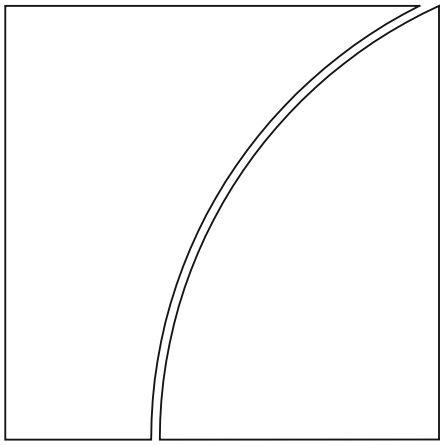

BIS Working Papers
No 1231

Monetary policy in the news: communication pass-through and inflation expectations

by Fiorella De Fiore, Alexis Maurin, Andrej Mijakovic and Damiano Sandri

Monetary and Economic Department

December 2024

JEL classification: E43, E50, E52, E58.

Keywords: Central bank communication, media coverage, large language models, households' expectations.

BIS Working Papers are written by members of the Monetary and Economic Department of the Bank for International Settlements, and from time to time by other economists, and are published by the Bank. The papers are on subjects of topical interest and are technical in character. The views expressed in them are those of their authors and not necessarily the views of the BIS.

This publication is available on the BIS website (www.bis.org).

ISSN 1020-0959 (print)
ISSN 1682-7678 (online)

Monetary policy in the news: communication pass-through and inflation expectations*

Fiorella De Fiore — BIS and CEPR
Alexis Maurin — STX Group
Andrej Mijakovic — EUI and LUISS
Damiano Sandri — BIS and CEPR

December 2024

Abstract

We analyse the media's role in channelling information about the Fed's monetary policy stance to the public. Using LLMs, we find a tight correspondence between FOMC communication and media coverage, although with significant variation over time. The communication pass-through weakened during the ZLB period and improved with the introduction of press conferences, which now exert strong influence on the media. Media coverage effects households' inflation expectations, particularly when inflation is high and volatile, while we do not detect a direct impact of FOMC communication. This underscores the media's crucial function in channelling central banks' communication to the public.

Keywords: Central bank communication, media coverage, large language models, households' expectations.

JEL Codes: E50, E52, E58.

*Fiorella De Fiore, Fiorella.DeFiore@bis.org; Alexis Maurin, Alexis.Maurin@stxgroup.com; Andrej Mijakovic andrej.mijakovic@eui.eu; Damiano Sandri Damiano.Sandri@bis.org. The views expressed in this paper are our own and should not be interpreted as reflecting the views of the Bank for International Settlements. Alexis Maurin worked on this paper when he was a member of the Monetary and Economic Department of the BIS. We thank Yuriy Gorodnichenko, Benoit Mojon, Andrei Zlate, Riccardo Correia, Taejin Park, Olatz Roman and participants at the 2024 CEBRA Annual Meetings and EUI for useful comments and discussions.

"You know that the FOMC pays very close attention to what it says [...]. Getting this message out to the public depends a good deal on the work you do in reporting on the FOMC, analyzing its statements and actions, and explaining its role and objectives."— Janet Yellen (2012), remarks at the Society of American Business Editors and Writers

1 Introduction

Over the past few decades, central banks have undergone a dramatic change in their approach to public communication (Blinder et al., 2008; Yellen, 2012). Prior to the 1990s, the prevailing strategy for central banks was to stir the economy by altering the short-term rate and disclose as little information as possible about monetary policy decisions. Explaining the rationale behind them or providing indications about future interest rates was considered inappropriate because it could trigger market overreactions or constrain future policy decisions. Much has changed since then. Central banks and academics have increasingly recognized that consumption and investment decisions depend not only on current borrowing costs or saving returns but are also strongly influenced by expectations regarding future interest rates and inflation. Hence, monetary policy can considerably enhance its efficacy by leveraging public communication to influence expectations.

The importance of communication became particularly evident during the period after the 2007-08 global financial crisis (GFC) when many central banks in advanced economies found themselves constrained by the zero lower bound (ZLB), preventing deeper interest rate cuts. In those circumstances, central banks extensively relied on forward guidance to provide additional monetary stimulus, indicating their determination to keep interest rates low for a long period. Public communication also played a critical role during the post-pandemic inflation surge. To prevent a possible de-anchoring of inflation expectations, central banks went to great lengths to communicate their determination to regain control of inflation by keeping rates elevated for as long as needed to bring inflation firmly back to target.

Recent years have seen a flurry of research to understand how central bank communication affects financial markets and professional forecasters. Yet, far less is known about central banks' ability to influence the broader public. As discussed in Blinder et al. (2023), this is a critical area for research since effective public communication is essential to improve monetary policy transmission and support public confidence in central banks. In this paper, we contribute to this research agenda by examining the degree of consistency between the official communication of the Federal Reserve about monetary policy

decisions and the media coverage, which we also refer to as the pass-through of FOMC communication to the media. Despite the rise of social media, traditional media remain the most important vehicle through which the general public acquires information about central banks (Hayo and Neuenkirch, 2015, 2018; Gardt et al., 2021; Conrad, Enders and Glas, 2022). Hence, for central banks to successfully influence public perceptions, it is critical that their communication is properly understood and portrayed by the media.

To assess the pass-through of FOMC communication to the media, we leverage the recent advancements in textual analysis enabled by the introduction of large language models (LLMs). We compare the performance of state-of-the-art LLMs (GPT-4, BERT, and RoBERTa) in evaluating the monetary policy stance conveyed during post-FOMC press conferences against human classification. For BERT and RoBERTa, we use model versions that have been specifically trained on FOMC communication. Despite this, we find that GPT-4 performs similarly to RoBERTa and considerably better than BERT in assessing the sentiment of FOMC communication. Moreover, GPT-4 shows higher accuracy in assessing longer texts on which RoBERTa and BERT have not been trained (Gambacorta et al., 2024). Therefore, we select GPT-4 to perform our analysis on newspaper articles.

We analyse around 14,000 articles from eight major news outlets, covering 224 FOMC announcements from 1994—when the FOMC released a first press statement—to 2023. We use GPT-4 to assess the sentiment about the monetary policy stance, on a scale from hawkish to dovish. We then examine the pass-through over the entire sample as well as variations over time. We also explore differences across the FOMC communication tools, including press statements, introductory remarks to press conferences, and the answers provided by the Fed chair to journalists. Furthermore, we examine whether the communication pass-through varies across phases of the Fed chairs' tenure. Finally, we evaluate whether the media coverage of FOMC decisions affects households' inflation expectations.

We find, in general, a high degree of consistency between the monetary policy sentiment expressed in FOMC communication and media coverage. The correspondence is especially strong when the FOMC expresses a particularly hawkish or dovish message. Furthermore, we document that FOMC communication influences the media coverage especially in the days immediately after the FOMC meetings, when most newspaper articles covering monetary policy are published.

Regarding how the consistency between the FOMC and media sentiment has varied over time, we find that the communication pass-through became considerably impaired after the 2008 financial crisis. In those years, monetary policy became constrained by

the ZLB and the Fed had to rely on new tools to provide accommodation. Our results underscore the severe communication challenges associated with explaining how new tools will impact the economy and how they will be deployed.

The communication pass-through began to improve around 2011 with the introduction of the post-FOMC press conferences. These events provided journalists with the opportunity to clarify their understanding of FOMC communication by directly posing questions to the Fed chair. The analysis shows that the Fed chairs' answers have become the most effective tool in shaping media coverage, However, the analysis also provides a cautionary message about these tools. The sentiment expressed in the press conference answers is at times not fully aligned with the message conveyed in FOMC written communication. This discrepancy may partly stem from the Fed chair's efforts to correct misinterpretations of the FOMC decisions that arise from journalists' questions. But it may also reflect the inherent communication challenges of addressing a wide-range of questions under difficult circumstances. The press conferences are thus high-stakes events, which can strongly shape media coverage but may also provide a message not entirely consistent with written FOMC communication.

We also document that the pass-through of FOMC communication to media coverage is significantly weaker in the first few months of a Fed chair's tenure. This underscores the challenges faced by new chairs in introducing their communication style and establishing credibility with the media.

Finally, we also provide evidence that media communication tends to influence household inflation expectations. Using data from the NY Fed's Survey of Consumer Expectations, we show that when the media sentiment becomes more hawkish after an FOMC meeting, households tend to reduce their medium-term inflation expectations. These effects are particularly pronounced during the post-pandemic period, when the inflation surge likely increased households' attention to monetary policy news. Interestingly, we do not find evidence that households directly respond to the sentiment expressed in FOMC communication. These results corroborate the notion motivating the analysis that the media coverage of monetary policy decisions plays a crucial role in channelling FOMC communication to the public.

The paper is structured as follows. In Section 2, we provide a brief review of the literature. In section 3, we present information about FOMC communication and media coverage. In Section 4, we describe the procedure to extract our sentiment indicators, and in Section 5, we examine the consistency between FOMC communication and the media. In Section 6, we examine the influence of media coverage on household inflation expectations. Section 7 concludes.

2 Literature review

The literature on central bank communication is vast. Most studies have focused on the impact of central bank communication on financial markets. A highly influential line of research has used event-study approaches to examine the reactions of asset prices in narrow time windows around central bank communication, including monetary policy statements, press conferences, and governors' speeches (Kuttner, 2001; Cochrane and Piazzesi, 2002; Bernanke and Kuttner, 2005; Gürkaynak, Sack and Swanson, 2005; Cieslak and Schrimpf, 2019; Andrade and Ferroni, 2021; Swanson, 2021). These papers document that central bank communication can strongly affect financial markets, impacting bond yields, stock prices, and exchange rates. By exploiting financial market reactions across different assets and maturities, these studies also offer important insights regarding the extent to which central bank communication provides elements of forward guidance and/or information about the economic outlook.

More recently, the literature has exploited technological advancements in language processing tools to directly examine the content of central bank communication, rather than inferring it from market reactions. Much of this work has focused again on the impact of central bank communication on financial markets (Lucca and Trebbi, 2009; Hansen and McMahon, 2016; Neuhierl and Weber, 2019; Hansen, McMahon and Tong, 2019; Gardner, Scotti and Vega, 2022; Doh, Song and Yang, 2022; Cieslak and McMahon, 2023; Cieslak, McMahon and Pang, 2024). Leveraging computer-based techniques to analyze both textual and nonverbal aspects of communication, recent studies have documented that central banks' governors also influence financial markets via their tone of voice and facial expressions (Gorodnichenko, Pham and Talavera, 2023; Curti and Kazinnik, 2023; Alexopoulos et al., 2024).

Rather than focusing on the impact of central bank communication on financial markets, our analysis examines the potential for central bank communication to reach the broader public via the media. This is a necessary condition for central banks to influence the public given that households still heavily rely on traditional media to learn about monetary policy (Hayo and Neuenkirch, 2018; Gardt et al., 2021; Conrad, Enders and Glas, 2022). However, two other conditions must also be satisfied.

First, people must be responsive to monetary policy news when provided with such information. Survey-based analyses using information provision experiments provide encouraging supportive evidence. For example, households tend to adjust their inflation expectations when informed about central banks' inflation targets and interest rate decisions (Coibion, Gorodnichenko and Weber, 2022; Coibion et al., 2023), especially if

central banks use a simple narrative (Haldane and McMahon, 2018).

Second, people must be attentive to monetary policy announcements. The available evidence on this point is more mixed. Based on data from Australia, Claus and Nguyen (2020) document that households' expectations of economic conditions adjust following monetary policy shocks in line with theoretical predictions. Lewis, Makridis and Mertens (2019) find that US policy rate shocks affect households' confidence in the state of the economy. Using surveys conducted right before and after FOMC press conferences, Lamla and Vinogradov (2019) document an increase in the proportion of households that have heard about monetary policy news after the press conferences. However, they detect very little effect on households' expectations. Similarly, De Fiore, Lombardi and Schuffels (2021) examine the NY Fed Survey of Consumer Expectations over the pre-pandemic period and find that households do not materially change inflation expectations around monetary policy announcements, suggesting that the general public pays limited attention to the FOMC communication of monetary policy decisions. Binder, Campbell and Ryngaert (2024) also document that, in general, FOMC communication does not trigger changes in inflation expectations. However, they do find instances, especially during the Covid-19 pandemic, when households adjusted inflation expectations after FOMC announcements.

Our study provides new evidence that, although households tend to be unresponsive to FOMC communication alone, they are sensitive to the media coverage. Specifically, households tend to lower medium-term inflation expectations when the media coverage turns more hawkish. A few other papers also examine the role of media coverage in channelling information about monetary policy. Pinter and Kočenda (2023) manually assess the articles of the major French newspapers covering ECB policy decisions. They find that monetary policy shocks affect French firms' and consumers' expectations only when they are properly reported in the media. Schmanski et al. (2023) extract textual sentiment about monetary policy from FOMC communication and Dow Jones newswire articles using dictionary-based techniques. They find that media sentiment correlates more tightly with changes in bond yields and economic forecasts by professional analysis than FOMC sentiment. Our paper extracts textual sentiment by leveraging recent developments in LLMs. We study media coverage across a large set of newspapers and focus on its influence on households' inflation expectations.

Our paper is also related to several studies that examine the media coverage of central bank communication. Neuenkirch (2014) and Munday and Brookes (2021) investigate the textual features that might solicit larger news coverage. Berger, Ehrmann and Fratzscher (2011) examine how favorably the media report about the monetary policy

decisions by the ECB. Hayo and Neuenkirch (2012) and Hendry (2012) show that media coverage of communication by the Bank of Canada can influence equity markets and interest rates. Ter Ellen, Larsen and Thorsrud (2022) construct monetary policy shocks for Norway based on differences between the policy announcements and prior media coverage and show that they predict macroeconomic responses consistent with an information channel. Ehrmann and Wabitsch (2022) examines the role of Twitter in relaying information about the ECB's monetary policy communication.

3 FOMC communication and media coverage

We systematically examine FOMC communication and media coverage since February 1994, when the FOMC issued for the first time a press statement to announce a change in the monetary policy stance. Before that, monetary policy decisions were implemented via market operations by the Fed, without being publicly announced. Press statements became more frequent in the second half of the 1990s and have been regularly released after every FOMC meeting since 2000. In April 2011, the Fed also started to use press conferences to provide additional information regarding monetary policy decisions and the future policy stance. The Fed chair opens press conferences by reading introductory remarks and then answers questions from journalists. Press conferences were initially held only after FOMC meetings when updated economic projections were released but have become regular events after every FOMC meeting since 2019.

In total, our analysis covers all press statements, introductory remarks at press conferences, and answers to journalists' questions for all the 224 FOMC meetings that took place between February 1994 and the end of 2023. Figure 1 illustrates several features of these data. The left-side panel shows the rising frequency of FOMC communication, captured by an increasing use of press statements and press conferences.[1] The growing emphasis on public communication is also reflected in the lengthening of communication, as illustrated in the middle panel. The average number of words in the statements increased from 110 in 1994 up to 811 in 2014. From that point onward, statements have become shorter but have been complemented with considerably longer introductory remarks (1326 words on average) and especially extensive answers provided during press conferences (4980 words on average).

Besides becoming more frequent and lengthy, FOMC communication has also adopted a simpler language over time, as measured by the Flesch-Kincaid readability score and

[1] FOMC meetings are generally held eight times per year. However, meetings can occur more frequently during crisis episodes, such as during the global financial crisis and the COVID-19 pandemic.

illustrated in Figure 1c.[2] The readability score of press statements improved (as shown by an increase in the index) between 1994 and the eve of the 2008 financial crisis. Statements became more complex during the ensuing ZLB period as the Fed started to deploy new monetary policy tools, such as large-scale asset purchases and forward guidance. However, the introduction of press conferences allowed the Fed chair to communicate with the public using simpler language, especially when answering questions.

Figure 1: FOMC communication

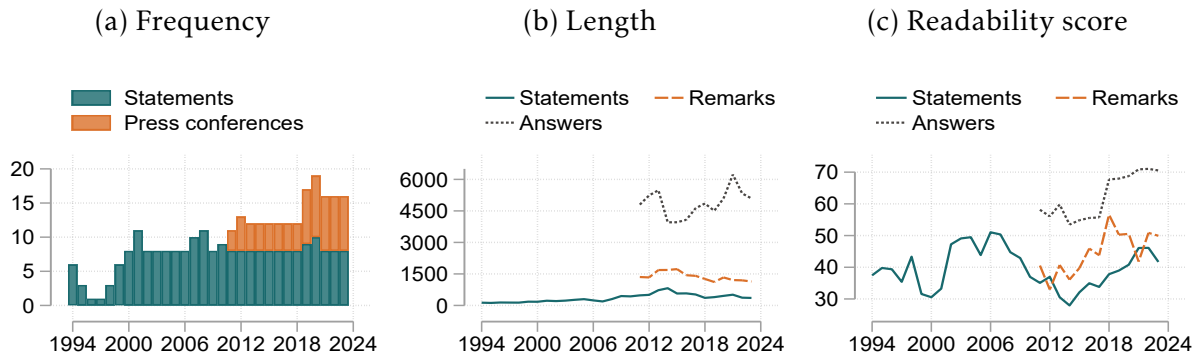

Notes: The left panel shows the yearly average number of press statements and press conferences following FOMC meetings. The middle panel shows the yearly average number of words embedded in press statements, introductory remarks and in the answers during press conferences. The right panel shows the Flesch–Kincaid readability score for each communication tool.

Turning to the media coverage of FOMC meetings, we collect articles published between two days before and seven days after the meeting from eight major media outlets: the Financial Times, Wall Street Journal, Washington Post, Los Angeles Times, CNN, USA Today, New York Post, and Fox News.[3] We select articles that include references to the FOMC or the Fed and to interest rates or monetary policy decisions.[4] Our search criteria return 40,683 articles.

We then follow a 3-step data cleaning procedure to remove articles that are not relevant for our analysis. First, we eliminate transcripts of FOMC communication, for example distributed via newsletters and selected daily market articles. Second, we drop duplicate articles or very similar ones, for example, online articles that underwent minor

[2]The Flesch-Kincaid readability score captures complexity along two main dimensions, word length and sentence length.

[3]The selection of media outlets is based on readership in the US and considerations related to licensing agreements required to process the articles using LLMs. Coverage of FT articles starts in 2005 because of licensing limitations. We retrieve articles from the Financial Times using its dedicated application programming interface (API) and we use the Factiva API for all other newspapers.

[4]More specifically, we select articles that include either the expressions "FOMC" or "Federal Open Market Committee" or "FED" or "Federal Reserve," as well as the word "interest rate" or "interest rates" or "monetary" or "federal funds rate" or "fed funds rate".

changes. Third, we leverage the capabilities of GPT-4 to select articles that are sufficiently focused on monetary policy decisions by the Federal Reserve. Details of the data cleaning procedure are provided in Appendix A. The final sample of analysis includes 14,021 articles whose distribution across media outlets is reported in Appendix Table A.1.

As illustrated in Figure 2a, the media coverage of the FOMC decisions remained broadly stable between 1994 and the 2008 financial crisis. Based on our sample of analysis–excluding the Financial Times for which we have data starting only in 2005—FOMC meetings were covered on average in about 33 articles between 2 days prior and 7 days after the meeting. The media coverage of FOMC meetings has considerably increased after the financial crisis, reaching about 55 articles per meeting.[5] Therefore, the increased frequency and length of FOMC communication were matched with larger media coverage. Regarding the distribution of articles around FOMC meetings, Figure 2b shows that the bulk of them are published the day of the meeting or the subsequent one.[6] Finally, in Figure 2c we observe that the readability score of media articles is considerably higher than the average readability of FOMC communication across all communication tools. This points to an important role played by the media in interpreting and simplifying Fed communication to reach the broad public.

Figure 2: Media coverage of FOMC meetings

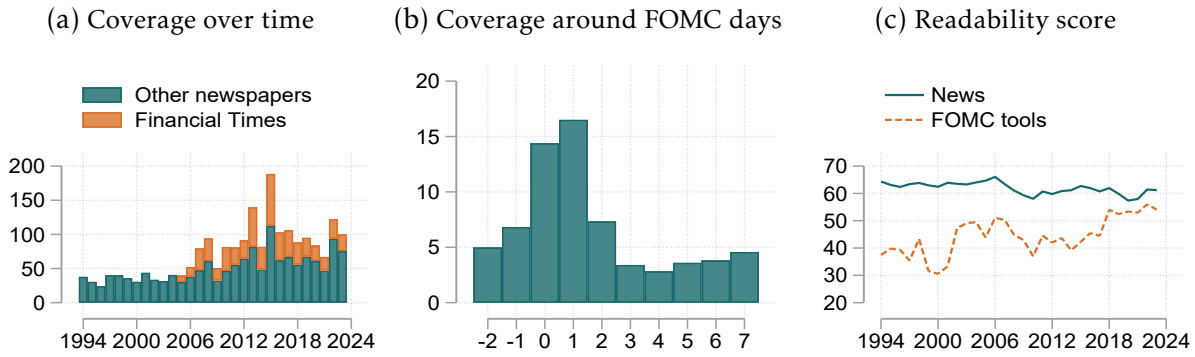

Notes: The left panel shows the average number of articles per FOMC meeting per year. The middle panel shows the average number of articles published around the day of the FOMC meeting. The right panel shows the average Flesch–Kincaid readability score for the articles and FOMC communication tools.

[5]The same figures including the FT are 39 for the period January 2005 - September 2008, and 81 for the period September 2008 - December 2023.

[6]To determine whether an article was published before or after the FOMC press statement, we used the timestamp of the article when available. When the timestamp was not available, we read and classified the articles as published pre- or post-FOMC based on their content.

4 Sentiment extraction

The main focus of the paper is to examine the degree of consistency between FOMC communication regarding the stance of monetary policy and media coverage. To assess the monetary policy sentiment, we rely on state-of-the-art LLMs.

4.1 Comparison of alternative LLMs

We consider three LLMs that have been used in the literature to analyze Fed communication: GPT-4, RoBERTa, and BERT. Regarding GPT, we use the most advanced version available at the time of our study, GPT-4. For RoBERTa, we use the version developed by Shah, Paturi and Chava (2023), which was pre-trained on extracts from FOMC minutes, press conferences and speeches by Federal Reserve officials. We refer to this version as RoBERTa-SPC23. Finally, for BERT, we use the version used by Gorodnichenko, Pham and Talavera (2023) (GoPhTA23) which was pre-trained based on FOMC statements. We refer to this model as BERT-GoPhTa23.

To select the most suitable LLM for the purpose of our analysis, we compare their performance in assessing the Fed Chair's answers during the press conference against the interpretation of human research assistants made available by GoPhTA23 for all FOMC press conferences until June 2019. We compare the performance of LLMs based on the press conference answers rather than statements or introductory remarks because the answers use a less scripted language that is arguably harder to decipher. For that reason, the analysis of answers is also likely more representative of LLMs' performance in assessing the rich narrative presented in newspaper articles.

GoPhTA23 asked research assistants to classify each press conference answer on a scale from -10 (very hawkish) to +10 (very dovish). They then averaged scores across the research assistants to classify each answer as dovish (≥ 0.5), neutral (between -0.5 and 0.5), or hawkish (≤ -0.5). BERT-GoPhTa23 and RoBERTa-SPC23 have been pre-trained to assess FOMC text on a three-point scale, including dovish, neutral or hawkish. In the case of GPT-4, we use a finer 5-point scale to capture richer nuances in the text. Specifically, we provide GPT-4 with the following prompt:[7]

[7]We tested different versions of this prompt and found that the results are similar. To ensure greater replicability of the results, we set the temperature in GPT (i.e. the degree of randomness in the model's output) to 0. In a few instances—about 0.3% of cases—GPT returns words that do not correspond to the dovish/hawkish categories. We re-classify these statements as neutral.

You're a research assistant working at the Fed. You have a Master degree in Economics. Your task is to understand and classify the monetary policy stance described in the text given to you into one of the following five categories: dovish, mostly dovish, neutral, mostly hawkish, hawkish. Dovish means that the central bank is more lenient towards higher inflation to support economic activity and employment, and hence is more likely to lower interest rates or keep them low for a period of time. Hawkish means the opposite. Neutral is neither one nor the other.

To map GPT-4's 5-point scores to the 3-point scale used in GoPhTA23, we code "mostly dovish" as dovish and "mostly hawkish" as hawkish. This limits the number of neutral assessments in line with GoPhTA23's mapping of human scores into the 3-point scale which treats as neutral only those answers with scores very close to zero (between -0.5 and 0.5) on a scale from -10 and 10.

Table 1 compares the accuracy of the LLMs, measured as the percentage of correct matches between the model and the human classification.[8] GPT-4 and and RoBERTa-SPC23 perform considerably better than BERT-GoPhTa23. On average, they correctly predict the sentiment in 60 and 61 percent of cases respectively, compared to 42 percent by BERT-GoPhTa23. Looking at individual sentiment categories, BERT-GoPhTa23 scores relatively well for "hawkish" and "dovish" answers but performs poorly for "neutral" answers. The weaker performance of BERT-GoPhTa23 relative to RoBERTa-SPC23 is likely explained by the fact that BERT-GoPhTa23 was trained using only FOMC statements. In contrast, RoBERTa-SPC23 was also trained based on press conference transcripts. Despite a comparable average performance, the accuracy of GPT-4 and RoBERTa varies across categories of sentiment. GPT-4 more accurately captures "neutral" answers, while detecting less accurately the "hawkish" or "dovish" categories. That said, GPT-4 has an overall more balanced performance as it is less prone to make severe mistakes, defined as the instances when it detects "hawkish" while the sentiment is "dovish" or viceversa.[9]

The strong performance of GPT-4 is particularly remarkable considering that GPT-4 was not subject to any specific training based on FOMC communication. Hence, GTP-4 is likely to perform even better relative to BERT-GoPhTa23 and RoBERTa-SPC23 in the analysis of newspapers since neither BERT-GoPhTa23 nor RoBERTa-SPC23 have been trained on media sources. Another important advantage of using GPT-4 rather than RoBERTa-SPC23 for our analysis is that the assessment of media sentiment is performed based on full articles rather individual paragraphs, since Factiva API does not provide

[8] Appendix Table B.1 reports the associated confusion matrices.

[9] The occurrence of such mistakes is shown in the confusion matrices reported in Appendix B. For instance, RoBERTa assesses 24% (7%) of statements that are hawkish (dovish) as dovish (hawkish), while this happens only for 13% (0.8%) of the statements with GPT-4.

Table 1: Accuracy scores for press conference answers across LLMs

	Human's classification			
Model	All	Hawkish	Neutral	Dovish
GPT-4	0.60	0.16	0.91	0.38
RoBERTa-SPC23	0.61	0.34	0.75	0.54
BERT-GoPhTa23	0.42	0.54	0.26	0.60
Number of answers	692	117	336	239

Notes: The accuracy scores report the share of correct matches between the model's and human's classification of press conference answers.

paragraph splits. As shown in Gambacorta et al. (2024), GPT-4 performs considerably better than RoBERTa-SPC23 in assessing longer texts. In light of these considerations, we conduct our analysis using GPT-4.

4.2 FOMC and media sentiment

Using GPT-4, we extract the sentiment of FOMC communication as expressed in press statements, introductory remarks to the press conference, and the Fed Chair's answers to journalists. The sentiment of FOMC statements and remarks is extracted based on the full text of these documents. This is consistent with the treatment of media articles that will also be analyzed at the full-text level, as described later on. We convert the sentiment classification into numerical values by assigning the dovish, mostly dovish, neutral, mostly hawkish, hawkish categories to $\{-1, -0.5, 0, 0.5, 1\}$ values, respectively. Regarding the press conference answers, we extract the sentiment of each answer and convert it to numerical values. We then construct an aggregate sentiment for the press conference answers as the average sentiment across the answers.[10]

We also construct an aggregate sentiment of the FOMC communication for each meeting by taking a simple average across the sentiment expressed in the press statement, the introductory remarks and the press conference answers. The sentiment of FOMC meetings that did not feature a press conference simply corresponds to the sentiment expressed in the press statement. Note that all sentiment indexes range between [-1,1], with

[10]Note that we compute the average sentiment by excluding answers that are classified as neutral. This is because most answers are classified as neutral, hence their inclusion would mechanically reduce the variance of the press conference answers' sentiment relative to the sentiment extracted from statements and remarks. The standard deviation of the aggregate answer sentiment without neutral answers is 0.51, compared to standard deviations of 0.62 and 0.72 for statements and remarks, respectively. If neutral answers are included, the standard deviation of the aggregate answer sentiment falls to 0.17. The key results of the analysis are, in any case, robust to including neutral answers.

higher values denoting more hawkish sentiment.

Figure 3a illustrates the evolution of the FOMC sentiment between 1994 and 2023, together with the dynamics of the federal funds rate. Several observations stand out. First, the FOMC sentiment displays considerable variation over time, with sharp dovish shifts around economic downturns. For example, the FOMC sentiment turns rapidly dovish during the 2001, 2008, and 2020 recessions. We also see a rapid shift toward a hawkish sentiment in mid-2021, with the onset of the post-pandemic inflation surge. Second, sentiment dynamics well align with prevailing narratives about monetary policy phases. For example, the FOMC sentiment is characterized as persistently dovish in the years after the 2008 financial crisis, when the Fed engaged in several rounds of quantitative easing and used forward guidance to signal its intention to keep interest rates low for long. Third, fluctuations in FOMC sentiment correlate closely with the federal funds rates. We also see a tendency for changes in FOMC sentiment to anticipate interest rate adjustments, especially hiking cycles. For example, FOMC sentiment turned more hawkish in 2004 and 2015, right before interest rate hikes. Communication also turned sharply hawkish with the inflation surge in 2022, right before one of the sharpest tightening cycles in the history of the Fed.

Figure 3: FOMC sentiment

(a) Aggregate

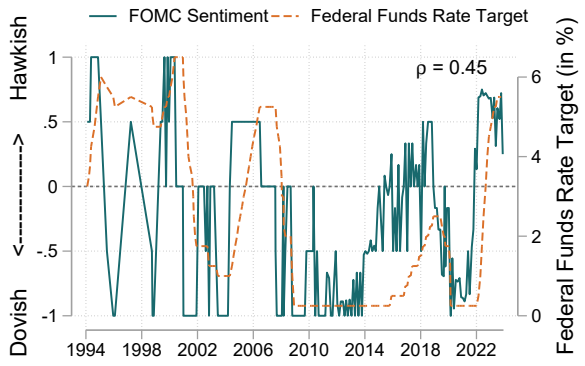

(b) By communication tool

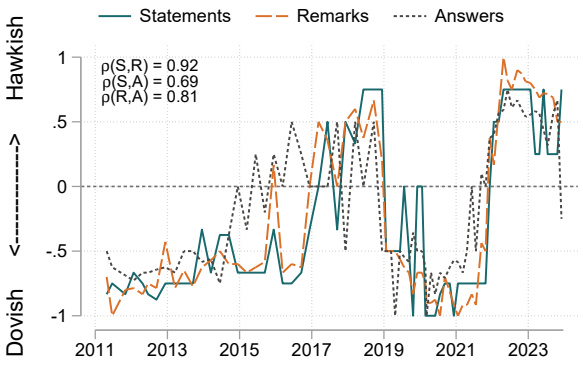

Notes: Panel (a) compares the evolution of the sentiment of the aggregate FOMC official communication to that of the Federal Funds Rate target. ρ is the correlation coefficient between the two variables shown. Panel (b) plots the sentiment of each individual communication tool, i.e. statements (S), introductory remarks (R), and answers at press conferences (A). ρ is the correlation coefficient between the two indicated tools.

Figure 3b illustrates the sentiments of individual communication tools. We see a very tight comovement between the sentiment expressed in press statements (S) and in the introductory remarks (R) to the press conference, with a correlation coefficient of 0.92. This underscores how written texts are carefully drafted by the Fed to provide a consis-

tent message. The sentiment expressed in the answers during the press conference (A) is instead less tightly connected with the sentiment in written communication. The correlation between the sentiment of the answers and remarks is 0.81. The one between answers and statements is even lower, at 0.69. The limited correlation between the sentiment of the answers and written statements may partly reflect the communication challenges arising during the press conferences, when the Fed Chair must address a broad range of questions from journalists on the spot. Discrepancies between these two communication tools may also result from attempts by the Fed Chair to push back against misinterpretations of FOMC decisions that may emerge during the press conference. For example, journalists may ask whether a dovish passage in the post-FOMC press statement rules out interest rate hikes in the near future, prompting the chair to clarify that the FOMC retains discretion to adjust the policy stance based on upcoming data.

Turning to the sentiment expressed in the media, we follow the same approach used to construct the sentiment of FOMC communication. Specifically, we ask the LLMs to assess the full text of each newspaper article along the five hawkish/dovish categories.[11] We then convert this classification into numerical values and average the scores across articles to obtain an aggregate media sentiment for each FOMC meeting.

5 Consistency between FOMC and media sentiment

We now move on to examining the consistency between the sentiment extracted from Fed communication and its media coverage. In the next section, we start by assessing the average degree of consistency over the entire sample period and across all communication tools. We then examine in the subsequent section how the communication pass-through has varied over time and how it differs across communication tools.

5.1 Average degree of consistency

The analysis uncovers, on average, a strong degree of consistency between the FOMC communication and the media coverage. Figure 4 illustrates the evolution of the aggregate FOMC sentiment and the aggregate media sentiment based on all articles published up to 7 days after the FOMC meeting. We see, in general, a strong comovement between these two series, with a correlation of 0.83.

To examine more systemically how changes in FOMC sentiment are reflected in the

[11] Note that since the Factiva API does not provide the split of newspaper articles in individual paragraphs, the analysis cannot be conducted at the paragraph level.

Figure 4: FOMC and media sentiment

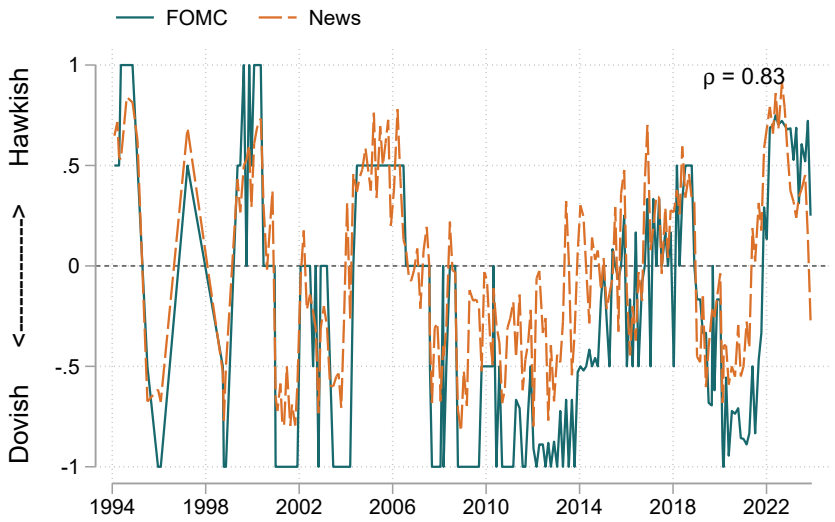

Notes: Green solid line denotes the aggregate FOMC sentiment. Orange dashed line denotes media sentiment. ρ is the correlation coefficient.

media coverage, we estimate the following regression:

$$S_{a,j,m}^{Media} = \alpha_j + \beta S_m^{FOMC} + u_{a,j,m} \tag{1}$$

where $S_{a,j,m}^{Media}$ is the sentiment of article a, published in newspaper j, up to 7 days after the FOMC meeting m. The variable S_m^{FOMC} is the aggregate sentiment of the FOMC communication for meeting m. The regression includes newspaper-fixed effects, α_j, and standard errors are clustered at the newspaper level. The focus of the analysis is on the coefficient β, which captures the strength of the pass-through from FOMC communication to the media coverage. The first column in Table 2 reports the regression estimates. The coefficient β is positive and highly statistically significant, equal to 0.65, confirming a strong degree of consistency between FOMC communication and the media. Note that this result is not driven by specific media outlets. In Appendix Table C.1, we re-estimate equation (1) for each media outlet separately and find a consistently strong degree of pass-through.

Besides assessing the average degree of pass-through, we also examine if the strength of the pass-through is influenced by the intensity or direction of the FOMC sentiment. To this end, we extend the baseline regression equation (1) as follows:

$$S_{a,j,m}^{Media} = \alpha_j + \beta_1 S_m^{FOMC} + \beta_2 D_m + \beta_3 S_m^{FOMC} \times D_m + u_{a,j,m} \tag{2}$$

Table 2: Sentiment pass-through and media coverage

	Sentiment			Number of articles		
	(1) Base	(2) Size	(3) Sign	(4) Base	(5) Size	(6) Sign
FOMC sentiment	0.651***	0.422***	0.662***	1.541		
	(0.031)	(0.040)	(0.056)	(2.147)		
Strong sentiment		0.101***			-0.652	
		(0.014)			(2.679)	
FOMC X Strong sentiment		0.261***				
		(0.030)				
Dovish sentiment			-0.176**			-3.457
			(0.056)			(3.244)
FOMC X Dovish sentiment			-0.203*			
			(0.096)			
Constant	0.121***	0.048***	0.154**	34.178***	34.116***	38.333***
	(0.008)	(0.013)	(0.047)	(1.456)	(2.061)	(2.690)
No. of meetings	224	224	184	224	224	184
Observations	10472	10472	9202	224	224	184
R2	0.38	0.39	0.41	0.00	0.00	0.01

Notes: Strong sentiment denotes meetings in which FOMC sentiment is above the 75th percentile or below the 25th percentile of the sentiment distribution. Dovish sentiment denotes meetings in which FOMC sentiment is below 0. Standard errors are clustered at the newspaper level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

where D_m is a dummy variable that takes value 1 if a certain criterion is satisfied. More specifically, we first differentiate the results based on the strength of the FOMC sentiment, namely whether it is particularly hawkish or dovish. In this case, the dummy D_m captures FOMC meetings where the sentiment is within the top or bottom quartile of the sentiment distribution. The estimation results are reported in column (2) of Table 2. The coefficient β_3 is positive and statistically significant, showing that meetings with more extreme FOMC sentiment tend to influence more strongly the media coverage.

We then examine if the strength of pass-through differs depending on whether the FOMC sentiment is hawkish or dovish. To explore this issue, the dummy D_t takes value 1 to denote meetings with dovish sentiment, i.e. meetings in which $S^{FOMC} < 0$, and value 0 for hawkish meetings. Column (3) shows that the interaction coefficient β_3 is negative and statistically significant, although only at the 10 percent level. Thus, the strength of pass-through appears to be somewhat weaker when the FOMC sentiment is dovish. As later discussed, this result reflects weaker pass-through during the post-GFC years when the ZLB became binding.

Consistency between the FOMC and media sentiment is critical for an accurate transmission of monetary policy communication to the public. However, it is also important for the media to provide a balanced coverage of FOMC decisions. For example, if the

FOMC and media sentiments are generally consistent but the media covers much more extensively hawkish FOMC meetings, this may create an overly hawkish perception about the FOMC. To assess this aspect, we regress the number of media articles for each FOMC meeting over the FOMC sentiment. As shown in column (4) in Table 2, we do not detect any correlation, implying that media coverage does not vary with the FOMC sentiment. Furthermore, columns (5) and (6) show that the media coverage does not differ depending on whether the FOMC sentiment is strong or weak or on whether it is dovish or hawkish.

The analysis so far has focused on the average pass-through of the FOMC sentiment to articles published in the 7 days following the policy meeting. We now explore how the FOMC sentiment propagates over time in the media by differentiating articles based on their publication dates. To this end, we estimate the following specification

$$S^{Media}_{a,j,m,h} = \alpha_j + \beta_h S^{FOMC}_m + \gamma_h S^{Media,pre}_m + u_{a,j,m,h} \qquad \text{for} \quad h = 0, 1, ..., 7 \tag{3}$$

where the subscript h captures the day of publication of each article relative to the FOMC meeting m. More precisely, $h = 0$ denotes the day of the meeting, $h = 1$ the day after and so on. We also control for the average media sentiment in the two days preceding the meeting $S^{Media,pre}_m$. Figure 5 shows the estimates of the β_h coefficients. We see that the pass-through of the FOMC sentiment to the media coverage is particularly strong in the day following the meeting. As shown in Figure 2b, this is also the day with the highest number of articles covering the FOMC, thus arguably the most important one to shape public views. The strength of the pass-through declines modestly over the subsequent days as media coverage of FOMC decisions decreases.

5.2 Differences across time and communication tools

The previous section has documented that the FOMC sentiment is, in general, well reflected in the media. We now explore whether the strength of the pass-through has varied over time and which types of FOMC communication tools tend to influence the media coverage more strongly.

A closer inspection of Figure 4 already suggests that the correspondence between the FOMC and media sentiment varies over time. The two series comoved tightly during the pre-GFC period and during the hawkish turn triggered by the post-pandemic inflation surge. In contrast, the media sentiment did not track the FOMC sentiment that closely in the post-GFC period, when monetary policy was constrained by the zero lower bound. To examine these patterns more clearly, we assess the evolution of the pass-through over time by re-estimating our baseline equation (1) over rolling (overlapping) windows in-

Figure 5: Dynamic response of media sentiment

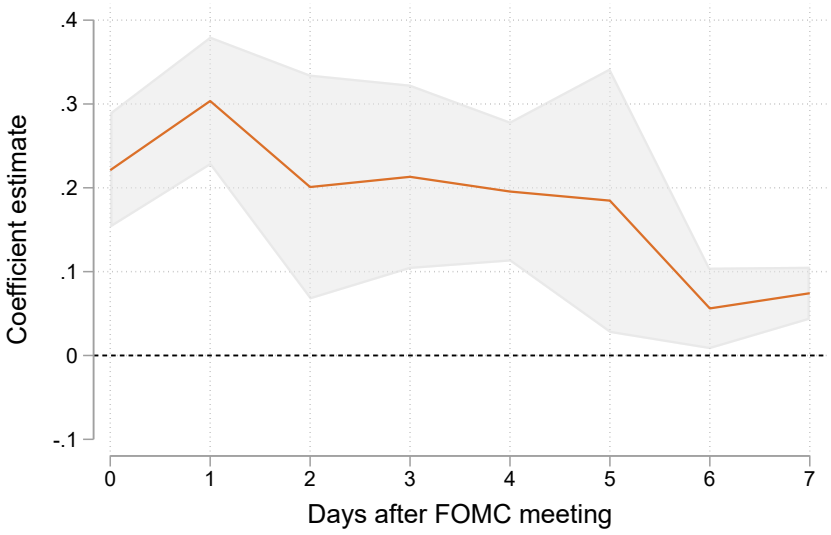

Notes: Standard errors are clustered at the newspaper level. 90 percent confidence bands in grey. Controls include pre-meeting news sentiment.

cluding 70 FOMC meetings, each covering about 10 years.[12] Figure 6 reports the estimates of the β coefficient for each rolling window, with the horizontal axis denoting the date of the FOMC meeting in the middle of the window. We see a pronounced decline in the communication pass-through when monetary policy became constrained by the zero lower bound after the GFC. This underscores the communication challenges faced by the Fed at that time, reflecting the new uncharted conduct of unconventional monetary policy. The communication pass-through started to improve with the introduction of press conferences in 2011, suggesting that this communication tool proved effective in clarifying the stance of monetary policy. Finally, the strength of the pass-through has reached historic highs in the most recent period, when the rolling regression sample includes the sharp communication turns by the FOMC during COVID and the subsequent inflation surge.

To further explore the role of press conferences in influencing media coverage, we assess whether the strength of pass-through differs depending on whether FOMC meetings are followed or not by a press conference. As previously discussed, press conferences were introduced in April 2011 and were held until 2019 only when updated economic projections were released. By focusing on the period between April 2011 and 2019 which saw FOMC meetings with and without press conferences, we estimate equation (2) where

[12] Although FOMC meetings are now typically 8 per year, meetings were less frequent in the early part of our sample, resulting on average in around 7 meetings per year. Appendix Figure C.1 reports results for other window lengths. The overall pattern remains similar.

Figure 6: Strength of the communication pass-through over time

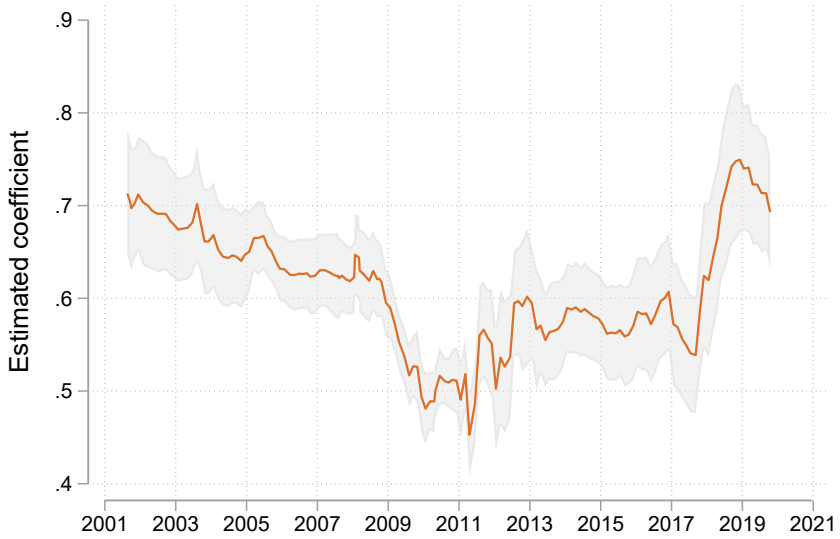

Notes: Estimated over overlapping 70-meeting windows. Each point on the x-axis corresponds to the date of the FOMC occurring in the middle of the rolling window. The grey area indicates the 90 percent confidence interval. Standard errors are clustered at the newspaper level.

the dummy D_t denotes FOMC meetings followed by a press conference. The results reported in column (1) of Table 3 show that the pass-through tends to be considerably stronger for FOMC meetings with a press conference.

An additional approach to assess the importance of press conferences in affecting the media coverage is to examine the pass-through of each FOMC communication tool separately. To this end, we re-estimate our baseline equation by including as separate regressors the sentiment conveyed in the press statement, in the introductory remarks to the press conference, and in the answers provided by the Fed during the press conference. In this case, the regression is run including all FOMC meetings featuring a press conference. The results are reported in column (2) of Table 3. We see that the sentiment of the answers plays a considerably larger effect in influencing the media coverage, both relative to the introductory remarks to the press conference and to the initial press statement. These results highlight the critical role that press conferences—and especially the answers by the Fed chair—play in influencing the media coverage.

Finally, we also examine if the pass-through varies depending on the phase of the chair tenure. Column (3) shows that the strength of pass-through is considerably weaker under a new chair, that is, during the first 6 months of the appointment. The effect is quantitatively large, reducing the transmission of FOMC communication to media by about 40 percent. These results suggest that the first few months of a new chair are

Table 3: Pass-through differences across time, tools and chair tenure

	(1) Meetings between April 2011 and 2018	(2) Only meetings with press conference	(3) All meetings
FOMC sentiment	0.534***		0.651***
	(0.038)		(0.032)
PressConf=1	-0.126***		
	(0.021)		
PressConf=1 × FOMC sentiment	0.118**		
	(0.036)		
Statements sentiment		0.187***	
		(0.025)	
Remarks sentiment		0.140***	
		(0.024)	
Answers sentiment		0.406***	
		(0.037)	
NewChair=1			0.173***
			(0.026)
NewChair=1 × FOMC sentiment			-0.255***
			(0.033)
Constant	0.253***	0.102***	0.116***
	(0.012)	(0.004)	(0.009)
No. of meetings	62	73	224
No. of articles	3982	5179	10472
R2	0.22	0.35	0.39

Notes: New chair is a dummy that takes on value 1 in periods of six months after the inauguration of a new chair. Standard errors are clustered at the newspaper level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

particularly delicate to establish clear communication with the media and, hence, the public.

6 Media sentiment and households' expectations

The analysis has documented a high degree of consistency between the sentiment expressed in FOMC communication and the media coverage. But does the media sentiment in turn influence households, especially regarding their inflation expectations? A key goal of central bank communication with the public is indeed to manage inflation expectations.

To address this question, we examine the inflation expectations of US households collected by the New York Fed in the Survey of Consumer Expectations. The survey has been conducted since June 2013 and involves the participation of about 1,300 respondents each month. We use data up to September 2023, the last vintage available at the

time of our study.

A useful feature of this survey for the purposes of our analysis is that participants submit their answers on different days. Building on Binder, Campbell and Ryngaert (2024) and De Fiore, Lombardi and Schuffels (2021), the empirical approach involves testing whether the inflation expectations collected a few days after FOMC meetings differ from those collected a few days before the meetings, and whether these differences are correlated with changes in media sentiment. Formally, we estimate the following regression

$$\pi^e_{it} = \alpha_i + \gamma_m + \left(\beta_1 + \beta_2 \Delta S^{News}_m\right) Post_{im} + u_{it} \tag{4}$$

The dependent variable π^e_{it} denotes the inflation expectation of survey participant i on the day t. The regression controls for participants' fixed effects, α_i, as well as for FOMC meetings' fixed effects, γ_m. The γ_m fixed effects identify inflation expectations collected between w days before and after the FOMC meeting m. The dummy variable $Post_{im}$ identifies inflation expectations collected within w days after the FOMC meeting m. Finally, the variable ΔS^{News}_m captures the change in media sentiment around the FOMC meeting m. This is computed as the difference in the average media sentiment between two days after and two days before the meeting m.

The coefficient β_1 captures whether inflation expectations submitted within w days after the meeting differ from those submitted within w days prior to the meeting. The focus of the analysis is on the coefficient β_2. This coefficient captures if changes in inflation expectations before and after the meeting are correlated with changes in the media coverage. Specifically, the β_2 coefficient should be negative if a hawkish turn in media sentiment reduces inflation expectations, as we would expect from a theoretical standpoint. We estimate equation (4) using alternative windows w surrounding FOMC meetings. We start by setting $w = 3$, hence including only inflation expectations collected between 3 days before and 3 days after the meeting. And then increase this window up to 12 days. Shorter windows provide a cleaner identification of the effect of media sentiment on inflation expectations since inflation expectations are less likely to be contaminated by events unrelated to the FOMC. However, shorter windows reduce the regression sample.[13] We consider both short and medium-term inflation expectations, defined respectively as expected inflation over the next 12 months and expected inflation between 24 and 36 months from the time of the survey. We refer to these expectations as 1 and 3-year-ahead inflation expectations.[14] The regression is estimated accounting for the sur-

[13] In a few instances, households fall within two FOMC meeting windows for longer window lengths. We exclude these responses from the analysis.

[14] For each month, we winsorize the tails of the distribution of inflation expectations at the 10 percent

vey weights and is based on the entire sample of inflation expectations from June 2013 to September 2023.

Figure 7 shows the estimates of the β_2 coefficient for 1 and 3-year ahead inflation expectations, based on the alternative time windows surrounding FOMC meetings.[15] The left chart shows that media sentiment has no material effect on short-term inflation expectations. The point estimates for β are negative, especially for windows between 3 and 7 days, but they are not statistically significant. This result is not surprising given that short-term inflation expectations tend to be highly affected by recent inflation developments. Central bank communication is indeed more geared at influencing medium-term inflation expectations rather than near-term expectations, also in light of the lags associated with monetary policy transmission.

The media sentiment exercises instead a stronger impact on medium-term inflation expectations. The right chart shows that when the media sentiment becomes more hawkish, households tend to revise down their expectations about the 3-year ahead inflation rate. The quantitative effects are sizeable. Using the average estimate across the different windows, a one-standard-deviation hawkish shift in media sentiment tends to lower 3-year ahead inflation expectations by about 0.18 percentage points. These results provide encouraging evidence about central banks' ability to affect households' inflation expectations via the media.

Figure 7: The impact of media sentiment on households' inflation expectations

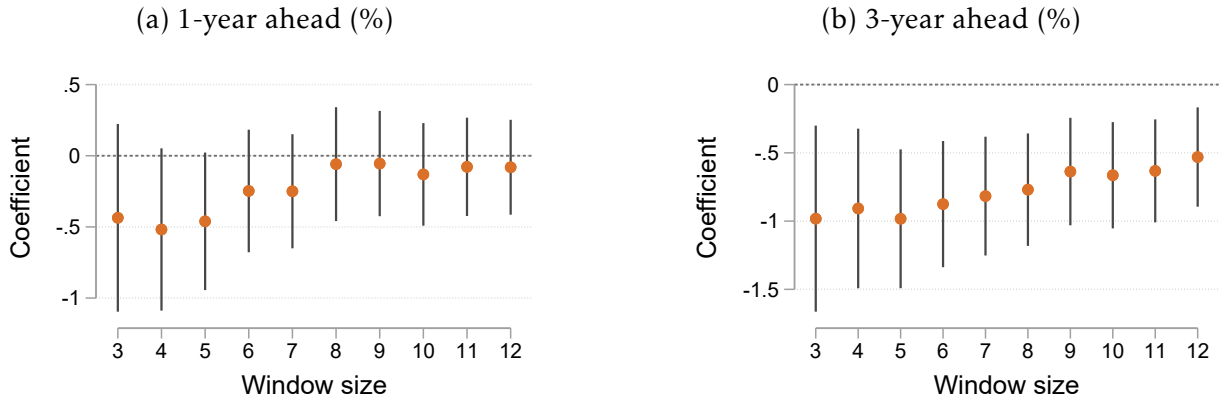

Notes: The figure reports estimates of the coefficient β_2 in equation (4) for different windows surrounding FOMC meetings, with 90% confidence intervals. The left and right panels consider inflation expectations over the next 12 months and between 24 and 36 months, respectively.

level to remove outliers. The regression results are similar if data are winsorized at the 5 percent level, although less precisely estimated, as shown in Appendix Figure D.1. Note that we exclude the first 6 interview rounds for each respondent since inflation expectations tend to be higher and more uncertain at the beginning of the survey participation (Kim and Binder, 2023).

[15] See Appendix Tables D.1 and D.2 for the regression estimates.

It is also interesting to examine whether the impact of media sentiment on households' expectations has varied over time. During the period from June 2013 (the starting period of the analysis) until the onset of the COVID-19 pandemic, inflation rates remained persistently low and stable. As a result, households might have paid less attention to monetary policy news and have been less likely to revise inflation expectations based on such information. In contrast, the post-pandemic surge in inflation and central banks' explicit efforts to manage inflation expectations through public communications, may have heightened households' attention to monetary policy news, thereby amplifying its influence on inflation expectations. The econometric results support this hypothesis. Figure 8 shows estimates of the β_2 coefficient based on the sample before and after 2020. The coefficient estimates for the post-2020 period are considerably larger than those for the pre-2020 period. This indicates that media coverage of monetary policy tends to have much stronger effects on households' inflation expectations during times when inflation is high and volatile.

Figure 8: The impact of media sentiment on 3-year-ahead inflation expectations

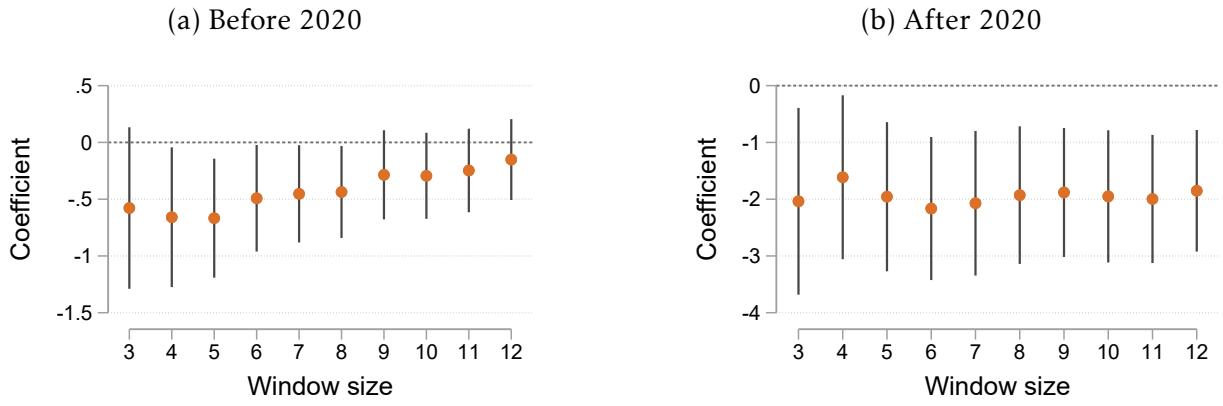

Notes: The figure reports estimates of the coefficient β_2 in equation (4) for different windows surrounding FOMC meetings, with 90% confidence intervals, over different periods. The left panel considers the period up to 2020, while the right panel covers the period from 2020 onwards.

A possible concern with the results presented so far is that changes in inflation expectations around FOMC meetings might be directly influenced by the FOMC communication rather than by the media coverage. To explore this issue, we estimate the following expanded version of equation 4:

$$\pi^e_{it} = \alpha_i + \gamma_m + \left(\beta_1 + \beta_2 \Delta S^{News}_m + \beta_3 \Delta S^{FOMC}_m\right) Post_{im} + u_{it} \tag{5}$$

which now also controls for the change in the FOMC sentiment ΔS^{FOMC}_m. Note that since we do not have a measure of FOMC sentiment right before the meeting, we construct

ΔS_m^{FOMC} as the difference between the FOMC sentiment of meeting m and the media sentiment in the 2 days before the meeting, mirroring the procedure to construct the change in the media sentiment ΔS_m^{News}.

Table 4 reports the estimates for the regression coefficients β_2 and β_3. As shown earlier in the paper, the FOMC and media sentiment tend to be highly correlated, making it challenging to disentangle their effects. However, the results still paint a consistent picture across different regression windows. We continue to see that a hawkish turn in media sentiment tends to reduce medium-term inflation expectations. In contrast, we do not detect any effect of the FOMC sentiment on inflation expectations, at both horizons. These results are consistent with the premise of the analysis, namely that households do not directly read FOMC statements or listen to the press conferences. Therefore, the media plays a critical role in conveying monetary policy information to the public and, in turn, influencing households' inflation expectations.

Table 4: Sensitivity of inflation expectations to media and FOMC sentiment

	1-year ahead (%)				3-year ahead (%)			
	(1) 3-day	(2) 5-day	(3) 7-day	(4) 9-day	(5) 3-day	(6) 5-day	(7) 7-day	(8) 9-day
Post-FOMC X Δ News Sentiment	-0.519	-0.532*	-0.280	-0.057	-0.810*	-0.914***	-0.775***	-0.530**
	(0.409)	(0.299)	(0.250)	(0.229)	(0.445)	(0.325)	(0.276)	(0.250)
Post-FOMC X Δ FOMC sentiment	0.142	0.151	0.068	0.006	-0.295	-0.146	-0.095	-0.249
	(0.306)	(0.228)	(0.203)	(0.175)	(0.324)	(0.260)	(0.246)	(0.204)
Meetings	82	82	82	82	82	82	82	82
Respondents	11739	11737	11732	11731	11734	11732	11727	11726
Observations	60902	60826	60725	60536	60907	60831	60731	60542
R2	0.74	0.74	0.74	0.75	0.70	0.70	0.70	0.70

Notes: Standard errors are clustered at the respondent level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. For each window length, each column corresponds to a different measure of inflation expectations. Columns (1), (2), (3) and (4) refer to the 12m-ahead point estimate. Columns (5), (6), (7) and (8) refer to the point estimate for the inflation rate in 24/36 months.

7 Conclusion

Over the last decades, central banks have placed increasing emphasis on public communication, by explaining the rationale of policy decisions and providing indications about the future stance. A key goal underpinning these efforts is to enhance the effectiveness of monetary policy and strengthen the anchoring of inflation expectations. A large body of literature has documented that central bank communication can strongly influence financial markets. However, much less is known about the impact on the broader public, especially on households.

This paper contributes to shedding light on central banks' ability to influence the public by examining the degree of consistency between FOMC communication and the media coverage. We focus on media coverage because survey evidence shows that households rely on the media to gather information on monetary policy and central banks' decisions. Hence, a strong pass-through of FOMC communication to the media coverage is a key pre-condition for central banks' ability to reach the public.

The analysis shows that FOMC communication is generally well portrayed by the media, especially when the FOMC sentiment is particularly dovish or hawkish. However, there have been notable differences across time. The communication pass-through deteriorated considerably post-GFC, when the ZLB became binding and the Fed had to conduct monetary policy by deploying new tools. The pass-through then strengthened with the introduction of the post-FOMC press conferences. The answers provided by the Fed Chair during press conferences are now the most influential communication tool to shape media coverage. The analysis also underscores the communication challenges faced by new chairs, as it documents a lower pass-through to media coverage in the first few months of a chair's tenure. Finally, we show that when media coverage turns more hawkish, households tend to lower medium-term inflation expectations. This effect was particularly strong during the post-pandemic inflation surge, hinting that households may pay more attention to monetary policy news when inflation is high and volatile, exactly when central bank communication becomes more critical. We, instead, do not detect any association between the sentiment directly expressed in FOMC communication and households' expectations, underscoring the critical role played by the media in channelling central banks' communication to the public.

Looking ahead, future research could expand the analysis along various dimensions. First, it would be valuable to assess the role of governors' speeches in influencing media coverage, as these speeches also convey important indications about the monetary policy stance (Swanson, 2023). Second, the work could be extended to examine the consistency between central communication and media coverage concerning the economic outlook. In this context, it would be interesting to investigate potential differences in the coverage of the economic outlook across media outlets, which may reflect different editorial or political leanings. Finally, it would be worthwhile to examine the sentiment regarding monetary policy and the economic outlook portrayed on social media, given the growing use of these platforms among large segments of the population.

References

Alexopoulos, Michelle, Xinfen Han, Oleksiy Kryvtsov, and Xu Zhang. 2024. "More than words: Fed Chairs' communication during congressional testimonies." *Journal of Monetary Economics*, 142: 103515.

Andrade, Philippe, and Filippo Ferroni. 2021. "Delphic and odyssean monetary policy shocks: Evidence from the euro area." *Journal of Monetary Economics*, 117: 816–832.

Berger, Helge, Michael Ehrmann, and Marcel Fratzscher. 2011. "Monetary policy in the media." *Journal of Money, Credit and Banking*, 43(4): 689–709.

Bernanke, Ben S, and Kenneth N Kuttner. 2005. "What explains the stock market's reaction to Federal Reserve policy?" *The Journal of Finance*, 60(3): 1221–1257.

Binder, Carola Conces, Jeffrey R Campbell, and Jane M Ryngaert. 2024. "Consumer inflation expectations: Daily dynamics." *Journal of Monetary Economics*, 103613.

Blinder, Alan S, Michael Ehrmann, Jakob de Haan, and David-Jan Jansen. 2023. "Central Bank Communication with the General Public: Promise or False Hope?" *Journal of Economic Literature*.

Blinder, Alan S, Michael Ehrmann, Marcel Fratzscher, Jakob De Haan, and David-Jan Jansen. 2008. "Central bank communication and monetary policy: A survey of theory and evidence." *Journal of Economic Literature*, 46(4): 910–945.

Cieslak, Anna, and Andreas Schrimpf. 2019. "Non-monetary news in central bank communication." *Journal of International Economics*, 118: 293–315.

Cieslak, Anna, and Michael McMahon. 2023. "Tough Talk: The Fed and the Risk Premium." Available from https://ssrn.com/abstract=4560220.

Cieslak, Anna, Michael McMahon, and Hao Pang. 2024. "Did I Make Myself Clear? The Fed and the Market in the Post-2020 Framework Period." Available from https://ssrn.com/abstract=4920858.

Claus, Edda, and Viet Hoang Nguyen. 2020. "Monetary policy shocks from the consumer perspective." *Journal of Monetary Economics*, 114: 159–173.

Cochrane, John H, and Monika Piazzesi. 2002. "The Fed and interest rates-a high-frequency identification." *American Economic Review*, 92(2): 90–95.

Coibion, Olivier, Dimitris Georgarakos, Yuriy Gorodnichenko, and Michael Weber. 2023. "Forward Guidance and Household Expectations." *Journal of the European Economic Association*.

Coibion, Olivier, Yuriy Gorodnichenko, and Michael Weber. 2022. "Monetary policy communications and their effects on household inflation expectations." *Journal of Political Economy*, 130(6): 1427–1716.

Conrad, Christian, Zeno Enders, and Alexander Glas. 2022. "The role of information and experience for households' inflation expectations." *European Economic Review*, 143: 104015.

Curti, Filippo, and Sophia Kazinnik. 2023. "Let's face it: Quantifying the impact of nonverbal communication in FOMC press conferences." *Journal of Monetary Economics*, 139: 110–126.

De Fiore, Fiorella, Marco J Lombardi, and Johannes Schuffels. 2021. "FOMC Announcements and Household Inflation Expectations." *BIS Working Paper No 956*.

Doh, Taeyoung, Dongho Song, and Shu-Kuei Yang. 2022. "Deciphering federal reserve communication via text analysis of alternative fomc statements." *Federal Reserve Bank of Kansas City Working Paper Forthcoming*.

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

Gambacorta, Leonardo, Byeungchun Kwon, Taejin Park, Pietro Patelli, and Sonya Zhu. 2024. "CB-LMs: language models for central banking." *BIS Working Paper 1215*.

Gardner, Ben, Chiara Scotti, and Clara Vega. 2022. "Words speak as loudly as actions: Central bank communication and the response of equity prices to macroeconomic announcements." *Journal of Econometrics*, 231(2): 387–409. Special Issue: The Econometrics of Macroeconomic and Financial Data.

Gardt, Marius, Siria Angino, Simon Mee, and Gabriel Gl'ockler. 2021. "ECB Communication with the Wider Public." *ECB Economic Bulletin 8/2021*.

Gorodnichenko, Yuriy, Tho Pham, and Oleksandr Talavera. 2023. "The voice of monetary policy." *American Economic Review*, 113(2): 548–584.

Gürkaynak, Refet S, Brian P Sack, and Eric T Swanson. 2005. "Do actions speak louder than words? The response of asset prices to monetary policy actions and statements." *International Journal of Central Banking*, 1(1): 55–93.

Haldane, Andrew, and Michael McMahon. 2018. "Central bank communications and the general public." 108: 578–583.

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

Hansen, Stephen, Michael McMahon, and Matthew Tong. 2019. "The long-run information effect of central bank communication." *Journal of Monetary Economics*, 108: 185–202.

Hayo, Bernd, and Edith Neuenkirch. 2018. "The influence of media use on layperson monetary policy knowledge in Germany." *Scottish Journal of Political Economy*, 65(1): 1–26.

Hayo, Bernd, and Matthias Neuenkirch. 2012. "Bank of Canada communication, media coverage, and financial market reactions." *Economics Letters*, 115(3): 369–372.

Hayo, Bernd, and Matthias Neuenkirch. 2015. "Self-monitoring or reliance on media reporting: How do financial market participants process central bank news?" *Journal of Banking & Finance*, 59: 27–37.

Hendry, Scott. 2012. "Central Bank Communication or the Media's Interpretation: What Moves Markets?" *Bank of Canada Working Paper 2012-9*.

Kim, Gwangmin, and Carola Binder. 2023. "Learning-through-Survey in Inflation Expectations." *American Economic Journal: Macroeconomics*, 15(2): 254–78.

Kuttner, Kenneth N. 2001. "Monetary policy surprises and interest rates: Evidence from the Fed funds futures market." *Journal of Monetary Economics*, 47(3): 523–544.

Lamla, Michael J, and Dmitri V Vinogradov. 2019. "Central bank announcements: Big news for little people?" *Journal of Monetary Economics*, 108: 21–38.

Lewis, Daniel J, Christos Makridis, and Karel Mertens. 2019. "Do monetary policy announcements shift household expectations?" *FRB of New York Staff Report*, , (897).

Lucca, David O, and Francesco Trebbi. 2009. "Measuring Central Bank Communication: An Automated Approach with Application to FOMC Statements." National Bureau of Economic Research Working Paper 15367.

Munday, Tim, and James Brookes. 2021. "Mark my words: the transmission of central bank communication to the general public via the print media." *Staff Working Paper No. 944*.

Neuenkirch, Matthias. 2014. "Federal Reserve communications and newswire coverage." *Applied Economics*, 46(25): 3119–3129.

Neuhierl, Andreas, and Michael Weber. 2019. "Monetary policy communication, policy slope, and the stock market." *Journal of Monetary Economics*, 108: 140–155.

Pinter, Julien, and Evžen Kočenda. 2023. "Media Treatment of Monetary Policy Surprises and Their Impact on Firms' and Consumers' Expectations." *Journal of Money, Credit and Banking*.

Schmanski, Bennett, Chiara Scotti, Clara Vega, and Hedi Benamar. 2023. "Fed Communication, News, and Echo Chambers." *Finance and Economics Discussion Series 2023-036*.

Shah, Agam, Suvan Paturi, and Sudheer Chava. 2023. "Trillion Dollar Words: A New Financial Dataset, Task Market Analysis." *arXiv preprint arXiv:2305.07972*.

Swanson, Eric T. 2021. "Measuring the effects of federal reserve forward guidance and asset purchases on financial markets." *Journal of Monetary Economics*, 118: 32–53.

Swanson, Eric T. 2023. "The importance of fed chair speeches as a monetary policy tool." *AEA Papers and Proceedings*, 113: 394–400.

Ter Ellen, Saskia, Vegard H Larsen, and Leif Anders Thorsrud. 2022. "Narrative monetary policy surprises and the media." *Journal of Money, Credit and Banking*, 54(5): 1525–1549.

Yellen, Janet L. 2012. "Revolution and Evolution in Central Bank Communications." Speech delivered at the Haas School of Business, University of California, Berkeley, November 13.

Appendix

A Selection procedure for media articles

As described in the paper, we implement a 3-step cleaning procedure to eliminate articles that are not relevant for the analysis. First, we drop transcripts of FOMC communication, for example as distributed via newsletters. Second, we eliminate articles that are very similar to each other, for example as in the case of small revisions to articles posted online. We perform this step by measuring the "cosine similarity" between articles from the same newspaper covering the same FOMC meeting. This index quantifies the similarity between two texts by computing the cosine of the angle between their vector representations in a multi-dimensional space. For our purpose, these vectors represent the frequency of each word in each text. Articles with a similarity score exceeding 0.9 and with a length ratio below 1.5 are filtered out. If these articles are published on the same day, we keep the longest one. Otherwise, we keep the first one published. Note that we retain articles with a high similarity score but a length ratio exceeding 1.5 because these are more likely to reflect substantial revisions. Third, we use GPT-4 to select articles that are sufficiently focused on monetary policy decision by the Federal Reserve. We do so by providing GPT-4 with the following prompt: "You're a research assistant working for the Fed. You have a Master's degree in Economics. Your task is to classify the following articles on a scale from 0 to 5 depending on the extent to which they focus on a monetary policy decision by the Federal Reserve. Specifically, 0 means that the article is unrelated to a monetary policy decision by the Federal Reserve. On the opposite end, 5 means that the article is squarely focused on a monetary policy decision by the Federal Reserve." We keep the articles that have a score of 4 or 5. This cleaning procedure leaves us with a final sample of 14,021 articles. Table A.1 shows the distribution of articles across media outlets.

Table A.1: Distribution of articles covering FOMC meetings across media outlets

Newspaper	Number of articles	Percent
Wall Street Journal	5,118	36.50
Financial Times	3,929	28.02
Washington Post	1,498	10.68
USA Today	1,083	7.72
Los Angeles Times	1,049	7.48
CNN	850	6.06
New York Post	456	3.25
Fox News	38	0.27
Total	14,021	100

B Additional results on LLM comparison

Table B.1: Confusion matrices for press conference answers across LLMs

Actual	Predictions					
	GPT-4			RoBERTa-SPC23		
	Hawkish	Neutral	Dovish	Hawkish	Neutral	Dovish
Hawkish	16.2	70.94	12.8	34.2	41.9	23.9
Neutral	1.5	91.4	7.1	6.8	74.7	18.5
Dovish	0.8	61.1	38.1	7.1	38.9	54.0

Notes: The confusion matrix rows represent the actual classes, while the columns represent the predicted classes. Each entry is the percentage of data points classified into a given class.

C Additional results regarding sentiment pass-through

Table C.1: Sentiment pass-through by media outlet

	(1) WSJ	(2) FT	(3) WP	(4) LAT	(5) CNN	(6) USAT	(7) NYP	(8) FOX
FOMC sentiment	0.598***	0.586***	0.665***	0.729***	0.751***	0.798***	0.737***	0.770***
	(0.014)	(0.017)	(0.020)	(0.021)	(0.029)	(0.023)	(0.037)	(0.171)
Constant	0.107***	0.108***	0.110***	0.064***	0.149***	0.083***	0.093***	0.180
	(0.008)	(0.011)	(0.013)	(0.016)	(0.019)	(0.017)	(0.029)	(0.106)
No. of meetings	224	161	223	210	104	213	173	13
Observations	3792	2903	1178	810	642	769	355	23
R2	0.32	0.27	0.48	0.57	0.48	0.56	0.45	0.48

Notes: Heteroskedasticity robust standard errors. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Figure C.1: Trailing rolling window estimation - alternative windows

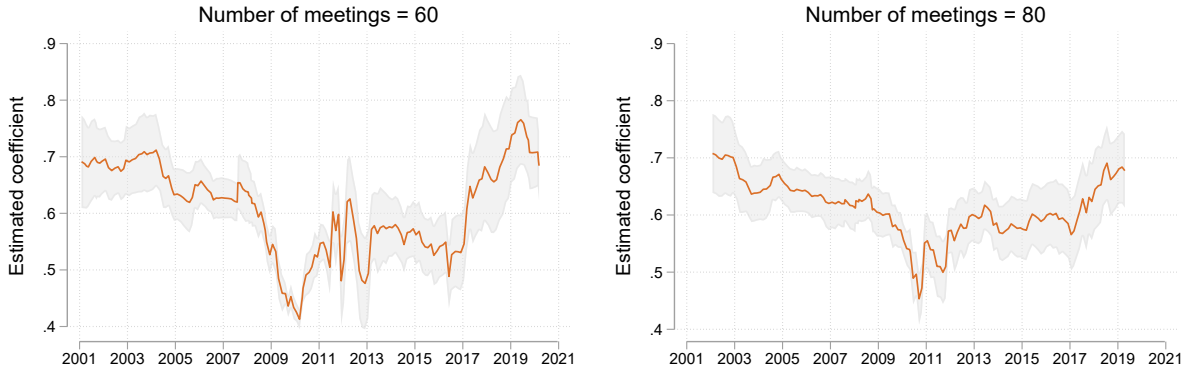

Notes: Estimated over overlapping windows. The grey area indicates the 90 percent confidence interval. Standard errors are clustered at the newspaper level.

D The effects of media sentiment on inflation expectations

Table D.1: Media sentiment and 1-year ahead inflation expectations

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Window length (days)	3	4	5	6	7	8	9	10	11	12
Post-FOMC X Δ News Sentiment	-0.436	-0.518	-0.461	-0.247	-0.250	-0.059	-0.055	-0.131	-0.078	-0.081
	(0.401)	(0.347)	(0.294)	(0.262)	(0.244)	(0.243)	(0.225)	(0.219)	(0.210)	(0.203)
Meetings	82	82	82	82	82	82	82	82	82	82
Respondents	11739	11739	11737	11736	11732	11731	11731	11730	11729	11729
Observations	60902	60875	60826	60780	60725	60631	60536	60394	60337	60251
R2	0.74	0.74	0.74	0.74	0.74	0.75	0.75	0.75	0.75	0.75

Notes: Standard errors are clustered at the respondent level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Each column corresponds to a different window length.

Table D.2: Media sentiment and 3-year ahead inflation expectations

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Window length (days)	3	4	5	6	7	8	9	10	11	12
Post-FOMC X Δ News Sentiment	-0.982**	-0.907**	-0.983***	-0.876***	-0.817***	-0.770***	-0.637***	-0.664***	-0.632***	-0.531**
	(0.414)	(0.355)	(0.309)	(0.281)	(0.265)	(0.250)	(0.239)	(0.237)	(0.229)	(0.221)
Meetings	82	82	82	82	82	82	82	82	82	82
Respondents	11734	11734	11732	11731	11727	11726	11726	11725	11724	11724
Observations	60907	60880	60831	60786	60731	60637	60542	60400	60343	60257
R2	0.70	0.70	0.70	0.70	0.70	0.70	0.70	0.70	0.70	0.70

Notes: Standard errors are clustered at the respondent level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Each column corresponds to a different window length.

Figure D.1: Media sentiment on households' inflation expectations - 5% winsorization

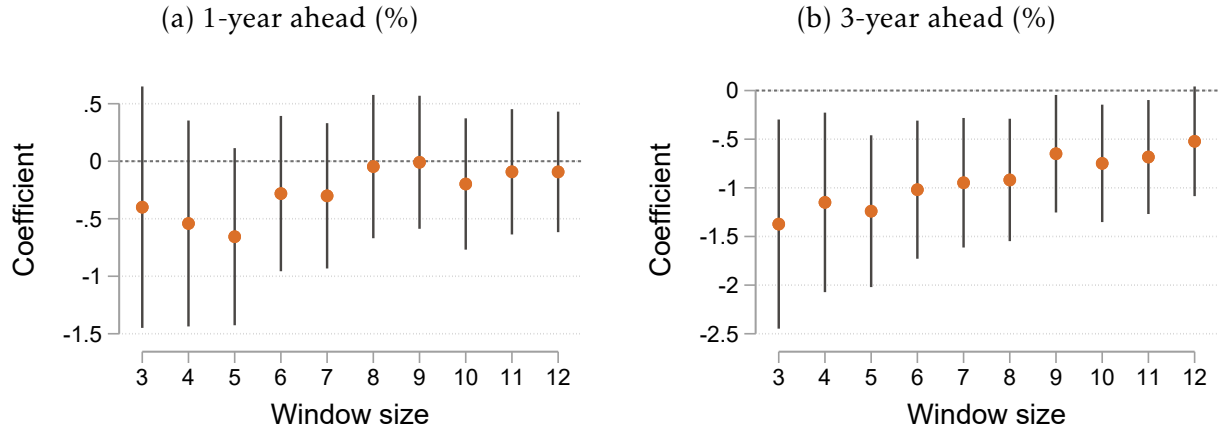

Notes: The figure reports estimates of the coefficient β_2 in equation (4) for different windows surrounding FOMC meetings, with 90% confidence intervals. The left and right panels consider inflation expectations over the next 12 months and between 24 and 36 months, respectively. Inflation expectations are winsorized at the 5% level.

Previous volumes in this series

1230 December 2024	Whither inflation targeting as a global monetary standard?	Claudio Borio
1229 November 2024	Through stormy seas: how fragile is liquidity across asset classes and time?	Nihad Aliyev, Matteo Aquilina, Khaladdin Rzayev and Sonya Zhu
1228 November 2024	Retail fast payment systems as a catalyst for digital finance	Giulio Cornelli, Jon Frost, Jonathan Warren, Clair Yang and Carolina Velásquez
1227 November 2024	Decentralised dealers? Examining liquidity provision in decentralised exchanges	Matteo Aquilina, Sean Foley, Leonardo Gambacorta and William Krekel
1226 November 2024	How Does Fiscal Policy affect the Transmission of Monetary Policy into Cross-border Bank Lending? Cross-country Evidence	Swapan-Kumar Pradhan, Előd Takáts and Judit Temesvary
1225 October 2024	Trade fragmentation, inflationary pressures and monetary policy	Ludovica Ambrosino, Jenny Chan and Silvana Tenreyro
1224 October 2024	A theory of economic coercion and fragmentation	Chris Clayton, Matteo Maggiori and Jesse Schreger
1223 October 2024	Carbon prices and reforestation in tropical forests	José A Scheinkman
1222 October 2024	Artificial intelligence and big holdings data: opportunities for central banks	Xavier Gabaix, Ralph S J Koijen, Robert Richmond and Motohiro Yogo
1221 October 2024	Bank geographic diversification and funding stability	Sebastian Doerr
1220 October 2024	The road to net zero: a fund flow investigation	Louisa Chen and Koji Takahashi
1219 October 2024	Stablecoins, money market funds and monetary policy	Iñaki Aldasoro, Giulio Cornelli, Massimo Ferrari Minesso, Leonardo Gambacorta, Maurizio Michael Habib
1218 October 2024	Bank specialisation and corporate innovation	Hans Degryse, Olivier De Jonghe, Leonardo Gambacorta and Cédric Huylebroek

All volumes are available on our website www.bis.org.