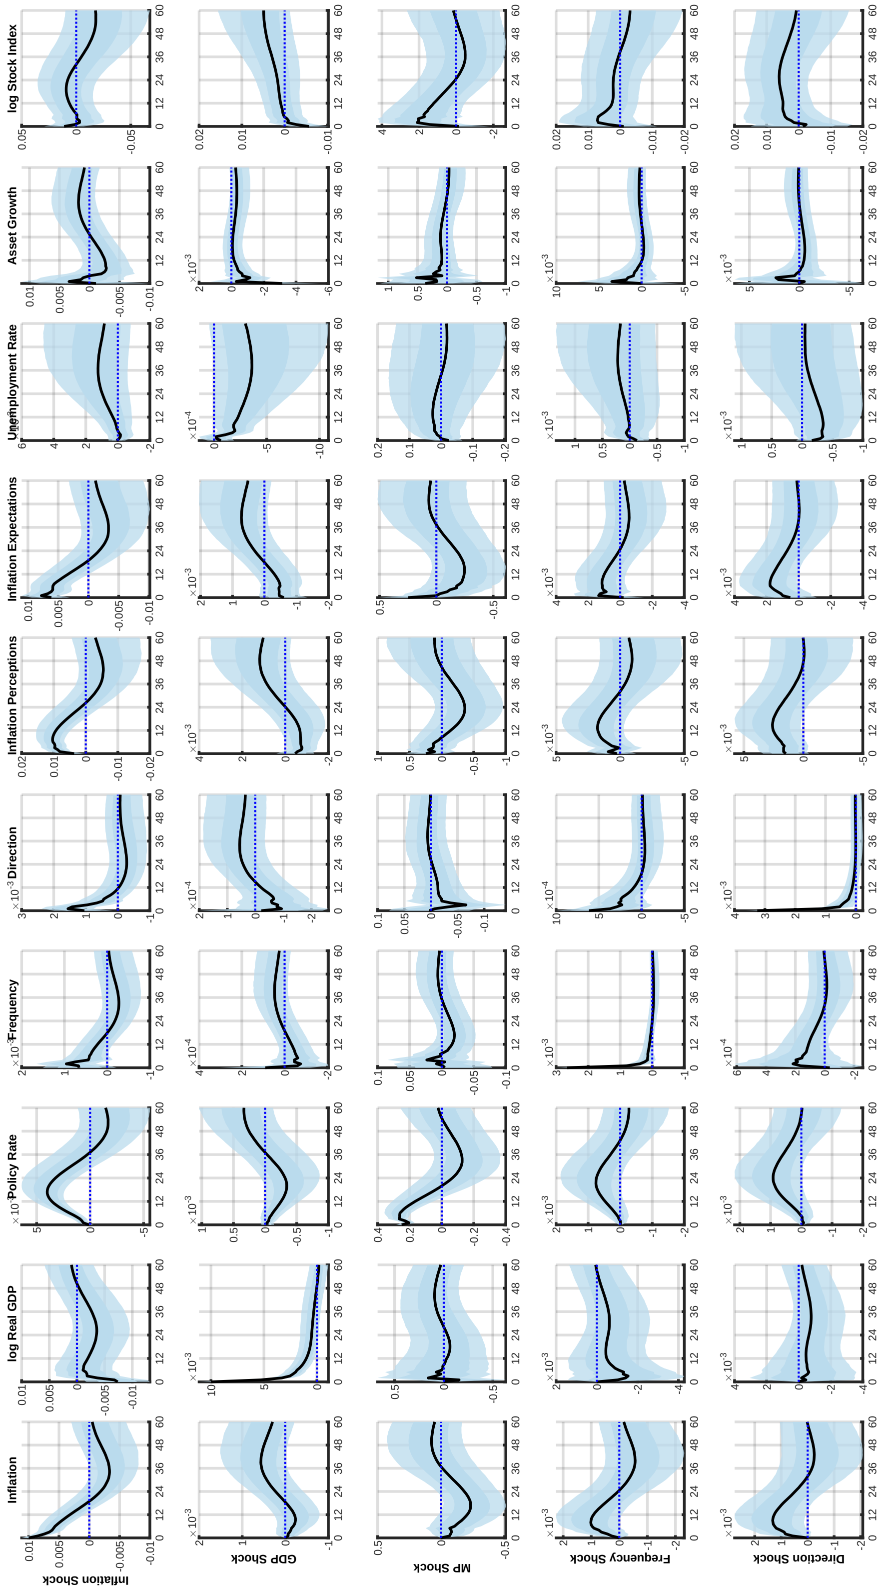

Note: The (dark) blue shaded areas represent the 5–95 (16-84) percentiles of the impulse responses identified using inertial restrictions. The horizontal axis shows the time horizon in months.

Figure 19: Impulse Responses to all Shocks in Italy

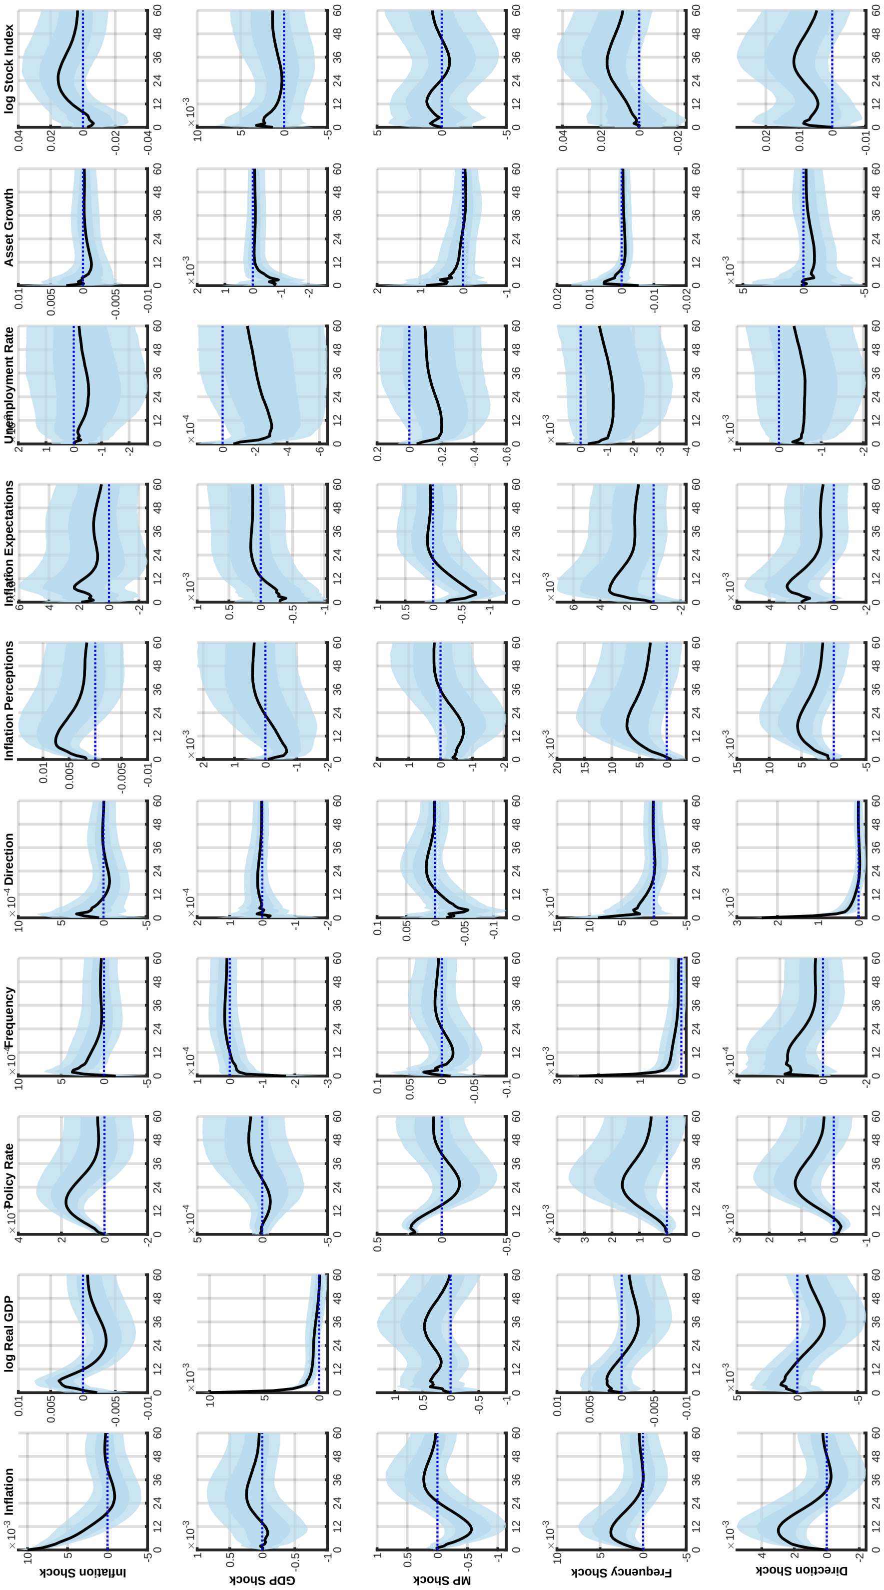

Note: The (dark) blue shaded areas represent the 5–95 (16-84) percentiles of the impulse responses identified using inertial restrictions. The horizontal axis shows the time horizon in months.

Figure 20: Impulse Response Functions of Qualitative Inflation Beliefs from the Innovation to Media Frequency

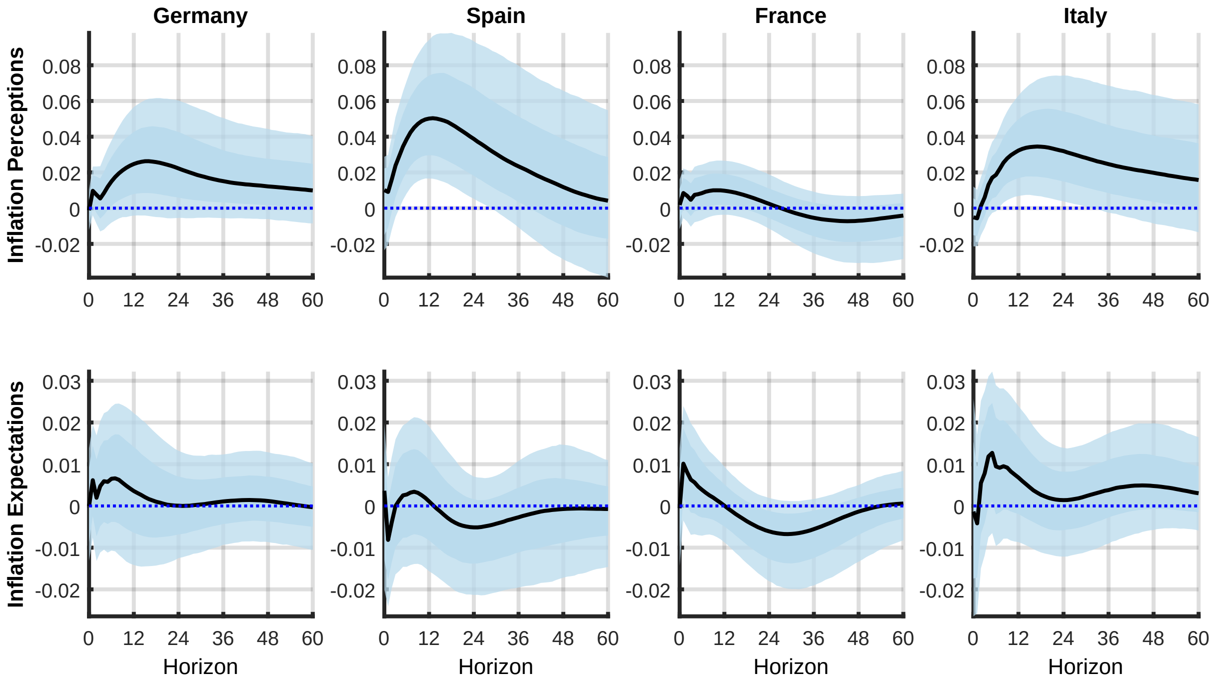

Note: The (dark) blue shaded areas represent the 5–95 (16-84) percentiles of the impulse responses identified with the innovation to media frequency about inflation. The impulse responses are normalized a standard deviation increase in the average number of mentions of the topic inflation per article per month. The figures display the responses of inflation perceptions and inflation expectations in the four countries: Germany, Spain, France, and Italy. The horizontal axis shows the time horizon in months.

G Contribution of Information Signals

Figure 22 presents the contributions of prior beliefs, media reports, and private observations to model-implied inflation expectations, together with the observed quantitative inflation expectations. Overall, the model-implied series track observed expectations closely, although they differ slightly in magnitude after 2021. In all countries, the estimated peak in inflation expectations lies somewhat below the observed peak, and the subsequent decline in model-implied expectations is slightly volatile than in the data.

Regarding the relative contributions of each component, prior beliefs remain consistently important across countries, reinforcing earlier evidence that they play a central role in expectation updating. As noted previously, media signals become more influential after 2021, coinciding with an increase in both the volume and intensity of inflation-related coverage.

Figure 21: Impulse Response Functions of Qualitative Inflation Beliefs from the Innovation to Media Direction

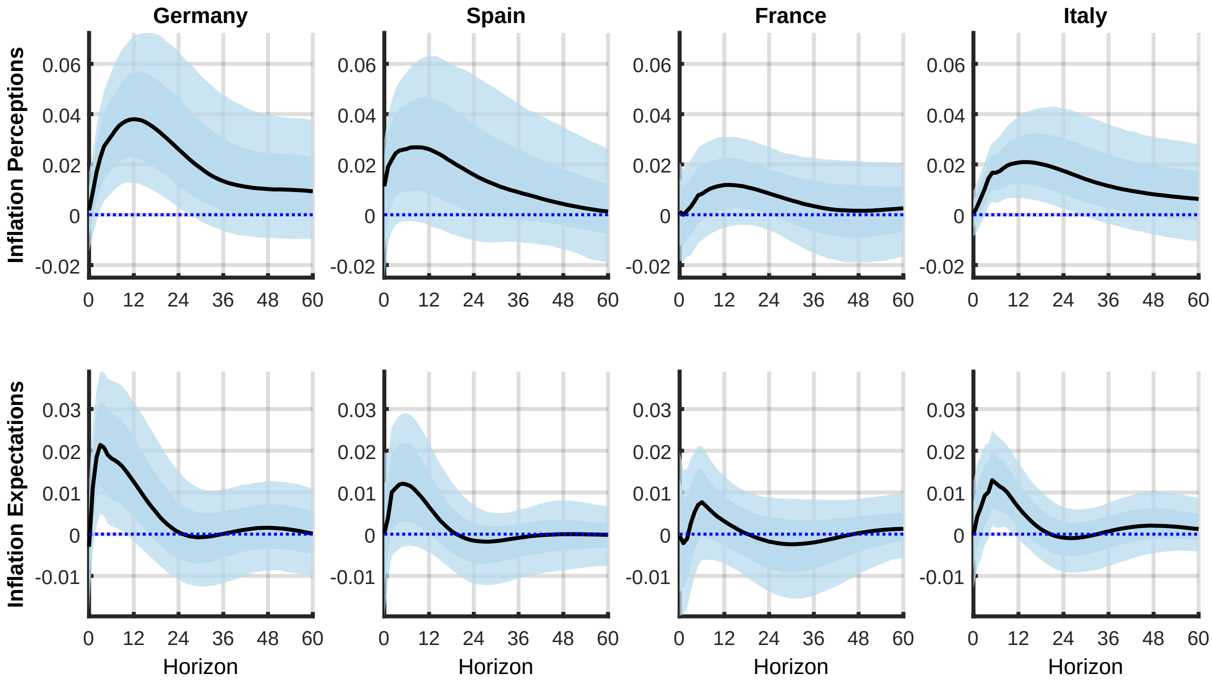

Note: The (dark) blue shaded areas represent the 5–95 (16-84) percentiles of the impulse responses identified with the innovation to media frequency about inflation. The impulse responses are normalized a standard deviation increase in the average direction surrounding the topic inflation per article per month. The figures display the responses of inflation perceptions and inflation expectations in the four countries: Germany, Spain, France, and Italy. The horizontal axis shows the time horizon in months.

H Microfoundation of the New Keynesian Model

The following derivations of the New Keynesian Model are based on the rational expectations open-economy model by Galí and Monacelli (2005, 2008) and on the behavioral currency union model by Bertasiute et al. (2020). First, I will derive the core equations 7 and 8 of the currency union model used in Section 6. Second, I provide a general-equilibrium consistent framework for deriving the expectation formation process presented in Section 5. Finally, I will derive the loss function of the central bank, which provides the groundwork for the central bank loss function used in Section 6.

H.1 Currency Union Model

To derive the IS and Phillips curves (7 and 8), I proceed in three steps. I first obtain the household's optimality conditions, the relationship between domestic and CPI inflation, and the risk-sharing condition within the union; I then turn to firms' technology and price

Figure 22: Contribution Towards Estimated Consumer's Inflation Expectations using Estimated Parameter Weights

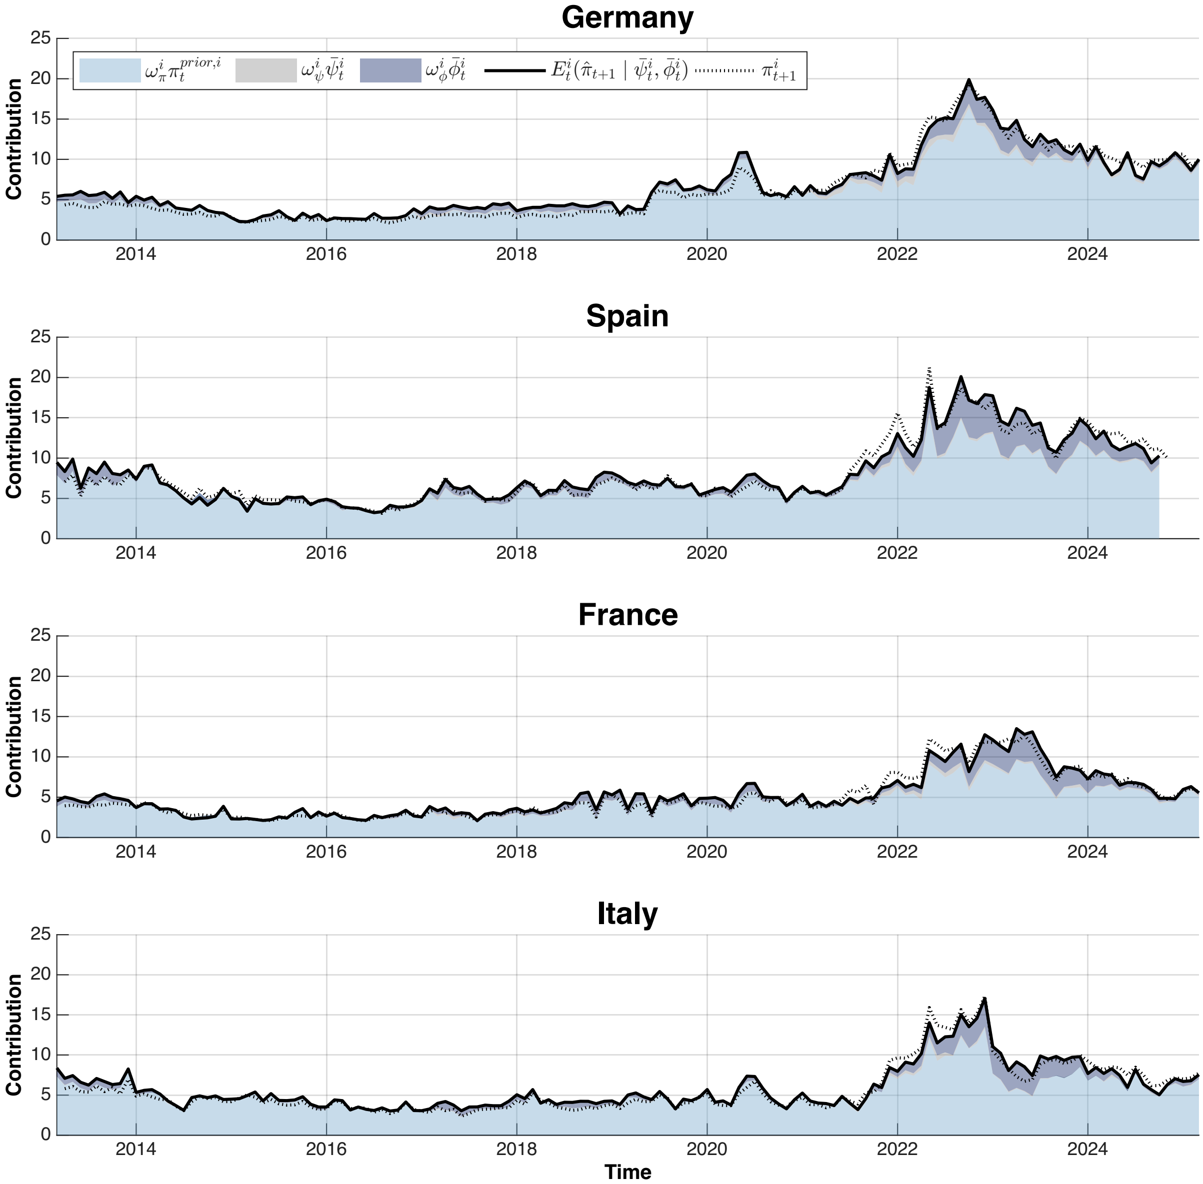

Note: The figure shows the contribution of the model estimates towards estimated consumer's inflation expectations ($\mathrm{E}^i_t(\pi_{t+1} \mid \bar{\psi}^i_t, \bar{\phi}^i_t)$, black line) given the prior belief ($\rho^i_\pi \pi^{prior,t}_t$; blue area), the media report ($\rho^i_\psi \bar{\psi}^i_t$; light grey area) and her private observation of inflation ($\rho^i_\phi \bar{\phi}^i_t$; dark grey area) over time together with the observed inflation expectations (π^i_{t+1}, black dotted line). Observations are demeaned.

setting; finally, I combine aggregate demand with marginal-cost and inflation dynamics to obtain the two core equations.

H.1.1 Households

Consider a generic country i part of a currency union. The representative household in country i chooses consumption C_t^i, labor N_t^i and bond holdings B_t^i to maximize

$$\mathrm{E}_t^i \sum_{\tau=0}^{\infty} \beta^\tau \left[\frac{(C_{t+\tau}^i)^{1-\sigma}}{1-\sigma} - \frac{(N_{t+\tau}^i)^{1+\zeta}}{1+\zeta} \right]. \tag{H.1}$$

Variable C_t^i is a composite consumption index defined by

$$C_t^i \equiv \frac{(C_{i,t}^i)^{1-\alpha}(C_{F,t}^i)^\alpha}{(1-\alpha)^{1-\alpha}\alpha^\alpha}.$$

The variable $C_{i,t}^i$ denotes consumption of the representative household in country i (superscript), of goods produced in country i (subscript), given by the CES aggregator

$$C_{i,t}^i = \left(\int_0^1 C_{i,t}^i(j)^{\frac{\epsilon-1}{\epsilon}} \, \mathrm{d}j \right)^{\frac{\epsilon}{\epsilon-1}},$$

where $j \in [0,1]$ indexes the type of goods, within the set produced in country i. In the preferences described above, $\alpha \in [0,1]$ is the weight of imported goods in the utility of private consumption. Equivalently, α can be seen as an index of openness. Following Galí and Monacelli (2005), I assume that each country produces a continuum of differentiated goods, that each good is produced by a distinct firm, and that no good is produced in more than one country.

Variable $C_{F,t}^i$ is an index of country i's consumption of imported goods given by

$$C_{F,t}^i = \exp \int_0^1 c_{f,t}^i \, \mathrm{d}f,$$

where $c_{f,t}^i = \log C_{f,t}^i$ is the log of an index of the quantity of goods consumed in country i, and produced in country f. In what follows, lower case letters will denote logs of the respective variables.

The index $C_{f,t}^i$ is defined symmetrically to $C_{i,t}^i$, namely

$$C_{f,t}^i = \left(\int_0^1 C_{f,t}^i(j)^{\frac{\epsilon-1}{\epsilon}} \, \mathrm{d}j \right)^{\frac{\epsilon}{\epsilon-1}}.$$

The parameter $\epsilon > 1$ denotes the elasticity of substitution between varieties produced within any given country independently of the producing country.

Maximization of Equation H.1 is subject to a sequence of budget constraints of the form:

$$\int_0^1 P_t^i(j) C_{i,t}^i(j) dj + \int_0^1 \int_0^1 P_t^f(j) C_{f,t}^i(j) dj df + B_t^i = B_{t-1}^i R_{t-1} + W_t^i N_t^i - T_t^i$$

where W_t^i is the gross wage and T_t^i are lump sum transfers including profits from firms. R_{t-1} is the interest rate for households. $P_t^f(j)$ is the price of good j produced in country f expressed in units of the single currency. The optimal allocation of expenditure on the goods produced in a certain country gives the following demand functions:

$$C_{i,t}^i(j) = \left(\frac{P_t^i(j)}{P_t^i}\right)^{-\epsilon} C_{i,t}^i \qquad \text{and} \qquad \text{(H.2)}$$
$$C_{f,t}^i(j) = \left(\frac{P_t^f(j)}{P_t^f}\right)^{-\epsilon} C_{f,t}^i$$

for all $i, j, f \in [0,1]$. Note that $P_t^i = \left(\int_0^1 P_t^i(j)^{1-\epsilon}\,\mathrm{d}j\right)^{\frac{1}{1-\epsilon}}$ denotes the index of prices of *domestically produced* goods, and $P_t^f = \left(\int_0^1 P_t^f(j)^{1-\epsilon}\,\mathrm{d}j\right)^{\frac{1}{1-\epsilon}}$ denotes the index of prices of goods *imported from country* f. Therefore, $\int_0^1 P_t^i(j) C_{i,t}^i(j)\,\mathrm{d}j = P_t^i C_{i,t}^i$ and $\int_0^1 P_t^f(j) C_{f,t}^i(j)\,\mathrm{d}j = P_t^f C_{f,t}^i$. Moreover, optimal allocation of expenditures on imported goods by country of origin yields

$$P_t^f C_{f,t}^i = P_t^* C_{F,t}^i,$$

where $P_t^* = \exp\int_0^1 p_t^f\,\mathrm{d}f$ is the union-wide price index. As there is a continuum of countries, $P_t^* = P_t^{*i}$, which implies that the union-wide price index is identical to the price index abroad, excluding country i. Finally, define the *consumer price index* (CPI) for country i as

$$P_{c,t}^i = (P_t^i)^{1-\alpha}(P_t^*)^{\alpha}. \qquad \text{(H.3)}$$

Then, the optimal allocation of expenditure between imported and domestically produced goods is

$$P_t^i C_{i,t}^i = (1-\alpha) P_{c,t}^i C_t^i \qquad \text{and} \qquad \text{(H.4)}$$
$$P_t^* C_{F,t}^i = \alpha P_{c,t}^i C_t^i.$$

Note that the total consumption expenditure by country i's household as $P_t^i C_{i,t}^i + P_t^* C_{F,t}^i = P_{c,t}^i C_t^i$. Conditional on an optimal allocation of expenditures, the period budget constraint can be rewritten as:

$$P_{c,t}^i C_t^i + B_t^i = B_{t-1}^i R_{t-1} + W_t^i N_t^i - T_t^i.$$

The first order conditions for country i's household optimization problem are given by

$$(C_t^i)^{-\sigma} = \beta R_t \, \mathrm{E}_t^i \left((C_{t+1}^i)^{-\sigma} \frac{P_{c,t}^i}{P_{c,t+1}^i} \right) \tag{H.5}$$

$$(C_t^i)^{-\sigma} \frac{W_t^i}{P_{c,t}^i} = (N_t^i)^{\zeta} \tag{H.6}$$

As standard, I require that agents' subjective transversality condition is satisfied, $\lim_{\tau \to \infty} \mathrm{E}_t^i \beta^{t+\tau} (C_{t+\tau}^i)^{-\sigma} \frac{B_{t+\tau}^i}{P_{t+\tau}^i} \leq 0$.

H.1.2 Domestic and CPI Inflation

Define the *bilateral terms of trade* between countries i and f as $S_{f,t}^i \equiv \frac{P_t^f}{P_t^i}$. The *effective terms of trade* can be written as

$$S_t^i = \frac{P_t^*}{P_t^i} = \exp \int_0^1 (p_t^f - p_t^i) \, \mathrm{d}f = \exp \int_0^1 s_{f,t}^i \, \mathrm{d}f.$$

Therefore, taking logs yields $s_t^i = \int_0^1 s_{f,t}^i \, \mathrm{d}f$.

Using the definition of the CPI, Equation H.3, one can derive the following relationship between CPI and domestic price levels

$$P_{c,t}^i = P_t^i (S_t^i)^{\alpha}. \tag{H.7}$$

The same equation in logs yields

$$p_{c,t}^i = p_t^i + \alpha s_t^i.$$

From the latter equation, it follows that the relationship between *CPI inflation*, $\pi_{c,t}^i \equiv p_{c,t}^i - p_{c,t-1}^i$, and *domestic inflation*, $\pi_t^i \equiv p_t^i - p_{t-1}^i$, is given by

$$\pi_{c,t}^i = \pi_t^i + \alpha \Delta s_t^i. \tag{H.8}$$

H.1.3 Risk-Sharing

Under the common interest rate regime, the first order condition analogous to Equation H.5 will hold for the representative household in country f. Under the assumption of complete markets, the two first order conditions can be combined as follows:

$$\frac{1}{R_t} = \beta \, \mathrm{E}_t^i \left[\left(\frac{C_t^i}{C_{t+1}^i} \right)^{\sigma} \left(\frac{P_{c,t}^i}{P_{c,t+1}^i} \right) \right] = \beta \, \mathrm{E}_t^f \left[\left(\frac{C_t^f}{C_{t+1}^f} \right)^{\sigma} \left(\frac{P_{c,t}^f}{P_{c,t+1}^f} \right) \right]$$

As in Bertasiute et al. (2020), I assume

$$\vartheta_i = \frac{\int_0^1 \mathrm{E}_t^i(C_{t+1}^i(P_{c,t+1}^i)^{1/\sigma})}{\int_0^1 \mathrm{E}_t^f(C_{t+1}^f(P_{c,t+1}^f)^{1/\sigma})} = 1.$$

Galí and Monacelli (2005) explain that ϑ is a constant, which generally depends on initial conditions. Without loss of generality, the authors assume symmetric initial conditions, in which case $\vartheta_i = \vartheta = 1$, as stated above. This assumption yields

$$C_t^i = C_t^f \left(\frac{P_{c,t}^f}{P_{c,t}^i}\right)^{1/\sigma} = C_t^f \left(S_{f,t}^i\right)^{\frac{1-\alpha}{\sigma}}, \tag{H.9}$$

where the last equality comes from the definition of the CPI in Equation H.3. Integrating over f and taking logs yields

$$c_t^i = c_t^* + \frac{1-\alpha}{\sigma} s_t^i.$$

H.1.4 Firms' Technology

Each country has a continuum of firms represented by the interval $[0, 1]$. Each firm produces a differentiated good with a linear technology:

$$Y_t^i(j) = A_t^i N_t^i(j)$$

for all $i, j \in [0, 1]$, where A_t^i is a country-specific productivity shifter. The latter is assumed to follow an AR(1) process in logs: $a_t^i = \rho_a a_{t-1}^i + \varepsilon_t^i$, where $a_t^i \equiv \log A_t^i$, $\rho_a \in [0, 1]$ and ε_t^i is white noise.

Real marginal costs are common across firms in country i and given in logs by

$$mc_t^i = w_t^i - p_t^i - a_t^i. \tag{H.10}$$

Denoting the aggregate output index for country i as $Y_t^i = (\int_0^1 Y_t^i(j)^{\frac{\epsilon-1}{\epsilon}} \mathrm{d}j)^{\frac{\epsilon}{\epsilon-1}}$, I can derive the log-linearized relationship between aggregate output and aggregate employment in country i (see Galí and Monacelli (2005) for details)

$$y_t^i = a_t^i + n_t^i. \tag{H.11}$$

H.1.5 Firms' Price Setting

I assume a staggered price setting, as in Calvo (1983). Therefore, a measure $1 - \theta$ of randomly selected firms sets new prices each period, with an individual firm's probability of re-optimizing in any given period being independent of the time elapsed since it last reset its prices. Let $Q_t^{i,opt}(j) = P_t^{i,opt}(j)/P_t^i$ denote the optimal price set by firm producing goods j relative to the aggregate price level. The optimal price-setting strategy for a firm resetting its price in period t log-linearized around steady state is given by

$$\hat{q}_t^{i,opt}(j) = (1-\beta\theta)\,\mathrm{E}_t^i \sum_{\tau=0}^{\infty} (\beta\theta)^\tau (\widehat{mc}_{t+\tau}^i + \pi_{t+\tau}^i),$$

where $\hat{x}_t = x_t - x$ denotes log deviations from steady state for a generic variable X_t.

H.1.6 Aggregate Demand and Output Determination

The clearing of market for good j produced in country i requires

$$\begin{aligned}
Y_t^i(j) &= C_{i,t}^i(j) + \int_0^1 C_{i,t}^f(j)\,\mathrm{d}f && \text{sub Equ. H.2 \& H.4} \\
&= \left(\frac{P_t^i(j)}{P_t^i}\right)^{-\epsilon} \left[(1-\alpha)\left(\frac{P_{c,t}^i}{P_t^i}\right) C_t^i + \alpha \int_0^1 \left(\frac{P_{c,t}^f}{P_t^i}\right) C_t^f\,\mathrm{d}f\right] && \text{sub Equ. H.7} \\
&= \left(\frac{P_t^i(j)}{P_t^i}\right)^{-\epsilon} \left[(1-\alpha)\left(S_t^i\right)^\alpha C_t^i + \alpha \left(S_t^i\right)^\alpha \int_0^1 \left(\frac{P_{c,t}^f}{P_{c,t}^i}\right) C_t^f\,\mathrm{d}f\right] && \text{sub Equ. H.9} \\
&= \left(\frac{P_t^i(j)}{P_t^i}\right)^{-\epsilon} \left(S_t^i\right)^\alpha \left[(1-\alpha) C_t^i + \alpha \int_0^1 \left(\frac{P_{c,t}^f}{P_{c,t}^i}\right)^{\frac{\sigma-1}{\sigma}} C_t^i\,\mathrm{d}f\right] && \text{sub Equ. H.3} \\
&= \left(\frac{P_t^i(j)}{P_t^i}\right)^{-\epsilon} \left(S_t^i\right)^\alpha C_t^i \left[(1-\alpha) + \alpha \int_0^1 \left(S_{f,t}^i\right)^{\frac{(1-\alpha)(\sigma-1)}{\sigma}}\,\mathrm{d}f\right].
\end{aligned}$$

Plugging in the latter the definition of country i's aggregate output $Y_t^i = (\int_0^1 Y_t^i(j)^{\frac{\epsilon-1}{\epsilon}}\,\mathrm{d}j)^{\frac{\epsilon}{\epsilon-1}}$, I can obtain the aggregate goods market clearing condition for country i

$$Y_t^i = C_t^i \left(S_t^i\right)^\alpha \left[(1-\alpha) + \alpha \int_0^1 \left(S_{f,t}^i\right)^{\frac{(1-\alpha)(\sigma-1)}{\sigma}}\,\mathrm{d}f\right].$$

This last equation can be log-linearized around a symmetric steady state using the fact that a symmetric steady state $S^i = 1$ and therefore $s^i = 0$ as

$$\hat{y}_t^i = \hat{c}_t^i + \frac{\alpha\iota}{\sigma} s_t^i, \tag{H.12}$$

where $\iota = \sigma + (1-\alpha)(\sigma-1)$. Log-linearize the Euler Equation, Equation H.5, yields

$$\hat{c}_t^i = \mathrm{E}_t^i \hat{c}_{t+1}^i - \frac{1}{\sigma}(r_t - \rho_\beta - \mathrm{E}_t^i \pi_{c,t+1}^i),$$

where $\rho_\beta = -\log\beta$ is the time discount rate. Substituting the last equation into Equation H.12 yields

$$\hat{y}_t^i = \mathrm{E}_t^i \hat{y}_{t+1}^i - \frac{1}{\sigma}(r_t - \rho_\beta - \mathrm{E}_t^i \pi_{c,t+1}^i) - \frac{\alpha\iota}{\sigma} \mathrm{E}_t^i \Delta s_{t+1}^i,$$

which can be rewritten using Equation H.8

$$\hat{y}_t^i = \mathrm{E}_t^i \hat{y}_{t+1}^i - \frac{1}{\sigma}(r_t - \rho_\beta - \mathrm{E}_t^i \pi_{t+1}^i) + \frac{\alpha(1-\iota)}{\sigma} \mathrm{E}_t^i \Delta s_{t+1}^i. \tag{H.13}$$

H.1.7 Marginal Cost and Inflation Dynamics

Given the Calvo pricing scheme, the dynamics of *domestic* inflation in terms of real marginal cost in each individual country are described by the difference equation

$$\pi_t^i = \beta\, \mathrm{E}_t^i \pi_{t+1}^i + \lambda \widehat{mc}_t^i, \tag{H.14}$$

where $\lambda = \frac{(1-\beta\theta)(1-\theta)}{\theta}$. Log-linearizing Equation H.6, and using Equation H.10 and H.11 combined with the market clearing result in Equation H.12, real marginal costs can be expressed as a function of output, productivity and the effective terms of trade

$$\widehat{mc}_t^i = (\sigma + \zeta)\hat{y}_t^i - (1+\zeta)a_t^i + \alpha(1-\iota)s_t^i. \tag{H.15}$$

Log-linearization of Equation H.9 gives

$$\hat{c}_t^i = \hat{c}_t^f + \frac{1-\alpha}{\sigma} s_{f,t}^i$$

and integrating over f yields

$$\hat{c}_t^i = \hat{c}_t^{cu} + \frac{1-\alpha}{\sigma} s_t^i,$$

where $\hat{c}_t^{cu}$ is consumption at the currency union level. Integrating the market clearing results Equation H.12 over i yields $\hat{y}_t^{cu} = \hat{c}_t^{cu}$ as $\int_0^1 s_t^i \,\mathrm{d}i = 0$. Therefore, substituting Equation H.12 in the above equation gives

$$\hat{y}_t^i = \hat{y}_t^{cu} + \frac{1+\alpha(\iota-1)}{\sigma} s_t^i.$$

Combining this with Equation H.15 yields

$$\widehat{mc}_t^i = (\sigma + \zeta)\hat{y}_t^i - (1+\zeta)a_t^i + \alpha(1-\iota)\frac{\sigma}{1+\alpha(\iota-1)}(\hat{y}_t^i - \hat{y}_t^{cu}). \tag{H.16}$$

The above equation implies the following natural level of output

$$\hat{y}_t^{i,n} = \frac{1+\zeta}{\sigma+\zeta+\alpha(1-\iota)\frac{\sigma}{1+\alpha(\iota-1)}} a_t^i + \frac{\alpha(1-\iota)\frac{\sigma}{1+\alpha(\iota-1)}}{\sigma+\zeta+\alpha(1-\iota)\frac{\sigma}{1+\alpha(\iota-1)}} \hat{y}_t^{CU}.$$

Defining the output gap as $\tilde{y}_t^i = \hat{y}_t^i - \hat{y}_t^{i,n}$, I can rewrite Equation H.16 as

$$\widehat{mc}_t^i = \left(\frac{\sigma}{1+\alpha(\iota-1)} + \zeta\right) \tilde{y}_t^i.$$

Substituting this last equation in Equation H.14 gives

$$\pi_t^i = \beta \, \mathrm{E}_t^i \pi_{t+1}^i + \kappa \tilde{y}_t^i,$$

where $\kappa = \frac{(1-\beta\theta)(1-\theta)\left(\frac{\sigma}{1+\alpha(\iota-1)}+\zeta\right)}{\theta}$.

Adding a cost-push shock, ξ_t^i, yields Equation 8 in Section 6

$$\pi_t^i = \beta \, \mathrm{E}_t^i \pi_{t+1}^i + \kappa \tilde{y}_t^i + \xi_t^i. \tag{H.17}$$

Rewriting Equation H.13 in terms of the output gap yields

$$\tilde{y}_t^i = \mathrm{E}_t^i \tilde{y}_{t+1}^i - \frac{1}{\sigma}(r_t - \rho_\beta - \mathrm{E}_t^i \pi_{t+1}^i) + \gamma \, \mathrm{E}_t^i \Delta s_{t+1}^i + \nu_t^i, \tag{H.18}$$

where $\gamma = \frac{\alpha(1-\iota)}{\sigma}$ and $\nu_t^i = \frac{1+\zeta}{\sigma+\zeta+\alpha(1-\iota)\frac{\sigma}{1+\alpha(\iota-1)}} \mathrm{E}_t^i(a_{t+1}^i - a_t^i) + \frac{\alpha(1-\iota)\frac{\sigma}{1+\alpha(\iota-1)}}{\sigma+\zeta+\alpha(1-\iota)\frac{\sigma}{1+\alpha(\iota-1)}} \mathrm{E}_t^i(\hat{y}_{t+1}^{cu} - \hat{y}_t^{cu})$. As in Bertasiute et al. (2020), I assume that expectations of changes in productivity and currency union level output gap, ν_t^i, are white noise. The last equation is identical to Equation 7 in Section 6.

H.2 Expectation Formation

The expectation formation rules in Section 6 are microfounded as the posterior mean of a Gaussian signal-extraction problem, following Drenik and Perez (2020) and the framework developed in Section 5.

H.2.1 Inflation Expectations

By Equation H.17, expected inflation satisfies $\mathrm{E}_t^i \pi_{t+1}^i = \beta^{-1}(\pi_t^i - \kappa \tilde{y}_t^i - \xi_t^i)$. To keep the signal-extraction problem tractable, I approximate the resulting law of motion of inflation by the stationary AR(1) process

$$\pi_{t+1}^i = \rho_\pi \pi_t^i + \varepsilon_t^i, \qquad \varepsilon_t^i \sim \mathcal{N}(0, \sigma_{\varepsilon,\pi}^{2,i}), \tag{H.19}$$

with $\rho_\pi \in [0, 1)$. The AR(1) coefficient ρ_π and innovation variance $\sigma_{\varepsilon,\pi}^{2,i}$ are treated as reduced-form parameters estimated from the data. At the beginning of period t, the household observes three signals about π_{t+1}^i.

Prior. The household's prior is its lagged expectation, which under the AR(1) law of motion satisfies

$$\pi_t^{\text{prior},i} = \rho_\pi \, \mathrm{E}_{t-1}^i \pi_t^i. \tag{H.20}$$

This is treated as a noisy signal of π_{t+1}^i with noise variance $\tau_\pi^i > 0$, capturing uncertainty about the persistence of inflation and past forecast errors.

Media signals. The household receives V_t^i independent media reports, each of the form

$$\psi_{v,t}^i = \pi_{t+1}^i + \varepsilon_{v,t}^{\psi,i}, \qquad \varepsilon_{v,t}^{\psi,i} \overset{\text{iid}}{\sim} \mathcal{N}(0, \frac{1}{\tau_\psi^i}), \qquad v = 1, \ldots, V_t^i, \tag{H.21}$$

where $1/\tau_\psi^i$ is the noise variance of a single media report. This signal-extraction step isolates the precision of household beliefs and therefore takes the mean of coverage to be π_{t+1}^i. The direction shock $\varepsilon_t^{\psi,i}$, which shifts this mean, is introduced in the media-content equation of Section 6 and does not enter the precision derived here.[20] The average media signal $\bar{\psi}_t^i = \frac{1}{V_t^i} \sum_{v=1}^{V_t^i} \psi_{v,t}^i$ is therefore distributed as

$$\bar{\psi}_t^i \sim \mathcal{N}\left(\pi_{t+1}^i, \, \frac{1}{\tau_\psi^i V_t^i} \right),$$

so that a higher volume of media reports V_t^i reduces the effective noise on the media signal by averaging, consistent with the media frequency mechanism in Section 6.

Private signal. The household makes one private observation of future inflation. Its own reading of price conditions, which it cannot observe perfectly. The private signal is therefore

$$\bar{\phi}_t^i = \pi_{t+1}^i + \varepsilon_t^{\phi,i}, \qquad \varepsilon_t^{\phi,i} \sim \mathcal{N}(0, \frac{1}{\tau_\phi^i}),$$

where $1/\tau_\phi^i$ reflects the noise in the household's individual price observations, as in Section 5.

[20] The two disturbances are distinct: $\varepsilon_t^{\psi,i}$ is a common shock to the mean of coverage $\bar{\psi}_t^i$ that shifts every report in the same direction and corresponds to the media-direction innovation of Section 4, whereas $\varepsilon_{v,t}^{\psi,i}$ is idiosyncratic, mean-zero noise in each of the V_t^i reports whose effect on the mean signal shrinks with V_t^i. The calibration shock is applied to $\varepsilon_t^{\psi,i}$.

Posterior mean. All three signals are conditionally independent given π^i_{t+1}. Applying Bayes' rule under Gaussianity, the posterior distribution of π^i_{t+1} is normal with variance

$$\tau^i_{\text{post},t} = \tau^i_\pi + V^i_t \tau^i_\psi + \tau^i_\phi$$

and mean

$$\mathrm{E}^i_t \pi^i_{t+1} = \omega^i_{\pi,t}\, \pi^{\text{prior},i}_t + \omega^i_{\psi,t}\, \bar{\psi}^i_t + \omega^i_{\phi,t}\, \bar{\phi}^i_t, \tag{H.22}$$

where

$$\omega^i_{\pi,t} = \frac{\tau^i_\pi}{\Lambda^i_t}, \qquad \omega^i_{\psi,t} = \frac{V^i_t \tau^i_\psi}{\Lambda^i_t}, \qquad \omega^i_{\phi,t} = \frac{\tau^i_\phi}{\Lambda^i_t}, \qquad \Lambda^i_t \equiv \tau^i_\pi + V^i_t \tau^i_\psi + \tau^i_\phi,$$

and $\omega^i_{\pi,t} + \omega^i_{\psi,t} + \omega^i_{\phi,t} = 1$. Equation H.22 is identical to the inflation expectation rule in Section 6.

H.2.2 Output Gap Expectations

Output gap expectations are formed analogously. The output gap is approximated by the AR(1) process

$$\tilde{y}^i_{t+1} = \rho_y \tilde{y}^i_t + \varepsilon^{\tilde{y},i}_t, \qquad \varepsilon^{\tilde{y},i}_t \sim \mathcal{N}(0, \sigma^{2,i}_{\varepsilon,\tilde{y}}), \tag{H.23}$$

with $\rho_y \in [0,1)$. The household observes: a prior $\tilde{y}^{\text{prior},i}_t = \rho_y\, \mathrm{E}^i_{t-1}\tilde{y}^i_t$ with noise variance $1/\tau^i_{\tilde{y}}$; V^i_t media signals $\varphi^i_{v,t} = \tilde{y}^i_{t+1} + \varepsilon^{\varphi,i}_{v,t}$ with $\varepsilon^{\varphi,i}_{v,t} \sim \mathcal{N}(0, \frac{1}{\tau^i_\varphi})$; and a private observation $\bar{\varsigma}^i_t$ with noise variance $1/\tau^i_\varsigma$. Applying Bayes' rule as above yields

$$\mathrm{E}^i_t \tilde{y}^i_{t+1} = \omega^i_{\tilde{y},t}\, \tilde{y}^{\text{prior},i}_t + \omega^i_{\varphi,t}\, \bar{\varphi}^i_t + \omega^i_{\varsigma,t}\, \bar{\varsigma}^i_t, \tag{H.24}$$

where

$$\omega^i_{\tilde{y},t} = \frac{\tau^i_{\tilde{y}}}{\Gamma^i_t}, \qquad \omega^i_{\varphi,t} = \frac{V^i_t \tau^i_\varphi}{\Gamma^i_t}, \qquad \omega^i_{\varsigma,t} = \frac{\tau^i_\varsigma}{\Gamma^i_t},$$

with $\Gamma^i_t \equiv \tau^i_{\tilde{y}} + V^i_t \tau^i_\varphi + \tau^i_\varsigma$ and $\omega^i_{\tilde{y},t} + \omega^i_{\varphi,t} + \omega^i_{\varsigma,t} = 1$. Equation H.24 is identical to the output gap expectation rule in Section 6.

H.3 Monetary Policy

To define the objectives of the central bank, I derive the central bank loss functions in Section 6 from a second-order approximation to household welfare, following the approach in Woodford (2003) and Galí and Monacelli (2008). I will first discuss the baseline loss function and the alternative loss function with country heterogeneity.

H.3.1 Baseline Central Bank Loss Function

Start from the lifetime utility of the representative household in country i, Equation H.1. A second-order Taylor expansion of period utility around the zero-inflation, efficient steady state ($\bar{\pi} = 0$ and $\bar{y} = 0$) yields a period welfare loss proportional to

$$\mathcal{W}_t^i \approx -\frac{1-\alpha}{2}\left[\frac{\epsilon}{\lambda}(\pi_t^i)^2 + (\sigma + \zeta)(\tilde{y}_t^i)^2\right] + \text{t.i.p.}, \tag{H.25}$$

where $\lambda = \frac{(1-\beta\theta)(1-\theta)}{\theta}$ is the Calvo adjustment frequency parameter and t.i.p. denotes terms independent of policy. The first term captures the welfare cost of price dispersion induced by inflation: under Calvo pricing, cross-sectional price dispersion is proportional to $(\pi_t^i)^2$ to second order (see Woodford (2003), Chapter 6). The second term captures the cost of inefficient resource allocation when output deviates from its natural level.

Aggregating across countries using weights $w(i)$ and discounting gives the expected discounted welfare loss for the currency union

$$\mathcal{W}^{CU} = -\,\mathrm{E}_0 \sum_{t=0}^{\infty} \beta^t \frac{1}{\sum_i w(i)} \sum_{i=1}^{I} w(i)\, \mathcal{W}_t^i. \tag{H.26}$$

Minimising the negative of Equation H.26 subject to the equilibrium conditions is equivalent to minimising

$$\mathcal{L}^B = \lambda_\pi(\pi_t^{CU})^2 + \lambda_{\tilde{y}}(\tilde{y}_t^{CU})^2,$$

where $\lambda_\pi = \epsilon/\lambda$, $\lambda_{\tilde{y}} = \sigma + \zeta$, and deviations are taken from the inflation target $\bar{\pi}$ and steady-state output gap $\bar{y}$.

When the central bank also internalizes within-union heterogeneity, Equation H.25 implies an additional welfare cost from cross-country dispersion. Decompose country-level inflation as $\pi_t^i = \pi_t^{CU} + (\pi_t^i - \pi_t^{CU})$ and similarly for the output gap. Substituting into Equation H.26 and using the fact that the cross terms vanish under the currency-union average yields

$$\begin{aligned}\mathcal{L}^B = {}& \lambda_\pi(\pi_t^{CU})^2 + \lambda_{\tilde{y}}(\tilde{y}_t^{CU})^2 \\ & + \lambda_{\pi^i}\sum_{i=1}^{I}(\pi_t^i - \pi_t^{CU})^2 + \lambda_{\tilde{y}^i}\sum_{i=1}^{I}(\tilde{y}_t^i - \tilde{y}_t^{CU})^2\end{aligned} \tag{H.27}$$

where λ_{π^i} and $\lambda_{\tilde{y}^i}$ are the welfare weights on cross-country dispersion, inherited from Equation H.25. This is the baseline loss function $\mathcal{L}^B$ as Equation 13.

H.3.2 Alternative Central Bank Loss Function

As an additional tool, the central bank can also communicate via the media. However, communication is costly and therefore the alternative loss function is given by

$$\begin{aligned}\mathcal{L}^A = {} & \lambda_\pi(\pi_t^{CU})^2 + \lambda_{\tilde{y}}(\tilde{y}_t^{CU})^2 \\ & + \lambda_{\pi^i}\sum_{i=1}^{I}(\pi_t^i - \pi_t^{CU})^2 + \lambda_{\tilde{y}^i}\sum_{i=1}^{I}(\tilde{y}_t^i - \tilde{y}_t^{CU})^2 \\ & + \lambda_z\sum_{i=1}^{I}(z_t^i)^2 + \lambda_{\Delta z}\sum_{i=1}^{I}(z_t^i - z_{t-1}^i)^2. \end{aligned} \tag{H.28}$$

The last two terms penalize the level and the change of central bank communication z_t^i. These communication costs can be motivated as a reduced-form representation of reputational or operational costs of country-specific communication, analogous to the instrument smoothing term in standard Taylor rules (Woodford, 2003). Equation H.28 is identical to the alternative loss function $\mathcal{L}^A$, Equation 14.

H.4 Robustness

H.4.1 Robustness to Country Weights

The baseline calibration sets country weights to each country's share of private final consumption in the union: 0.37 for Germany, 0.15 for Spain, 0.27 for France, and 0.21 for Italy. These weights confound the results: a country with a larger share draws more of the central bank's attention mechanically, independently of its media transmission. To separate the two forces, I re-solve the model with equal weights ($w^1 = w^2 = w^3 = w^4 = 0.25$) and examine how optimal communication changes. These results are displayed in Figure 23. Under equal weights, Spain receives the strongest communication signal, reflecting its low reporting frequency combined with high media precision, and France receives an almost equally strong signal, reflecting both low frequency and low precision. Germany and Italy receive weaker signals. As in the baseline, the central bank communicates over a longer horizon where reporting frequency is lower. Stripping out country size therefore leaves media transmission as the force ordering optimal communication: the weakest-transmission countries receive the strongest and most persistent signals.

Figure 23: Impulse Response Functions from Currency-Union Supply Disturbance using Equal Country Weights

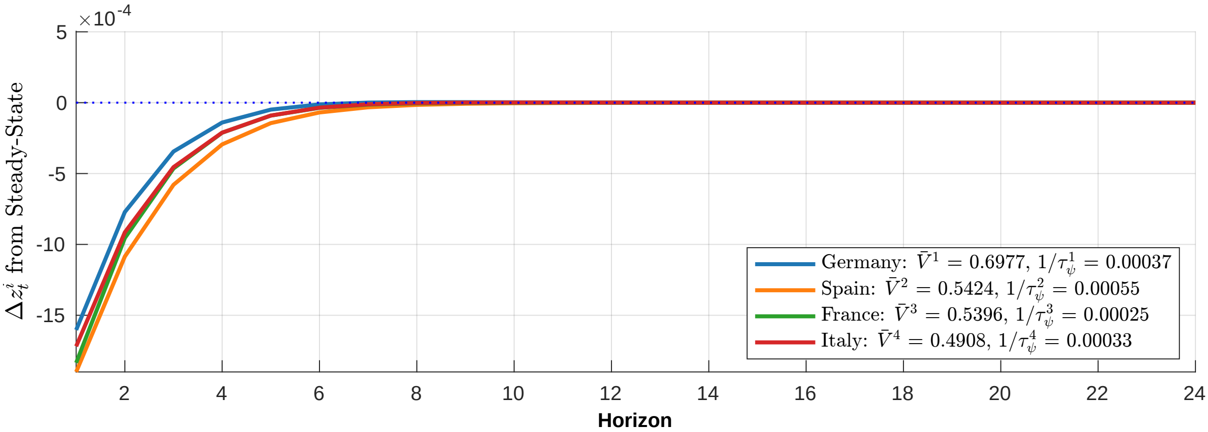

Note: The figure shows the impulse response functions of central bank communication, z_t^i, in response to a currency-union supply disturbance, ξ_t^i, for the countries Germany (blue line), Spain (yellow line), France (green line) and Italy (red line). Instead of country-specific weights, equal weights were used: $w^1 = w^2 = w^3 = w^4 = 0.25$.

H.4.2 Robustness to Welfare Gains

The welfare comparison in Section 6.6 relies on two parameters that are assigned rather than estimated: the pass-through of central-bank communication into media content, b_z, and the communication cost weights, $\lambda_z = \lambda_{\Delta z}$. Because the size of the welfare gain from communication depends on both, I assess how the comparison behaves across a range of plausible values. Table 8 reports the consumption-equivalent welfare gain of the communication regime $\mathcal{L}^A$ relative to the baseline $\mathcal{L}^B$ for a grid of pass-through and cost combinations. The gain is positive in every cell, so the welfare ranking favors communication throughout the grid: allowing the central bank to use country-specific communication reduces the welfare loss relative to the policy rate alone regardless of the particular values chosen for b_z and $\lambda_z = \lambda_{\Delta z}$. The magnitude, however, varies systematically along both dimensions. The gain increases in b_z, since stronger pass-through means a larger share of the central bank's signal reaches households through the media, making communication a more effective instrument. It decreases in λ_z, since higher costs discourage the central bank from using communication, shrinking the improvement it can deliver. At the extremes, the gain ranges from 0.007%—when pass-through is weak and communication is costly—to 1.14%—when pass-through is strong and communication is cheap. The baseline calibration ($b_z = 0.3, \lambda_z = \lambda_{\Delta z} = 10$) lies in the interior of this grid and yields a gain of 0.068%. The headline result therefore rests neither on an optimistic pass-through nor on a negligible communication cost, and the qualitative conclusion, that communication complements the policy rate and raises welfare, is robust across the entire range considered.

Table 8: Welfare Gains from Communication across Pass-Through and Cost Parameters

	Pass-through of communication, b_z				
$\lambda_z = \lambda_{\Delta z}$	0.1	0.3	0.5	0.7	0.9
1	0.074%	0.310%	0.507%	0.661%	0.780%
10	0.010%	0.068%	0.143%	0.219%	0.292%
30	0.003%	0.027%	0.064%	0.106%	0.151%
50	0.002%	0.017%	0.042%	0.073%	0.106%
70	0.001%	0.013%	0.031%	0.055%	0.082%

Notes: Each entry is the consumption-equivalent welfare gain (in percent) of the communication regime $\mathcal{L}^A$ relative to the baseline $\mathcal{L}^B$, for combinations of the pass-through b_z (columns) and the communication cost weights $\lambda_z = \lambda_{\Delta z}$ (rows). Positive values indicate that country-specific communication reduces the welfare loss relative to the policy rate alone.