The Messenger Matters*

Alena Wabitsch
University of Oxford

First Version: October 2024
This Version: February 2025 (Latest)

Abstract

Does the messenger matter in communication with the public? For policy communication on monetary, climate, fiscal, or other issues to be impactful, it must successfully reach people and influence their beliefs. I combine novel empirical evidence and theoretical methods to study how messenger characteristics, particularly nationality, influence central bank communication in the Euro area. First, I construct a multilingual dataset of over 8 million tweets, and document three novel stylized facts for ingroup audiences: higher reporting on policymakers, greater information availability, and stronger belief updating. Second, I design an inflation forecasting experiment, identifying causal evidence that ingroup messengers significantly increase the use of information. Third, I incorporate these effects into a stylized coordination model and demonstrate that delegating communication to multiple heterogeneous messengers maximizes welfare when public information quality is high, while centralizing communication is preferable when it is low. These findings identify the strategic selection of messengers as a powerful and novel policy tool, complementing traditional disclosure policies.

Keywords: Central Bank Communication, Inflation Expectations, Ingroup Effects, Coordination Games
JEL: E52, E58, D82, D83, D84

*Ethics approval from the University of Oxford, Reference: ECONCIA21-22-24. AEA RCT Registry ID: AEARCTR-0010727. Funded by: the Austrian Economic Association (NOeG), St Catherine's College (Oxford), and the Department of Economics (Oxford). For valuable comments I am grateful to Miguel Ballester, Tobias Cagala, Michael Ehrmann, Basil Halperin, Christian Hellwig, Alexandre Kohlhas, Seung Lee, Alistair Macaulay, Michael McMahon, Matthew Naylor, Lukas Nord, Luba Petersen, Ryan Rholes, Peter Rickards, Hannah Roemer, Séverine Toussaert, as well as for participants at the following seminars and conferences: *2024*: CEBRA Annual Meeting, EEA-ESEM, EABCN: New Challenges in Monetary Economics at the University of Mannheim, Esade Workshop in Behavioral and Experimental Economics, European Association of Young Economists (EAYE) Conference (Paris School of Economics), Annual MMF PhD Conference (University of Surrey), CEPR Central Bank Communication RPN Seminar Series, London School of Economics (Behavioural Finance Group), Oxford-CEPR Central Bank Communication Workshop, University of Oxford (Macroeconomics Workshop); *2023*: Austrian Economic Association (NOeG) Winter Workshop, University of Oxford (Behavioural Lab), the Annual Meeting of the NOeG, Sverige's Riksbank, Conference on Expectations in Dynamic Macroeconomic Models (TU Vienna). Email: alena.wabitsch@economics.ox.ac.uk

1 Introduction

Effective communication of a policy to the public is critical for its impact: people will only respond if they receive the information and believe it. This includes health campaigns promoting vaccine uptake during pandemics, subsidies and taxes incentivizing climate-friendly behavior, and monetary policy aimed at influencing consumption and savings decisions. How can policies—whether on climate, health, fiscal, monetary, and other issues—be communicated effectively to the public?

One important dimension that can be varied is the messenger. Communicating with the general public means that the audience is heterogeneous across various dimensions—nationality, gender, education, and more. Literature on social identity has identified "ingroup effects": Receivers respond more strongly to messages delivered by someone who shares characteristics with them (i.e., someone from an "ingroup"), as formalized for economic decision-making by Malmendier and Veldkamp (2022). Since policy communication occurs in institutional context, the identity of messengers has widely been neglected by policymakers and the policy communication literature. However, should ingroup effects also exist for policy communication, choosing diverse messengers could make this communication more effective.

In this paper, I combine novel empirical evidence and theoretical methods to address two questions. First, positively, how do messenger characteristics impact policy communication with the public? Second, normatively, how should policy be communicated to the public in light of these messenger effects? I explore this in the context of central bank communications, as central banks have made communication a core component of their monetary policy toolkit, recognizing its power (see Blinder et al. (2024) for a recent overview). Specifically, to investigate ingroup dynamics, I study how nationality—a salient yet understudied characteristic—influences central bank communication in the Euro area. First, I create a new multilingual dataset of tweets that provides novel evidence of ingroup effects along two dimensions: information availability and information processing (belief updating). Second, I design a unique inflation forecasting experiment, in which participants of four different nationalities forecast inflation subject to receiving signals from various messengers. The treatment is an isolated variation in messenger nationality, allowing me to precisely estimate the pure causal effect of messenger nationality on inflation expectations. With this design, I further confirm the efficacy of two concrete examples of delegating communication: to other board members of the same institution, or to other institutions. Finally, accounting for these reduced-form insights, I develop a stylized coordination model to derive optimal, welfare-maximizing communication policies. I find that strategic delegation of communication is a new and powerful policy tool that complements the well-studied disclosure policy. When

the quality of public information is high, delegating to heterogeneous messengers maximizes welfare; conversely, centralizing communication can be more socially desirable when public information quality is low.

Three novel stylized facts based on big data from Twitter motivate that the messenger's nationality matters in the monetary union of the Euro. First, the reporting on the messenger (here, the president of the European Central Bank (ECB)) rises strongly for the ingroup. Second, information availability—measured as the volume of tweets—is higher for the ingroup. And third, after the ECB's press conference, Twitter users update their beliefs more strongly when they are in the ingroup. These facts are based on over 8 million multilingual tweets posted between 2016 and 2022, which give meaningful insight into real-world phenomena such as modern information supply and responses to communication events.

To establish causal mechanisms, I design a pre-registered inflation forecasting experiment tailored to assess ingroup effects on inflation expectations, underlying mechanisms, and messenger-driven attention to information. In this experiment, participants from Germany, Spain, France, and Italy forecast inflation under six different scenarios. Each time, they receive a signal about future inflation from varying messengers, and then get the chance to update their forecasts. The treatment consists of isolated variation in the nationality of otherwise generic messengers. Using a standard Bayesian belief updating framework, I estimate the causal nationality-based ingroup effect on updating inflation expectations. The within-subject design allows me to further distinguish between two drivers of ingroup effects: homophily (preference for similar others) and heterophobia (aversion to different others). In addition, I provide causal reduced-form evidence for two concrete policy recommendations: delegating communication to diverse messengers within an institution, and delegating communication to other institutions.

Across all messenger treatments, participants under-use signals by approximately ten percentage points compared to a Bayesian benchmark. Signals from ingroup messengers are used 5.2 percentage points more than those from outgroup messengers—roughly half of the average information under-use. This constitutes the pure nationality-based ingroup effect on information processing related to inflation expectations. Delegating communication to ECB experts of the same nationality, or to national central banks (NCBs), successfully replicates these ingroup effects, albeit to a lesser extent. The positive effects on belief updating are primarily driven by trust, largely associated with reported perceived quality, and homophily rather than heterophobia.

After the experiment, participants face a survey, which provides more precise insights on information availability, complementing Twitter findings. Individuals indeed receive more news about their respective ingroup policymakers, confirming better reach beyond information

availability. However, while messengers influence real-world information availability, they do not causally affect attention to information in the experiment. Instead, attention responds causally to the level of inflation and its uncertainty, in line with attention to information being endogenous to the inflationary environment.

The evidence from both Twitter and the experiment shows that messenger effects positively impact the extent to which the ingroup receives and uses information. This makes central bank communication more effective for ingroup audiences. Reduced-form evidence further underscores benefits of delegating communication to board members of certain nationality or to the Eurosystem of NCBs. Leveraging these dynamics, what should optimal public communication policy look like?

To formally assess the normative predictions for optimal communication policy that maximizes social welfare, I develop a model in the spirit of a "beauty contest" (Morris and Shin, 2002). Agents are either in the in- or outgroup with the messenger, and they care about aligning their actions with an unknown fundamental as well as with all other agents' actions. While linked to empirical findings based on nationality, the model remains general, allowing findings to extend to public communication with heterogeneous messengers and receivers across various dimensions (such as gender, ethnicity, socio-economic background, or economic expertise), applicable to policy communication beyond monetary issues. Two communication policies are evaluated with the model: a disclosure policy and a delegation policy. The disclosure policy involves setting the precision of public information, while the delegation policy entails choosing the messenger(s) to deliver public information, thereby influencing the share of agents classified as ingroup.

I find that when the coordination motive is weak, increasing the share of ingroup agents through delegation is always beneficial—particularly so when public information is relatively precise. However, when coordination is strong and disclosing public information harms welfare because agents over-rely on imprecise public information and neglect their better private information, the presence of outgroup agents mitigates such welfare losses. This is because outgroup agents receive and use less public information. Additionally, as the quality of public information improves in these circumstances, disclosure raises welfare, while delegation may still be harmful. Thus, strategic delegation serves as a powerful additional policy tool to maximize welfare beyond disclosure.

The messenger matters. For communication to be an effective policy tool, the interplay between messengers' characteristics and heterogeneous receivers is important. By strategically delegating communication to messengers with diverse characteristics, policymakers can influence how messages are received and processed, and ultimately maximize social welfare through disclosing public information.

1.1 Related Literature

Central bank communication with the general public While communication with financial market participants has been an acknowledged powerful policy tool for a while now, targeting the non-expert public is a more recent endeavor of central banks. Although deemed a potentially powerful policy tool affecting institutional trust, inflation expectations, and economic volatility after monetary policy decisions, doubts remain whether central banks actually reach the public. Blinder et al. (2024) summarizes this literature comprehensively, building on work by e.g., Blinder et al. (2008); Ehrmann et al. (2013); Binder (2017); Blinder (2018); Haldane and McMahon (2018); Coibion et al. (2020). Although central bank communication successfully affects media coverage (Ter Ellen et al., 2022; Lamla and Vinogradov, 2019), and non-experts do respond to some communication events, they are generally not as responsive as experts (Ehrmann and Wabitsch, 2022). I emphasize the critical role of the messenger, offering new perspectives and actionable policy recommendations, supported by both positive and normative evidence. I offer new perspectives and policy recommendations to address these challenges, focusing on the role of the messenger and related ingroup effects. I add to this literature by evaluating the importance of the messenger, and by providing novel positive and normative evidence on concrete policy recommendations.

Ingroup effects and social identity Social identity influences various domains, including economics, education, healthcare, and finance. In the context of central banking, recent evidence suggests policymakers' gender and ethnicity affect their perception, trust, and influence (D'Acunto et al., 2022; Bodea et al., 2021; Bodea and Kerner, 2022).[1] Malmendier and Veldkamp (2022) formalize ingroup-dependent belief formation in economic decision-making more generally. I am the first to incorporate this into a model of strategic complementarities, after empirically estimating it in the context of central bank communication. My estimation of the nationality-related ingroup effect on belief updating represents novel causal evidence attributed to a single, isolated characteristic. With my unique experimental design, I can precisely identify underlying mechanisms, and provide novel evidence on concrete policy recommendations.

Experiments in Macroeconomics I design a new inflation forecasting experiment to causally infer the effects of messengers on inflation forecasts. It is the first forecasting

[1] In education, studies show that Black teachers positively impact Black students' educational outcomes, with stronger effects when gender also matches (Gershenson et al., 2022; Price, 2010; Carrell et al., 2010). In healthcare, Alsan et al. (2019) find that Black men are more likely to seek preventative care when interacting with Black male doctors. Similarly, in financial advising, Stolper and Walter (2019) find that homophily significantly increases the likelihood of clients following financial advice.

experiment that also causally measures attention to information through interactive design features, enabling novel causal inferences about how both messengers but also inflationary conditions impact attention paid to information. The experiment is also unique in its reflection of credible economic dynamics by using real historical inflation data in a fully controlled within-subject design, where each participant experiences all treatments. The within-subject design most closely relates to McMahon and Rholes (2023), who examine central bank credibility, and offers more flexibility than between-subject randomized control trials (RCTs): For instance, while the thematically closely related RCT by D'Acunto et al. (2022) makes important contributions in highlighting the significance of ingroup effects for central bank communication, my design allows for attributing ingroup effects causally to a single characteristic by varying only the nationality of generic experts in writing. It is also the first to causally estimate the influence of homophily and heterophobia in this context and provides novel causal evidence for policy recommendations. Presenting multiple inflation scenarios further limits the impact of the contemporaneous economic environment, which can affect effect sizes in information provision RCTs (Weber et al., 2025).

Optimal disclosure in coordination games The social value of public information in games with strategic complementarities, first explored by Morris and Shin (2002), has evolved significantly since.[2] I am the first to consider characteristics of the public signal's messenger, and to differentiate between heterogeneous agent types based on the alignment with these characteristics in such a model. The heterogeneity in agents determines agents' information availability and processing. Cornand and Heinemann (2008) discusses partial information availability but my model differs in introducing a belief updating bias and endogenizing signal availability to signal size, extending an idea by Nimark and Pitschner (2019) whereby larger signals are more likely reported. The belief updating bias that I introduce is a resonance weight to expectation formation (Malmendier and Veldkamp, 2022), complementing behavioral distortions such as myopia, or extra discounting of the future, explored by Angeletos and Huo (2021). My model set-up allows me to consider a new communication policy: the delegation of communication. This policy complements the widely explored disclosure policy (e.g., James and Lawler (2011); Angeletos and Lian (2018); Bassetto (2019); Kohlhas (2020, 2022)).

[2]They show that public information can disproportionately influence agents' actions, potentially leading to welfare-reducing overreactions to noisy public signals. Developments since Morris and Shin (2002)'s seminal work include Svensson (2006) arguing that increased precision of public information is likely beneficial for reasonable parameter values, and Amador and Weill (2010) examining a dynamic setting, highlighting varying welfare outcomes.

2 Three Stylized Facts

This section presents novel motivating observational evidence showing how the messenger's nationality matters in a monetary union like the Euro area (EA). I derive three stylized facts based on big data of over 8 million multilingual tweets posted between 2016 and 2022: First, the reporting on the messenger (i.e., the ECB president) rises strongly for the ingroup, indicating that the messenger is indeed noticed. Second, information availability—measured as the volume of tweets—is relatively higher for the ingroup. And third, after the ECB's press conference, Twitter users update their beliefs more strongly when they are in the ingroup.

2.1 Twitter Data

To capture real-world patterns of information availability and belief updating, I create a multilingual dataset of publicly available tweets posted between 1 November 2016 and 31 October 2022, using the Academic Twitter API. This time span covers three full years of Mario Draghi's presidency (2016-2019) until the switchover in ECB presidency on November 1, 2019, and equally three full years of Christine Lagarde's presidency (2019-2022). To cover the "Big Four" Euro area nations (Germany, France, Italy and Spain), tweets are either in German, French, Italian or Spanish. Tweet language is used to proxy individuals' nationality.[3] To capture general dynamics on Twitter, I also include tweets in English, the language of ECB communication, which unsurprisingly makes up for half of the tweet sample. However, because English language is a poor proxy for nationality, English tweets are excluded for analyses where nationality is key. Collected tweets contain the abbreviation "ecb" or the words "european central bank" or their respective equivalents in any of the mentioned languages.[4] This gives just over 8 million (8,031,937) tweets after thoroughly cleaning and translating the dataset. Table A5 provides an overview of the number of tweets by language in the final sample. Details of the cleaning steps and descriptive evidence, including some idiosyncrasies in tweet samples across different languages, are in Appendix A.

The Ingroup Individuals are considered to be in the ingroup if their nationality, proxied by tweet language, matches the nationality of the ECB's president. This means Italian are in

[3]Appendix A.2 demonstrates that tweet language is a valid proxy for nationality. The experiment in this paper will not rely on this assumption.

[4]These terms for selection follow Ehrmann and Wabitsch (2022), except that I exclude "draghi" and "lagarde" because both presidents have/had careers beyond the ECB presidency that are/were of public interest. To avoid picking up tweets about Draghi's political career in Italy post-ECB, or Lagarde's IMF presidency pre-ECB, I exclude tweets that mention these two presidents without referring to the ECB. Tweets are only collected if the keywords occur in the original part of the tweet. This includes tweets, retweets, replies or quote retweets.

the ingroup between November 2016 and October 2019 of the tweet sample, and French are in the ingroup between November 2019 and October 2022 respectively.[5]

Why Tweets? Compared to surveys, data from Twitter have several advantages: the sample size is much larger as it captures the full population of relevant tweets, allowing for the depiction of real-world phenomena. The high frequency of tweets, along with the ability to observe the same users multiple times, enables real-time analysis of reactions to central bank communication events and the study of associated belief updates. Unlike surveys, which capture expectations less frequently or only once, Twitter captures spontaneous expressions of opinion from a broad audience, reducing concerns about observer effects (e.g., Hawthorne or experimenter demand effects). While the frequency of tweets is comparable to market data typically used to study financial market reactions to central bank disclosures, Twitter uniquely allows for the study of the wider population rather than just market participants (Ehrmann and Wabitsch, 2022). However, Twitter data also presents some limitations: the evidence focuses on specific real policymakers with various attributes, is influenced by other trends and economic events, and uses imperfect control groups (the outgroup). Although these limitations prevent causal interpretation, the emerging real-world trends highlight key dimensions of messenger effects. In addition, all limitations of Twitter data are addressed with the experiment later in the paper, including answering whether higher tweet volume indicates better reach, whether tweet sentiment translates into inflation expectations, and what drives these findings.

2.2 Stylized Fact 1: Higher Reporting on Ingroup Policymakers

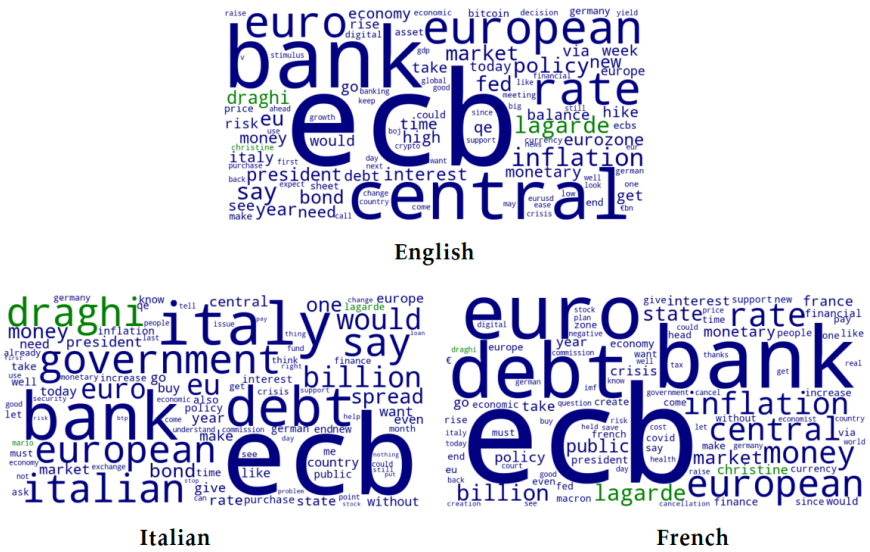
The first stylized fact that emerges from analyzing tweets is that the reporting on the main ECB policymakers varies strongly by language: ECB presidents are mentioned much more in tweets if they are in a language's ingroup than outgroup. This fact is important to establish that ECB presidents are indeed talked about, and their nationalities play a role for how information gets reported. Thereby supporting that any further identified differences between languages could indeed be due to being in- and outgroup with the ECB policymaker.

To highlight the different focus on policymakers, Figure 1 shows how frequently ECB presidents are talked about in the various language samples. Recall that tweet samples are only selected based on ECB-related key words, not based on ECB presidents. Notably, while "draghi" ranks among the top 10 most frequent words in the Italian sample, "lagarde" only ranks around the top 50. Similarly, in the French sample, "lagarde" ranks among the

[5]Note that the experiment in this paper will abstract from specific policymakers, defining the ingroup as participants who share the same nationality as the generic messengers.

top 10, while "draghi" is around the 60th most frequent word. In the English sample, the two policymakers are within 6 ranks, and similarly in the German and Spanish sample, the rankings of the two policymakers are much closer.

Figure 1: Reporting on Policymakers by Language



Notes: Word clouds show the 100 most frequent words in English (top), in Italian (bottom left) and French (bottom right) between 2016 and 2022. Word size indicates word frequency. "Christine", "Lagarde", "Mario" and "Draghi" are highlighted in green for better visibility. Word clouds are based on cleaned, translated and lemmatized tweets. Stopwords are removed. No ngrams are used. Word clouds in German and Spanish are found in Figure A10 in Appendix A.

Thus, while across languages ECB presidents are talked about in connection with the ECB, the reporting on these policymakers is heterogeneous, with a stronger focus on ingroup policymakers. This highlights that the nationality of ECB presidents alone might be enough to impact the availability and content of information, thereby impacting how effective central bank communication can be.

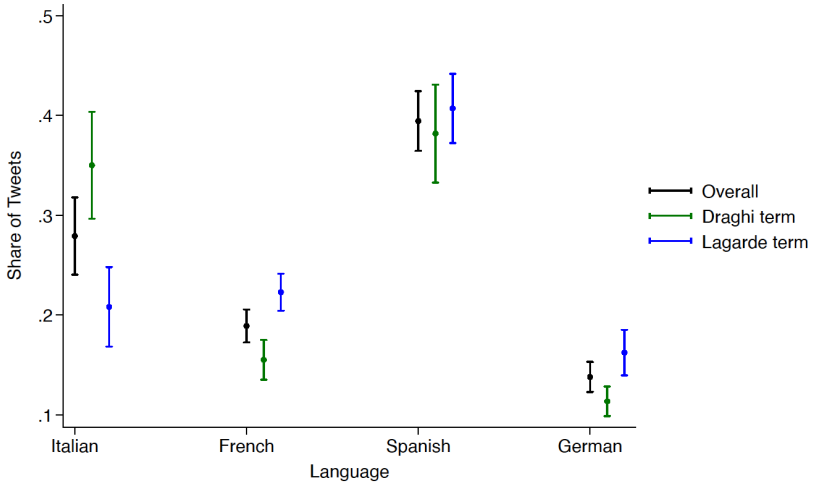
2.3 Stylized Fact 2: Higher Information Availability

The second stylized fact that emerges from the data is that information availability about the ECB increases for the ingroup, indicating a better chance of reaching a wider audience.

To show this, I analyse tweet volume by language and find that tweet volume about the ECB is relatively higher when a nationality is in the in- rather than the outgroup. Figure 2 plots the shares of tweets by language in six-week intervals corresponding to the ECB's press conference cycle. Shares are shown for the entire duration of the sample (in black) and split into the sample years covering Draghi's presidency (in green) and Lagarde's presidency

(in blue). Using shares of 6-week periods is beneficial to reduce concerns about time trends driving the results (e.g., avoiding a rise in Twitter volume over time because of user adoption being wrongfully associated with Lagarde-related tweet behavior), but also to avoid a few events driving the results. Overall, the Spanish tweet the most (well over a third of tweets), followed by the Italians (just under 30%). The French and Germans are much less active on Twitter for this topic. Comparing the differences in tweet shares between Draghi's term and Lagarde's term, Italians tweet substantially more during Draghi's presidency, while the opposite is true for the French: their share in tweets rises by around 50% during Lagarde's term. The Spanish and German samples serve as a type of control group.[6]

Figure 2: Share of Tweets by Language



Notes: The plot shows the share of tweets by language per 6-week cycle between ECB Press Conferences (with 95%-CI). The sample considered includes tweets about the ECB in German, Spanish, French and Italian between 2016 and 2022. The color of dots and error bars indicate the time horizon considered for the share of a language's tweets: black uses the entire sample, while green (blue) uses the sub-sample during Mario Draghi's (Christine Lagarde's) presidency.

A simple regression reveals a significant positive ingroup effect on information availability via Twitter: for the ingroup, information availability increases by 10.5 percentage points. These ingroup effects are positive across all levels of monetary expertise, with the strongest impact observed among users who are neither strict experts nor complete non-experts, where

[6]Indeed, the Spanish tweet volume does not change much between presidencies. However, German tweets does increase slightly during Lagarde's term. While this may be attributed to Lagarde, this might also be driven by Isabel Schnabel—who is very active on Twitter—joining the ECB Executive Board just two months after Lagarde. Further, this might even merely be a composition effect, assuming that (for whatever reason) German Twitter was growing relatively more than the average during this time. Appendix A discusses some of the most important observed trends by language. In sum, while the patterns strongly suggest that being in the ingroup increases information availability on Twitter, the presence of confounding factors prevents any causal interpretation, as often is the drawback of observational data.

the (non-)expert classification follows the benchmark in Ehrmann and Wabitsch (2022) (see Table A7).

I confirm the information availability patterns found on Twitter with traditional media, using the volume of printed newspaper articles about the ECB. The ECB president's nationality significantly increases the share of their corresponding national newspaper volume by 6.1 percentage points (or by roughly 272 articles). Appendix A.4 shows the full newspaper analysis.

Later in this paper, I show that information is not only more available, but individuals are also better reached by ingroup policymakers (see 3.5 and Appendix B.4).[7]

In sum, information availability about the ECB increases for the ingroup, making central bank communication more effective at reaching this group. Next, I will assess whether the received information is also more effective at shifting beliefs of the ingroup.

2.4 Stylized Fact 3: Stronger Belief Updating

The final key takeaway from Twitter is on information processing. Following central bank communication, the ingroup updates their beliefs more strongly. For this, I analyze the content of tweets using common natural language processing (NLP) methods.

The Signal The most prominent communication event of the ECB is its press conference, which takes place every six weeks.[8] It is particularly suitable to study, as it is the ECB's main communication event where any rate changes are announced, a statement on the economy is given, and questions from the press are taken. It is well-anticipated by observers and communication is carefully designed by the central bank. The final sample contains 48 press conferences, where the first took place on December 8, 2016 and the last on October 27, 2022.

The Beliefs I assess the sentiment of tweet content as a broad indicator of individuals' beliefs. There are various NLP approaches to estimating text sentiment, each with its own advantages and disadvantages. While the results are robust to more complex NLP methods, I show results based on a dictionary-based sentiment indicator due to its transparency and simplicity. The dictionary-based method provides a continuous estimate of how positive or negative a text snippet is.[9] The distribution of sentiment is highly similar across all languages,

[7]This finding is based on a survey about receiving news about central bank policymakers that is embedded in my experiment.

[8]The monetary policy announcement changed from 13:45 to 14:15 CET after July 21, 2022. Similarly, the starting time of the press conference shifted from 14.30 to 14:45 CET (see here).

[9]I use the polarity measure integrated in the Python package TextBlob, which is based on the Princeton University's WordNet lexicon (Loria, 2018).

precluding any idiosyncratic differences in tweet sentiment between them. This, alongside details of measuring sentiment, is shown in detail in Appendix A.

Relative Belief Updates To assess individual belief updates, I compare the sentiment of tweets before and after ECB press conferences. Belief updates are calculated as the absolute change in sentiment between a user i's last tweet during the ECB's quiet period—7 days before the conference—$\text{Sentiment}_{i,t-1}$ and the sentiment of the same user's first post-conference tweet posted within 24 hours of the conference $\text{Sentiment}_{i,t}$, relative to the average update around each press conference to account for differences in signals.[10] The absolute difference captures the magnitude of belief change, regardless of direction, and subtracting the average belief update of all individuals at conference j controls for variations in information content or timing, focusing on deviations from the common trend. This approach captures responses to the press conference and avoids picking up other ECB communications due to the quiet period's restricted communication. Figure A12 in the Appendix shows the distribution of sentiment before and after press conferences.

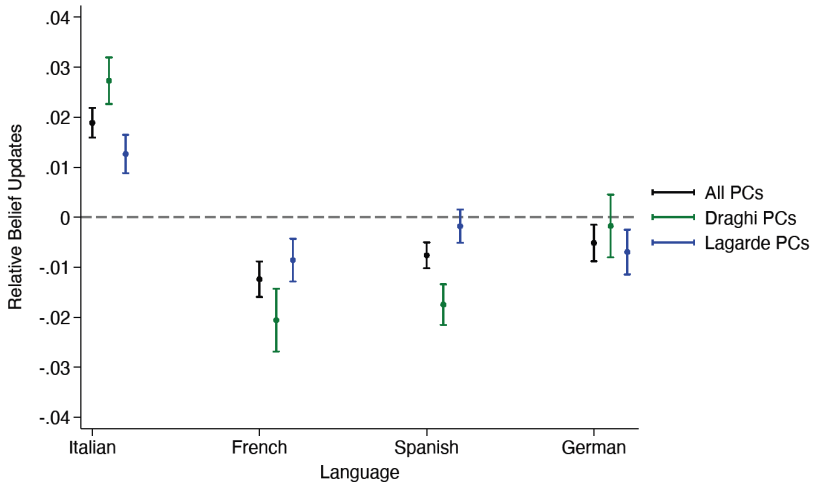
Figure 3 plots these relative belief updates by nationality for the entire duration of the sample (in black) and split into press conferences led by president Draghi (in green) and president Lagarde (in blue). The following conclusions can be drawn: First, Italians generally update more strongly than others. Second, Italians and French update their beliefs more strongly when they are in the ingroup—the Italians under Draghi and the French under Lagarde—rather than the outgroup. Again, the Spanish and German sample act as a control.[11] Based on these four nations, the overall estimated ingroup effect on belief updating estimated by a simple OLS is 0.014^{***}.

Finding that the ingroup exhibits larger absolute changes in sentiment suggests that they update their beliefs more strongly in response to central bank communication. But do these larger belief updates indicate that the ingroup shifts their beliefs more towards the information communicated by the central bank? While the presented descriptive evidence does not allow for a definitive conclusion, a detailed analysis accounting for the actual communicated signal in Appendix A.5 offers further insights: While all individuals' beliefs are responsive to novel information from the central bank, ingroup agents are more susceptible than outgroup agents. Ingroup agents rely less on their existing beliefs when updating beliefs and instead place more weight on new information from the central bank. They use new information over twice

[10] Formally : $\text{Relative Update}_{ij} = |\text{Sentiment}_{i,t} - \text{Sentiment}_{i,t-1}| - \frac{1}{N_j}\sum_{k=1}^{N_j}|\text{Sentiment}_{k,t} - \text{Sentiment}_{k,t-1}|$

[11] While the Germans show no difference in updating behavior between presidents, the Spanish are updating more strongly during Lagarde's presidency. Instead of this representing a Lagarde effect, however, this might instead be driven by the Spanish de Guindos. As the vice president of the ECB, he is also present at press conferences and his term time largely coincides with Lagarde's term.

Figure 3: Belief Updates by Tweet Language



Notes: Plot shows the absolute size of belief updates relative to the average during a given ECB Press Conference (with 95%-CI). The sample considered includes tweets about the ECB in German, Spanish, French and Italian between 2016 and 2022. The color of dots and error bars indicate the time horizon considered for the share of a language's tweets: black uses the entire sample, while green (blue) uses the sub-sample during Mario Draghi's (Christine Lagarde's) presidency.

as much as outgroup agents, while relying on existing beliefs about half as much (see Table A10). These ingroup effects persist regardless of individual monetary expertise, though they are slightly more pronounced among non-experts than experts as visualized by Figure A13. Overall, the heightened responsiveness implies that communication is more effective within the ingroup.

2.5 Discussion

The evidence I provide demonstrates that nationality of the messenger matters in a multi-national monetary union like the Euro area. When individuals are in the ingroup with the messenger, information is more available and beliefs respond more strongly. These ingroup effects are present for both experts and non-experts. Thus, it is not solely journalists reporting differently about the ECB that drives these results. While this evidence provides valuable insights into real-world phenomena, it is affected by the aforementioned data limitations and cannot inform on the underlying mechanisms. To address this, and to establish the causal effect of the messenger on inflation expectations, I design an experiment.

3 Inflation Forecasting Experiment

This section describes the individual-choice inflation forecasting experiment and its findings. In this pre-registered experiment, participants in Germany, Spain, France and Italy forecast inflation in six different inflation scenarios, being given information by six different messengers (the treatments). The experiment is explicitly designed to provide causal effects of how messenger characteristics affect belief updating about future inflation, the mechanism that drives differences in belief updating, and messenger-driven attention to information. The design allows me to identify the size of the nationality-based ingroup effect on belief updating, and whether such effects are rooted in homophily rather than heterophobia. Finally, findings inform on two potential policy recommendations: First, can the ECB communicate more effectively by strategically communicating using different board members more prominently? And second, in a similar vein, can beneficial ingroup effects be reaped by communicating more through the network of the Eurosystem's national central banks instead of centralized communication via the ECB?

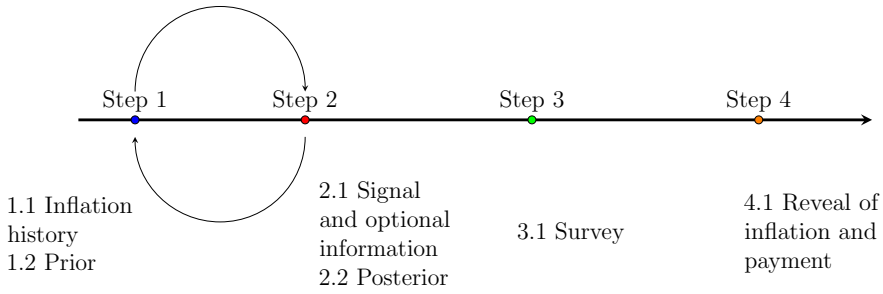
3.1 Experimental Design

Methodologically, this inflation forecasting experiment is individual-choice and follows primarily a within-subject design. It is incentivized, fully controls the environment in which participants forecast, and most closely relates to McMahon and Rholes (2023) who study credibility varying the forecast performance of central banks. In my experiment, two decisions that participants make are key: forecasting inflation and requesting additional information (or not). After participants complete all forecasting tasks, they are surveyed on their perceived ability of messengers, on their trust in institutions and policymakers, and on their news consumption of various policymakers and monetary policy. This allows me to draw conclusions on ingroup effects for inflation expectations, the driving mechanism, whether the ingroup is reached more and causally pays more attention to information. The timeline of the experiment is shown in Figure 4 and each aspect is explained in detail below.[12]

Inflation Forecasting Tasks At the core of this experiment are inflation forecasting tasks, which measure how participants update their inflation expectations following information provided by randomly varying messengers. Updates of inflation expectations are later evaluated through a standard Bayesian belief updating framework. What this means is that posterior beliefs should be determined by the value and the precision of both the prior and

[12]The experimental interface design facilitates task comprehension. It incorporates interactive tools and visuals, and is coded in oTree (Chen et al., 2016). See Appendix C for screenshots of the experimental interface.

Figure 4: Timeline of Experiment



Notes: The timeline goes from left to right. Steps 1-2 in repeat for all 6 forecasting tasks before going on to step 3 and 4.

the signal. Thus, it is important that the experimental design not only provides values of the signals and their precision, but also collects data on participants' priors and their precision.

Each forecasting task follows the same structure: First, participants observe 10 periods of inflation history before submitting their own inflation point and range forecasts for the next period (period 11).[13] This point forecast represents their prior, and the corresponding range forecast their inverted prior precision, which indicates participants' uncertainty associated with their prior. Participants will then receive a professional forecast for period 11 (i.e., the signal). This signal is provided as text but also appears in the interactive chart (see A26). The forecast is linked to randomly varying messengers, and comes together with the given messenger's forecast history of the previous 10 periods. This forecast history allows participants to assess the signal precision via the messenger's historical forecast error (comparing the historical forecasts to realized inflation).[14] While this requires the assumption that participants perceive historical forecast errors as the signal noise—or at least do not change their interpretation of provided signal noise between treatments—, this reflects how forecast quality and credibility is typically assessed in the real world. The historical forecast performance is visually displayed in a chart, with the exact values appearing when participants hover over data points. I refer to each combination of inflation and forecast values as a scenario. To preclude any differences in the six used scenarios from impacting effects, scenarios are randomized between messenger treatments. Finally, participants are given the chance to update their inflation forecast (point + range), i.e., they provide their posterior.[15]

[13] Range forecasts are participants' expected upper and lower bounds of inflation. Input values are restricted to ensure that participants' point forecasts are within their provided range. The experimental software interactively displays the forecasts and inflation history of a given sequence. For an example of this, see Figure A25 in Appendix C.

[14] 10 data points of forecasted and realized inflation allow participants to assess the precision of the signal fairly well. Signal precision is proxied by the inverse average forecasting error of the 10 periods. Table A11 gives an overview of the values used for signal precision.

[15] Figures A25 and A26 in Appendix C show the interface for these forecasting tasks. Instructions emphasize

Messenger Treatments The forecasting tasks are designed with ten messenger treatments—out of which each participant faces six—to precisely test four hypotheses (H1-H4). Table 1 gives an overview of all treatments and their alignment with the hypotheses.[16] Treatments are grouped into three categories: *Generic Expert Treatments*, *ECB Expert Treatments*, and *Institutional Expert Treatments*.

Table 1: Overview of Messenger Treatments

Treatment	Hypotheses	Messenger
1	H1	Expert from France
2	H1	Expert from Italy
3	H1	Expert from Germany
4	H1	Expert from Spain
5	H2, H3	Expert from France who represents the **ECB**
6	H2, H3	Expert from Italy who represents the **ECB**
7	H2, H3	Expert from Germany who represents the **ECB**
8	H2, H3	Expert from Spain represents the **ECB**
9	H3, H4	Expert who represents the European Central Bank (**ECB**)
10	H4	Expert who represents the National Central Bank (**NCB**)

Notes: The order of treatments is randomized at the participant-level. Additionally, the forecasting sequences are randomly allocated to treatments for each participant. In the experiment, font color is uniformly black and font colors displayed here simply aim to provide a better overview of the treatments to the reader. H1-H4 refer to the hypotheses that the treatments will be testing, which are explained in more detail in Subsection 3.4. For treatments 1–4, and 5–8 participants face only two out of the four treatments: the in- and one outgroup messenger (where the outgroup messengers' nationality is held constant within a participant). All participants face treatments 9 and 10.

1. *Generic Expert Treatments:* These quantify the pure ingroup effect by assessing how individuals use information from generic experts distinguished only by nationality. Participants encounter two messengers: one by an expert of their nationality (*generic ingroup expert*) and one by an expert of a randomly selected nationality out of the three other nationalities (*generic outgroup expert*).

2. *ECB Expert Treatments:* These add institutional context by introducing messengers representing the ECB with specified nationalities. As for the generic expert treatments, one ECB expert matches the participant's nationality (*ECB ingroup expert)*, while the other ECB expert does not (*ECB outgroup expert*). The latter messenger's nationality matches the nationality of the generic outgroup expert.

3. *Institutional Expert Treatments:* These involve generic experts without specified nationalities representing either the *ECB* or the participant's national central bank (*NCB*).

that participants may either use completely new values for their updated forecasts, some new values, or stick entirely to the values of their initial forecast, of which they are reminded (see Figures A22 and A23 in Appendix C for the full instructions).

[16] Section 3.4 discusses the hypotheses and findings and Appendix B provides more details on how treatments correspond to hypotheses.

To avoid fatigue, each participant faces only six out of the ten treatments—two from each group. Treatment order and inflation scenarios are randomized at the participant-level to ensure valid comparisons.

Inflation and Forecast Data Data for the forecasting tasks are based on actual annual inflation realizations in the Euro area between 1990Q1 and 2023Q1 (considered at quarterly frequency), and real-world quarterly forecasts of Euro area annual inflation provided by the ECB Survey of Professional Forecasters (SPF).[17] Using real data ensures that economic dynamics are credible. To emphasize this, participants are explicitly told that the dynamics of inflation and its forecasts always follow the real Euro area economy and that forecasts are real historical expert forecasts. The large amount of data points (97 quarters) make it unlikely that participants would recognize a historical pattern, and successfully predict the next period. Six sequences are randomly selected—one for each of the 6 forecasting tasks that participants complete. Appendix B shows the quarterly data series considered for the experiment (Figure A14) and the forecasted and realized inflation for selected sequences (Figure A15), as well as more details on the forecasting precision averages, realized inflation averages, and standard deviations of realized inflation for 6 randomly selected sequences (Table A11).

Incentivizing Forecasts Participants are incentivized to minimize forecast errors. Applying common incentive structures from the experimental literature ensures decisions are theory-consistent and of high quality (Charness et al., 2016). Similar to McMahon and Rholes (2023), I incentivize priors, posteriors, and their respective precisions, using methods from the learning-to-forecast literature (e.g., Rholes and Petersen (2021), Mokhtarzadeh and Petersen (2021)).

For *point forecasts*, participants receive a bonus payment based on their forecasting score, which decreases with the absolute forecast error. A perfect forecast yields the maximum score, and the score diminishes as the forecast error increases, reducing by a consistent proportion for each percentage point of error.

For *range forecasts*, participants earn a payoff only if the realized inflation falls within their specified forecast range. The payoff decreases as the width of the forecast range increases. This scoring rule incentivizes participants to keep their forecast ranges as narrow as possible while still covering inflation values they consider probable.

[17] The SPF is a quarterly survey of expectations for the rates of inflation and other variables. Respondents of the SPF are experts affiliated with financial or non-financial institutions based within the European Union. More information on the SPF can be found here.

More details on bonus payments, including how incentivizing bonus payments are calculated and randomly selected are shown in Appendix B.

Attention to Information In addition to the passively endowed forecast of a messenger (i.e., the signal), participants may actively choose more information by a given messenger. This provides a measure of attention to central bank communication. Participants are offered the opportunity to read more about the messenger's economic analyses underlying the given forecast. These analyses are summaries of historical ECB communication that were released roughly at the same time as the forecasts from the SPF.[18] This aims to mimic real central bank communication while ensuring that the numerical signal effectively summarizes the accompanying qualitative information. In the experiment, there are three "Read More" buttons, which participants have to click on to reveal the optional information. This measures attention by tracking participants revealing information piecewise, similar to Bartoš et al. (2016) who count clicks on "learn more" buttons to measure acquired information about job applicants' résumés. One can think of the pieces of information as a flow of information, for which participants decide when to "stop", leaving them with their demanded level of acquired information. An overview of all additional information pieces and their corresponding inflation scenario is given in Table A12 in Appendix B. A widely used alternative measure of attention is the amount of time spent on making a decision. However, this proxy of attention suffers from substantial drawbacks, such as being unable to differentiate causes of short decision times (e.g., fast thinking vs. low attention or unrelated interruptions).

Post-Experimental Survey After completing the forecasting tasks, participants fill out a survey aimed at understanding the mechanisms behind their decisions.[19]

Participants first rated the *perceived ability* of the treatment messengers to forecast inflation and provide economic analyses on a 7-point Likert scale. This measure identifies any prejudice or experience that might have influenced their perceived signal precision of the messengers. Accounting for these perceptions allows to distinguish between perceived signal quality, and remaining discrimination beyond perceived signal quality, which participants are not aware of and may be viewed as a bias.

To account for pre-existing perceptions of institutions and policymakers that could influence decision-making, participants report their *trust* and *exposure* (the extent to which they know of and follow news) regarding monetary institutions (the ECB and participants'

[18] This ensures the given additional information links to the corresponding forecast scenario.

[19] All survey questions are answered after the forecasting tasks to avoid priming but before revealing the realized values, preventing task performance from influencing responses about similar policymakers. Details of the survey as well as the entire survey are shown in Figures A28-A33 in Appendix C.

respective NCB) and representative real-world policymakers from messenger treatments. The representative policymakers are NCB governors and ECB board members of the same four nationalities at the time of the experiment, Table A13 provides an overview of chosen representative policymakers.

Finally, they were asked about their *monetary policy expertise* and were tested on their attention during the experiment.

3.2 Experimental Sample

The experiment is run in four different countries (the big four nations in the EA) to ensure that any identified effects that may be culturally driven are fairly representative of the EA.[20] While the random sample is not perfectly representative of the populations in Germany, Spain, Italy or France, it captures the populations fairly well regarding employment status, education and income. Although participants are slightly younger than the population average, the sample is much more diverse compared to common samples in lab experiments, as these commonly include students only. Appendix B.3 gives a detailed overview of participants in the sample. Participants' nationality is required to also coincide with the country of residence to preclude the inclusion of participants with mixed or potentially weaker cultural identification due to living abroad. In total, 400 participants (100 per nationality) do the experiment online via Prolific. One outlier is removed from the Spanish sample, leaving the final sample at 399.[21] A power analysis can be found in Appendix B. To ensure participants understand how the experiment and their incentives work, they must pass a comprehension quiz on the instructions to be allowed to proceed.[22] The vast majority of data is collected on October 31, with a few responses being added within 48 hours. On average, participants earned £3.58 in bonus payments and thus £7.58 overall.[23]

3.3 Estimating Belief Updating

Framework to Assess Belief Updating I use a standard Bayesian belief updating framework to assess how agents update their inflation expectations following central bank

[20] This is a between-subject element of the experiment.

[21] The outlier is dropped due to forecasted inflation values being consistently above 20, going up as high as 70, which indicates complete disregard of how forecasts are incentivized, and thus casts doubt as to whether instructions have been carefully read.

[22] Overall, 553 participants start the experiment, out of which 153 drop out because they either fail the quiz on instructions, voluntarily return the experiment before passing the quiz, or time out (only very few do so).

[23] Participants earn £4 for completing the experiment and may earn up to £6 in bonus payment, depending on the accuracy of their point and range forecasts.

communication.[24] The posterior belief is a weighted average of the prior and the signal, based on their relative precision (e.g., Veldkamp (2011)). Appendix B.2 provides more explanations.

Experimental evidence suggests that people are not perfectly Bayesian and tend to under-infer from signals.[25] Whether this is the case here, I will now test empirically.

Empirical Estimation This subsection shows the empirical estimates of how participants use their priors and signals. According to the theoretical framework, posterior beliefs are a weighted combination of prior beliefs and received signals, with weights determined by their relative precision.[26] I test this empirically using the following regression:

$$Posterior_i = \gamma \underbrace{\left(\frac{\alpha_i}{\alpha_i + \beta} A_i\right)}_{\text{weighted Prior}} + \delta \underbrace{\left(\frac{\beta}{\alpha_i + \beta} B\right)}_{\text{weighted Signal}} + \epsilon_i, \tag{1}$$

where the coefficients γ and δ indicate whether agents act Bayesian. If $\gamma = \delta = 1$, agents update their beliefs according to Bayes' rule. The parameter γ captures the influence of the prior: $\gamma > 1$ implies over-reliance on priors, while $\gamma < 1$ indicates under-use (base-rate neglect). Similarly, δ reflects responsiveness to new information: $\delta < 1$ means under-use of signals, and $\delta > 1$ implies over-use.

Estimating Eq. 1 by OLS without a constant yields $\gamma = 1.21$ and $\delta = 0.90$. This indicates that compared to a Bayesian benchmark, the signal is substantially under-used, regardless of nationality (see Figure 5 and Table A14 in the Appendix). Participants under-use the signal by about 10 percentage points on average, aligning with findings in existing literature.[27] This observed signal under-use suggests that increasing signal use is beneficial for closing the gap with a Bayesian benchmark. Can certain messengers achieve this?

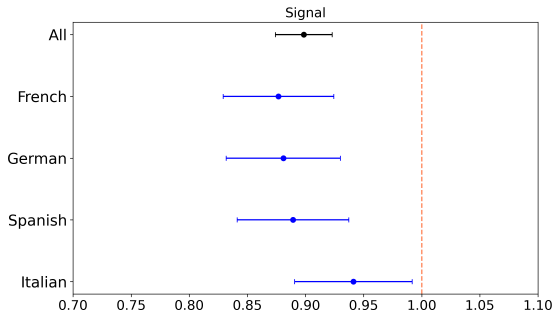
[24] Inflation data and professional forecasts may not perfectly align with the theory, which is reflected in the interpretation of results.

[25] Benjamin (2019) provides an overview of biases in belief updating. This under-inference may be due to factors like extreme-belief aversion (Benjamin et al. (2016)) or conservatism bias, where beliefs are updated too slowly, as discussed by Phillips and Edwards (1966).

[26] As outlined in Subsection B.2, the Bayesian benchmark assumes inflation is normally distributed and the signal is independently and identically distributed (i.i.d.), requiring forecast errors to be serially uncorrelated. However, historical quarterly EA inflation is non-normally distributed, and inflation forecasts show slight autocorrelation, invalidating the signal's i.i.d. assumption. Thus, I refer to a Bayesian *benchmark* that assumes i.i.d. signals and normally distributed inflation. Deviations from this benchmark do not necessarily imply non-Bayesian behavior; it serves as an imperfect reference point. Importantly, I focus on how participants' updates differ between messenger treatments, which are equally affected by these data limitations.

[27] For example, Benjamin (2019) find a δ coefficient of 0.86 in a meta-analysis of incentivized tasks with a single signal.

Figure 5: Signal Use Across Treatments by Nationality



Notes: The plot shows the regression coefficients of signal use with 95% confidence intervals. Results for the entire sample is shown in black and results fore samples split by nationality are shown in blue. The dashed vertical line at 1 indicates the Bayesian benchmark for signal use.

To test how different messengers affect signal use, I expand Eq. (1):

$$Posterior_{ij} = \gamma \left(\frac{\alpha_{ij}}{\alpha_{ij} + \beta_j} A_{ij} \right) + \sum_{j=1}^{J} \delta_j T_j \left(\frac{\beta_j}{\alpha_{ij} + \beta_j} B_j \right) + \epsilon_{ij}, \tag{2}$$

where T_j is a dummy variable that equals 1 for each messenger treatment $j \in J = \{1, 2, \ldots, 6\}$. For each treatment j, I estimate a separate δ_j coefficient to test hypotheses.

One caveat is that, due to the limited sample size, the randomized matching of messenger treatments with task order and inflation scenarios is imperfectly balanced.[28] To ensure that I identify variation in messenger treatments rather than inflation scenarios, I control for inflation scenarios and the order of messenger treatments in the analysis. While including these controls properly identifies the treatment effects, the Bayesian interpretation of the absolute size of coefficients is no longer valid. Therefore, I focus on the relative signal uptake between treatments to test my hypotheses.

3.4 Causal Findings on Updating Inflation Expectations

This subsection presents how different messengers impact belief updating, providing the pure ingroup effect and causal evidence for two concrete policy examples of delegating communication, and explores the mechanism behind the observed ingroup effects.

To obtain a clean causal nationality-based ingroup effect on updating inflation expectations that addresses a core research question of this paper, I test the first hypothesis:

[28] With a sufficiently large sample, randomization of inflation sequences and messenger treatment order (as described in Section 3.1) should balance out any effects. However, perfect balancing is not achieved here due to sample size limitations.

Hypothesis 1: *There exists a pure ingroup effect to updating inflation expectations, when signals come from experts of varying nationality.*

To test this, I take the *Generic Expert* treatments (1-4), and compare signal use between in- and outgroup experts. These experts are generic with the only information known about them being their nationality, which is made salient to participants (i.e., "an expert from Germany/France/Italy/Spain"). Except for the attached nationality, messaging experts are identical, allowing for a clean identification of a pure causal effect by avoiding having to account for any features related to a real person or institutions.

I find support for Hypothesis 1: National ingroup effects are significant in central bank communication, with ingroup members showing stronger belief updates. The size of this causal nationality-driven ingroup effect is 0.052^{***} (see column (3) in Table 2).[29] Meaning, while participants under-use signals of both in- and outgroup experts, they use the signal of the ingroup expert roughly 5 percentage points more, when forming their posterior beliefs about inflation. Putting this effect size in perspective, the ingroup effect is akin to closing half of the sample average gap to the Bayesian benchmark (as shown in subsection 3.3). Thus, this ingroup effect to updating inflation expectations is sizable and positively impacts the effectiveness of a signal.

Since central banks do not communicate without institutional context, it is relevant to consider how institutional reputation may affect policymakers' perception and credibility. Does the ingroup effect persist within the institutional context?

Hypothesis 2: *There exists an ingroup effect when signals come from ECB experts of varying nationality.*

This expands the pure ingroup effect from hypothesis 1 by adding the ECB as the institution of the expert. I compare signal use of messenger treatments 5-8 (ingroup vs. outgroup ECB experts).

I find support for Hypothesis 2, albeit somewhat weak: The ingroup effect diminishes to roughly a quarter of the average gap to the Bayesian benchmark (0.028^{*}) once the ECB is introduced as the institution of the expert (see Table 2). This effect is the causal ingroup effect within the ECB context but adding institutional dimension may introduce noise. Two sources of noise seem particularly likely: First, the institutional information may trigger associations beyond nationality that may cause other in- or outgroup effects, weakening the nationality

[29] This effect is based on the estimation that controls for inflation scenario, order of treatments and individual fixed effects.

Table 2: Main Effects from Experiment (H1, H2, H3, H4)

	(1)	(2)	(3)
Pure Ingroup Effect (H1):	0.064**	0.047***	0.052***
	(0.029)	(0.017)	(0.017)
R-squared	0.960	0.986	0.994
N	795	795	795
Ingroup Effect for ECB Experts (H2):	-0.014	0.022	0.028*
	(0.029)	(0.018)	(0.017)
R-squared	0.960	0.985	0.994
N	795	795	795
Homophily - Ingroup ECB Expert vs. Generic ECB Expert (H3):	0.006	0.034*	0.035**
	(0.029)	(0.018)	(0.018)
R-squared	0.961	0.985	0.993
N	794	794	794
Heterophobia - Outgroup ECB Expert vs. Generic ECB Expert (H3):	0.020	0.014	0.013
	(0.027)	(0.017)	(0.017)
R-squared	0.965	0.986	0.994
N	795	795	795
NCB vs ECB: Institutions Effect (H4):	0.027	0.031*	0.034**
	(0.028)	(0.017)	(0.017)
R-squared	0.965	0.987	0.994
N	795	795	795
Inflation Scenario		✓	✓
Individual-FE			✓

Notes: Effects result from comparing coefficients of interest of the same regression. *Inflation Scenario* refers to controlling for the underlying data of forecasting tasks and the order in which they appeared. *N* refers to the number of observations (forecasting tasks) in the regression. N may slightly vary across treatments due to instances of infinite prior precision. Column (3) represents the main results. Stars correspond to the following p-values: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

treatment (e.g., along the dimensions of level of education, field, perceived elitism, geographic location, etc.). Stolper and Walter (2019) show that various attributes can simultaneously cause ingroup effects. Second, even though the effect holds the institutional context constant across treatments, it is now plausible that the combination of nationality and ECB triggers associations with real policymakers. Participants might then distort the provided signal precision by what they know about the forecast ability of these associated policymakers, as suggested by Esponda et al. (2023). Besides messenger ability, also trust could then play a role, which is an important factor in the estimates as I will show in Subsection 3.4.1. Considering these caveats, my diminished results compare somewhat to D'Acunto et al. (2022) who do not find significant results for inflation expectations when varying messengers of

the same institutions. However, their study differs from my experiment in an important aspect: The use of various real policymakers as messengers, which means multiple messenger characteristics vary simultaneously between treatments. As my experimental design only varies a single characteristic at a time, any potential noise from other characteristics is avoided, which could further explain why I find significant effects for inflation expectations, while D'Acunto et al. (2022) do not.

An alternative interpretation of the diminished effect is that ECB as an institution carries enough weight to dwarf the relevance of its policymakers' nationalities. Given the ingroup's greater endogenous exposure to news from relevant policymakers (as shown by survey data in subsection 3.6), even a modest ingroup effect on information processing may result in significant real-world differences in belief updating, as indicated by the belief updating patterns on Twitter.

The positive messenger effects on information availability and processing within the ECB context should be reflected in French households' inflation expectations, given the French nationality of the current ECB president. I find corroborative correlations of this pattern using cross-country representative survey data from the ECB's Consumer Expectations Survey (CES) between the start of the survey in 2020 and up until 2024. Compared to other nations, the French reported inflation expectations that are most aligned with ECB forecasts (see Figure A20).[30] The detailed analysis can be found in Subsection B.5 in the Appendix.

What is the root cause of these ingroup effects within the ECB context? Differential signal use between two ECB experts may stem from a relative over-use of the ingroup expert's signal (homophily), a relative under-use of the outgroup expert's signal (heterophobia), or a combination of both. To understand which one it is, I now test:

Hypothesis 3: *Higher signal use for ingroup experts (homophily), rather than lower signal use for outgroup experts (heterophobia), drives the ingroup effect within the ECB context.*

I take advantage of having ECB expert messenger treatments with and without nationalities attached. Homophily compares participants' updates to the signal from ECB ingroup expert vs. the generic ECB expert, and heterophobia compares the signal use of ECB outgroup experts vs. generic ECB expert (treatments 5-8 vs. 9).

I find support for Hypothesis 3. Participants update significantly more towards the ingroup ECB expert than the generic ECB expert (0.035^{**}), while they do not update significantly less to the outgroup ECB expert compared to the generic ECB expert (0.013). Thereby, ECB

[30] ECB forecasts are taken from the ECB's Macroeconomic Projection Database (MPD), which is published four times a year and contains information on the outlook for the EA and contributes to the ECB Governing Council's assessment of economics developments and risks to price stability.

ingroup effects are driven by homophily, not heterophobia (see Table 2).

Instead of varying messengers within an institution, could varying institutions return similarly affect belief updating? If NCBs are seen as an "ingroup institution", while the ECB is seen as an "outgroup institution", the ingroup effect might translate to institutions (without attaching an explicit nationality to messengers). The final hypothesis tests:

Hypothesis 4: *There exists an ingroup effect when signals come from national central banks (NCBs) rather than the ECB.*

I find support for Hypothesis 4 by comparing the signal use between treatments 9 and 10 across all participants: Participants update more towards the NCB signal than the ECB signal (0.034^{**})—a third of the average gap to the Bayesian benchmark). Effects are smaller than the pure ingroup effect but larger than the ECB ingroup effect. A possible explanation for the former is that institutions are not strictly viewed as "in- vs. outgroup institutions"; after all, the ECB is the supranational institution. Another explanation is that institutional reputation plays a big role, which is likely not constant across institutions. Finally, adding institutions to messengers treatments potentially adds noise, as discussed above for Hypothesis 2. However, for the NCBs in the experiment, the found effect is causal and meaningful.

This positive effect on belief updating in response to NCB communication suggests that delegating communication to national banks within the Eurosystem could be more effective than centralizing it through the ECB. Yet, this might come at costs and risks such as the cacophony of voices (Blinder, 2007). It is vital that the central bank carefully assesses the number of speeches given, and ensures focused and cohesive messaging (Do Hwang et al., 2021; Tutino, 2016). Overall, the benefits of delegating communication to NCBs highlighted here should be viewed as an upper bound, assuming that the cacophony of voices is avoided by ensuring identical messaging.

3.4.1 Mechanism: Perceived Quality and Trust

How conscious are participants of their discrimination between messengers? Are ingroup effects fully reflecting subjective perceptions of messenger ability (i.e., signal quality), or trust in real-world policymakers and institutions? Indeed, much of the effects correlate with these factors.

Observed differences in information use may reflect perceived signal quality, beyond the provided signal precision in the experiment. Participants might perceive their ingroup messengers as more able (i.e., perceiving their signals as more precise) based on subjective experience, disregarding the provided forecast history. To understand the extent to which

participants' information use corresponds to the subjectively perceived information quality, I interact messenger-dependent signal use with the self-reported perceived ability of a messenger to forecast inflation and provide economic analyses. The findings are in column (3) of Table A15 in Appendix B.3. Approximately 80% of the pure ingroup effect can be attributed to reported perceived ability. Since perceived ability is reported ex-post, it may also reflect participants' justification for their signal use beyond their genuine perceptions of ability. Either way, the majority of the pure ingroup effect can be attributed to discrimination that participants are aware of, while a fifth of the original size (0.010^{**}) represents discrimination beyond perceived quality.[31] Similar reductions in effect sizes are seen across hypotheses. The homophily-driven positive effect from ECB experts becomes more significant (5% level instead of 10%) but drops to roughly a third of its original size. Similarly, the increased signal use of NCB experts compared to ECB experts falls to slightly less than a third, while remaining significant at the 5% level. Therefore, while reported perceived quality explains much of the increased signal uptake, a remainder of ingroup effects persists.

Trust could be another driver of signal use and is in fact important. Columns (4) and (5) show that for ECB policymakers (homophily effect), accounting for trust does not fully explain differences in belief updating, though it explains more than half of the effect.[32] In contrast, the increased signal use from the NCB expert can be entirely explained by trust; when interacting signal uptake coefficients with trust, there is no significant difference between the NCB and ECB experts. Thus, trust fully explains the differences in signal use between institutional experts.

3.5 Ingroup Policymakers Raise Information Reach But Not Attention

Beyond belief updating, the other dimension of messenger effects identified on Twitter is information availability. Does the increased volume of reporting translate to better reach? I find that it does: Participants in the experiment are better informed of ECB policymakers of their in- rather than their outgroup. Their probability of knowing these policymakers and receiving news about them increases by 27% and 29% respectively (See Appendix B.4 for details of the analysis). Thus, the ingroup is reached better.

Is the ingroup better informed because they pay more attention to information or simply

[31] This is a conservative estimate, imposing that participants fully incorporate their own perceived messenger ability beyond the provided signal precision.

[32] Note that the ingroup effect for ECB experts (H2) is insignificant when restricting the sample to only individuals where trust in representative policymakers is available. Hence, the comparison made here is based on the homophily effect (H3), where more observations are available.

by the higher availability? I causally evaluate paid attention within my experiment and find that attention to information, measured by clicking on "read more" buttons, remains unchanged across messenger treatments. Participants click on average on around 1.8 out of 3 buttons, independent of the messenger treatment (see Figure A18). Therefore, the increased informedness on ingroup policymakers seems to be due to higher information availability, not increased attention to information.

3.5.1 Attention Is Endogenous

While attention to information does not vary between messengers, it instead responds to the inflationary environment, confirming that the experimental feature used to measure attention to information is successful. The randomization of messengers with inflation scenarios renders inflation scenarios' inflation level and uncertainty exogenous. Therefore, causal effects of inflation and uncertainty on attention to information can be obtained. This is similar to estimating the extent of attention that is endogenous to the inflationary environment.

Higher inflation levels and greater inflation uncertainty significantly increase attention to information (see Table 3 or Figure A18). A one percentage point increase in average historical inflation raises the likelihood of participants requesting any piece of information by 32.6 percentage points and raises the clicked buttons by almost 1 (0.921^{***}).

Similarly, a one-unit increase in inflation uncertainty raises the probability of requesting any information by 24.3 percentage points and raises clicked buttons by 0.701^{***}. These findings align with Cavallo et al. (2017), who find that individuals have weaker priors about inflation in low-inflation contexts, and with Weber et al. (2025), who conclude that attention to inflation is endogenous to its level. Besides confirming these results on the level of inflation, I further show that attention to information is endogenous to the uncertainty of inflation.[33]

3.5.2 Attention to Information Moderates Effects

While variation in the messenger does not impact attention to central bank communication, attentive participants—who request additional information—exhibit stronger ingroup effects across hypotheses (see A16). Notably, these effects disappear when limiting the sample to forecasting tasks where no buttons are revealed (i.e., "no attention"). Conversely, the effects intensify as individuals pay more attention (i.e., reveal more buttons). This underscores the importance of participant compliance with information treatments in experiments, as

[33]Endogenous attention to information is rationalized by the costs of being mistaken about inflation increasing with the level of inflation in the real world. In my controlled setting, where participants are only incentivized to minimize forecasting errors, this is not the case: Expected forecasting errors should not depend on the level of inflation, only on the uncertainty.

Table 3: Endogenous Attention

Dep.:	OLS Attention (Continuous)	Probit Attention (Binary)	Marginal Effects Attention (Binary)
Inflation Level	0.921***	0.864***	0.326***
	(0.090)	(0.136)	(0.050)
Inflation Uncertainty	0.701***	0.644***	0.243***
	(0.077)	(0.114)	(0.042)
Signal Precision	0.179***	0.164***	0.062***
	(0.033)	(0.048)	(0.018)
Cons	-1.117**	-2.193***	
	(0.467)	(0.439)	
Messenger Treatments			
Outgroup Expert	0.059	0.047	
	(0.063)	(0.091)	
ECB Ingroup Expert	-0.053	-0.063	
	(0.063)	(0.091)	
ECB Outgroup Expert	-0.029	-0.042	
	(0.063)	(0.091)	
ECB Expert	-0.064	-0.050	
	(0.063)	(0.091)	
NCB Expert	-0.055	-0.040	
	(0.063)	(0.091)	
Treatment Order	✓	✓	✓
Individual-FE	✓		
R-squared	0.692		
Pseudo R-squared		0.023	
N	2,394	2,394	2,394

Notes: The dependent variable in column (1) is the number of additional information pieces requested (i.e., buttons clicked), in column (2) it is a binary variable, equaling 1 if at least some additional information piece is requested and 0 otherwise. Column (3) shows marginal effects of probit shown in column (2). Inflation level refers to the average inflation level of an inflation scenario, inflation uncertainty refers to the standard deviation of inflation of the 10 periods making up an inflation scenario. Reference category for messenger treatments is the Generic Ingroup Expert treatment. Stars correspond to the following p-values: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

highlighted by Knotek et al. (2024). Details are in Appendix B.3.

3.6 Discussion of Experimental Findings and Policy Implications

The experimental findings demonstrate positive ingroup effects on information processing, which somewhat diminish within the ECB context. This might raise questions about how these results align with real-world evidence from Twitter and what they imply for policymaking.

Policymakers embody a multitude of characteristics beyond nationality, making real-world evidence difficult to square with the clean results from the experimental setting. A moderating factor for the pronounced results on belief updating seen on Twitter is certainly the increased availability of signals (i.e., news supply), as shown in Subsections 2.3 and 3.5. Ingroup policymakers are significantly better known, but also significantly followed more in the news

than outgroup policymakers.

In sum, sharing nationalities with the communicating policymaker positively impacts the likelihood of receiving information and the extent to which novel information is used in belief updating. Overall, this makes central bank communication more effective for ingroup agents. So how can policymaking benefit from this? Reduced-form experimental evidence highlights that if central bank communication wants to target certain groups of the population, it can be beneficial to delegate communication to selected board members, due to their higher perceived technical abilities. Similarly, because of higher trust, communicating through the Eurosystem of National Central Banks improves the effectiveness of central bank communication. Even though attention to information does not causally increase for the ingroup, they are more likely to receive news about ingroup policymakers. To evaluate optimal communication policy, I develop a stylized coordination game that incorporates messenger effects on information availability and processing. This model provides normative recommendations for communication strategies in contexts with heterogeneous messengers and receivers.

4 Modeling Optimal Communication

Building on the empirical findings, this section develops a model with strategic complementarity to evaluate optimal public communication that maximizes social welfare. Extending Morris and Shin (2002), a "beauty contest" with private and public signals, the model incorporates the messenger's role in public communication. Based on shared characteristics with the messenger, agents are classified as either in- or outgroup agents, which affects their availability and processing of the public signal. The social value of public communication is modeled as a "beauty contest", reflecting the dual impact of public information disclosure: Agents respond to public signals due to both the information provided about a fundamental and their expectations of others' reactions, creating strategic complementarity as they align their actions with both the fundamental and others' choices (coordination motive).[34]

Two communication policies are evaluated: a disclosure policy, affecting the precision of public signals, and a novel delegation policy, which involves choosing messengers to influence the composition of in- and outgroup agents, while keeping the single public signal constant.

Interestingly, I find that delegating communication to harness ingroup effects is not always socially desirable. Increasing the share of ingroup agents raises social welfare, when public information is precise and coordination motives are moderate. By contrast, when private

[34]Examples for this coordination motive are firms setting prices considering their competitors' prices, or agents negotiating wages similar to others.

information is better and coordination is strong, the presence of outgroup agents is indeed desirable as their availability and processing of public information mitigates potential welfare losses. Thus, adjusting the composition of ingroup and outgroup agents through delegation should be used strategically. Thereby, delegation serves as a powerful additional policy tool to maximize welfare (conditional on messenger expertise), complementing disclosure policies.

The model is kept general and is applicable to any public communication with heterogeneous messengers and receivers, including monetary policy as well as other policies, such as climate, health, fiscal or education. While nationality within the Euro area is chosen as an application to build intuition, the model extends to any dimension of heterogeneity, such as gender, ethnicity, socio-economic background, or economic expertise.

4.1 The Model Set-Up

A continuum of agents $i \in [0,1]$ receive private and public signals and aim to align their actions with both an unknown fundamental and the actions of others (the coordination motive inspired by "beauty contest"). My key innovation is distinguishing between two *agent types* $h \in \{g, o\}$—ingroup (g) and outgroup (o)—based on their alignment with the messenger's characteristics ($\theta_i = \theta_m$ for ingroup, $\theta_i \neq \theta_m$ for outgroup). Reflecting my empirical findings, these types differ in the availability and processing of public information.

4.1.1 Information Structure

The fundamental $x \sim \mathcal{N}(\mu, \tau_x^{-1})$ represents an unobserved variable, such as inflation. Each agent $i \in [0,1]$ predicts x based on private and public signals.

$$\text{Private signal:} \quad y_i = x + \epsilon_{iy}, \quad \epsilon_{iy} \sim \mathcal{N}(0, \tau_y^{-1}) \tag{3}$$

$$\text{Public signal:} \quad Y = x + \epsilon_Y = x + \epsilon_z + \epsilon_V, \quad \epsilon_z \sim \mathcal{N}(0, \tau_z^{-1}), \quad \epsilon_V \sim \mathcal{N}(0, \tau_V^{-1}) \tag{4}$$

The central bank's own information about the fundamental ($z = x + \epsilon_z$) has a fixed noise component ϵ_z. As part of its *Disclosure Policy*, the central bank decides how much of this information to disclose by controlling the additional noise term ϵ_V. This affects the overall precision of the public signal τ_Y beyond τ_z. Full disclosure corresponds to $\tau_V \to \infty$, while complete opacity is $\tau_V \to 0$, and partial disclosure is any interim case, $\tau_V \in \mathbb{R}^+$. Accounting for both disclosure and noise in the public information, the observed public signal can be written as $Y \sim \mathcal{N}(x, \tau_Y^{-1})$, where $\tau_Y^{-1} = \tau_V^{-1} + \tau_z^{-1}$.[35]

[35]Note that ϵ_V and ϵ_z are independent of each other, x, and of ϵ_{iy}.

The second decision the central bank makes is on its *Delegation Policy*. By choosing the messenger(s) of communication, the central bank sets the share of ingroup agents in the economy, denoted α. Delegating communication assumes a single public signal Y and, unlike the disclosure policy, does not aim at changing the precision τ_Y but influences whether different agents receive and how they process the public signal.

Public Signal Availability: The public signal is not always available to all agent types—it is not strictly common knowledge. This deviates from coordination games in the literature, except for Cornand and Heinemann (2008) who restrict availability of public signals, but who do not let this depend on agent type and signal size.

Ingroup agents always receive the public signal and are therefore always informed. Outgroup agents, however, may or may not receive the public signal depending on whether its magnitude $|Y|$ exceeds their individual threshold d_o. Each outgroup agent $o \in [0,1]$ has a threshold d_o drawn from a truncated normal distribution on the interval $[0,\infty)$, denoted $d_o \sim \mathcal{N}_+(0,1)$ (see Appendix D.1 for details). An outgroup agent observes and acts on Y only if $|Y| \geq d_o$; otherwise, the agent relies solely on their private signal. This captures the idea that agents obtain public signals indirectly, for example, through media outlets that report only sufficiently "newsworthy" signals (Nimark and Pitschner, 2019). Larger signals are more likely to be reported, increasing the likelihood that outgroup agents receive them.

The probability that an outgroup agent observes a public signal Y depends on its realized value ($\mathbb{P}(|Y| \geq d_o)$). Consequently, the fraction of informed outgroup agents is:

$$A = 2\Phi(|Y|) - 1, \tag{5}$$

where Φ is the standard normal cumulative distribution function. The fraction of outgroup agents who do not observe the signal is $1 - A$. As $|Y|$ becomes very large, A approaches 1, meaning nearly all outgroup agents observe large signals, whether positive or negative.

Belief Updating: All agents have a diffuse common prior that x is normally distributed with mean zero and zero precision. Agents follow Bayesian belief updating for private signals, which they receive regardless of the size of the signal. For ingroup agents, this is also the case for the public signal. However, informed outgroup agents incorporate *Information Resonance* in their belief updating (Malmendier and Veldkamp, 2022). The overall weight applied to a public signal is called the *Relevance Weight* ω_{im} and is defined as:

$$\omega_{im} = \rho_{im}\tau_Y = \left(2 - 2\Phi(\chi \,||\, \theta_i, \theta_m \,||)\right)\tau_Y, \tag{6}$$

where ρ_{im} is the *Resonance Weight*, which defines the resonance that a signal coming from m has for receiver i. $\|\theta_i, \theta_m\|$ is the Euclidean distance between the agent's characteristics θ_i and the messenger's characteristics θ_m, and χ captures the sensitivity to mismatching characteristics.[36] The normal cumulative distribution function Φ equals 0.5 at 0, and hence the expression $2 - 2\Phi(\cdot)$ takes on a value of 1 at zero and then declines quickly towards zero.

Scaling the resonance weight ρ_{im} with the (true) accuracy of the signal τ_Y reflects the idea that agents will not weight signals that have no information content at all.

The information sets are summarized as:

$$\Omega_i = \begin{cases} \{y_i, Y\} & \text{if } h = g, \\ \{y_i, Y\} & \text{if } h = o \text{ and } |Y| \geq d_o, \\ \{y_i\} & \text{if } h = o \text{ and } |Y| < d_o. \end{cases} \tag{7}$$

In addition to the signals, all agents know the share of ingroup agents α, and the fraction of informed outgroup agents A. Each informed agent believes that all agents—both ingroup and outgroup—process any received public signals in the same way they do.

Agents' expected values of the fundamental x are:

$$\text{Ingroup Agents:} \quad E_{ig}[x \mid y_i, Y] = \frac{\tau_y y_i + \tau_Y Y}{\tau_y + \tau_Y}, \tag{8}$$

$$\text{Outgroup Agents:} \quad E_{io}[x \mid y_i, Y] = \frac{\tau_y y_i + \rho_{im}\tau_Y Y}{\tau_y + \rho_{im}\tau_Y} \quad \text{if } |Y| \geq d_o, \tag{9}$$

$$E_{io}[x \mid y_i] = y_i \quad \text{if } |Y| < d_o. \tag{10}$$

4.1.2 Payoffs

Agents $i \in [0, 1]$ choose action $a_i \in \mathbb{R}$ to maximize their payoff $u_i \in \mathbb{R}$. Such an action could reflect decisions about consumption, savings or investments. Their payoff depends on their own action, strategic complementarities, and the state of the exogenous fundamental $x \in \mathbb{R}$:

$$u_i = -(1 - r)(a_i - x)^2 - r(a_i - \bar{a})^2, \tag{11}$$

where $\bar{a} = \int_0^1 a_i \, di$ is the average action. The term $-(a_i - x)^2$ represents the quadratic loss between the agent's action a_i and the exogenous fundamental state x. It reflects the agent's desire to align their action with the fundamental state. $-(a_i - \bar{a})^2$ captures the deviation of agent i's action from the actions of all other agents a_j. This represents the "beauty contest"

[36] If $\chi = 0$: agents learns from everyone without discounting anyone's information $\rho_{im} = 1$. If χ is very large: agents disregard any information from messengers that do not exactly match their characteristics $\rho_{im} < 1$. $0 < \chi < 1$ represents cases in-between.

term: agents care about coordinating their actions with the average action of others.[37] r gives the weight on the coordinating, second-guessing motive.

4.1.3 Social Welfare

Social welfare is defined as the (normalized) average of individual utilities:

$$W(a,x) \equiv \frac{1}{1-r}\int_0^1 u_i(a,x)\,di = -\int_0^1 (a_i - x)^2\,di. \tag{12}$$

The social planner cares only about agents' distance to the fundamental x since the coordination motive introduces a social inefficiency, where agents' desire to align with others distorts individual actions without improving outcomes based on the true fundamental state.

4.1.4 Agents' Actions

While the social planner only cares about keeping all agents' actions close to the state x, the agent's action is determined by the information available to her and the first-order condition:

$$a_{ih} = (1-r)E[x|\Omega_i] + rE[\bar{a}|\Omega_i], \tag{13}$$

Neither the central bank nor the agents observe other agents' actions. Instead, agents form expectations about the average action, based on their information sets and requiring the determination of equilibrium signal extraction weights. To start with, the expectation of others' private signals is their expected state of the economy. For ingroup agents this is:

$$E_{ig}[y_j \mid y_i, Y] = E_{ig}[x \mid y_i, Y] = \frac{\tau_y y_i + \tau_Y Y}{\tau_y + \tau_Y}, \tag{14}$$

where $j \neq i$, hence y_j represents other agents' private signals. For informed outgroup agents this is:

$$E_{io}[y_j \mid y_i, Y] = E_{io}[x \mid y_i, Y] = \frac{\tau_y y_i + \rho_{im}\tau_Y Y}{\tau_y + \rho_{im}\tau_Y}. \tag{15}$$

For uninformed outgroup agents this is simply their private signal, which coincides with their expected value of the fundamental:

[37] The former term is akin to an agent's aim to forecast inflation correctly so as to avoid over- or under-consumption. The latter term of aligning forecasts with others can be viewed as to avoid disruptions or mismatches in wages, contracts, and prices, even if the forecast is not perfect. This coordination reduces uncertainty in economic interactions like wage-setting and price adjustments, ensuring smoother participation in the economy.

$$E_{io}[y_j \mid y_i] = E_{io}[x \mid y_i] = y_i. \tag{16}$$

4.1.5 Timeline

The game consists of two stages. First, the central bank communicates, having decided on their policy to disclose (τ_Y) and their policy to delegate (α) in order to maximize expected welfare. Second, agents receive their signals and choose their actions a_i to maximize expected utility based on their information sets.

4.2 Equilibrium

An equilibrium of the game consists of strategies for the central bank and the agents such that no player has an incentive to deviate.

4.2.1 Derivation of Equilibrium Actions

Suppose that all agents of type h follow a linear strategy:

$$a_{ih} = \kappa_h y_i + (1 - \kappa_h)Y. \tag{17}$$

Ingroup and Informed Outgroup Agents All informed agents of type h choose the same optimal weight κ_h in equilibrium. The optimal weight depends on an agent's expectations about the behavior of other agents. The shares of the population in the ingroup (α), outgroup ($1 - \alpha$), and informed outgroup (A) are common knowledge. However, agents assume all informed agents process information identically, unaware of any biases in others' or their own updating.[38]

Lemma 1. *The optimal signal extraction weights κ_h for ingroup and informed outgroup agents are given by:*

$$\kappa_g^* = \frac{\tau_y q}{\tau_Y + \tau_y q} \quad \text{and} \quad \kappa_o^* = \frac{\tau_y q}{\rho_{im}\tau_Y + \tau_y q}, \tag{18}$$

where $q = (1 - r + r(1 - \alpha)(1 - A))$. Proof in Appendix D.2.1.

Equilibrium signal extraction weights κ_o^* and κ_g^* differ only by how the public signal's precision is being perceived (i.e., ingroup agents use τ_Y instead of $\rho_{im}\tau_Y$). The equilibrium

[38] Assuming a lack of awareness of biases aligns with the inherent definition of bias and is consistent with its usage in the literature (e.g., Angeletos and Huo (2021); Gust and Lopez-Salido (2024).

actions of informed agents thus are:

$$a_{ig}(y_i, Y) = \frac{\tau_Y Y + \tau_y q y_i}{\tau_Y + \tau_y q} \quad \text{and} \quad a_{io}(y_i, Y) = \frac{\rho_{im}\tau_Y Y + \tau_y q y_i}{\rho_{im}\tau_Y + \tau_y q}. \tag{19}$$

Uninformed Outgroup Agents An uninformed outgroup agent can solely rely on their private signal:

$$a_{io}(y_i) = y_i. \tag{20}$$

Average Action The average action in equilibrium is:

$$\begin{aligned} \bar{a} &= \alpha a_{ig} + (1-\alpha)A a_{io} + (1-\alpha)(1-A)a_{io}(y_i) \\ &= x\left[\alpha\kappa_g^* + (1-\alpha)A\kappa_o^* + (1-\alpha)(1-A)\right] + Y\left[\alpha(1-\kappa_g^*) + (1-\alpha)A(1-\kappa_o^*)\right]. \end{aligned} \tag{21}$$

Note that my model nests Morris and Shin (2002). This is the case in two special cases:

1. $\alpha = 1$: All agents are ingroup agents, leading to homogeneous agents who update without bias.

2. As $|Y| \to \infty$ and $\chi = 0$: All outgroup agents receive the signal ($A \to 1$) and update without bias ($\rho_{im} = 1$), eliminating differences between agent types.

4.2.2 Uniqueness of Equilibrium

Proposition 1. *The equilibrium strategies derived constitute the unique linear equilibrium of the game. Proof in Appendix D.2.2.*

4.3 Welfare Implications and Policy Analysis

Having derived the equilibrium, I now assess the welfare implications of the central bank's communication policies.

4.3.1 Social Welfare Computation

Expected social welfare, accounting for ingroup and outgroup agents, is:

$$\begin{aligned}
E[W(a,x)|x] &= -E\left[\int_0^1 (a_i - x)^2\, di\right] \\
&= -\alpha \int_0^1 E\left[(a_{ig}(y_i, Y) - x)^2\right] di - (1-\alpha)A \int_0^1 E\left[(a_{io}(y_i, Y) - x)^2\right] di \\
&\quad - (1-\alpha)(1-A) \int_0^1 E\left[(a_{io}(y_i) - x)^2\right] di \\
&= -\alpha \left[\kappa_g^2 E({\epsilon_{iy}}^2) + (1-\kappa_g)^2 E(\epsilon_Y^2)\right] \\
&\quad - (1-\alpha)A \left[\kappa_o^2 E({\epsilon_{iy}}^2) + (1-\kappa_o)^2 E(\epsilon_Y^2)\right] - (1-\alpha)(1-A) E({\epsilon_{iy}}^2) \\
&= -\alpha \frac{\tau_Y + \tau_y q^2}{\left[\tau_Y + \tau_y q\right]^2} - (1-\alpha)A \frac{\rho_{im}^2 \tau_Y + \tau_y q^2}{\left[\rho_{im}\tau_Y + \tau_y q\right]^2} - \frac{(1-\alpha)(1-A)}{\tau_y}.
\end{aligned} \tag{22}$$

How does the precision of information affect expected equilibrium social welfare? Conducting comparative statics on Eq. (22) highlights the implications of raising the precision of public signals—akin to the disclosure policy—for welfare and policy design.[39]

4.3.2 The Precision of Public Information—The Disclosure Policy

Recall that the central bank can fully disclose its information, thereby increasing the precision τ_Y of the public signal. The upper bound of Y's precision, achieved under full disclosure ($\tau_V \to \infty$), is τ_z, which represents the precision in the central bank's information itself. Conversely, under complete opacity, $\tau_Y = 0$ (since $\tau_V \to 0$), meaning that Y is entirely noise and carries no meaningful information about the fundamental.

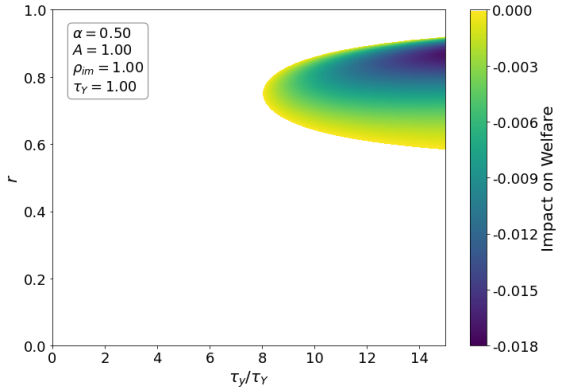
Proposition 2. *Increasing the precision of the public signal (τ_Y) improves social welfare only if the public signal is sufficiently precise relative to private signals and if the coordination motive r is not too high. Derivation in Appendix D.2.4.*

The canonical finding by Morris and Shin (2002) is that welfare is not necessarily increasing in the precision of the public signal—only if τ_Y is sufficiently large, especially compared to τ_y. This is because disclosing noisy public information causes agents with a desire to coordinate to neglect their private signals, which diminishes welfare when private information is better. Since my model nests Morris and Shin (2002), I use this as a benchmark in Figure 6a. Here, disclosure harms welfare when the coordination motive is sufficiently high ($r > 0.5$), and

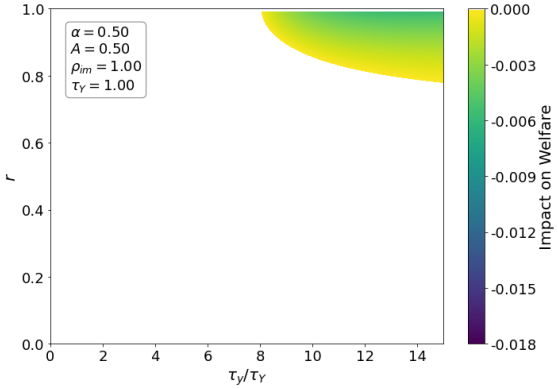
[39] For completeness, the impact on welfare of raising the precision of private signals is discussed in Appendixppendix D.2.3.

private signals are relatively precise compared to the public signal. The other two panels in Figure 6 highlight what happens to welfare if messenger effects are introduced one by one.

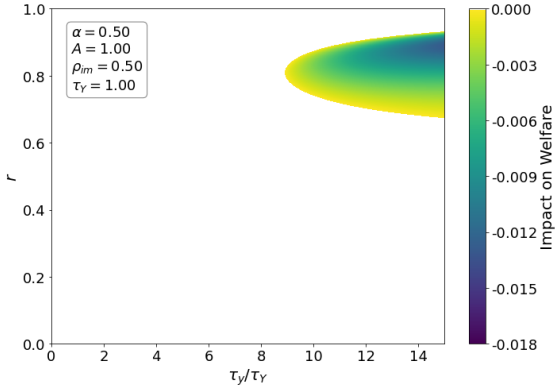
Figure 6: Disclosure's Effect on Social Welfare



(a) Benchmark



(b) Restricted Signal Availability



(c) Updating Bias

Notes: The plots visualize Eq. (41) by plotting the impact of increasing public signal precision (disclosure) on social welfare for different value combinations of parameters r and $\frac{\tau_y}{\tau_Y} = \tau_y$ since τ_Y is set to 1. White areas indicate strictly non-negative impact on welfare. Any other color indicates negative marginal effects on welfare. Panel (a) serves as a the benchmark of all agents receiving the public signal and updating without a bias (equivalent to only ingroup agents as in Morris and Shin (2002)). Panel (b) shows the case where outgroup agents do not receive all signals but update without bias. Panel (c) shows the case where outgroup agents receive all signals but update with a bias.

First, the outgroup's restricted signal availability reduces the parameter combinations for which disclosure harms welfare. Higher coordination is now required for disclosure to lower welfare (see Figure 6b). This is because not all public signals are universally observed. Instead, signal availability depends on signal size, where the outgroup is less likely to observe smaller signals (lower A). This mitigates welfare loss from disclosure by limiting the over-reliance on relatively imprecise public information.

Second, the biased updating of outgroup agents further eliminates parameter combinations

that render disclosure welfare-harming, as illustrated in Figure 6c. This occurs because outgroup agents downweight the public signal, which mitigates the over-reliance on imprecise public information in high coordination environments.

Thus, outgroup agents mitigate potential welfare losses from disclosure. Their presence acts as a buffer against scenarios where disclosure of public information could have detrimental welfare effects. Considering this, whenever the central bank has noisy information, it could delegate communication to optimally set the level of in- an outgroup agents (adjusting α), in tandem with choosing its optimal disclosure policy (τ_V). This raises the question of how the central bank should optimally balance delegation and disclosure policies to maximize welfare.

Corollary 1. *The presence of outgroup agents mitigates potential welfare losses from increased disclosure when public information is noisy.*

4.3.3 The Composition of In- and Outgroup Agents—The Delegation Policy

The central bank may decide to delegate communication of the single public signal from the default messenger (typically the president) to one or more messengers of varying characteristics. Delegation controls the share of ingroup agents (α); either by selecting multiple messengers communicating the same signal, or by selecting another messenger whose characteristics resonate with a larger share of the economy. Reduced-form evidence of this paper suggests the positive ingroup effect on belief updating carry over to two concrete examples of delegation policy: through other board members with varying characteristics, or through governors of national institutions.

For the central bank to decide on the policy tool of delegating communication, it needs to determine the welfare-maximizing share of ingroup agents, $0 \leq \alpha^* \leq 1$. This optimal share of ingroup agents depends on the coordination motive r, the availability of the public signal to the outgroup A (thus, the size of the signal $|Y|$), the size of outgroup agents' updating bias ρ_{im} and the precisions of signals τ_y and τ_Y.

Proposition 3. *An optimal share of ingroup agents (α^*) exists and depends on r, signal precision, signal size, and bias in information processing. Derivation in Appendix D.2.5.*

An explicit analytical solution for α^* is not feasible. Instead, I use numerical solutions to illustrate how α^* varies with parameters and discuss special cases.

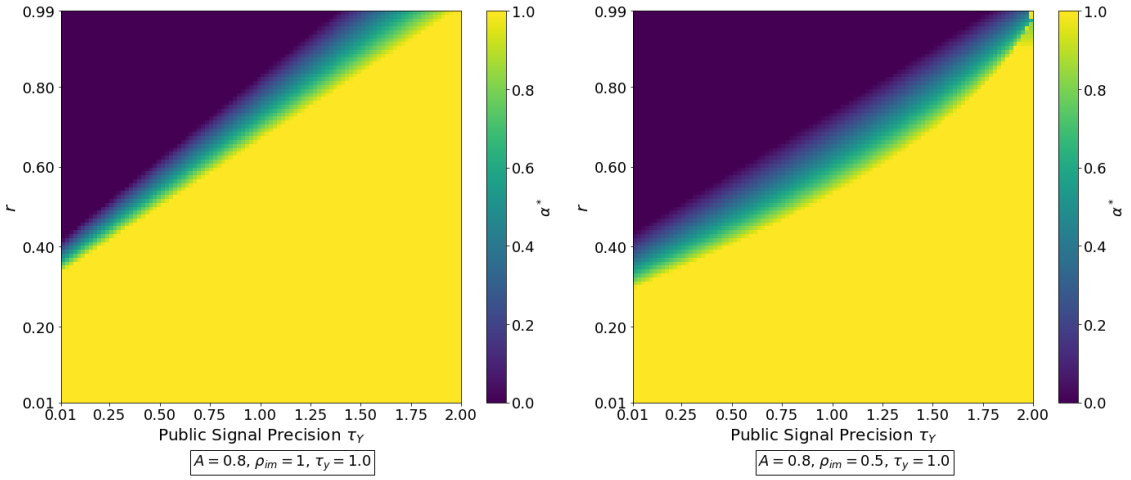
For the special case, where the outgroup faces restricted signal availability but updates without bias, $\alpha^* = \min\left\{1, \frac{\tau_Y+\tau_y}{3r\tau_y}\right\}$. Meaning, the optimal share of ingroup agents α^* is less than one so long as the private signal is sufficiently precise compared to the public signal

and there is sufficiently strong strategic complementarity (formally, $\frac{\tau_Y}{\tau_y} = 3r - 1$). $r < \frac{1}{3}$ always results in the corner solution of all agents optimally being ingroup agents ($\alpha^* = 1$). This result is in line with the optimal publicity result of Cornand and Heinemann (2008). Intuitively, over-reaction to a relatively noisy public signal is alleviated by outgroup agents not receiving smaller signals. The presence of both ingroup and outgroup agents creates a dynamic where the negative impact of ingroup agents' overreaction is outweighed by the positive effect of enhanced coordination on x.

Introducing an updating bias complicates the expression for α^*. Figure 7 shows α^* numerically, illustrating optimal delegation across different levels of coordination and public signal precision. The left column (Panels 7a and 7c) represents cases of outgroup agents updating without bias but not receiving all signals. The right column (Panels 7b and 7d) adds biased updating ($\rho_{im} = 0.5$). Top and bottom rows compare large signals (received by 80% of outgroup agents) to moderate signals (received by only 20%). Regardless of signal size, the more precise public information is relative to private information, the higher the optimal share of ingroup agents. Intuitively, this is because if public information is very precise, it is beneficial if agents receive this information, and fully incorporate it in their expectations. By contrast, when public information is relatively noisy, it is beneficial to have fewer ingroup agents, so α^* drops. This is especially pronounced for large signals, where the outgroup's bias to belief updating play a substantial role, lowering the optimal share of ingroup agents for more parameter combinations than in the absence of bias. Meaning, given the bias, fewer ingroup agents may be socially desirable, even in cases of weaker strategic complementarities and noisier public signals than if outgroup agents updated without bias. When signals are moderate, the optimal share of ingroup agents drops to a lesser extent. In this case, the belief updating bias makes little difference to the optimal share of ingroup agents.

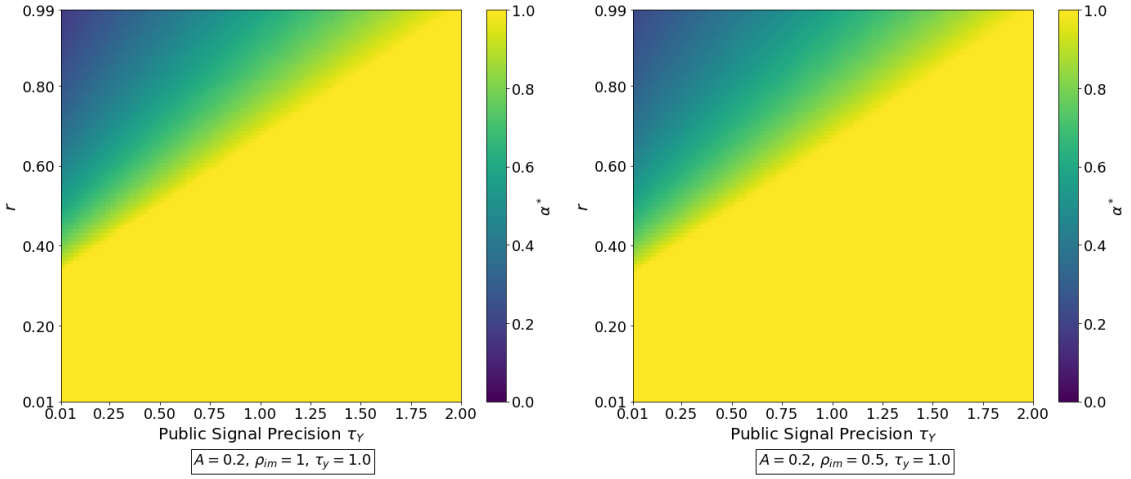
A Potential Caveat of Delegating Communication: Loss in Expertise Delegating communication away from a single president, who is presumably be the best-suited candidate to communicate, might come at the cost of public signal precision. Such precision losses arise if the messenger conveys less expertise than the original messenger. This potential trade-off must be accounted for when deciding on whether to delegate communication. Appendix D.3 discusses this formally.

Figure 7: Optimal Share of Ingroup Agents



(a) Large Signal - No Bias

(b) Large Signal - Updating Bias



(c) Moderate Signal - No Bias

(d) Moderate Signal - Updating Bias

Notes: Plots show the welfare-maximizing level of ingroup agents α^* for different combinations of coordination (r) and public signal precision (τ_Y). The top row plots show α^* for large signals (corresponding to $A = 0.8$), and bottom row plots show the same for moderate signals ($A = 0.2$). Severe updating bias corresponds to $\rho_{im} = 0.5$. Since $\tau_y = 1$, the public signal is more precise than the private signal whenever $\tau_Y > 1$ and noisier for $\tau_Y < 1$.

4.4 Policy Implications: When to Delegate?

In economies with heterogeneous agents and messengers, both disclosure and delegation policies contribute to social welfare. How should a central bank choose between these policies, or a mix of them? I evaluate disclosure (raising τ_V) against delegation (raising α) and illustrate the social welfare contours for environments of high and low coordination (Figure 8). The share of informed outgroup agents is set to 50%, and their updating bias is set to

$\rho_{im} = 0.95$, reflecting the pure ingroup effect this paper finds empirically. To showcase the role of coordination, r is set to take up rather extreme values, which are $r = 0.90$ and 0.10, respectively.

1. *When coordination is low and public signals are precise: Full disclosure and maximizing the share of ingroup agents ($\alpha^* = 1$) is optimal.*

   When public information is relatively noisier than private information ($\tau_y > \tau_Y$), delegating communication is always optimal ($\alpha^* = 1$), and so is full disclosure ($\tau_Y^* \to \infty$). Therefore, the central bank should delegate communication to maximize the share of ingroup agents, and disclose all its information (see Figure A36a). When public information is noisier than private signals, welfare improvements of raising α are rather small. Larger increases in the share of ingroup agents are required for notable welfare improvements.

   When private information is noisier ($\tau_y < \tau_Y$), delegating communication is optimal, as the corner solution of $\alpha^* = 1$ still applies. When the public signal is very precise, even small increases in ingroup agents improve welfare substantially.

2. *When coordination is high and public signals are noisy: Limiting the share of ingroup agents can be beneficial to prevent over-reliance on imprecise public information.*

   When public information is relatively noisier than private information ($\tau_y > \tau_Y$), the relative precision of the public signal must be sufficiently high for welfare to increase with the share of ingroup agents (see Figure A36b). If this is not achieved, delegating communication may be harmful to social welfare. With sufficiently high strategic complementarity it may be optimal to limit ingroup agents when public signals are noisy.[40] Such low levels of public signal precision may occur when the central bank's information is noisy (compared to private information), or when it decides not to disclose all or parts of its information. Note that when social welfare contours are downwards-sloping, both delegation and disclosure harm welfare (see Figure A36b). The latter reflects the canonical finding if Morris and Shin (2002), while the former now identifies an additional policy tool for these circumstances: Alongside limiting disclosure, welfare can be raised by limiting delegation (e.g., through centralized communication or selecting homogeneous messengers), as this lowers the share of ingroup agents α.

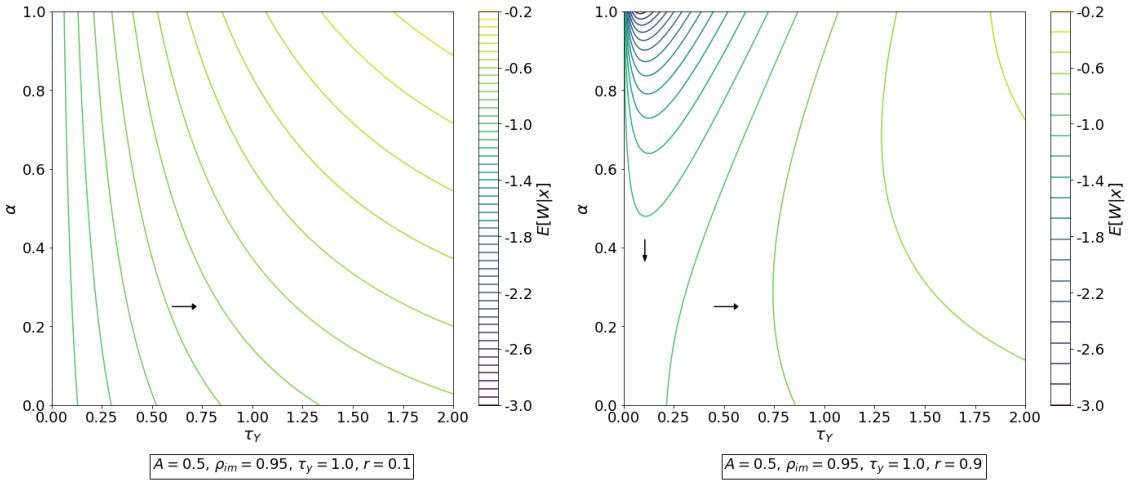
[40] The necessary level of strategic complementarity for this depends on the updating bias and the availability of the signal to outgroup agents. It is $r > \frac{1}{3}$ when outgroup agents update without a bias (irrespective of signal availability), but lower than that when there is biased updating and signals are available to outgroup agents. For example, when signals are large ($A \to 1$) and there is a slight updating bias of $\rho_{im} = 0.95$, $r > 0.156$ is sufficiently high strategic complementarity. Larger signals generally return lower sufficient r, while stronger updating biases raise sufficient r.

Furthermore, note that there exist circumstances, in which disclosure raises welfare, while delegation does not, as indicated by upwards-sloping welfare contours. Thus, delegation is a distinct policy tool, different from traditional disclosure policy. Still, delegation should be decided on in conjunction with disclosure considerations.

When private information is noisier ($\tau_y < \tau_Y$), delegating communication proves beneficial as ingroup agents use the precise public signal more effectively. Consequently, to align actions and maximize overall social welfare, raising the share of ingroup agents (α^*) as much as possible is optimal.

When coordination is sufficiently low, delegating communication is beneficial provided that any precision losses from lower expertise are sufficiently small. Precision losses are less concerning when private signals are relatively noisy but can be detrimental when public information is relatively noisier ($\tau_Y < \tau_y$). When coordination is high communication should only be delegated if the public signal is ensured to be sufficiently precise. In such cases, even with full disclosure, perceived losses in expertise of a delegated messenger may harm welfare due to the strong coordination motive causing over-reliance on the noisy public signal.

Figure 8: Social Welfare Contours



(a) Low Coordination (r=0.10) (b) High Coordination (r=0.90)

Notes: Plots show social welfare contours in (τ_Y, α)-space. τ_y is set to 1. Arrows indicate the direction in which welfare is increasing. Panel (a) shows a case of insufficient strategic complementarity (r=0.10). Panel (b) shows a case of sufficient strategic complementarity (r=0.90).

In conclusion, the central bank's communication strategy plays a pivotal role in shaping social welfare, particularly through the dual lenses of disclosure and delegation policies. While increased disclosure of public information can enhance coordination among agents, it is not

always welfare-improving, especially when the public signal is noisy. Outgroup agents buffer such potential welfare losses, as they rely less on public signals. Delegation, by increasing the share of ingroup agents, can improve coordination when public information is sufficiently precise, but must be carefully balanced against possible reductions in expertise and thus public signal quality. By jointly setting delegation and disclosure policies, while carefully considering the size of public signals, the precision of private information, and potential loss in messenger expertise from delegation, central banks can optimize their communication strategies to maximize social welfare.

5 Conclusion

The messenger matters. For communication to function effectively as a policy tool, the alignment between the messenger's characteristics and the audience's heterogeneity is critical. This paper, through a combination of empirical and theoretical insights, demonstrates that the ingroup effect—when the characteristics of the messenger and receiver match—significantly influences both the reach and the impact of public information.

There are ingroup-rooted benefits to two concrete ways of implementing a delegation policy: delegation of communication to regional central banks or to other board members with different characteristics. These ingroup effects are driven by homophily rather than heterophobia, meaning that ingroup messengers are relied upon more, without outgroup messengers being disproportionately disregarded. The effects are largely explained by perceived quality and trust, and while messengers influence information availability, they do not affect attention to information, which instead responds to the inflationary environment, including uncertainty.

To determine what optimal communication policy should look like, the paper develops a stylized coordination model, mirroring the empirical findings. It introduces both restricted availability of public signals and biased information processing, assessing how delegation and disclosure policies can be used strategically to maximize social welfare. While increasing the share of ingroup agents through delegation improves welfare when public information is precise, the presence of outgroup agents serves as a buffer against potential welfare losses from disclosing noisy public information. A dual approach of delegation and disclosure offers a nuanced framework for maximizing the social value of public information.

The generality of the model allows for broad applicability beyond monetary policy, extending to any public communication with heterogeneous audiences, including communication on climate or fiscal issues. Whether addressing dimensions like nationality, gender, ethnicity, socio-economic background, or expertise, the insights from this research provide a powerful framework for optimizing communication strategies to maximize social welfare.

References

Alsan, M., Garrick, O., and Graziani, G. (2019). Does diversity matter for health? experimental evidence from oakland. *American Economic Review*, 109(12):4071–4111.

Amador, M. and Weill, P.-O. (2010). Learning from prices: Public communication and welfare. *Journal of Political Economy*, 118(5):866–907.

Angeletos, G.-M. and Huo, Z. (2021). Myopia and anchoring. *American Economic Review*, 111(4):1166–1200.

Angeletos, G.-M. and Lian, C. (2018). Forward guidance without common knowledge. *American Economic Review*, 108(9):2477–2512.

Angeletos, G.-M. and Pavan, A. (2007). Efficient use of information and social value of information. *Econometrica*, 75(4):1103–1142.

Bartoš, V., Bauer, M., Chytilová, J., and Matějka, F. (2016). Attention discrimination: Theory and field experiments with monitoring information acquisition. *American Economic Review*, 106(6):1437–1475.

Bassetto, M. (2019). Forward guidance: Communication, commitment, or both? *Journal of Monetary Economics*, 108:69–86.

Benjamin, D. J. (2019). Errors in probabilistic reasoning and judgment biases. *Handbook of Behavioral Economics: Applications and Foundations 1*, 2:69–186.

Benjamin, D. J., Rabin, M., and Raymond, C. (2016). A model of nonbelief in the law of large numbers. *Journal of the European Economic Association*, 14(2):515–544.

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

Blinder, A. S. (2007). Monetary policy by committee: Why and how? *European Journal of Political Economy*, 23(1):106–123.

Blinder, A. S. (2018). Through a crystal ball darkly: The future of monetary policy communication. In *AEA Papers and Proceedings*, volume 108, pages 567–571. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Blinder, A. S., Ehrmann, M., De Haan, J., and Jansen, D.-J. (2024). Central bank communication with the general public: Promise or false hope? *Journal of Economic Literature*, 62(2):425–457.

Blinder, A. S., Ehrmann, M., Fratzscher, M., De Haan, J., and Jansen, D.-J. (2008). Central bank communication and monetary policy: A survey of theory and evidence. *Journal of economic literature*, 46(4):910–45.

Bodea, C., Ferrara, F. M., Kerner, A., and Sattler, T. (2021). Gender and economic policy: When do women speak with authority on economic issues? evidence from the euro area. *Evidence from the Euro Area (July 2, 2021)*.

Bodea, C. and Kerner, A. (2022). Fear of inflation and gender representation in central banking. *European Journal of Political Economy*, 74:102192.

Carrell, S. E., Page, M. E., and West, J. E. (2010). Sex and science: How professor gender perpetuates the gender gap. *The Quarterly journal of economics*, 125(3):1101–1144.

Cavallo, A., Cruces, G., and Perez-Truglia, R. (2017). Inflation expectations, learning, and supermarket prices: Evidence from survey experiments. *American Economic Journal: Macroeconomics*, 9(3):1–35.

Charness, G., Gneezy, U., and Halladay, B. (2016). Experimental methods: Pay one or pay all. *Journal of Economic Behavior & Organization*, 131:141–150.

Chen, D. L., Schonger, M., and Wickens, C. (2016). otree—an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9:88–97.

Cohen, J. (2013). *Statistical power analysis for the behavioral sciences*. Academic press.

Coibion, O., Gorodnichenko, Y., Kumar, S., and Pedemonte, M. (2020). Inflation expectations as a policy tool? *Journal of International Economics*, 124:103297.

Coibion, O., Gorodnichenko, Y., Kumar, S., and Ryngaert, J. (2021). Do you know that i know that you know...? higher-order beliefs in survey data. *The Quarterly Journal of Economics*, 136(3):1387–1446.

Cornand, C. and Heinemann, F. (2008). Optimal degree of public information dissemination. *The Economic Journal*, 118(528):718–742.

D'Acunto, F., Fuster, A., and Weber, M. (2022). A diverse fed can reach underrepresented groups. *University of Chicago, Becker Friedman Institute for Economics Working Paper*, (2021-95):21–20.

Do Hwang, I., Lustenberger, T., and Rossi, E. (2021). Does communication influence executives' opinion of central bank policy?. *Journal of International Money and Finance*, 115:102393.

Ehrmann, M., Soudan, M., and Stracca, L. (2013). Explaining european union citizens' trust in the european central bank in normal and crisis times. *The Scandinavian Journal of Economics*, 115(3):781–807.

Ehrmann, M. and Wabitsch, A. (2022). Central bank communication with non-experts–a road to nowhere? *Journal of Monetary Economics*, 127:69–85.

Esponda, I., Oprea, R., and Yuksel, S. (2023). Seeing what is representative. *The Quarterly Journal of Economics*, page qjad020.

European Central Bank (2016). What is the 'quiet period'? URL: https://www.ecb.europa.eu/ecb/educational/explainers/tell-me/html/what-is-the_quiet_period.en.html. Accessed: June 28, 2023.

Gershenson, S., Hart, C. M., Hyman, J., Lindsay, C. A., and Papageorge, N. W. (2022). The long-run impacts of same-race teachers. *American Economic Journal: Economic Policy*, 14(4):300–342.

Gust, C. and Lopez-Salido, D. (2024). Optimal monetary policy with uncertain private sector foresight.

Haldane, A. and McMahon, M. (2018). Central bank communications and the general public. In *AEA papers and proceedings*, volume 108, pages 578–583. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Hellwig, C. (2005). Heterogeneous information and the welfare effects of public information disclosures. Technical report, UCLA mimeo.

James, J. G. and Lawler, P. (2011). Optimal policy intervention and the social value of public information. *American Economic Review*, 101(4):1561–1574.

Knotek, E. S., Mitchell, J., Pedemonte, M., and Shiroff, T. (2024). The effects of interest rate increases on consumers' inflation expectations: The roles of informedness and compliance.

Kohlhas, A. N. (2020). An informational rationale for action over disclosure. *Journal of Economic Theory*, 187:105023.

Kohlhas, A. N. (2022). Learning by sharing: Monetary policy and common knowledge. *American Economic Journal: Macroeconomics*, 14(3):324–364.

Lamla, M. J. and Vinogradov, D. V. (2019). Central bank announcements: Big news for little people? *Journal of Monetary Economics*, 108:21–38.

Loria, S. (2018). textblob. *Release 0.15*, 2:269.

Malmendier, U. and Veldkamp, L. (2022). Information resonance. Technical report, Working Paper.

McMahon, M. and Rholes, R. (2023). Building central bank credibility: The role of forecast performance.

Mokhtarzadeh, F. and Petersen, L. (2021). Coordinating expectations through central bank projections. *Experimental economics*, 24:883–918.

Morris, S. and Shin, H. S. (2002). Social value of public information. *american economic review*, 92(5):1521–1534.

Nimark, K. P. and Pitschner, S. (2019). News media and delegated information choice. *Journal of Economic Theory*, 181:160–196.

Petersen, L. and Rholes, R. (2022). Macroeconomic expectations, central bank communication, and background uncertainty: A covid-19 laboratory experiment. *Journal of Economic Dynamics and Control*, 143:104460.

Pfajfar, D. and Žakelj, B. (2016). Uncertainty in forecasting inflation and monetary policy design: Evidence from the laboratory. *International Journal of Forecasting*, 32(3):849–864.

Phillips, L. D. and Edwards, W. (1966). Conservatism in a simple probability inference task. *Journal of experimental psychology*, 72(3):346.

Price, J. (2010). The effect of instructor race and gender on student persistence in stem fields. *Economics of Education Review*, 29(6):901–910.

Rholes, R. and Petersen, L. (2021). Should central banks communicate uncertainty in their projections? *Journal of Economic Behavior & Organization*, 183:320–341.

Stolper, O. and Walter, A. (2019). Birds of a feather: The impact of homophily on the propensity to follow financial advice. *The Review of Financial Studies*, 32(2):524–563.

Svensson, L. E. (2006). Social value of public information: Morris and shin (2002) is actually pro-transparency, not con. *American Economic Review*, 96(1):448–452.

Ter Ellen, S., Larsen, V. H., and Thorsrud, L. A. (2022). Narrative monetary policy surprises and the media. *Journal of Money, Credit and Banking*, 54(5):1525–1549.

Tutino, A. (2016). Central bank communication must overcome the public's limited attention span. *Economic Letter*, 11(6):1–4.

Veldkamp, L. L. (2011). *Information choice in macroeconomics and finance*. Princeton University Press.

Weber, M., Candia, B., Afrouzi, H., Ropele, T., Lluberas, R., Frache, S., Meyer, B., Kumar, S., Gorodnichenko, Y., Georgarakos, D., et al. (2025). Tell me something i don't already know: Learning in low-and high-inflation settings. *Econometrica*, 93(1):229–264.

Appendix

A Details on Twitter Data and Additional Evidence

A.1 Cleaning Twitter Data

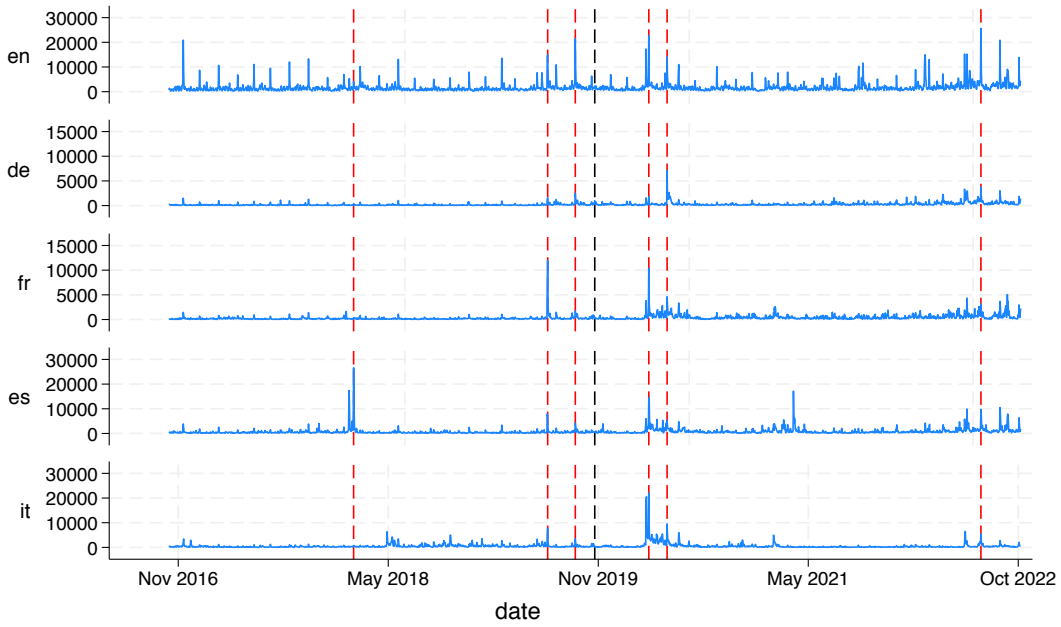
Tweets are collected between November 2022 and March 2023. Before any data cleaning, the dataset consists of just over 11 million tweets. First, I drop tweets that are unrelated to central banking. I start by manually identifying keywords that indicate unrelated tweets, such as any tweets about the English Cricket Board, which also abbreviates to ECB. This cleaning step is supported by two different word clouds. One visualizes the most frequent words for retained tweets, and may show any key words that are likely from unrelated tweets (e.g., "cricket", or names of cricket players), and the other visualizes most frequent words of dropped tweets, to inspire further keywords based on tweets are dropped for my final sample. Final word cloud for all languages of the cleaned sample are shown in Figure 1 and Figure A10. These word clouds indicate that the final sample indeed contains at a decisive majority tweets about central banking. This leaves just over 8 million (8,031,937) tweets, which I clean further, and translate to English using Google Translate's Neural Machine Translation. Specifically, I lower cases in each tweet and remove any special signs, emojis, digits, links to websites, in-tweet mentioned other users (indicated by preceding "@") except for "lagarde" and "ecb", and fully spell out contracted words (e.g., "can't" to "cannot" or "we'll" to "we will"). I translate all tweets, including English tweets, as this helps eradicate spelling mistakes, and ensures consistent, American English spelling of words, which is important for dictionary-based natural language processing (NLP) approaches.

A.2 Descriptive Analysis of Twitter Data

Tweet Volume The final sample contains just over 8 million tweets, whereby daily tweet volumes vary over time and by language (see Figure A9). The sample of English tweets stands out by its volume making up almost half of the entire sample (46%). This is not surprising as English is the official language of the ECB, and the lingua franca for policymakers and economists.

Idiosyncrasies of language samples and the validity of the proxy for nationality The daily tweet volume maxima (i.e., peaks) of tweet samples split by language coincides strongly with events of severe national interest. This is strong evidence in favor of the validity of proxying nationality with tweet language. Peaks in the time series across all

Figure A9: Daily tweet volumes by language.



Notes: Daily volume of tweets (using CET as this is the time zone corresponding to the headquarter of the ECB). The dashed grey vertical line indicates the changeover in presidency (November 1, 2019). Dashed red vertical lines indicate the following events (from left to right): Eurogroup's support for de Guindos (February 19, 2018), announcement of Lagarde as incoming ECB president (July 3, 2019), the day after PEPP announcement (March 19, 2020), German constitutional court ruling (May 5, 2020), ECB raising rates for the first time in 11 years (July 21, 2022).

samples generally refer to a decision by the ECB or an announcement of new ECB board members. Interestingly, the peaks of the time series are different for each language, and coincide with matters that have particular importance to different countries. For instance, the announcement of Lagarde as the incoming president of the ECB caused the peak in French data in July 2019, while the Eurogroup giving support to the candidacy of the Spanish vice president, de Guindos, in February 2018 caused the peak in the Spanish data. The peak in the German data coincides with German constitutional judges ruling that the ECB's asset purchasing programme partly violated the German constitution in May 2020. The peak in Italian data occurred in March 2020, one day after the announcement of the PEPP (Pandemic emergency purchase programme), which caused Italy's borrowing costs to fall sharply after the early and hard onset of Covid-19 in Italy. The peak in English data highlights how English tweets might be different from any other languages by marking the day on which the ECB raised interest rates for first time in 11 years (July 2022). This suggests that English is used relatively more for debating technicalities of EA-wide monetary policy than other languages that show a stronger focus on news concerning specific countries. At the same time, this

means that analyzing non-English tweets seems important to understand how particularly non-experts respond to central bank communication. Table A4 provides an overview of all peaks.

Table A4: Peaks in Tweet Volume by Language

Language	Peak Volume	Peak Date	Event
English	25,624	21 Jul 2022	ECB raising rates for first time in 11 years
German	7,098	05 May 2020	German constitutional court ruling
Spanish	26,599	19 Feb 2018	Eurogroup's support for Luis de Guindos
French	11,905	03 Jul 2019	Announcement of Lagarde as incoming ECB president
Italian	22,050	19 Mar 2020	Day after PEPP announcement

Notes: Tweet volume is defined as daily volume of tweets (using CET as this is the time zone corresponding to headquarter of the ECB).

Sentiment Measure Sentiment of tweet content is measured by a dictionary-based method that provides a continuous estimate of how positive or negative a text snippet is. A sentiment score is computed for each tweet using the polarity measure of the Python package TextBlob, which is based on the Princeton University's WordNet lexicon (Loria, 2018). Sentiment scores are estimated on tweets that are cleaned and translated (and thus spell-checked). This ensures that the same (English) dictionary can be used to create the sentiment scores. I normalize this sentiment indicator to assign values between 0 (most negative) and 1 (most positive) to each tweet, where neutral text gets a value of 0.5. The original polarity sentiment indicator returns a value between -1 and 1, where the most positive (negative) text obtains a value of 1 (-1) and neutral text gets a value of 0.

Sentiment across languages To be able to make meaningful comparisons of beliefs across languages, it is important that no language shows idiosyncratic differences in the measure of tweet sentiment. Given the strong similarity in the distribution tweet sentiment, this seems to be indeed the case. An overview of tweet sentiment across languages can be found in Table A5. The average and standard deviation of tweet sentiment is remarkable similar across languages.[41]

Time trends The sample covers six years; three entire years for each president. Many things happened in these six years, including digital advancements and increased computerization,

[41] Average tweet sentiment is close to neutral between 0.52 and 0.53, and the standard deviation is either 0.10 or 0.11. Even the 10^{th} and the 90^{th} percentiles of tweet sentiment do not vary considerably across countries, with values for the 10^{th} percentile ranging between 0.42 and 0.45, and for the 90^{th} between 0.63 and 0.65, respectively. It is noteworthy, however, that across languages most tweets have rather neutral sentiment around 0.5, and that tweets are slightly more positive than negative.

Table A5: Sentiment of Tweets by Language

Language	N	Mean	Std. dev.	p10	p90
English	3,697,934	0.53	0.10	0.45	0.65
German	565,854	0.52	0.10	0.42	0.63
Spanish	1,729,280	0.53	0.10	0.44	0.65
French	823,712	0.53	0.10	0.44	0.64
Italian	1,215,157	0.52	0.11	0.42	0.65

Notes: Tweet Sentiment ranges from 0 to 1, where more positive (negative) tweets are closer to 1 (0) and neutral tweets score 0.5. Means and standard deviations are based on all available individual tweets by language between 2016–2022.

all accelerated by the Covid-19 pandemic that moved much of communication online. Thus, it is not surprising that the number of tweets and the number of Twitter users has increased between the two three-year periods. The tweet volume increased by 53.5% between the three years of Draghi's presidency and the three years of Lagarde's. The total number of Twitter users increased even more: by 74.9%. Table A6 gives an overview of how the number of tweets and the number of Twitter users have changed between the two three-year periods by language. Blue font highlights when a nationality is in the ingroup with the ECB president. For the French sample, the increase in the number of tweets was by far the largest (145.3%). This coincides with their transition into the ingroup. Similarly, the number of French accounts grew substantially between the three years of Draghi's presidency and Lagarde's (by almost 93%). Only the number of German accounts grew more (by 127%), but from the lowest initial level. And even though the number of distinct users in the French sample is still below average (even when excluding English users), and the ranking of languages by the number of distinct users did not change between the two periods of presidency, it is noteworthy that the Italian sample—the Twitter users that fell out of the ingroup—grew the least (by only 30%). Similarly, Italian tweets saw the smallest growth (only 23.6%). Of course, there is a chance that this is a coincidence or merely reflecting heterogeneous levels of digitization across countries that have nothing to do with the ECB presidency, but this could also be another indicator of the ingroup behaving differently than the outgroup.

Table A6: Number of Tweets and distinct Twitter Users by Language and Presidency

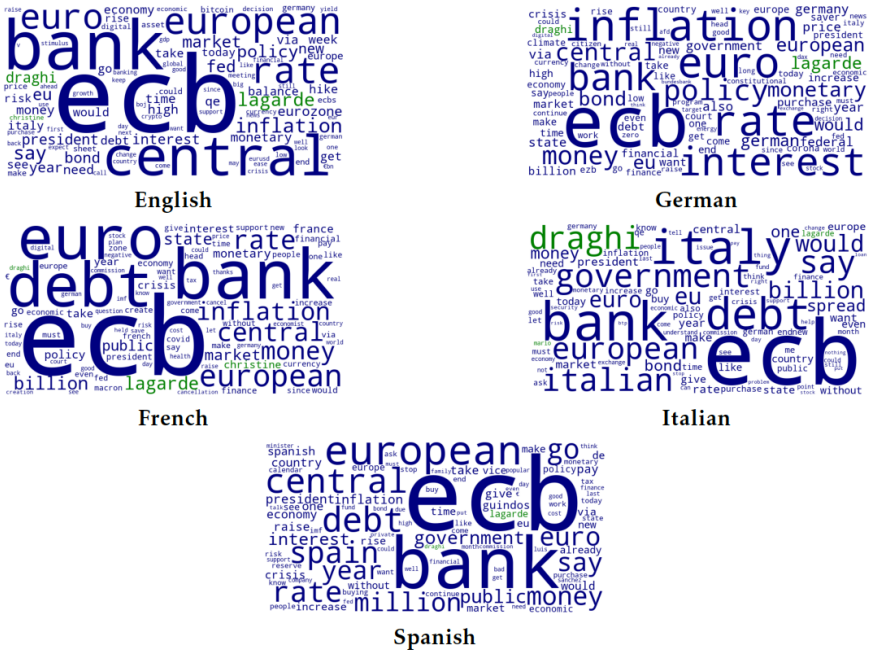
	Presidency		*Tweet Growth*
	Draghi	Lagarde	*(User Growth)*
English	1,595,810	2,102,124	31.7%
	(279,482)	(499,940)	(78.9%)
Spanish	623,046	1,106,234	77.6%
	(153,419)	(257,135)	(67.6%)
French	238,455	585,257	145.4%
	(62,737)	(120,942)	(92.8%)
Italian	543,504	671,653	23.6%
	(60,403)	(78,873)	(30.6%)
German	168,006	397,848	136.8%
	(29,903)	(67,829)	(126.8%)

Notes: The Table reports number of tweets and distinct Twitter users (in brackets) by language and presidency, where "Draghi" refers to tweets posted during 2016–2019 until the switchover in presidency in November 2019, and "Lagarde" refers to tweets posted afterwards (2019–2022). Languages are ranked by the number of distinct Twitter users (ranking remains the same across presidencies). The blue font indicates ingroup periods.

A.3 Additional Tables and Figures

The Sample The word clouds shown in Figure A10 indicate that the cleaning of tweets is successful in ensuring that tweets are indeed about central banking. In addition, the figure visualizes how different languages vary in their focus of discussions on Twitter, as discussed in the main part of the paper.

Figure A10: Word Clouds by Language



Notes: The figures show the 100 most frequent words between 2016-2022. Samples are split by original tweet language. Word size indicates word frequency. The following terms are highlighted in green for better visibility: "christine", "lagarde", "mario" and "draghi". Word cloud is based on cleaned, translated and lemmatized tweets. Stopwords are removed. No ngrams are included.

Table A7: Ingroup Effect on Information Supply on Twitter

	(1) All	(2) Experts	(3) Non-Experts
Ingroup=1	0.105***	0.0244**	0.0778***
	(0.0179)	(0.0103)	(0.0190)
ES	0.256***	0.129***	0.243***
	(0.0179)	(0.0103)	(0.0190)
FR	-0.00138	-0.0944***	0.0195
	(0.0200)	(0.0115)	(0.0212)
IT	0.0887***	0.0923***	0.0903***
	(0.0200)	(0.0115)	(0.0212)
Constant	0.138***	0.212***	0.142***
	(0.0127)	(0.00725)	(0.0134)
N	200	200	200
R-squared	0.582	0.738	0.505

Notes: The table shows OLS regression results of being in the ingroup with the ECB president on the share of tweets by language per 6-week press conference (PC) cycles, controlling for language. German acts as the baseline language. The number of observations reflects the 4 languages and 49 press conferences in the cycle, where one press conference cycle is split in partly being under Draghi's presidency and partly under Lagarde's, making it a total of 50 president-PC cycles combinations. The (non-)expert classification follows the benchmark in Ehrmann and Wabitsch (2022). Standard errors in parentheses. Significance level is indicated by stars: *** p<0.01, ** p<0.05, * p<0.1.

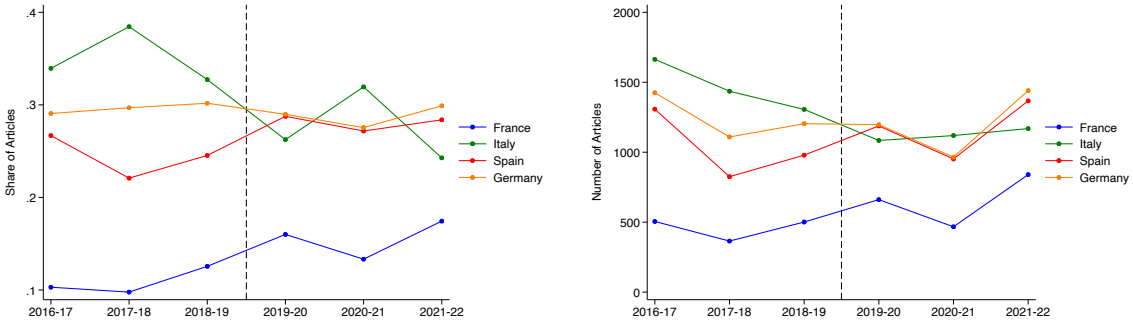
A.4 Does Print Media Confirm Increased Information Availability?

Twitter covers information availability beyond mere newspaper coverage, as it better reflects how exposed individuals are on average to given news story. However, to preclude doubts that information availability via print media provides a different picture to the results from Twitter, I analyse a sample of newspaper articles sourced from LexisNexis. The selection criteria mirror the Twitter sample: Articles are included in the sample if they contain the keywords "ECB", "European Central Bank" or the translated equivalents in the respective languages. For each nationality, the source of newspaper articles must be in the respective country and the language of the article must be the local language. The sample is restricted to only newspaper articles and the sample spans from November 1, 2016 to October 31, 2022, which reflects exactly three years of presidency for both Mario Draghi and Christine Lagarde. To avoid sample fluctuations over time that are due to LexisNexis-related data availability rather than reflecting actual changes in newspaper articles, I limit my sample to a few major newspapers for each country that are consistently available for the entirety of the sample, while ensuring political perspectives of selected newspapers are balanced. For Italian news, I include "Corriere della Sera" and "La Nazione". The French sample is comprised of "Le Figaro", "La Tribune", and "La Croix". The German sample contains "BILD" (incl. Sunday and regional editions), "Süddeutsche Zeitung" (incl. regional editions), and "Der Tagesspiegel". Finally, the included Spanish newspapers are "El Mundo" and "El País".

To assess newspaper article volume over time, the switchover in presidency acts as a pseudo-treatment, allowing to analyse the impact of the policymaker's nationality on national news with a difference-in-difference style analysis. While newspaper volume for Germany and Spain does not change much between the presidency of Draghi compared to Lagarde, Italian news decline during Lagarde's term time, while French news increase. Figure A11 visualizes these trends. Table A8 shows the estimated effect, controlling for time- and country-specific effects. The ECB president's nationality significantly increases the share of their corresponding national newspaper volume by 6.1 percentage points (or by roughly 272 articles).

In sum, also with traditional news, I find that the ECB president's nationality raises traditional national newspapers reporting. This confirms that there are positive ingroup effects on the availability of information, in line with the increased information availability seen on Twitter.

Figure A11: National Newspaper Articles Over Time



Notes: The left panel plots countries' share of articles each year, while the right panel shows the volume of national newspaper articles over time. Each point is based on the volume of newspaper articles over 12 months (from November 1 until October 31 of the following year). The vertical grey dashed line indicates the switch in presidents.

Table A8: Effect of ECB President's Nationality on Newspaper Volume

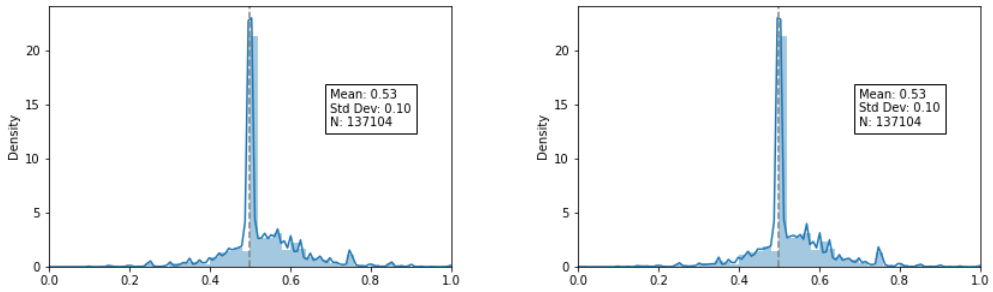
	(1) Share Of Articles	(2) Share Of Articles	(3) Number of Articles	(4) Number of Articles
Ingroup	0.061***	0.061**	271.8***	271.8**
	(0.018)	(0.021)	(68.1)	(71.6)
2017-2018	-0.000	0.020	-291.8***	-184.0
	(0.019)	(0.028)	(78.2)	(97.3)
2018-2019	-0.000	0.005	-228.0***	-181.0
	(0.012)	(0.025)	(67.2)	(155.4)
2019-2020	-0.000	-0.010	-193.0***	-212.0*
	(0.011)	(0.011)	(53.8)	(94.3)
2020-2021	-0.000	0.005	-349.8***	-291.5
	(0.017)	(0.041)	(66.0)	(144.3)
2021-2022	0.000	-0.013	-21.5	-80.0
	(0.015)	(0.031)	(73.2)	(124.0)
ES	-0.030**		-120.2*	
	(0.013)		(56.8)	
FR	-0.191***		-802.8***	
	(0.013)		(67.0)	
IT	-0.010	0.180***	-62.9	739.8***
	(0.020)	(0.021)	(76.6)	(71.6)
Constant	0.292***	0.100***	1404.0***	578.7***
	(0.007)	(0.013)	(42.2)	(104.7)
Observations	24	12	24	12
R^2	0.925	0.952	0.941	0.970

Notes: Column (1) shows the effect of being in the ingroup with the ECB president on the share of newspaper articles, controlling for time- and country-FE. Column (2) shows the same but limiting the sample to only nations that switch between the in- and outgroup (France and Italy). Similarly, columns (3) and (4) show the same for actual number of articles instead of shares. All regressions are estimated by OLS, and robust standard errors are shown in brackets. *Ingroup* is a dummy variable that equals 1 for the country of origin of the concurrent ECB president. Omitted baseline category for ingroup is the outgroup, for time periods it is 2016-2017, for countries it is Germany or France. Stars correspond to the following p-values: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

A.5 Empirical Estimation of Belief Updating

Priors and Posteriors To assess individual belief updating, I compare the sentiment of tweets before and after ECB press conferences. For users who tweet both before and after a conference, the posterior is defined as the sentiment of the first post-conference tweet: $Posterior_i = Sentiment_{i,t}$, where i represents an individual Twitter user and t is the time of the tweet. The prior is the sentiment of the user's last tweet during the ECB's quiet period (7 days before the conference) just before the press conference: $Prior_i = Sentiment_{i,t-1}$. This approach captures responses to the press conference and avoids picking up other ECB communications due to the quiet period's restricted communication. Figure A12 shows the distribution of priors and posteriors.

Figure A12: Distributions of Priors and Posteriors



Notes: The plots show the distribution of variables (priors and posteriors) by combining a histogram of each variable's values with a kernel density estimation (KDE) plot. The dotted vertical grey line goes through 0.5, which indicates neutral values across variables. Any values to the right (left) of the dotted line can be interpreted as positive (negative). The text boxes show the mean and standard deviation of the variables, as well as the number of observations. The left (right) panel shows normalized values of priors (posteriors) in the final sample of disaggregated analyses.

Empirical Strategy To empirically test the belief updating using Twitter data, I use an approach similar to Coibion et al. (2021). While the authors focus on the impact of the prior on the posterior in a single period, I analyze multiple updates of individuals over time, where t represents the time period corresponding to a specific press conference. Therefore, instead of including a constant, I include fixed effects for each updating event (i.e., press conferences). These fixed effects swallow any impacts the actual signals have on the updated posterior and simultaneously avoid making any assumptions about the signals.

$$Posterior_{i,t} = \beta_1 Prior_{i,t} + \delta' D_t + \epsilon_{i,t} \tag{23}$$

Coibion et al. (2021) assume that the distance between the prior's coefficient to 1 is explained by the effect of the missing signal. I first use this assumption to understand updating at different points in time first, and later incorporate actual signal values of press conferences for the main specifications. With the assumption, the prior coefficient is interpreted as follows: If the coefficient β_1 is equal to 0, the signals are interpreted as fully informative and the agents neglect their priors. Similarly, if $\beta_1 = 1$ the signals are fully uninformative, and if $0 < \beta_1 < 1$ the signals are partially informative. Thus, the larger the prior coefficient β_1 (i.e., closer to 1), the less agents respond to novel information from signals, and the more individuals rely on their prior beliefs when forming posterior beliefs. Note that a rational, Bayesian agent incorporates both the prior and the signal when forming their posterior belief, in proportion to their respective precision. Prior precision refers to the confidence or uncertainty assigned to the agent's prior belief, and signal precision is determined by signal accuracy. Since this information on priors and signals is not available in the observational data, no conclusions on whether updates are Bayesian or not can be made here.

To estimate ingroup effects, I include the interaction with a dummy ($Ingroup_{i,t}$) that indicates whether an individual i is in the ingroup. This shows whether and how individuals of the ingroup update differently than outgroup individuals. A negative coefficient of the interaction term suggests the ingroup relies less on their priors and more on new information more into their beliefs.

$$Posterior_{i,t} = \beta_1 Prior_{i,t} + \beta_2 Prior_{i,t} * Ingroup_{i,t} + \beta_3 Ingroup_{i,t} + \delta' D_t + \epsilon_{i,t} \quad (24)$$

Another advantage of avoiding the quantification of signals is that it allows me to flexibly define different updating horizons. Specifically, I first use Eq. (24) to assess how individuals update during quiet periods, which refer to the week before each press conference, where the ECB purposefully limits its communication with the public. Therefore, the expected coefficients for the prior should be higher than in the previous specification. This is because if there is less new information about central banking disclosed, beliefs should resemble their priors more. Second, I assess how individuals update outside quiet periods, excluding updates related to the press conferences. Here, the prior coefficient is expected to be smaller than during the quiet period, as the ECB communicates with the public during this time. The third specification restricts the sample to only updates to press conferences by enforcing two requirements: A posterior must be posted within 24 hours of the monetary policy announcement on a press conference day, and a prior must be posted in the most recent quiet period. The former ensures that the response is indeed due to information from the

press conference. The latter ensures that the update does not include information from other ECB communication. Whether updating around press conferences it is larger or smaller than outside quiet periods depends on whether the biggest shift in beliefs come from predominantly monetary policy related news, or other types of news that relate to the ECB (e.g., announcements of personnel changes).

The main results only focus on updates to press conferences. Instead of press conference fixed effects, I now include actual values for each press conference's signal $j \in [1, ..., J = 48]$ that allow to study how much individuals use novel information from the signal compared to their prior beliefs when updating their beliefs. To disentangle ingroup-driven effects on the signal, I also interact the signal with the ingroup dummy:

$$\begin{aligned} Posterior_{i,t} = \beta_1 Prior_{i,t} + \beta_2 Prior_{i,t} * Ingroup_{i,t} + \beta_3 Ingroup_{i,t} \\ + \beta_4 Signal_t + \beta_5 Signal_t * Ingroup_{i,t} + \epsilon_{i,t}. \end{aligned} \tag{25}$$

This allows to identify how much ingroup and outgroup individuals use the signal.

Results I run pooled regressions with all individuals by estimating Eq. (24) using OLS to show how beliefs are updated at different points in time: during quiet periods, outside quiet periods, and around press conferences.[42] Beliefs are updated the least during quiet periods (see column (1) in Table A9), more around press conferences (see column (3)), and most during the 5 weeks outside quiet periods (see column (2)). This makes sense intuitively, as the news about central banking are by definition limited during the quiet period, thus I expect posteriors to resemble their priors strongly. By contrast, press conferences communicate the most important monetary policy decisions, which is why seeing a larger shift in beliefs between the preceding quiet period and right after the press conference speaks for the credibility of any insights from analyses of these observational data. The biggest shift in beliefs occurs during the 5 weeks after a press conference and before the quiet period starts. During these weeks, many different kinds of news are released, which might explain the larger change in sentiment than when news are mostly about monetary policy.[43]

Across all three time periods, there exist ingroup effects, whereby ingroup individuals rely less on their priors (by around a third) than outgroup individuals (see Table A9). Lower prior reliance suggests higher responsiveness to new information; therefore, signals from the central bank are more effective in shaping beliefs.

[42] Regression results estimating individual coefficients from separate regressions for each agent confirm pooled regression findings and are available upon request.

[43] This might be even less surprising considering that news are expected at press conferences, and the timing of these press conferences is not a surprise.

Table A9: Belief Updates at Different Points in Time

	Quiet Period (QP)	*Outside QP (NQP)*	*Press Conferences (PCs)*
Dep. Var.: Posterior	**(1)**	**(2)**	**(3)**
Prior	0.511***	0.339***	0.492***
	(0.002)	(0.001)	(0.004)
*Ingroup*Prior*	-0.403***	-0.285***	-0.455***
	(0.005)	(0.002)	(0.010)
Ingroup	0.224***	0.154***	0.255***
	(0.003)	(0.001)	(0.005)
Outgroup Prior	0.511***	0.339***	0.492***
	(0.002)	(0.001)	(0.004)
Ingroup Prior	0.331***	0.208***	0.292***
	(0.002)	(0.001)	(0.005)
Fixed Effects	QP	NQP	PC
N	228,733	2,009,176	63,757
R-squared	0.958	0.958	0.954
Posterior Variable	most recent in same QP	most recent in same NQP	most recent from previous QP
Prior Variable	any in QP	any outside QP	within 24hrs after PC (closest to PC)

Notes: Regressions are estimated using OLS. Standard errors are shown in parentheses. *** indicates significance at the 1% level, ** at the 5% level, and * at the 10% level. N refers to the number of tweets. All regressions use individual tweets. The belief variables (prior and posterior) are determined by dictionary-based tweet sentiment. QP refers to "Quiet Period" and NQP to "Not Quiet Period" i.e., outside quiet periods. The quiet period refers to the 7 days leading up to a press conference, where there should not be any central bank signals (European Central Bank, 2016). To be included in the (N)QP regressions of columns (1) and (2) a user must tweet at least twice within the same (N)QP i.e., uninterrupted by a press conference (PC). To be included in the regression of column (3) a user must tweet in a quiet period (the prior) and within 24 hours of the following press conference, where the first tweet within this time window is selected as the posterior. Values for the "Ingroup Prior" are computed by taking the linear combination of relevant regression coefficients.

Table A10 shows belief updating around the press conference to better understand how individuals use their priors compared to the signal they receive. Column (1) includes fixed effects of individual press conferences (see Eq. (24)). While I do not obtain a coefficient for the signal this way, this specification requires fewer assumptions about the exact signals that individuals receive. The prior coefficient for outgroup individuals is at 0.49^{***}, compared to ingroup individuals for whom the prior coefficient is at 0.29^{***}. This may be interpreted as ingroup individuals relying almost a third less on their prior beliefs—thus being a third more open to novel information—when forming their beliefs following a press conference. This is even more pronounced when including values for the signals: Outgroup individuals' reliance on their priors is at 0.84^{***}, while this value is only 0.47^{***} for ingroup individuals (more than 40% less). Regressions shown in columns (2)-(5) are again estimated by OLS but now do not include time fixed effects (see Eq. (25)), but instead the actual signal of each press conference.

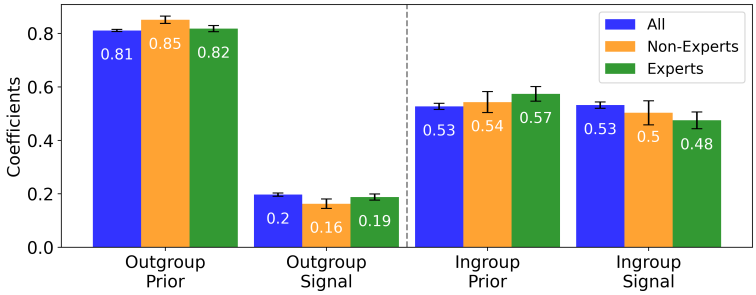
Columns (4) and (5) are the same regression as in column (3), but with the sample limited to non-experts and experts, respectively. Even when accounting for the actual signal that agents receive, ingroup effects are extremely persistent across specifications and subsamples. Ingroup agents hold on to their priors about a third less (even though both in- and outgroup agents get a larger coefficient for their posteriors compared to the regression specification without signal values). And similarly, ingroup agents use the signal more than twice as much as outgroup agents. In fact, their signal use becomes similar in size to their prior reliance. Figure A13 visualizes these results for the entire sample, and the non-(expert) subsamples.

Table A10: Belief Updating after Press Conferences

Dep. Var.: Posterior	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**
		All		*Non-Experts*	*Experts*
Prior	0.492***	0.835***	0.811***	0.851***	0.818***
	(0.004)	(0.002)	(0.002)	(0.007)	(0.006)
*Ingroup*Prior*	-0.455***	-0.806***	-0.782***	-0.795***	-0.701***
	(0.010)	(0.010)	(0.010)	(0.037)	(0.026)
Signal		0.166***	0.197***	0.163***	0.188***
		(0.003)	(0.003)	(0.009)	(0.006)
*Ingroup*Signal*			-0.164***	-0.147***	-0.169***
			(0.003)	(0.022)	(0.016)
Ingroup	0.255***	0.440***	0.498***	0.487***	0.457***
	(0.005)	(0.006)	(0.006)	(0.021)	(0.015)
Outgroup Prior	0.492***	0.835***	0.811***	0.851***	0.818***
	(0.004)	(0.002)	(0.002)	(0.007)	(0.006)
Ingroup Prior	0.292***	0.469***	0.527***	0.543***	0.574***
	(0.005)	(0.005)	(0.006)	(0.020)	(0.014)
Outgroup Signal			0.197***	0.163***	0.188***
			(0.003)	(0.009)	(0.006)
Ingroup Signal			0.532***	0.503***	0.475***
			(0.006)	(0.023)	(0.016)
Fixed Effects	PC	-	-	-	-
N	63,757	63,757	63,757	5,509	11,478
R-squared	0.954	0.946	0.946	0.944	0.957

Notes: Regressions are estimated using OLS. Regressions in columns (2)–(5) are estimated with no constant. Standard errors are shown in parentheses. *** indicates significance at the 1% level, ** at the 5% level, and * at the 10% level. N refers to the number of tweets. All regressions use individual tweets. The belief variables (prior and posterior) are determined by dictionary-based tweet sentiment. The signal is based on the OIS-2Y volatility indicator. Press Conference (PC) fixed effects refer the inclusion of dummies for each 6-week period starting at the moment of the press conference's monetary policy announcement. To be included in the regressions of columns (1)–(5) a user must tweet in a quiet period (the prior) and within 24 hours of the following press conference, where the first tweet within this time window is selected as the posterior. Columns (4) and (5) show the regression as shown in column (3), but restrict the sample to either non-experts or experts only. The (non-)expert classification follows the benchmark in Ehrmann and Wabitsch (2022). Values for the "Ingroup Prior" and "Ingroup Signal" are computed by taking the linear combination of relevant regression coefficients.

Figure A13: Belief Updating after Press Conferences By Expertise

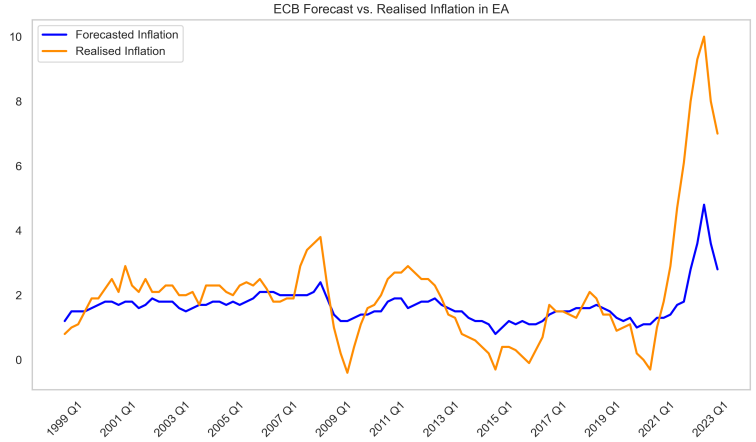


Notes: Bars show summed up OLS regression coefficients (base + interaction) with corresponding 95% confidence intervals of the priors (signals) in the left (right) panel for either individuals of the out- or the ingroup. The blue bars shows results for the entire sample, and the orange (green) bars for the subsample of non-experts (experts). The (non-)expert classification follows Ehrmann and Wabitsch (2022)

B Additional Information on the Experiment

Data in Forecasting Tasks The figure below shows the data used in the experiment for inflation and its forecasts.

Figure A14: Quarterly EA Inflation, SPF Forecasts, and SPF Forecast Errors

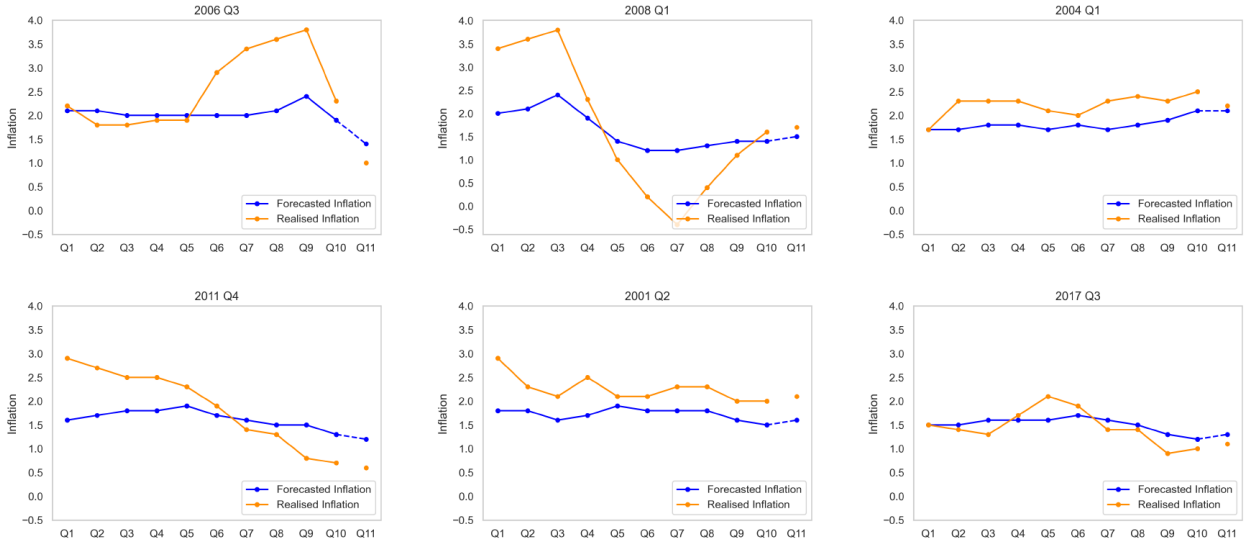


Notes: The figure shows realized quarterly EA HICP and forecasts of quarterly EA inflation (1-year-ahead). Data for inflation are taken from the ECB's Statistical Data Warehouse. Data for the inflation forecasts are based on the HICP forecast in the ECB Survey of Professional Forecasters (SPF), which can be found on the ECB website.

Details on Forecast Incentives Point forecasts for the prior and the posterior are incentivized with a bonus payment that is based on a participant's forecasting score $F_{i,t}$:

$$F_{i,t} = 6 * 3^{-|\mathbb{E}_{i,t-1}\{\pi_t\} - \pi_t|} \ , \tag{26}$$

Figure A15: Inflation Sequences



Notes: Sequences are randomly selected and based on real EA HICP and SPF 1-year-ahead forecasting data. The solid line represents the inflation (forecast) history that participants are shown initially. The dashed "predicted" line leads to the signal participants receive, and the unconnected final dot of the "realized" series is the value participants try to forecast. The first quarter of each sequence can be found at the top of each chart, however, note that date information will never be given to participants. Sequences are randomized across all messenger treatments.

where π_t is inflation at period t and $\mathbb{E}_{i,t-1}\{\pi_t\}$ is its forecast. A perfect forecast yields $F_{i,t} = 6$. This forecasting score is reduced by 2/3 of its value for each percentage point increase in the forecast error.

Range forecasts are incentivized using the scoring rule below, which follows the logic of e.g., Pfajfar and Žakelj (2016), Rholes and Petersen (2021) and Petersen and Rholes (2022).

$$U_{i,t}(r_{i,t}) = \left\{ \begin{array}{ll} 0 & \pi_{i,t} \notin \left[\underline{u_{i,t}}, \overline{u_{i,t}}\right] \\ \phi\left(\frac{1}{1+r_{i,t}}\right) & \pi_{i,t} \in \left[\underline{u_{i,t}}, \overline{u_{i,t}}\right] \end{array} \right\}, \tag{27}$$

where ϕ is a scalar that can be adjusted to scale average earnings, which are strictly increasing in ϕ. For this experiment, I set ϕ to 6. $\underline{u_{i,t}}$ represents the lower bound of a participant's forecast uncertainty, $\overline{u_{i,t}}$ the upper bound, and $r_{i,t} = ||\overline{u_{i,t}} - \underline{u_{i,t}}||$ represents a participant's forecast uncertainty. The intuition of this scoring rule is that if realized inflation falls out of a participant's forecast range, a participant earns nothing. But if instead realized inflation does fall within a participant's uncertainty bounds, the participant earns a payoff that is decreasing in the magnitude of her uncertainty. Thus, participants are incentivized to keep this range as narrow as possible, but also to ensure that it covers any inflation value they view as probable.

Table A11: Overview of Randomly Selected Sequences

10-Period Starting Date	10-Period Average of Forecasting Precision	10-Period Average of Inflation	10-Period STD of Inflation	Next-Period Signal	Next-Period Realisation
2006Q3	1.6	2.6	0.8	1.4	1
2008Q1	1.1	1.7	1.5	1.5	1.7
2004Q1	2.4	2.2	0.2	2.1	2.2
2011Q4	1.7	1.9	0.8	1.2	0.6
2001Q2	1.9	2.3	0.3	1.6	2.1
2017Q3	4.8	1.5	0.4	1.3	1.1

Notes: The table shows 10-period forecasting precision averages, realized inflation averages, and standard deviations of realized inflation for 6 randomly selected sequences. The indicated date marks the beginning of a 10-quarter sequence. Next-Period Signal and Realization refer to the forecast and realization of inflation, respectively, in the 11^{th} quarter of a given sequence. Data for realized EA inflation and for SPF's inflation forecasts are taken from the ECB's Statistical Data Warehouse.

After submitting their forecasts, participants are not immediately shown the realized inflation values. This approach avoids learning effects between forecasting tasks and ensures that their performance does not influence responses to subsequent questions about trust in institutions or policymakers similar to treatments. Instead, participants are informed of their point and range forecasting performance across all 12 forecasts (6 prior and 6 posterior forecasts) at the end of the survey. This is also when they learn which of the 24 decisions[44] has been selected randomly for their bonus payment. Randomly selecting one decision ensures that participants are incentivized to put high effort into each decision, as they cannot tell which of their forecasts will matter for their final payment.

Details on Treatments and Randomization Procedures In the *Generic Expert Treatments* (related to H1), participants encounter generic experts of four possible nationalities (Germany, France, Italy, and Spain), one of whom matches the participant's nationality (ingroup) and another randomly selected from the other three nationalities (outgroup).

The *ECB Expert Treatments* (related to H2 and H3) introduce ECB-affiliated experts with specified nationalities. For consistency, the outgroup expert in this treatment matches the outgroup nationality used in the Generic Expert Treatments, allowing a clearer comparison between ingroup and outgroup messengers across institutional and non-institutional contexts.

The *Institutional Expert Treatments* (related to H3 and H4) feature messengers from either the ECB or the participant's national central bank (NCB), such as the Deutsche Bundesbank or Banque de France, without specified nationalities. This treatment aims to test the effectiveness of delegating communication to NCBs, and whether homophily or heterophobia dominates ingroup effects with institutional context.

To minimize fatigue, each participant encounters only six out of the ten treatments—two

[44] 12 point and 12 range forecasts across 6 initial and 6 updated forecasts

from each group—with treatment order and inflation scenarios randomized to prevent learning effects and ensure valid comparisons.

Information Table A12 shows all information pieces revealed by clicking "Read more".

Table A12: Additional Information Pieces and Corresponding Inflation Scenarios

Inflation Scenario	**Information Pieces**
2006 Q3	1) Overall, the evidence that inflationary pressures are diminishing further has increased and, looking forward, inflation rates are expected to be in line with price stability, supporting the purchasing power of incomes and savings. 2) The recent decline in headline inflation mainly reflects the considerable easing in global commodity prices over the past few months, which more than offsets the impact of the sharp rise in unit labour costs in the first half of this year. Lower commodity prices and weakening demand should lead to inflationary pressures diminishing further. 3) The large declines in commodity prices and the impact of weakening demand on price developments strongly impacted this forecasts.
2008 Q1	1) Price developments are expected to remain moderate over the medium-term horizon. Global inflationary pressures may persist, while domestic price pressures are expected to remain low. 2) The still weak annual growth rate of bank loans to the private sector conceals the fact that monthly flows have now been positive for a number of months. At the same time, these aggregate developments continue to reflect mainly an ongoing strengthening in the annual growth of loans to households, while the annual growth of loans to non-financial corporations has remained negative. 3) Overall, price stability is expected to be maintained over the medium term, thereby supporting the purchasing power of Euro area households. Inflation expectations remain firmly anchored.
2004 Q1	1) The inflation rates are forecasted to remain elevated. Factors contributing to this include higher energy prices, indirect effects of past oil price increases, and potential stronger-than-expected wage developments. 2) The strength of monetary and credit growth, coupled with ample liquidity, pose inflationary risks over the medium to longer term. 3) The rapid rate of monetary growth is primarily driven by the stimulative impact of low-interest rates.
2011 Q4	1) The risks to the economic outlook for the Euro area are considered to be broadly balanced over the medium term, both for upside and downside risks to price developments. 2) The monetary analysis confirms subdued underlying growth in broad money and credit, and the annual growth rate of loans to the private sector continues to contract. 3) Progress has been made in improving the funding situation of banks, but further steps are needed to strengthen the resilience of banks and reduce the fragmentation of Euro area credit markets.
2001 Q2	1) Real economic activity in the Euro area was weak recently, but survey data and recent financial market developments indicate a gradual upturn in economic activity. 2) The expected pick-up in activity is supported by both external factors, such as global recovery, and domestic factors, such as ongoing adjustment efforts by companies to enhance competitiveness and profitability. 3) Downside risks to the main scenario for economic growth have declined, but they have not disappeared. Macroeconomic imbalances in some regions and high oil prices are mentioned as potential risks to economic activity in the Euro area.
2017 Q3	1) While there are some initial signs of stabilization in the growth slowdown and a mild increase in underlying inflation, overall inflation remains low. 2) Ongoing employment growth and increasing wages are underpinning the resilience of the Euro area economy. 3) Inflation expectations are at low levels, and measures of underlying inflation have remained generally muted. Inflation is expected to increase over the medium term, supported by monetary policy measures, economic expansion, and solid wage growth.

Notes: Information pieces are linked to inflation scenarios (each scenario's first quarter shown in the first column). Each inflation scenario includes three buttons, each displaying one piece of information. Information pieces are summarized statements based on ECB press conferences, which can be found in full here.

Details on Post-Experimental Survey To be able to account for any pre-existing perceptions of institutions and policymakers that could influence decision-making, participants are surveyed on monetary institutions (the ECB and the respective NCB) and policymakers that are representative of the messenger treatments 5–10. Participants are asked to report their *trust*[45] in these institutions and policymakers, and their respective *exposure* (i.e., the extent to which they know of and follow news about these institutions and policymakers). As representative ECB experts of the in- and outgroup, I take advantage of the fact that all four nationalities are represented among the six ECB board members at the time of the experiment.[46] As further in- and outgroup policymakers, NCB governors are included in the survey. The NCB governors at the time of the experiment are: François Villeroy de Galhau for France, Joachim Nagel for Germany, Ignazio Visco for Italy, and Pablo Hernández de Cos for Spain. Table A13 provides an overview of chosen representative policymakers.

Table A13: Representative Policymakers by nationality (at the time of experiment)

	ECB Board Member	**NCB Governor**
France	Christine Lagarde	François Villeroy de Galhau
Italy	Fabio Panetta	Ignazio Visco
Germany	Isabel Schnabel	Joachim Nagel
Spain	Luis de Guindos	Pablo Hernández de Cos

Notes: Overview of all representative in- and outgroup policymakers by nationality used in post-experimental survey. ECB Board Members represent in- and outgroup ECB experts: "Christine Lagarde" represents the French ingroup ECB expert, "Luis de Guindos" proxies the Spanish, "Isabel Schnabel" the German, and "Fabio Panetta" the Italian ECB ingroup expert. Similarly, the right column presents all four NCB governors at the time of the experiment. All participants face all policymakers, which are counted as ingroup (outgroup) policymakers whenever nationalities of participants and policymakers match (mismatch).

Participants are further asked about their *monetary policy expertise*, and they are tested on having paid attention during the experiment.

B.1 Power Analysis

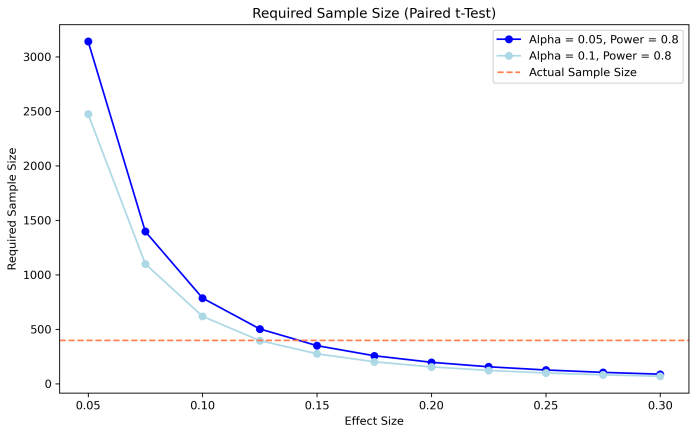
Power analysis is commonly used to determine the required number of observations to preclude insignificant results that are due to low statistical power (see e.g., Cohen (2013)). My choice of sample size for the experiment follows from such power analysis for comparisons of dependent

[45] Trust is only elicited for policymakers and institutions that the participant indicates to know of, and answers are given on a 7-point Likert scale (see for details Figure A32.

[46] "Christine Lagarde" is used to proxy for the French ingroup ECB expert, and for all others to proxy an outgroup ECB expert. Similarly, "Luis de Guindos" proxies the Spanish, "Isabel Schnabel" the German, and "Fabio Panetta" the Italian ECB ingroup expert. And respectively, these policymakers proxy an outgroup ECB expert, whenever the participant nationality does not match the policymaker's nationality.

means (since I use a within-subjects design). Based on the observational results from my Twitter analysis, pilot data, and findings from a comparable study (i.e., D'Acunto et al. (2022)), the expected effect size is approximately 0.1, ranging between 0.1 and 0.3. Figure A16 visualizes the required sample size for each of these expected effect sizes at various levels of required power and significance. The most commonly used power (alpha) specification is 0.8 (0.05), which is represented by the dark blue line. Considering this specification, my ideal sample size is just below 800, as this would allow me to statistically identify effect sizes of 0.1. However, due to budget constraints, the sample size of the experiment is limited to 400 participants (100 per nationality), which allows me to safely identify effect sizes of approximately 0.14 and above. This means that if the observed effects are smaller and statistically insignificant, their lack of significance might be due to insufficient sample size.

Figure A16: Power Analysis



Notes: The figure shows optimal sample size to ensure significance of effects for various power and alpha levels for expected effect sizes (paired t-test). The red horizontal dashed line indicates the actual sample size that was achievable with the available funding.

B.2 Bayesian Belief Updating Framework

To assess how agents update their inflation expectations following central bank communication, I use a standard Bayesian belief updating framework. This framework estimates how agents balance novel information (central bank communication) against their existing beliefs (priors). In the Bayesian case, agents weigh information and priors equally when both have the same precision (e.g., Veldkamp (2011)).

Assume macroeconomic expectations are normally distributed, where agent i believes inflation x has a mean of A_i and precision α_i. The prior belief is $x \sim \mathcal{N}(A_i, \alpha_i^{-1})$, and agents

receive a signal B about inflation x from the central bank: $B = x + e, \quad e \sim \mathcal{N}(0, \beta^{-1})$, where the signal has precision β. Applying Bayes' law, agents update their beliefs as:

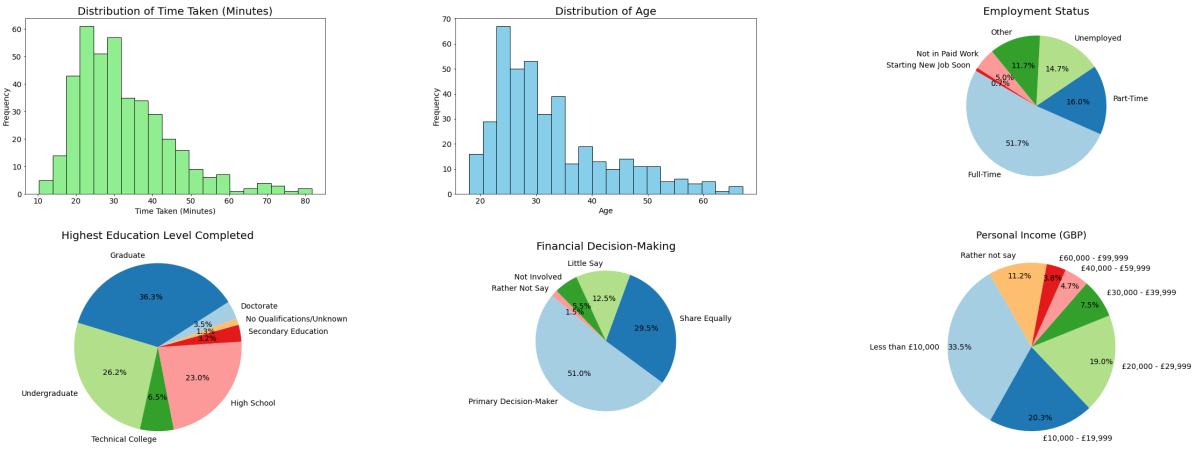
$$E_i[x|B] = \frac{\alpha_i A_i + \beta B}{\alpha_i + \beta}, \tag{28}$$

which shows that the posterior belief is a weighted average of the prior and the signal, based on their relative precision.

B.3 Additional Experimental Results, Tables and Figures

Participants in the Experiment The following background data of participants is provided by Prolific (as defined before running the study): age, sex, income bracket, marital status, employment status, student status, ethnicity, country of birth/residence, years lived in current country of residence, country spent most time in before turning 18, extent of financial decision-making, primary and fluent languages, highest obtained education, migration history, investments, property ownership, information on whether participant has lived abroad, is mono/multicultural, or has been raised monolingual. Figure A17 shows some of the participants' characteristics descriptively in more detail.

Figure A17: Descriptive Evidence of Experiment Participants



Notes: The figure shows descriptively various characteristics of experiment participants, where the characteristic of interest is shown on top of each panel. Only the histogram on the top left is not a participant characteristic but shows the amount of time (in minutes) participants needed to complete the experiment.

While the sample does not attempt to be representative of the populations in Germany, Spain, Italy or France, the participants in the collected sample are much more diverse compared to common samples in lab experiments, as these commonly include students only.

Participants' age ranges from 18 to 67 and averages at 33. More than half of the participants are employed full-time, only around 35% are students, more than half of participants earn more than 10,000 GBP per year, more than half of the participants are married or in a relationship, around a third of the sample is female (compared to two thirds being male), and the highest level of education obtained is distributed fairly evenly among high school, undergraduate, and graduate levels. Most participants complete the experiment in half an hour. A few outliers take around an hour. Prolific automatically disqualifies participants from submitting the experiment if they take longer than 90 minutes.

Average signal use The table below shows that participants uder-use the signal compared to a Bayesian benchmark across treatments.

Table A14: Regression Results By Nationality

	(1) **Overall**	(2) **Germany**	(3) **Spain**	(4) **France**	(5) **Italy**
Signal	0.899***	0.881***	0.889***	0.877***	0.941***
	(0.012)	(0.025)	(0.025)	(0.024)	(0.026)
Prior	1.206***	1.234***	1.198***	1.262***	1.139***
	(0.012)	(0.025)	(0.022)	(0.025)	(0.025)
R-squared	0.961	0.961	0.963	0.964	0.959
N	2,385	600	594	593	598

Notes: The table shows the regression results of estimating Eq. 1 by OLS without a constant. Standard errors are shown in brackets. Results for the entire sample are shown in column (1), while the other columns show results for samples split by nationality. Stars correspond to the following p-values: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Mechanism The table below shows the role of perceived signal quality and trust.

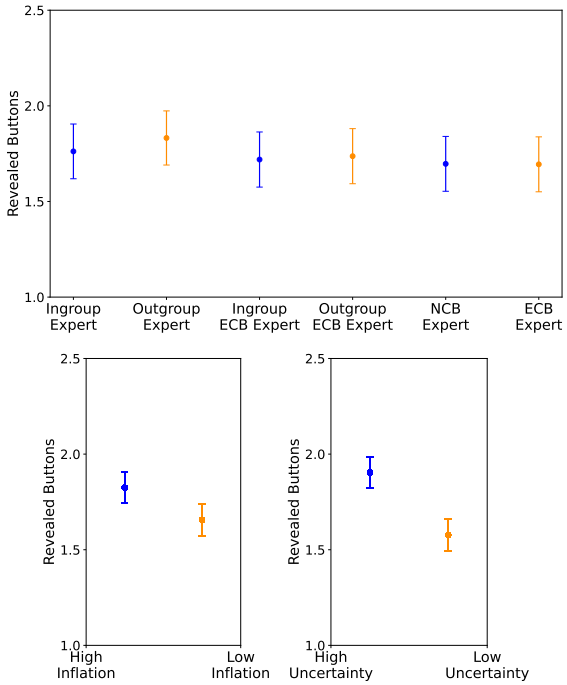
Table A15: Mechanism: Perceived Quality and Trust

	(1)	(2)	(3)	(4)
Pure Ingroup Effect (H1):	0.052***	0.010**		
	(0.017)	(0.005)		
R-squared	0.994	0.992		
N	795	795		
Ingroup Effect for ECB Experts (H2):	0.028*	0.005	0.053	0.006
	(0.017)	(0.004)	(0.044)	(0.015)
R-squared	0.994	0.993	0.998	0.997
N	795	795	315	315
Homophily - Ingroup ECB Expert vs. Generic ECB Expert (H3):	0.035**	0.009**	0.051**	0.021***
	(0.018)	(0.004)	(0.022)	(0.007)
R-squared	0.993	0.992	0.996	0.995
N	794	794	612	612
Heterophobia - Outgroup ECB Expert vs. Generic ECB Expert (H3):	0.013	0.008**	-0.004	0.012
	(0.017)	(0.004)	(0.033)	(0.011)
R-squared	0.994	0.993	0.998	0.998
N	795	795	477	477
NCB vs ECB: Institutions Effect (H4):	0.034**	0.008**	0.033*	0.005
	(0.017)	(0.004)	(0.017)	(0.004)
R-squared	0.994	0.994	0.994	0.993
N	795	795	779	779
Inflation Scenario	✓	✓	✓	✓
Individual-FE	✓	✓	✓	✓
Perceived Messenger Ability		✓		
Trust				✓
Knowing PMs/Inst			Yes	Yes

Notes: Effects result from comparing coefficients of interest of the same regression. *Inflation Scenario* refers to controlling for the underlying data of forecasting tasks and the order in which they appeared. *Perceived Messenger Ability* refers to the self-reported 7-point Likert scale indication of how able each messenger is thought to be at forecasting inflation and providing economic analyses. *Knowing PMs/Inst* refers to whether participants indicated to at least have heard of the institutions or real policymakers (see Section 3.1 for details on this) matching the messenger treatments. Conditional on knowing representative PMs and/or institutions, self-reported *Trust* is available (based on 7-point Likert scale). *N* refers to the number of observations (forecasting tasks) in the regression. N may slightly vary across treatments due to instances of infinite prior precision. Column (1) shows the main results again. Column (2) interacts signal use with perceived ability. Column (4) interacts this effect with trust in a representative policymaker or in the institution. Column (3) mirrors column (1) but on the restricted sample where policymakers and institutions are known in each treatment, which allows for better comparison with results in column (4). Stars correspond to the following p-values: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Attention to Information The top panel in Figure A18 shows that while attention to information across messenger treatments is fairly constant (around 1.8 revealed information buttons), the bottom panel of the figure shows that attention indeed responds to both the level of inflation and its uncertainty.

Figure A18: Requesting additional information pieces (with 95%-CI)



Notes: Top row shows the additional information pieces requested via the average number of buttons clicked (with 95%-CI) by messenger treatments. Bottom row shows the same by inflation level (left panel) or inflation uncertainty (right panel). Out of the six inflation scenarios, the three scenarios with the highest (lowest) inflation shown as "High (Low) Inflation". Similarly, the top (bottom) three scenarios in terms of the inflation history's standard deviation are shown as "High (Low) Inflation Uncertainty".

Attention and Ingroup Effects When participants request additional information, they show stronger ingroup effects across hypotheses. I attribute this to these participants being attentive. In fact, ingroup effects disappear when limiting the sample to forecasting tasks where no buttons are revealed (i.e., "no attention"), highlighting that treatments might not have been salient to these inattentive participants. In turn, effects become stronger the more attention individuals pay (i.e., the more buttons they reveal). Note that since attention is not exogenous in my experiment, these effects are not strictly causal. For instance, a lack of attention could also represent low effort exertion, resulting in missing the treatment variation.

Two robustness exercises show that it is indeed attention—not the additional information—that drives ingroup effects on signal use. Column (2) in Table A16 compares only forecasting

tasks where all information is retrieved and confirms effects (except that the ingroup effect within ECB context loses significance). Similarly, column (3) confirms robustness of all effects with the full sample, controlling for revealed buttons (for given inflation scenarios).

Table A16: The Role of Attention and Being Informed

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Pure Ingroup Effect (H1):	0.052***	0.066***	0.056***	0.057**	0.039			
	(0.017)	(0.025)	(0.018)	(0.024)	(0.036)			
R-squared	0.994	0.995	0.994	0.995	0.995			
N	795	463	795	488	307			
Ingroup Effect for ECB Experts (H2):	0.028*	0.029	0.029*	0.018	0.033	0.053	-0.006	-0.037
	(0.017)	(0.025)	(0.017)	(0.027)	(0.035)	(0.044)	(0.064)	(0.107)
R-squared	0.994	0.996	0.994	0.995	0.995	0.998	0.999	0.999
N	795	446	795	465	330	315	193	122
Homophily - Ingroup ECB Expert vs. Generic ECB Expert (H3):	0.035**	0.083***	0.040**	0.081***	-0.007	0.051**	0.110***	-0.030
	(0.018)	(0.026)	(0.018)	(0.025)	(0.032)	(0.022)	(0.031)	(0.042)
R-squared	0.993	0.995	0.994	0.995	0.995	0.996	0.997	0.998
N	794	437	794	462	332	612	361	251
Heterophobia - Outgroup ECB Expert vs. Generic ECB Expert (H3):	0.013	0.059**	0.007	0.057**	-0.032	-0.004	0.047	-0.251**
	(0.017)	(0.023)	(0.017)	(0.022)	(0.036)	(0.033)	(0.040)	(0.097)
R-squared	0.994	0.996	0.994	0.996	0.995	0.998	0.999	1.000
N	795	441	795	465	330	477	284	193
NCB vs ECB: Institutions Effect (H4):	0.034**	0.058**	0.038**	0.057**	0.005	0.033*	0.052**	0.006
	(0.017)	(0.025)	(0.018)	(0.024)	(0.034)	(0.017)	(0.024)	(0.035)
R-squared	0.994	0.995	0.994	0.995	0.995	0.994	0.996	0.995
N	795	433	795	464	331	779	455	324
Inflation Scenario	✓	✓	✓	✓	✓	✓	✓	✓
Individual-FE	✓	✓	✓	✓	✓	✓	✓	✓
Attention		Full	Controlled	Yes	No		Yes	No
Knowing PMs/Inst						Yes	Yes	Yes

Notes: Effects result from comparing coefficients of interest of the same regression. *Inflation Scenario* refers to controlling for the underlying data of forecasting tasks and the order in which they appear. *Attention - Full* refers to the subsample of inflation tasks, in which participants click on all information buttons to reveal additional information. *Attention - Controlled* refers to the full sample of inflation tasks, but controls for which buttons participants click on (for a given inflation scenario) to reveal additional information. *Attention - Yes* refers to the subsample of inflation tasks, in which participants click on at least one button to reveal additional information. Conversely, *Attention - No* reflects the subsample of inflation tasks, for which not a single button to reveal more information is clicked on. *Knowing PMs/Inst* refers to whether participants indicate to at least have heard of the institutions or real policymakers matching the messenger treatments. *N* refers to the number of observations (forecasting tasks) in the regression. N may slightly vary across treatments due to instances of infinite prior precision. Stars correspond to the following p-values: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Interestingly, attention further highlights another potential explanation of ingroup effects within the institutional context. When representative ECB policymakers are known but attention is low, there is substantial heterophobia (-0.251^{**}) This suggests that nationality may be used as a shortcut. Conversely, homophily is very large (0.110^{***}) when ECB policymakers are known and participants are attentive (see columns (6)-(8) in Table A16).

B.4 Information Availability: Exposure to News in the Media

The survey accompanying the experiment highlights that participants receive more news about in- than outgroup policymakers. The probabilities of having heard of, knowing, and following news about policymakers[47] are all larger for ingroup policymakers (see Table A17).

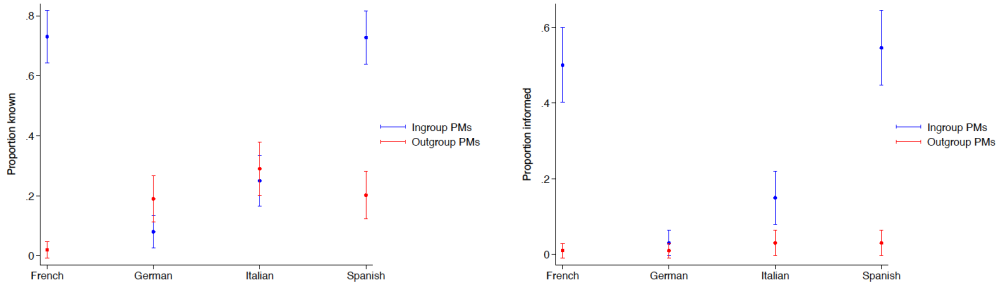
Table A17: Information Availability and Reach: Survey Evidence

	(1) Heard Of	(2) Know	(3) Informed
Ingroup=1	0.080**	0.271***	0.286***
	(0.035)	(0.031)	(0.024)
Constant	0.489***	0.175***	0.020
	(0.025)	(0.022)	(0.017)
R-squared	0.01	0.09	0.15
N	798	798	798

Notes: The representative in- and outgroup policymakers are all ECB board members (Lagarde, de Guindos, Schnabel, Panetta) but regressions for NCB governors (Villeroy de Galhau, de Cos, Nagel, Visco) are similar. Regressions are estimated by OLS. The Ingroup dummy indicates the respective probabilities for ingroup compared to outgroup policymakers. Data come from the survey question shown in Figure A31.

The likelihood of knowing policymakers increases by roughly a third if they are in the ingroup rather than outgroup. Participants are 29% more likely to follow news of a given policymaker if they are in the ingroup. Figure A19 plots the shares of respondents who indicate knowing or following news about in- and outgroup policymakers by participant nationality. Effects for ECB ingroup policymakers seem to be driven by the French and Spanish—coinciding with the nationalities of the ECB president and vice president. This suggests that nationality strongly matters, especially for these two ECB Board positions.

Figure A19: Knowing representative policymakers and receiving news



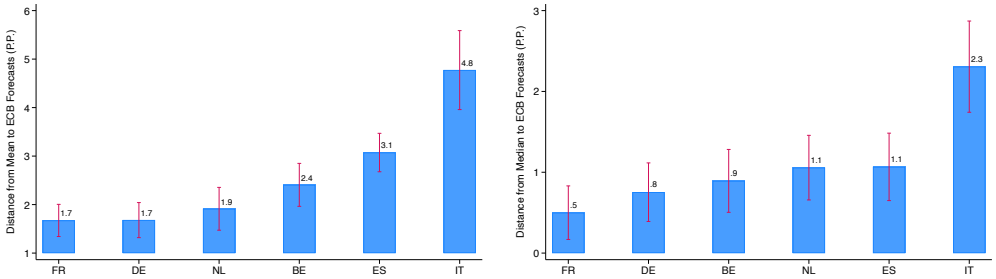
Notes: The left (right) panel shows the share of participants in the experiment who report knowing (receiving news about) representative in- and outgroup policymakers (PMs). The average response across all three outgroup PMs is used. These plots show results only for ECB board members (Lagarde, de Guindos, Schnabel, Panetta) but plots for NCB governors (Villeroy de Galhau, de Cos, Nagel, Visco) are very similar.

[47] As described earlier in the paper, the policymakers used to proxy for messengers are NCB governors (Villeroy de Galhau, de Cos, Nagel, Visco) and ECB board members (Lagarde, de Guindos, Schnabel, Panetta)

B.5 Real-world Forecasting

The experimental findings suggest that ingroup agents update more strongly to a signal; thus their inflation forecasts should—ceteris paribus—more closely resemble the forecasts by the ECB. To check whether this is indeed the case, I compare the distance between the ECB's 12 months-ahead inflation forecasts to forecasts of households in the EA. Values for the former are taken from the ECB's Macroeconomic Projection Database (MPD). The MPD is published four times a year (in March, June, September and December).[48] The latter come from the ECB's Consumer Expectations Survey (CES), which is a representative survey of consumers in selected EA countries started in 2020. I use their aggregated monthly mean and median indicators of quantitative household inflation expectations over the next 12 months.[49]

Figure A20: Comparing Household Inflation Forecasts to ECB Projections



Notes: The left (right) panel shows the difference between mean (median) household twelve months-ahead inflation forecasts and ECB projection (lagged by 1 month) with 95% confidence intervals. I consider fifteen survey waves for six countries, whereby data for Belgium are missing for the first wave (April 2020).

Figure A20 shows that forecasts of 1 year-ahead inflation by households in France indeed most closely resembled ECB's forecasts (lagged by 1 month).[50] Available data spans from April 2020 until December 2023, when the ECB's president was French (Christine Lagarde). During this time, the mean forecasts for French households are on average 1.7 percentage points above the ECB's projections, the lowest distance across countries (together with German households). Across countries, mean household forecasts are on average 2.6 percentage points above ECB forecasts. Similarly, French median household forecasts are most similar to ECB forecasts (0.5 percentage point distance), with a cross-country average distance of 1.1 percentage points. While these findings are correlations rather than causal effects, the trend is in line with ingroup agents updating more closely to the signal.

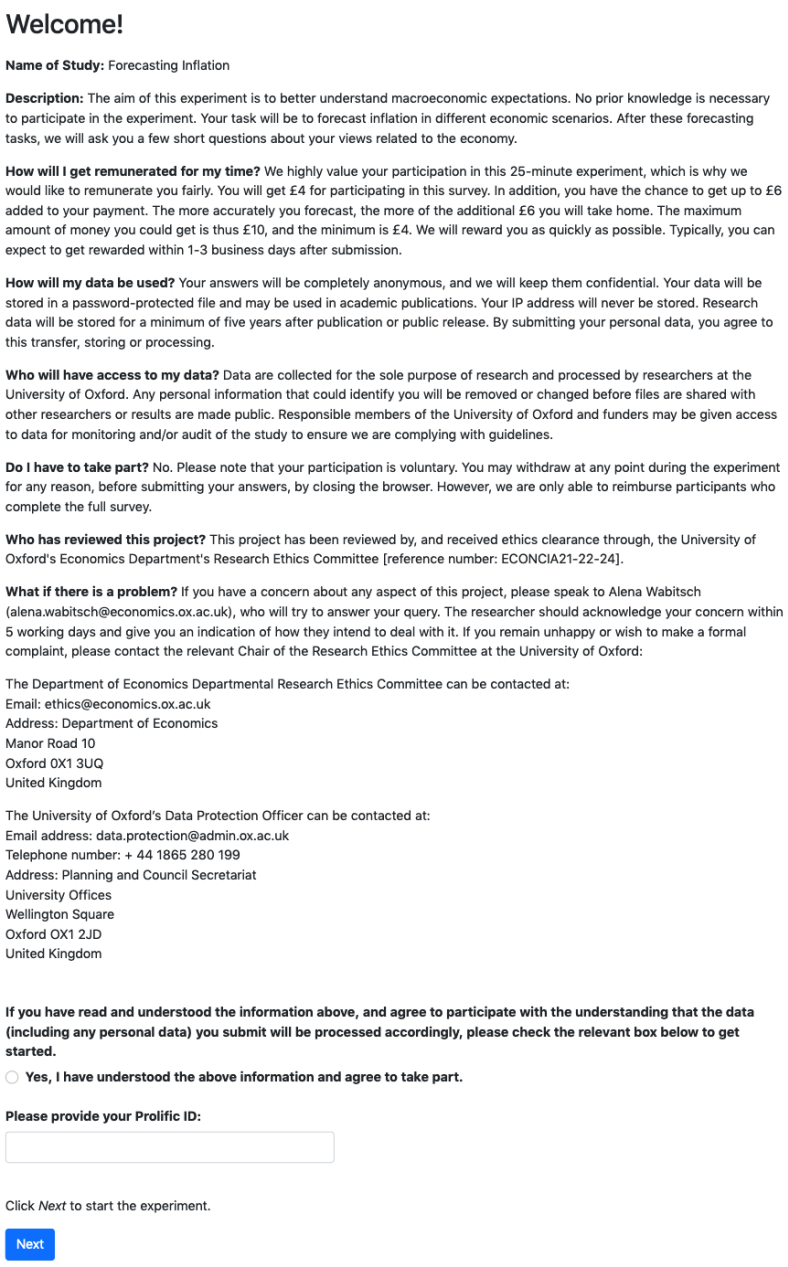
[48]The MPD contains information on the outlook for the EA and contributes to the ECB Governing Council's assessment of economic developments and risks to price stability. The June and December projections are conducted by Eurosystem staff, while the March and September projections are conducted by ECB staff.

[49]The survey provides mean and median values that are population-weighted and winsorised at the 2nd and 98th percentiles of the weighted distribution of responses for each survey wave and country.

[50]Not lagging ECB's forecast results in very similar findings.

C Interface of the Experiment

Figure A21: Welcome Page:



Welcome!

Name of Study: Forecasting Inflation

Description: The aim of this experiment is to better understand macroeconomic expectations. No prior knowledge is necessary to participate in the experiment. Your task will be to forecast inflation in different economic scenarios. After these forecasting tasks, we will ask you a few short questions about your views related to the economy.

How will I get remunerated for my time? We highly value your participation in this 25-minute experiment, which is why we would like to remunerate you fairly. You will get £4 for participating in this survey. In addition, you have the chance to get up to £6 added to your payment. The more accurately you forecast, the more of the additional £6 you will take home. The maximum amount of money you could get is thus £10, and the minimum is £4. We will reward you as quickly as possible. Typically, you can expect to get rewarded within 1-3 business days after submission.

How will my data be used? Your answers will be completely anonymous, and we will keep them confidential. Your data will be stored in a password-protected file and may be used in academic publications. Your IP address will never be stored. Research data will be stored for a minimum of five years after publication or public release. By submitting your personal data, you agree to this transfer, storing or processing.

Who will have access to my data? Data are collected for the sole purpose of research and processed by researchers at the University of Oxford. Any personal information that could identify you will be removed or changed before files are shared with other researchers or results are made public. Responsible members of the University of Oxford and funders may be given access to data for monitoring and/or audit of the study to ensure we are complying with guidelines.

Do I have to take part? No. Please note that your participation is voluntary. You may withdraw at any point during the experiment for any reason, before submitting your answers, by closing the browser. However, we are only able to reimburse participants who complete the full survey.

Who has reviewed this project? This project has been reviewed by, and received ethics clearance through, the University of Oxford's Economics Department's Research Ethics Committee [reference number: ECONCIA21-22-24].

What if there is a problem? If you have a concern about any aspect of this project, please speak to Alena Wabitsch (alena.wabitsch@economics.ox.ac.uk), who will try to answer your query. The researcher should acknowledge your concern within 5 working days and give you an indication of how they intend to deal with it. If you remain unhappy or wish to make a formal complaint, please contact the relevant Chair of the Research Ethics Committee at the University of Oxford:

The Department of Economics Departmental Research Ethics Committee can be contacted at:
Email: ethics@economics.ox.ac.uk
Address: Department of Economics
Manor Road 10
Oxford 0X1 3UQ
United Kingdom

The University of Oxford's Data Protection Officer can be contacted at:
Email address: data.protection@admin.ox.ac.uk
Telephone number: + 44 1865 280 199
Address: Planning and Council Secretariat
University Offices
Wellington Square
Oxford OX1 2JD
United Kingdom

If you have read and understood the information above, and agree to participate with the understanding that the data (including any personal data) you submit will be processed accordingly, please check the relevant box below to get started.

○ **Yes, I have understood the above information and agree to take part.**

Please provide your Prolific ID:

Click *Next* to start the experiment.

Next

Figure A22: Instructions (1):

Instructions

Please read these instructions carefully to make good decisions during the experiment. The quality of your decisions directly impact your bonus payment, which can be up to £6, in addition to the £4 participation payment. There will be a quiz on these instructions on the next page. If you answer at least one question incorrectly three times, the experiment will be ended early.

Throughout the experiment, you can always access these instructions by clicking the *"Show Instructions"* button at the bottom of any page.

Your Goal in the Experiment:
Your task is to **accurately forecast the annual inflation rate in the Euro area**. This means, you will be predicting the percentage increase or decrease in prices in the entire Euro area between a specific point in time and exactly 1 year earlier. Your goal is to **minimise forecast error**, which is the difference between your forecasted value of annual inflation and the realised annual inflation.

*Note: The more accurate your inflation forecasts, the lower your forecast error, the **more bonus payment** you get!*

Whenever we talk about inflation in this experiment, what we mean is the *annual* inflation rate (i.e., the percentage change in prices of goods and services over the course of exactly 1 year.) Data for inflation reflect real economic scenarios of the Euro area. Sometimes you will be shown professional forecasts, which are based on analyses of real European experts for the given, real economic scenarios.

The Experiment:

You will make two types of inflation forecasts:

- **Point Forecast:** Your 'Point Forecast' of inflation is your best guess of the exact value of inflation for the next unknown period.
- **Range Forecast:** Your 'Range Forecast' of inflation is the range of plausible values that you think will almost certainly contain the realised value of inflation in the next period (i.e. the upper and the lower inflation bounds of inflation).

The closer your point forecast is to the realised inflation, the more money you get. Similarly, the smaller your indicated range of inflation, the more money you get (as long as the realised inflation rate is inside your forecasted range).

What to Expect:
You will face 6 **distinct** forecasting tasks. Each task is independent and unrelated to the others. In each task, you will first make your inflation forecasts (*'Initial Forecasts'*). You will then receive additional information about the next period's inflation rate and have a chance to adjust your forecasts (*'Updated Forecasts'*). After all 6 forecasting tasks, we will ask you a few questions, and then unveil the realised inflation rates and your performance.

What to do in each Forecasting Task:

1. Examine the given 10 periods of inflation history.
2. Make your Initial Forecasts for the next period:
   - A point forecast of inflation
   - The corresponding range forecast of inflation
3. You will be provided professional inflation forecasts, which are based on real historic forecasts of European experts for the respective scenarios. *Note: The provided historic differences between the realised inflation and the forecasted inflation represent the experts' historic forecast errors.*
4. Make your Updated Forecasts for the next period:
   - As before, provide a point forecast, and the corresponding range forecast.
   - *Note: Your Updated Forecasts can be the same as your Initial Forecasts, use some of the same values, or use completely different values.*
   - We will remind you of the values of your Initial Forecasts, when you are making your Updated Forecasts.

Overview of all the steps of a forecasting task:

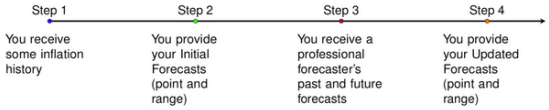


Figure A23: Instructions (2):

Interacting with the data:
Historical data, professional forecasting history, and your point and range forecasts will be visualised in an interactive chart.

- Hover over chart dots for values.
- Type all your forecasts in percentages (up to 1 decimal) into the clearly labelled input text boxes.
- You may submit positive values (prices are going up), negative values (prices are going down), or a value of zero (prices are staying the same).
- Your forecast range can be as big or small as you prefer.
- Your point forecast doesn't have to be the mid-point of your range forecast. Your point forecast may be closer or farther from the upper than the lower bound.
- Your point forecast must always lie within your forecasted range - the software will prevent other inputs.
- Your point inflation forecasts will be visualised with green squares, and the corresponding range forecasts with green triangles.
- Submit your forecasts by clicking the 'Next' button. Note: You won't be able to change your answers after submitting them.

How your decisions get you bonus payment:

You can earn up to £6 in bonus payment for the accuracy of your point and range forecasts. The full £6 bonus requires your point forecast to match realised inflation (down to 1 decimal), and your range forecast to have the upper and lower bound on the value of the realised inflation (also down to 1 decimal).

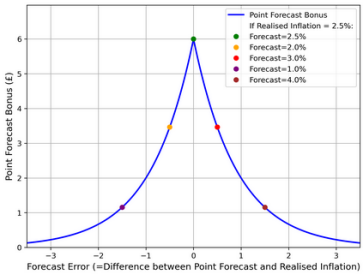
*Your bonus will be awarded for **either** your point forecast accuracy **or** your range forecast accuracy in **one** set of forecasts (Initial **or** Updated) in **one of the 6** forecasting tasks.* Here's how it works:

- Our software *randomly* selects one of your six forecasting tasks.
- For that forecasting task, the software *randomly* chooses either your Initial Forecasts or your Updated Forecasts.
- Within these selected Forecasts, the software *randomly* picks either your point forecast or your range forecast.

*Note: This means that to get a high bonus payment, you must take **both** the point and range forecasts in **all** Initial and Updated Forecasts across **all** forecasting tasks equally seriously. This is because you will only be paid for **one** randomly selected decision across the entire experiment.*

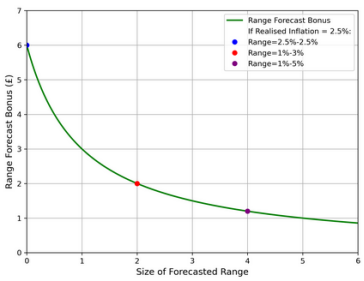
Point forecast bonus:

- Your exact forecast payment is computed like this: $P = 6 * 3^{-|\text{Forecast Error}|}$
- The larger your forecast error (distance of forecasted inflation above or below the realised inflation), the less you earn.
- The figure below shows how your point forecast bonus falls, the further away your forecast is from the realised inflation value.
- The colourful dots in the figure visualise the following example: Suppose that realised inflation turns out to be **2.5%**:
  - If your point forecast of inflation was 2.5%, you would earn $P = 6 * 3^{-0} =$ £6.00
  - If your point forecast of inflation was 2.0%, you would earn $P = 6 * 3^{-0.5} =$ £3.46
  - If your point forecast of inflation was 3.0%, you would earn $P = 6 * 3^{-0.5} =$ £3.46
  - If your point forecast of inflation was 1.0%, you would earn $P = 6 * 3^{-1.5} =$ £1.15
  - If your point forecast of inflation was 4.0%, you would earn $P = 6 * 3^{-1.5} =$ £1.15



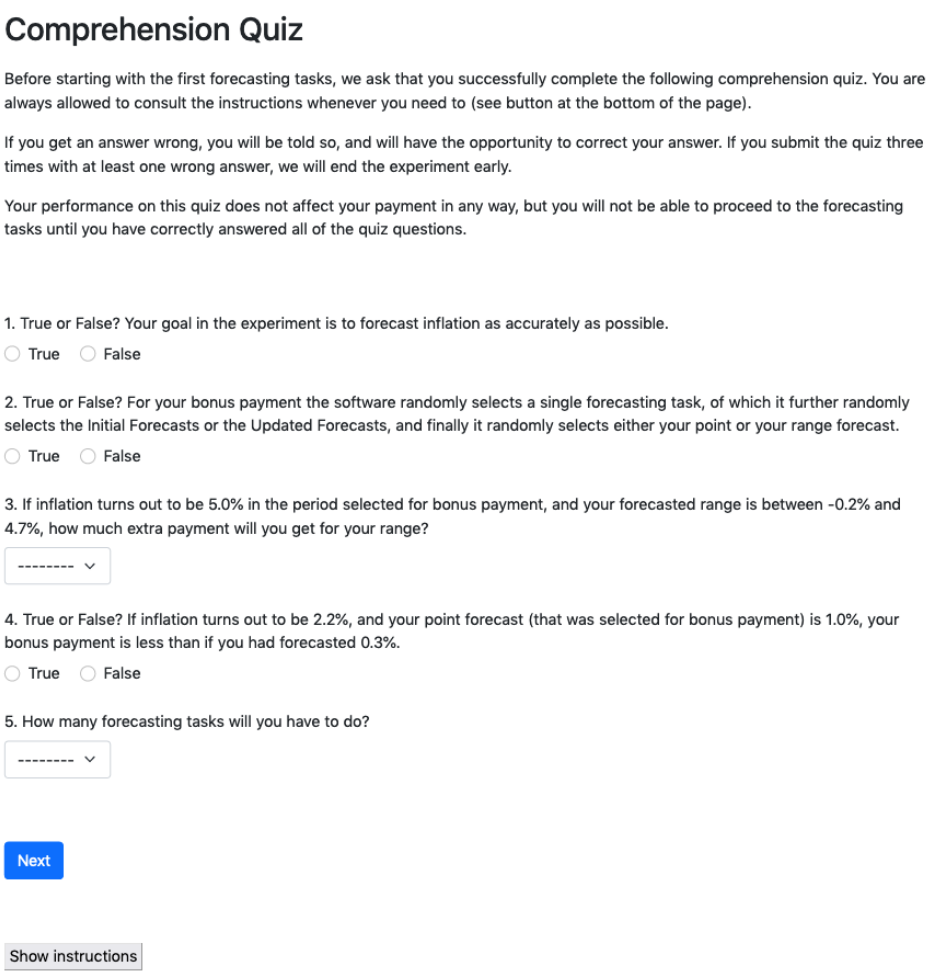
Range forecast bonus:

- If realised inflation is inside your range forecast, you score $R = 6 * [1 / (1 + \text{total range})]$
- The total range of your forecast is given by the gap between the upper bound of the range forecast and the lower bound of the range forecast.
- The larger the range you create, the less money you earn, when inflation is within your range forecast.
- If inflation falls **outside** your forecast range, you earn **nothing**
- The figure below shows how your range forecast bonus falls, the larger your forecasted range is (as long as the realised value of inflation falls within that range).
- The colourful dots in the figure visualise the following example: Suppose that realised inflation turns out to be **2.5%**:
  - If your forecasted range was between 2.5% and 2.5%, then you would earn $R = 6 * [1 / (1 + 0)] =$ £6.00
  - If your forecasted range was between 1% and 3%, you would earn $R = 6 * [1 / (1 + 2)] =$ £2.00
  - If your forecasted range was between 1% and 5%, you would earn $R = 6 * [1 / (1 + 4)] =$ £1.20
  - If your forecasted range was between 0% and 2%, you would earn nothing because realised inflation is not within your range.



Note that the *"Next"* button may appear with a delay on certain pages. Similarly, parts of certain pages will be displayed only after a few seconds. Go through the experiment at your own pace and don't rush any decision.

Figure A24: Comprehension Quiz:



Comprehension Quiz

Before starting with the first forecasting tasks, we ask that you successfully complete the following comprehension quiz. You are always allowed to consult the instructions whenever you need to (see button at the bottom of the page).

If you get an answer wrong, you will be told so, and will have the opportunity to correct your answer. If you submit the quiz three times with at least one wrong answer, we will end the experiment early.

Your performance on this quiz does not affect your payment in any way, but you will not be able to proceed to the forecasting tasks until you have correctly answered all of the quiz questions.

1. True or False? Your goal in the experiment is to forecast inflation as accurately as possible.

○ True ○ False

2. True or False? For your bonus payment the software randomly selects a single forecasting task, of which it further randomly selects the Initial Forecasts or the Updated Forecasts, and finally it randomly selects either your point or your range forecast.

○ True ○ False

3. If inflation turns out to be 5.0% in the period selected for bonus payment, and your forecasted range is between -0.2% and 4.7%, how much extra payment will you get for your range?

----

4. True or False? If inflation turns out to be 2.2%, and your point forecast (that was selected for bonus payment) is 1.0%, your bonus payment is less than if you had forecasted 0.3%.

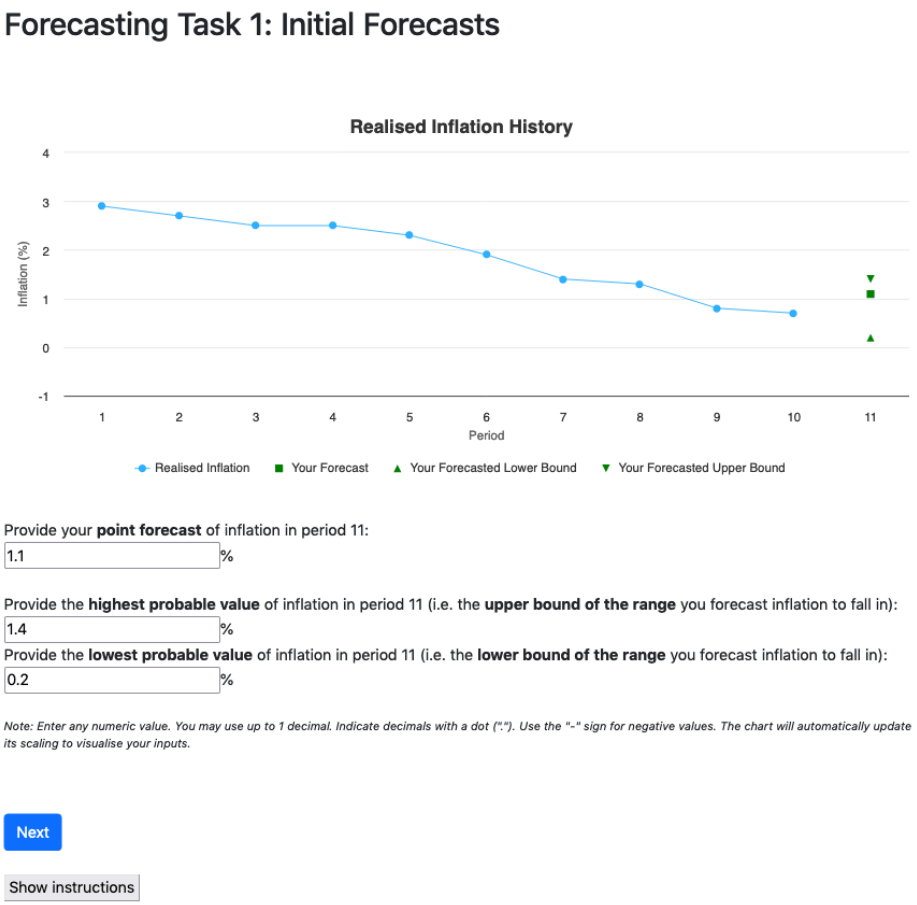
○ True ○ False

5. How many forecasting tasks will you have to do?

Next

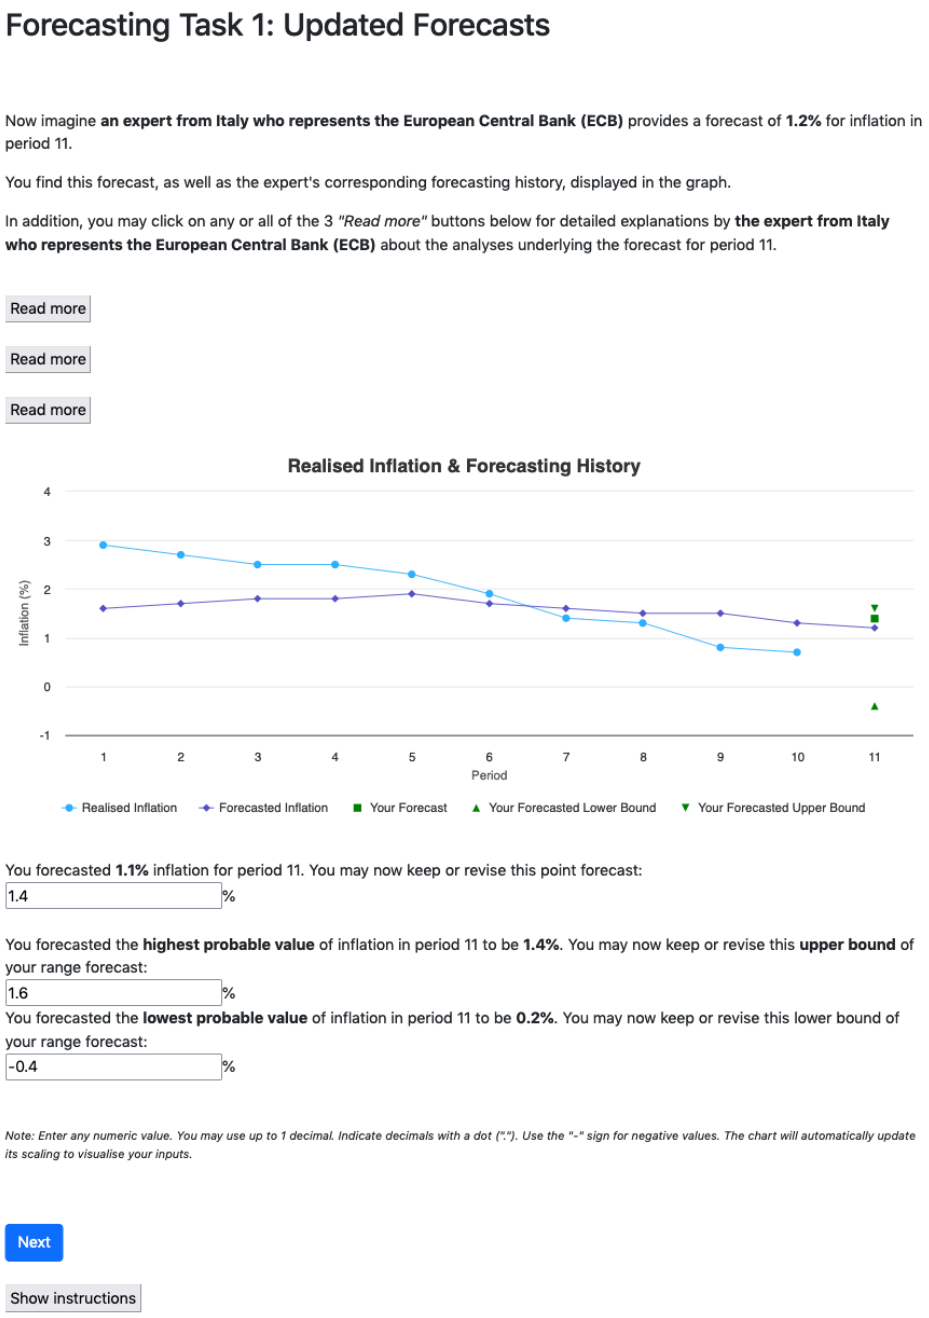
Show instructions

Figure A25: Prior elicitation:



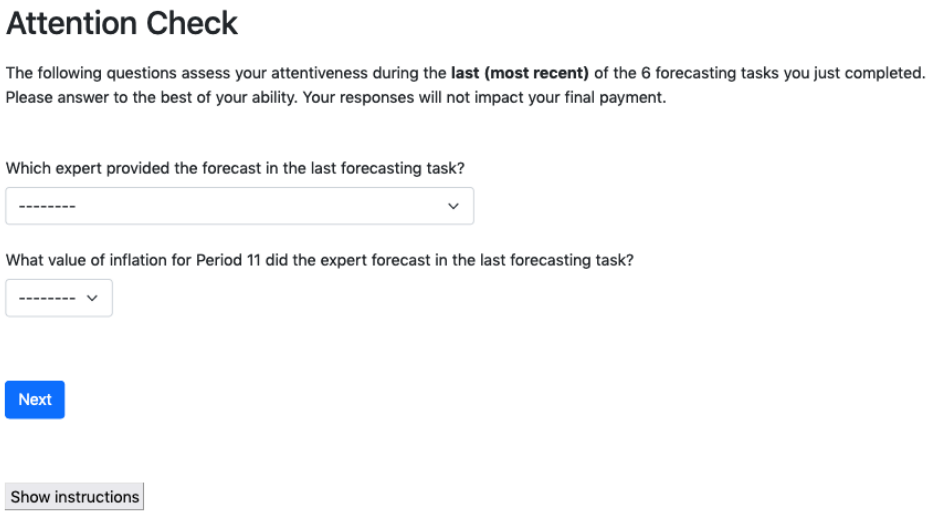
Notes: Participants are forced to spend at least 10 seconds on this page to reduce the risk of rushing decisions without paying attention to the provided information.

Figure A26: Posterior elicitation:



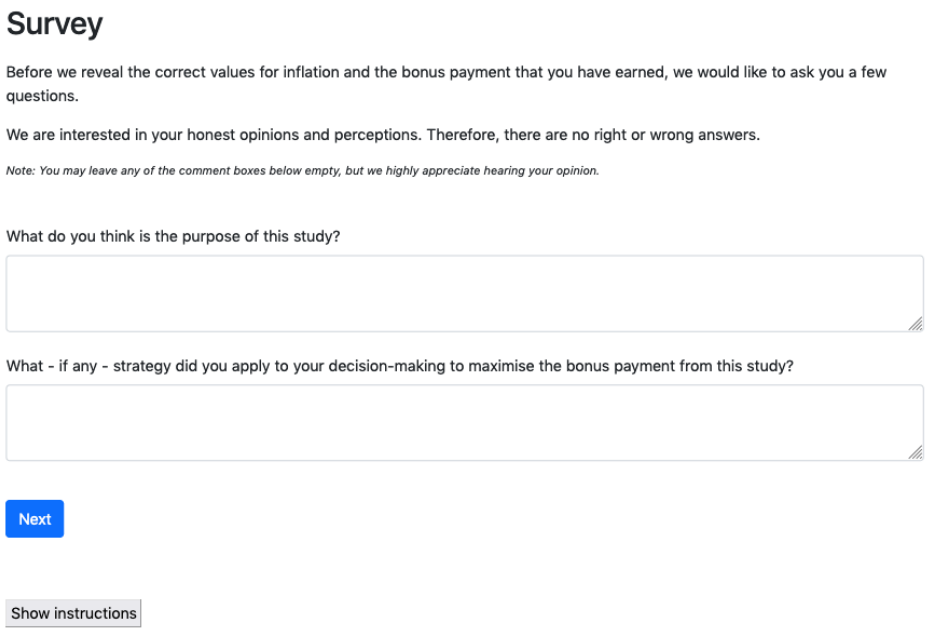
Notes: The first sentence of this page is shown in isolation for 5 seconds before the rest of the page appears to increase reporting on the messenger and signal. Participants are then forced to spend at least another 10 seconds on this page to reduce the risk of rushing decisions without paying attention to the provided information. Hovering over the dots in the graph reveals exact values of realized inflation and inflation forecasts.

Figure A27: Attention Check



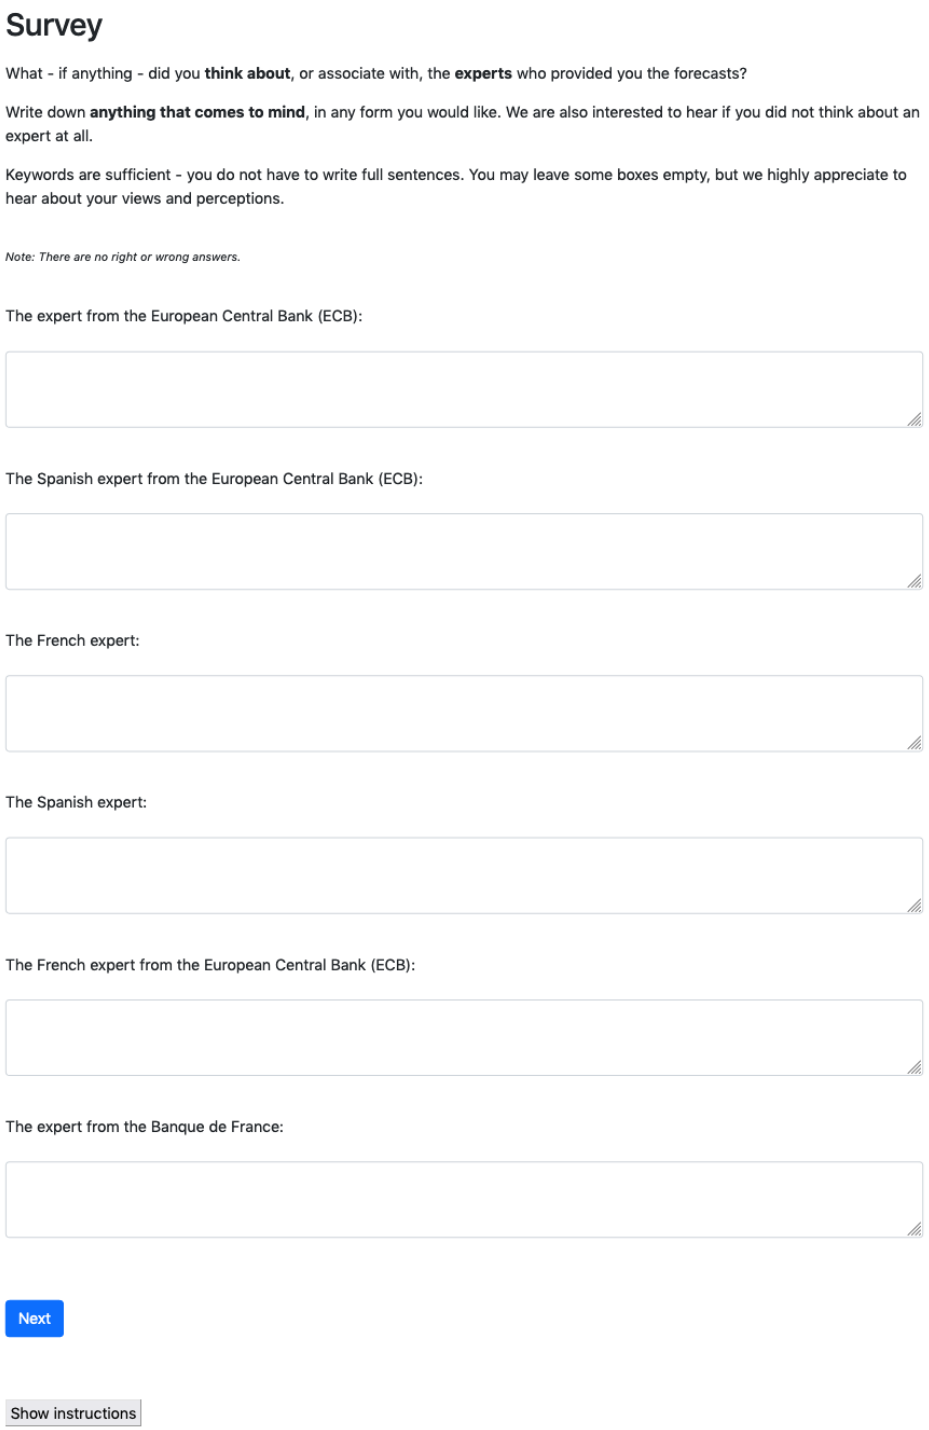
Notes: The order of answer options for experts is randomized. The options for period 11 values reflect actual signal values of the 6 forecasting tasks.

Figure A28: Survey: Purpose of Experiment & Strategy



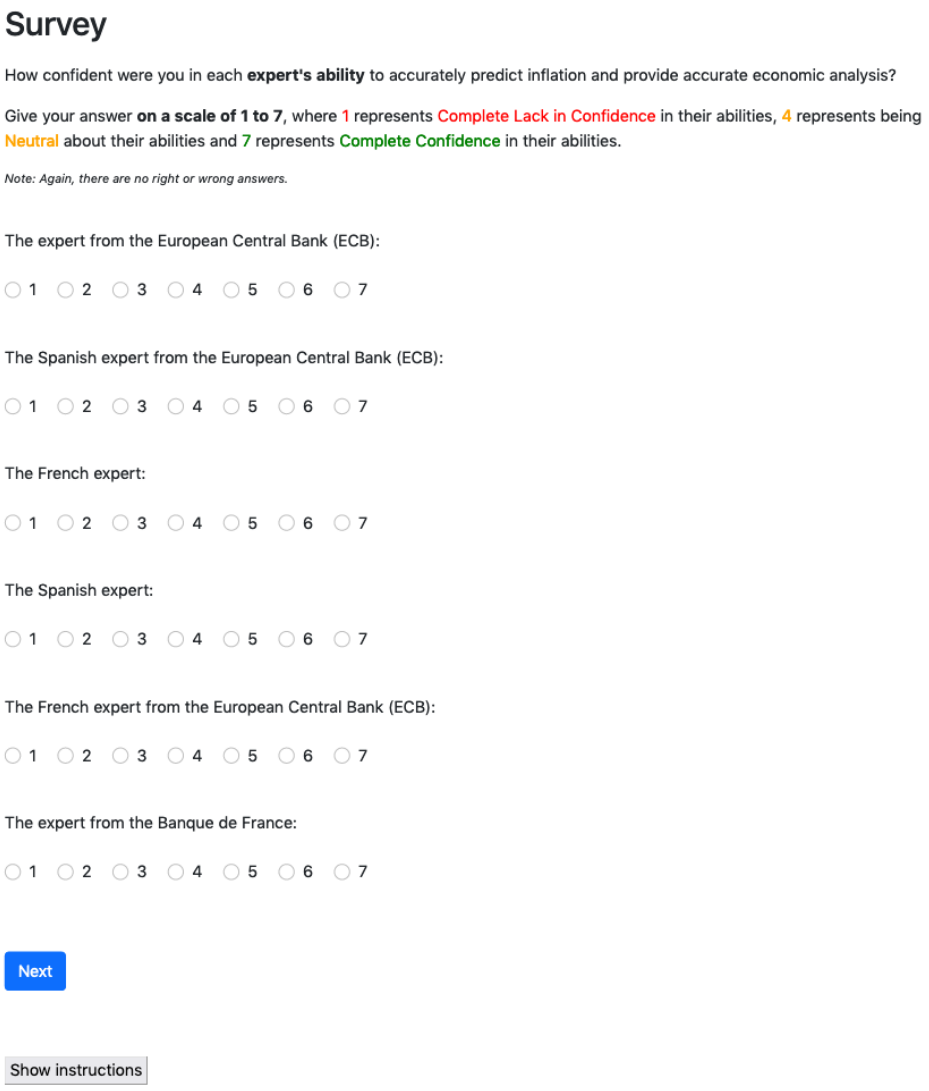
Notes: Participants are allowed to skip any or all of these questions entirely.

Figure A29: Survey: Association



Notes: Participants are allowed to skip any or all of these questions entirely. To make it easier for participants to recall, the order of experts represents the occurrence of experts in forecasting tasks (i.e., randomized across individuals).

Figure A30: Survey: Ability



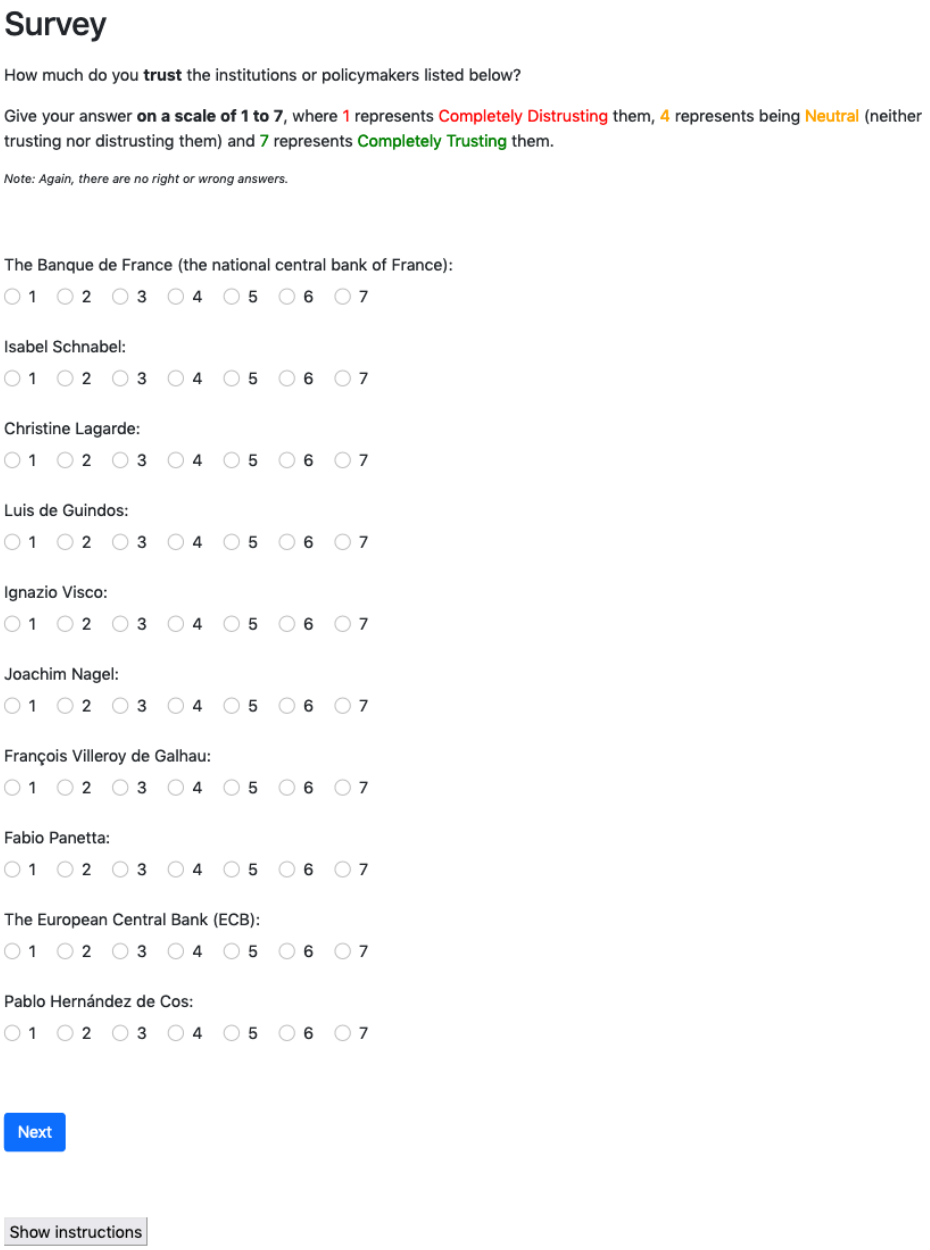
Notes: Order of experts represents occurrence of experts in forecasting tasks (i.e., randomized across individuals).

Figure A31: Survey: Exposure



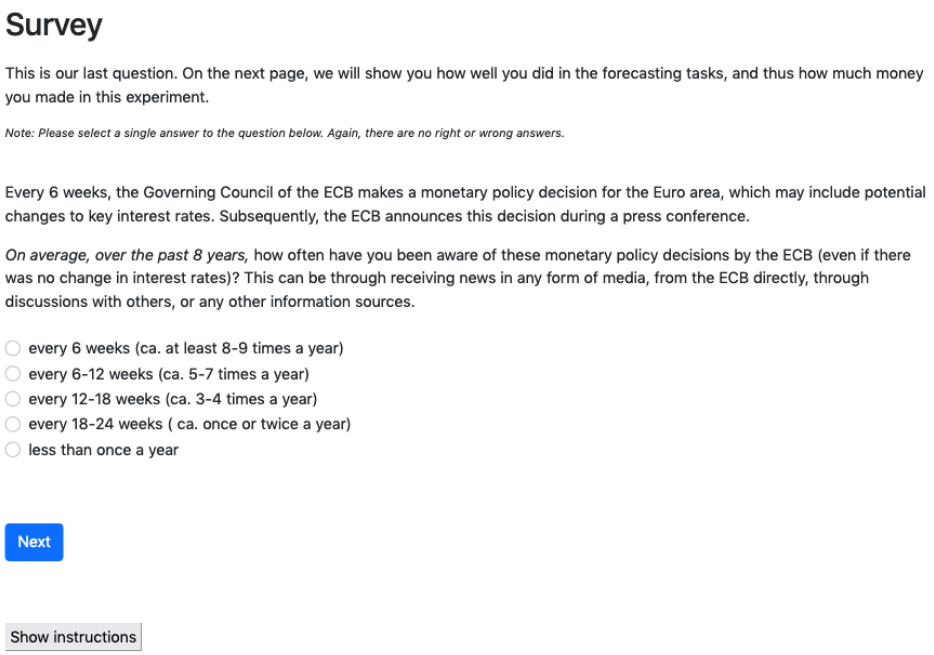
Notes: The order in which institutions and policymakers appear is randomized at the individual-level.

Figure A32: Survey: Trust



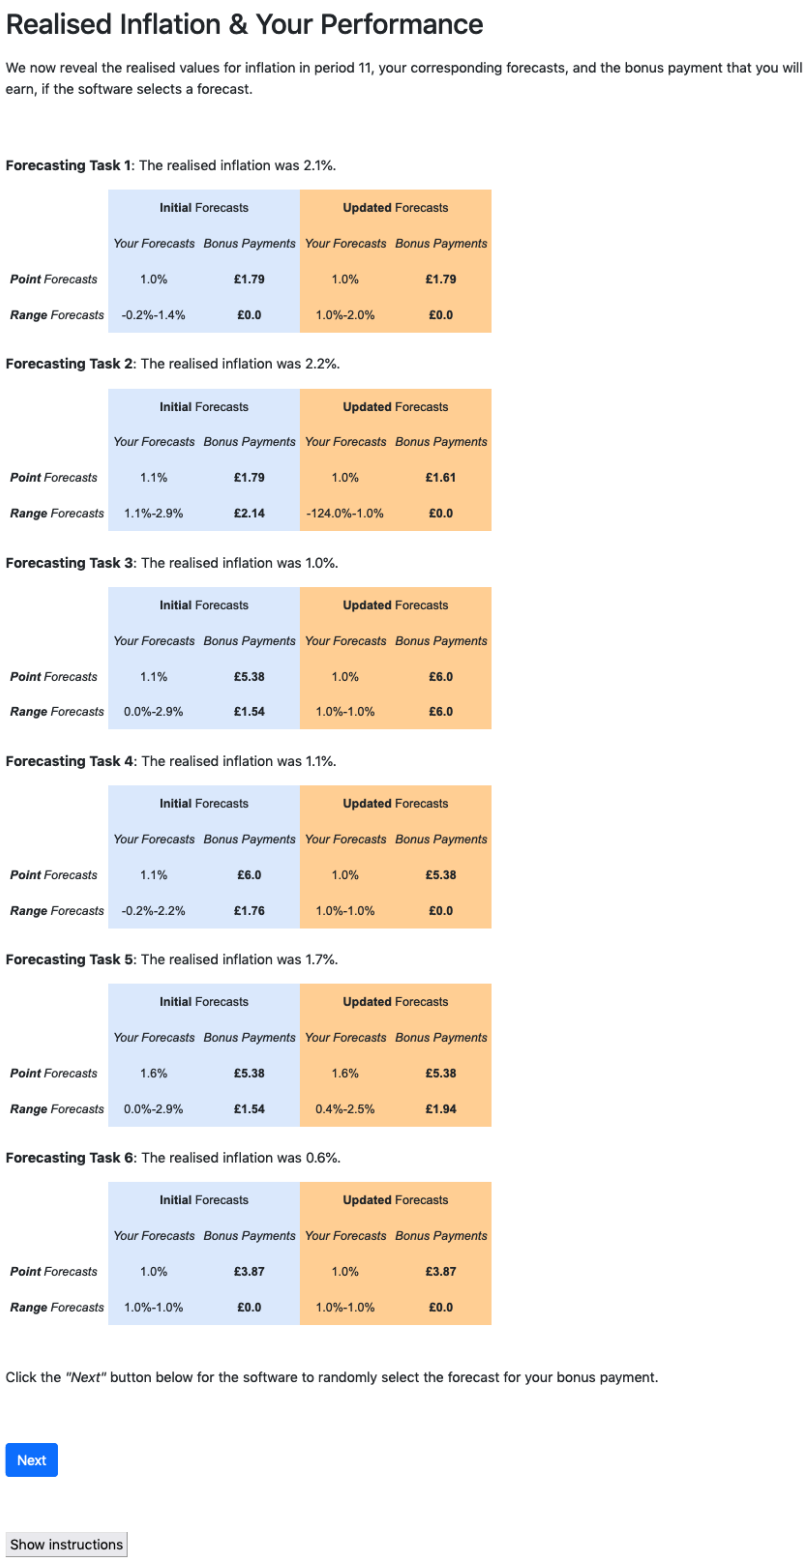
Notes: This page is only shown to participants who indicate knowing at least one institution of policymaker. Only institutions and policymakers that the participant indicates knowing appear in this list, as indicated trust is only meaningful if participants know the institutions and policymakers. The order in which institutions and policymakers appear is randomized at the individual-level.

Figure A33: Survey: Monetary Policy Expertise



Notes: The question is allows for potential comparison of experimental and Twitter results, distinguishing between experts and non-experts. With some caveats, answers to this questions can further be seen as an imperfect proxy for monetary policy expertise.

Figure A34: Payoff Overview



Realised Inflation & Your Performance

We now reveal the realised values for inflation in period 11, your corresponding forecasts, and the bonus payment that you will earn, if the software selects a forecast.

Forecasting Task 1: The realised inflation was 2.1%.

	Initial Forecasts		**Updated** Forecasts	
	Your Forecasts	*Bonus Payments*	*Your Forecasts*	*Bonus Payments*
***Point** Forecasts*	1.0%	**£1.79**	1.0%	**£1.79**
***Range** Forecasts*	-0.2%-1.4%	**£0.0**	1.0%-2.0%	**£0.0**

Forecasting Task 2: The realised inflation was 2.2%.

	Initial Forecasts		**Updated** Forecasts	
	Your Forecasts	*Bonus Payments*	*Your Forecasts*	*Bonus Payments*
***Point** Forecasts*	1.1%	**£1.79**	1.0%	**£1.61**
***Range** Forecasts*	1.1%-2.9%	**£2.14**	-124.0%-1.0%	**£0.0**

Forecasting Task 3: The realised inflation was 1.0%.

	Initial Forecasts		**Updated** Forecasts	
	Your Forecasts	*Bonus Payments*	*Your Forecasts*	*Bonus Payments*
***Point** Forecasts*	1.1%	**£5.38**	1.0%	**£6.0**
***Range** Forecasts*	0.0%-2.9%	**£1.54**	1.0%-1.0%	**£6.0**

Forecasting Task 4: The realised inflation was 1.1%.

	Initial Forecasts		**Updated** Forecasts	
	Your Forecasts	*Bonus Payments*	*Your Forecasts*	*Bonus Payments*
***Point** Forecasts*	1.1%	**£6.0**	1.0%	**£5.38**
***Range** Forecasts*	-0.2%-2.2%	**£1.76**	1.0%-1.0%	**£0.0**

Forecasting Task 5: The realised inflation was 1.7%.

	Initial Forecasts		**Updated** Forecasts	
	Your Forecasts	*Bonus Payments*	*Your Forecasts*	*Bonus Payments*
***Point** Forecasts*	1.6%	**£5.38**	1.6%	**£5.38**
***Range** Forecasts*	0.0%-2.9%	**£1.54**	0.4%-2.5%	**£1.94**

Forecasting Task 6: The realised inflation was 0.6%.

	Initial Forecasts		**Updated** Forecasts	
	Your Forecasts	*Bonus Payments*	*Your Forecasts*	*Bonus Payments*
***Point** Forecasts*	1.0%	**£3.87**	1.0%	**£3.87**
***Range** Forecasts*	1.0%-1.0%	**£0.0**	1.0%-1.0%	**£0.0**

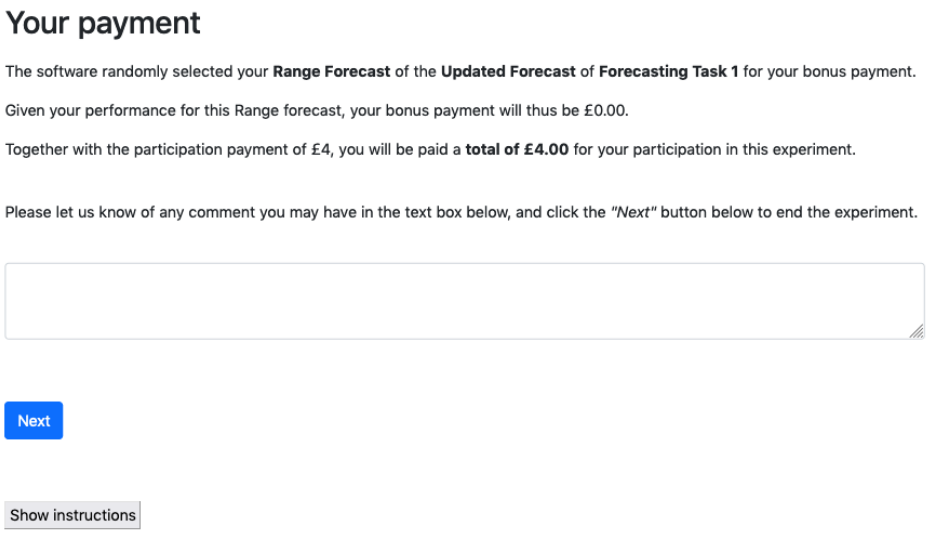
Click the *"Next"* button below for the software to randomly select the forecast for your bonus payment.

Next

Show instructions

Notes: Realized inflation values for period 11 and the corresponding performance of participants are only shown at the very end of the experiment and after all survey questions to avoid learning effects between forecasting tasks and performance-driven biases to survey answers asked after the experiment.

Figure A35: Payoff Reveal



Notes: The software reveals the randomly selected a decision for each participant's bonus payment.

D Model Details

D.1 Threshold for Public Signal Availability

To model the variability in agents' thresholds, I use the tractable and convenient truncated normal distribution. Each outgroup agent's threshold d_o follows a truncated normal distribution $\mathcal{N}_+(0,1)$ on $[0,\infty)$, which ensures thresholds are non-negative. The probability density function (PDF) of d_o is:

$$f_{d_o}(d) = \frac{\phi(d)}{\Phi(\infty) - \Phi(0)} = 2\phi(d), \quad \text{for } d \geq 0,$$

where $\phi(d)$ and $\Phi(d)$ are the standard normal PDF and cumulative distribution function (CDF), respectively. The cumulative distribution function (CDF) of d_o is:

$$F_{d_o}(d) = \int_0^d f_{d_o}(z)\, dz = 2\Phi(d) - 1.$$

This CDF determines the fraction of outgroup agents who observe Y based on their individual thresholds.

D.2 Proofs and Derivations

D.2.1 Proof of Lemma 1: Optimal Signal Extraction Weights

Lemma 1 claims that the optimal weights κ_h for ingroup and informed outgroup agents are given by:

$$\kappa_g^* = \frac{\tau_y q}{\tau_Y + \tau_y q} \quad \text{and} \quad \kappa_o^* = \frac{\tau_y q}{\rho_{im}\tau_Y + \tau_y q},$$

where $q = (1 - r + r(1-\alpha)(1-A))$.

Below outlines the corresponding proof of these optimal signal extraction weights.

Uninformed Outgroup Agents An uninformed outgroup agent simply chooses $a_{io}(y_i) = y_i$. Since Y is unobserved by this agent, he relies on $E[Y] = x$, so he uses $E_{io}[x \mid y_i] = y_i$ for any Y.

To derive the action of informed agents, suppose that the population, consisting of the different agent types h, follows a linear strategy of the form:

$$a_{jh} = \kappa_h y_j + (1 - \kappa_h)Y. \tag{29}$$

Ingroup Agents An ingroup agent g, who always receives the public signal, updates without bias and expects all agents in the economy to update fully Bayesian. She assumes that both ingroup agents and informed outgroup agents who receive the public signal Y use the same signal extraction weight κ_g. While she knows that uninformed outgroup agents rely solely on their private signals, she believes that informed outgroup agents process the public signal in the same unbiased way as she does. In other words, the ingroup agent is unaware of any biases or differences in information processing among other agents and expects all agents to update their beliefs and choose actions just like herself, given their information sets. Knowing the fraction of outgroup agents A that receive the public signal and the overall share of outgroup agents $(1-\alpha)$, her expected average action of the population is:

$$\begin{aligned}
E_{ig}[\overline{a} \mid y_i, Y] &= \left[\alpha\kappa_g + (1-\alpha)A\kappa_g + (1-\alpha)(1-A)\right] E_{ig}[y_j \mid y_i, Y] \\
&\quad + \left[\alpha(1-\kappa_g) + (1-\alpha)A(1-\kappa_g)\right] Y \\
&= \left[\alpha\kappa_g + (1-\alpha)A\kappa_g + (1-\alpha)(1-A)\right] \frac{\tau_y y_i + \tau_Y Y}{\tau_y + \tau_Y} \\
&\quad + \left[\alpha(1-\kappa_g) + (1-\alpha)A(1-\kappa_g)\right] Y \\
&= Y\left[(1-\kappa_g)\left[\alpha + (1-\alpha)A\right] + \frac{\tau_Y}{\tau_y+\tau_Y}\left[\kappa_g\left(\alpha + (1-\alpha)A\right) + (1-\alpha)(1-A)\right]\right] \\
&\quad + y_i\left[\frac{\tau_y}{\tau_y+\tau_Y}\left[\kappa_g\left(\alpha + (1-\alpha)A\right) + (1-\alpha)(1-A)\right]\right].
\end{aligned} \tag{30}$$

Thus, the ingroup agent g's optimal action is given by:

$$\begin{aligned}
a_{ig}(y_i, Y) &= (1-r)E_{ig}[x \mid y_i, Y] + rE_{ig}[\overline{a} \mid y_i, Y] \\
&= r\left\{Y\left[(1-\kappa_g)\left(\alpha + (1-\alpha)A\right) + \frac{\tau_Y}{\tau_y+\tau_Y}\left[\kappa_g\left(\alpha + (1-\alpha)A\right) + (1-\alpha)(1-A)\right]\right]\right. \\
&\quad \left. + y_i\left[\frac{\tau_y}{\tau_y+\tau_Y}\left[\kappa_g\left(\alpha + (1-\alpha)A\right) + (1-\alpha)(1-A)\right]\right]\right\} + (1-r)\frac{\tau_y y_i + \tau_Y Y}{\tau_y + \tau_Y} \\
&= \left[\frac{\tau_y}{\tau_y+\tau_Y}(1 - r + r[\kappa_g(\alpha + (1-\alpha)A) + (1-\alpha)(1-A)]\right] y_i \\
&\quad + \left[1 + \frac{\tau_y}{\tau_y+\tau_Y}\left[r[-\kappa_g(\alpha + (1-\alpha)A) + \alpha + (1-\alpha)A] - 1\right]\right] Y.
\end{aligned} \tag{31}$$

Comparing coefficients with the population's strategy (Eq. (29)) gives:

$$\kappa_g = \frac{\tau_y}{\tau_y + \tau_Y}\left[1 - r + r\left[\kappa_g\left(\alpha + (1-\alpha)A\right) + (1-\alpha)(1-A)\right]\right], \tag{32}$$

from which I can solve for κ_g:

$$\kappa_g* = \frac{\tau_y q}{\tau_Y + \tau_y q}. \tag{33}$$

Informed Outgroup Agents The informed outgroup agent o—who receives both signals but updates with a bias—looks very similar to the ingroup agent but down-weights the public signal by the resonance weight ρ_{im}. His expected average action of the population is:

$$\begin{aligned}
E_{io}[\overline{a} \mid y_i, Y] &= [\alpha\kappa_o + (1-\alpha)A\kappa_o + (1-\alpha)(1-A)]\, E_{io}[y_j \mid y_i, Y] \\
&\quad + [\alpha(1-\kappa_o) + (1-\alpha)A(1-\kappa_o)]\, Y \\
&= \alpha[\kappa_o E_{io}[y_j \mid y_i, Y] + (1-\kappa_o)Y] + (1-\alpha)A[\kappa_o E_{io}[y_j \mid y_i, Y] + (1-\kappa_o)Y] \\
&\quad + (1-\alpha)(1-A)E_{io}[y_j \mid y_i, Y] \\
&= Y\left[(1-\kappa_o)\left[\alpha + (1-\alpha)A\right] + \frac{\rho_{im}\tau_Y}{\tau_y + \rho_{im}\tau_Y}\left[\kappa_o\left(\alpha + (1-\alpha)A\right) + (1-\alpha)(1-A)\right]\right] \\
&\quad + y_i\left[\frac{\tau_y}{\tau_y + \rho_{im}\tau_Y}\left[\kappa_o\left(\alpha + (1-\alpha)A\right) + (1-\alpha)(1-A)\right]\right].
\end{aligned} \tag{34}$$

Thus, the informed outgroup agent o's optimal action is given by:

$$\begin{aligned}
a_{io}(y_i, Y) &= (1-r)E_{io}[x \mid y_i, Y] + rE_{io}[\overline{a} \mid y_i, Y] \\
&= \left[\frac{\tau_y}{\tau_y + \rho_{im}\tau_Y}(1 - r + r[\kappa_o(\alpha + (1-\alpha)A) + (1-\alpha)(1-A)]\right] y_i \\
&\quad + \left[1 + \frac{\tau_y}{\tau_y + \rho_{im}\tau_Y}\left[r[-\kappa_o(\alpha + (1-\alpha)A) + \alpha + (1-\alpha)A] - 1\right]\right] Y.
\end{aligned} \tag{35}$$

Comparing coefficients with the population's strategy (Eq. (29)), I get for the informed outgroup agent:

$$\kappa_o = \frac{\tau_y}{\tau_y + \rho_{im}\tau_Y}\left[1 - r + r\left[\kappa_o\left(\alpha + (1-\alpha)A\right) + (1-\alpha)(1-A)\right]\right], \tag{36}$$

from which I can solve for κ_o:

$$\kappa_o* = \frac{\tau_y q}{\rho_{im}\tau_Y + \tau_y q}. \tag{37}$$

Therefore, the optimal weights κ_h for ingroup and informed outgroup agents are as stated in Lemma 1.

D.2.2 Proof of Proposition 1: Uniqueness of Linear Equilibrium

Proposition 1 claims the equilibrium strategies derived constitute the unique linear equilibrium of the game.

Proof: To prove that the equilibrium strategy is indeed a linear combination of the available signals, and that the equilibrium is therefore unique, I extend the proof by Morris and Shin (2002) to accommodate heterogeneous agent types with biased information processing. The optimal action of agents is given by Eq. (13), and the average action is given in Eq. (21), which can also be written as:

$$\begin{aligned}\bar{a} &= \int_0^{\alpha} a_{ig}(y_i, Y)\, di + A \int_{\alpha}^{1} a_{io}(y_i, Y)\, di + (1-A) \int_{\alpha}^{1} a_{io}(y_i)\, di \\ &= \alpha \bar{a}_{ig}(x, Y) + (1-\alpha) A \bar{a}_{io}(x, Y) + (1-\alpha)(1-A)x.\end{aligned} \tag{38}$$

Integrating over the interval $[0, \alpha]$ for ingroup agents and $[\alpha, 1]$ for outgroup agents returns the average action as a weighted sum of the expected actions of each group, given the symmetry and independence of agents within each group. The integrals over outgroup agents are weighted by A and $(1-A)$, which represent the fractions of informed and uninformed agents, respectively. Remember that agents expect other informed agents to update and act just like themselves, as agents are naive about their own biases and those of others. This means that ingroup agents assume the average action of informed agents is $\bar{a}_i(x, Y) = \bar{a}_{ig}(x, Y)$, while informed outgroup agents assume $\bar{a}_i(x, Y) = \bar{a}_{io}(x, Y)$ across all informed agents. The average action of uninformed outgroup agents is known by all and is $\bar{a}_{io}(x) = x$, as defined earlier. Hence, the optimal strategy of any informed agent of type h is:

$$\begin{aligned}a_{ih}(y_i, Y) &= qE_{ih}[x|y_i, Y] + [r(\alpha + (1-\alpha)A)]E_{ih}[\bar{a}|y_i, Y] \\ &= qE_{ih}[x|y_i, Y] + [r(\alpha + (1-\alpha)A)]qE_{ih}[\bar{E}(x)] + [r(\alpha + (1-\alpha)A)]^2 E_{ih}[\bar{E}[\bar{a}(x, Y)]], \quad j \neq i \\ &= q \sum_{s=0}^{\infty} \left(r(\alpha + (1-\alpha)A)\right)^s E_{ih}\left[\bar{E}_i^s[x|y_i, Y]\right].\end{aligned} \tag{39}$$

$\bar{E}_i^s[x]$ denotes the average expectation of the average expectation (s-times) of x, and $\bar{E}_i^s[\bar{a}]$ is the average expectation of the average expectation (s-times) of the average action, where all averages concern informed agents only and are dependent on agent type h. The infinite sum captures the recursiveness of higher-order beliefs, where agents form expectations about the expectations of others in this strategic interaction setting. $E_{ih}\left[\bar{E}_i^s[x|y_i, Y]\right]$ is a linear combination of y_i and Y.[51] Thus, $a_{ih}(y_i, Y)$ is also a linear combination of these signals. By

[51] The proof by Morris and Shin (2002) is directly applicable to ingroup agents. It extends to expectation formation with the resonance weight ρ_{im} of informed outgroup agents as the recursive structure and the linearity of the expectation function (Eq. (9)) are preserved. The updating bias merely reduces the influence of the public signal, and thus shifts the coefficients in the explicit linear expression slightly.

establishing that the optimal actions of agents are linear combinations of their signals, and that these strategies lead to consistent expectations about others' actions, it is shown that the linear solution constitutes the unique equilibrium.

D.2.3 Private Signal Precision May Harm Welfare

Proposition 4. *The effect of increasing the precision of private signals (τ_y) on expected welfare is ambiguous and depends on the coordination parameter r and the presence of outgroup agents ($\frac{\partial E(W|x)}{\partial \tau_y} \gtrless 0$).*

Derivation: Taking the derivative of $E[W(a,x)|x]$ with respect to τ_y shows that $\frac{\partial E[W|x]}{\partial \tau_y} \gtrless 0$ and depends on the particular parameter values.

$$\begin{aligned}\frac{\partial E(W \mid x)}{\partial \tau_y} = &-\frac{\alpha q^2}{\left(\tau_Y + \tau_y q\right)^2} - \frac{(1-\alpha)Aq^2}{\left(\rho_{im}\tau_Y + \tau_y q\right)^2} \\ &+ \frac{2Aq(1-\alpha)\left(\tau_y q^2 + \rho_{im}^2 \tau_Y\right)}{\left(\rho_{im}\tau_Y + \tau_y q\right)^3} \\ &+ \frac{2q\alpha\left(\tau_y q\right)^2 + \tau_Y\right)}{\left(\tau_Y + \tau_y q\right)^3} + \frac{(1-\alpha)(1-A)}{\tau_y^2}\end{aligned} \tag{40}$$

At first glance, more precise private signals should enhance welfare by improving alignment with x. However, due to biased information processing and strategic complementarities, increasing τ_y can reduce welfare. Outgroup agents, unaware of their bias, overestimate others' reliance on private signals and overemphasize their own private information to coordinate. This leads to misalignment with the fundamental and increased action dispersion, especially when strategic complementarities are strong and the outgroup's updating bias is severe. The resonance weight required for negative welfare effects to materialize is much lower than what I find empirically.[52]

This finding contrasts with Morris and Shin (2002) and even Cornand and Heinemann (2008), who also restrict public signal availability, but aligns with other work such as Hellwig (2005) and Angeletos and Pavan (2007), finding that better private information increases reliance on individual signals, worsening coordination and increasing action dispersion. Specifically, Hellwig (2005) introduces heterogeneous private information, and Angeletos and Pavan (2007) show that welfare effects depend on the relationship between equilibrium and efficient coordination degrees.

[52]Specifically, $\rho_{im} \lesssim 0.3$ is needed. This occurs for $r \gtrsim 0.5$, equally precise public and private information, and ensuring at least half of agents are outgroup agents who are almost surely informed. For other combinations negative welfare effects are even less of a concern.

Corollary 2. *When there are significant biases in outgroup agents' information processing, increasing τ_y may reduce social welfare.*

D.2.4 Derivation for Proposition 2: Welfare Implications of the Public Signal Precision

Taking the derivative of $E[W(a,x)|x]$ with respect to τ_Y shows that $\frac{\partial E[W|x]}{\partial \tau_Y} \gtrless 0$ and depends on the particular parameter values.

$$\begin{aligned}\frac{\partial E(W \mid x)}{\partial \tau_Y} = &\frac{2\alpha\left(\tau_Y + \tau_y q^2\right)}{\left(\tau_Y + \tau_y q\right)^3} - \frac{\alpha}{\left(\tau_Y + \tau_y q\right)^2} \\ &+ \frac{2A(1-\alpha)\rho_{im}\left(\rho_{im}^2\tau_Y + \tau_y q^2\right)}{\left(\rho_{im}\tau_Y + \tau_y q\right)^3} - \frac{A(1-\alpha)\rho_{im}^2}{\left(\rho_{im}\tau_Y + \tau_y q\right)^2}\end{aligned} \tag{41}$$

D.2.5 Derivation for Proposition 3: The Optimal Share of Ingroup Agents

The optimal α^* is found by setting $\frac{\partial E[W|x]}{\partial \alpha} = 0$ and solving for α. Complex non-linearities makes an explicit analytical solution for α^* unfeasible. The Weierstrass Extreme Value Theorem guarantees the existence of a welfare-maximizing α^*, given that the expected welfare function $E[W(a,x)|x]$ is continuous in α and the domain $[0,1]$ is compact. α^* is either a corner solution (0 or 1) or is obtained by setting the first-order condition to zero:

$$\begin{aligned}\frac{\partial E(W \mid x)}{\partial \alpha} = &\frac{\alpha 2 r \tau_y (1-A) q}{\left(\tau_Y + \tau_y q\right)^2} - \frac{\alpha 2 r \tau_y (1-A)\left(\tau_Y + \tau_y q^2\right)}{\left(\tau_Y + \tau_y q\right)^3} + \frac{(1-\alpha) 2A (1-A) r \tau_y q}{\left(\rho_{im}\tau_Y + \tau_y q\right)^2} \\ &- \frac{(1-\alpha) 2A (1-A) r \tau_y \left(\tau_y q^2 + \rho_{im}^2 \tau_Y\right)}{\left(\rho_{im}\tau_Y + \tau_y q\right)^3} + \frac{A\left(\tau_y q^2 + \rho_{im}^2 \tau_Y\right)}{\left(\rho_{im}\tau_Y + \tau_y q\right)^2} - \frac{\tau_Y + \tau_y q^2}{\left(\tau_Y + \tau_y q\right)^2} + \frac{1-A}{\tau_y} = 0.\end{aligned}$$

D.3 Details on Expertise Loss

Proposition 5. *Delegating communication to increase α can be welfare-improving if the marginal benefit outweighs the marginal cost from a reduction in public signal precision due to the potential lower expertise of the delegated messengers: $\frac{\partial E[W|x]}{\partial \alpha} > k\frac{\partial E[W|x]}{\partial \tau_Y}$.*

For delegation to be welfare-improving, the marginal benefit must outweigh the marginal cost: $\frac{\partial E[W|x]}{\partial \alpha} > k\frac{\partial E[W|x]}{\partial \tau_Y}$, where k represents the rate at which precision decreases as α increases. This occurs when either the direct improvement in agents' alignment with the fundamental (due to more ingroup agents) is significant, or when the loss in public signal precision (τ_Y) per unit increase in α (represented by k) is small, minimizing the marginal cost.

D.3.1 Proof Sketch of Proposition 5: Delegating Communication vs. Loss in Expertise

The trade-off is formalized by comparing the marginal benefit of increasing α with the marginal cost of reduced τ_Y: Suppose that increasing α (the share of ingroup agents) comes at the cost of reducing τ_Y (the precision of the public signal). This is because delegating communication to messengers who better match the characteristics of more agents might involve using messengers who are less expert or less effective communicators, thereby reducing the quality (precision) of the public signal.

Assume that τ_Y is a decreasing function of α:

$$\tau_Y = \tau_Y(\alpha), \quad \text{with} \quad \frac{d\tau_Y}{d\alpha} \leq 0.$$

Recall that expected social welfare is given by:

$$E[W(a,x)|x] = -\alpha \frac{\tau_Y + \tau_y q^2}{\left[\tau_Y + \tau_y q\right]^2} - (1-\alpha)A \frac{\rho_{im}^2 \tau_Y + \tau_y q^2}{\left[\rho_{im}\tau_Y + \tau_y q\right]^2} - \frac{(1-\alpha)(1-A)}{\tau_y}.$$

To find the condition under which increasing α leads to an increase in expected welfare—considering the dependence of τ_Y on α—I take the total derivative of expected welfare with respect to α is:

$$\frac{dE[W|x]}{d\alpha} = \frac{\partial E[W|x]}{\partial \alpha} + \frac{\partial E[W|x]}{\partial \tau_Y} \cdot \frac{d\tau_Y}{d\alpha}.$$

The goal is to determine whether $\frac{dE[W|x]}{d\alpha} \leq 0$ (welfare increases as α increases, since welfare is defined as a negative value of losses).

Eq. (D.2.5) shows the expression for $\frac{\partial E[W|x]}{\partial \alpha}$. This partial derivative captures the direct effect of changing α on expected welfare, holding τ_Y constant, and represents the net change in welfare due to reallocating a marginal agent from the outgroup to the ingroup.

The expression for $\frac{\partial E[W|x]}{\partial \tau_Y}$ is shown in Eq. (41). This partial derivative captures how changes in the precision of the public signal affect expected welfare, holding α constant.

To compute $\frac{d\tau_Y}{d\alpha}$, assume that τ_Y decreases linearly with α:

$$\tau_Y = \tau_Y^0 - k\alpha, \quad \text{with} \quad k > 0 \quad \text{and} \quad 0 \leq \alpha \leq 1,$$

where τ_Y^0 is the initial precision of the public signal when $\alpha = 0$, and k represents the rate at which precision decreases as α increases. Then:

$$\frac{d\tau_Y}{d\alpha} = -k < 0.$$

Now, the total derivative is:

$$\frac{dE[W|x]}{d\alpha} = \frac{\partial E[W|x]}{\partial \alpha} - k\frac{\partial E[W|x]}{\partial \tau_Y}.$$

Under which conditions is the total derivative negative (since welfare is negative, a negative derivative means welfare increases)? The total effect of increasing α on welfare depends on:

1. *Direct Marginal Benefit* ($\frac{\partial E[W|x]}{\partial \alpha}$): Increasing α directly reduces the distance from the fundamental for ingroup agents $E\left[(\bar{a}_{ig}(x,Y)-x)^2\right]$ because these agents process the public signal without bias. It also reduces the weight on the higher distance from the fundamental of outgroup agents.

2. *Indirect Marginal Cost* ($-k\frac{\partial E[W|x]}{\partial \tau_Y}$): Increasing α reduces τ_Y, lowering the precision of the public signal. A lower τ_Y increases the distance from the fundamental $E\left[(\bar{a}_{ih}(x,Y)-x)^2\right]$ for both ingroup and any remaining informed outgroup agents.

The marginal benefit must outweigh the marginal cost for delegation to improve welfare:

$$\frac{\partial E[W|x]}{\partial \alpha} > k\frac{\partial E[W|x]}{\partial \tau_Y}.$$

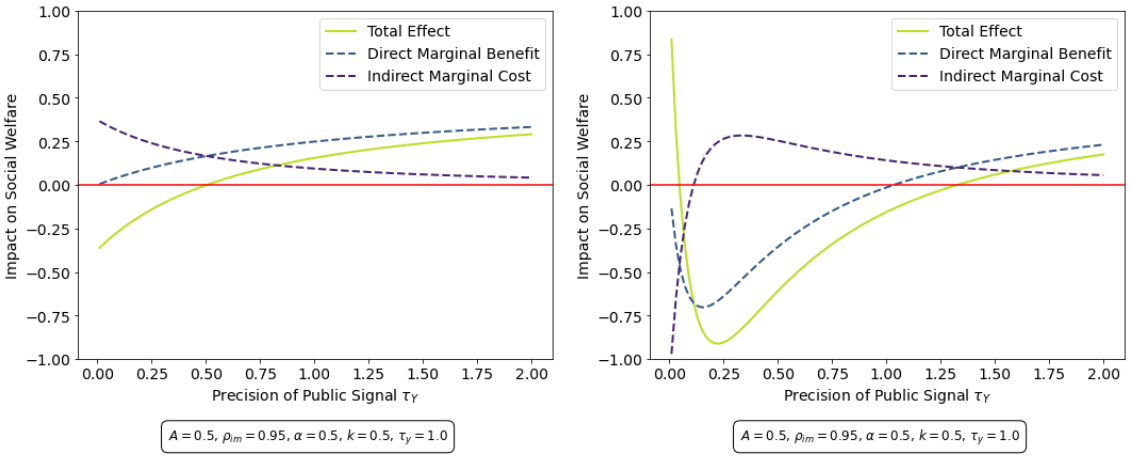
This condition formalizes the trade-off between the gain from enhanced coordination (marginal benefit) and the potential loss in public signal precision (marginal cost). If this inequality holds, increasing α improves expected welfare, proving Proposition 5. The specific conditions depend on the model parameters $(r, \tau_y, \tau_Y^0, k, \rho_{im}, A)$.

If $\frac{\partial E[W|x]}{\partial \alpha} > 0$ and $-k\frac{\partial E[W|x]}{\partial \tau_Y} < 0$, then the total derivative $\frac{dE[W|x]}{d\alpha} \gtrless 0$. Delegation is welfare-improving only when the marginal benefit of increasing α exceeds the marginal cost from reducing τ_Y. This either occurs when the direct improvement in agents' alignment with the fundamental is substantial, or when the marginal cost, which is the loss in public signal precision (τ_Y) per unit increase in α (represented by k), is sufficiently small.

Therefore, delegating communication to increase α is welfare-improving if the gain from enhanced coordination outweighs the potential loss in public signal precision due to less expertise of the delegated messengers.

Figure A36 visualizes the trade-off between direct marginal benefits and indirect marginal costs of delegating to messengers with lower (perceived) expertise in very low and very high coordination environments.

Figure A36: Social Welfare Impact of Delegation with Expertise Loss



(a) Low Coordination (r=0.10) (b) High Coordination (r=0.90)

Notes: Plots show impact on welfare by raising the share of ingroup agents via delegation, accounting for potential losses in the precision of the public signal due to lower expertise of the messenger. Green line shows the total impact, blue and purple dashed lines show the marginal benefit and marginal cost of delegating, respectively. τ_y is set to 1, A, k and initial α are set to 0.5 each. Arrows indicate the direction in which welfare is increasing. Panel (a) shows a case of insufficient strategic complementarity (r=0.10). Panel (b) shows a case of sufficient strategic complementarity (r=0.90).