




One Fed, Many Voices:

Coordinated Communication vs. Transparent Debate

**Milena Djourelova, Filippo Ferroni,
Leonardo Melosi, and Alessandro T. Villa**

REVISED
March 27, 2026

WP 2025-23

<https://doi.org/10.21033/wp-2025-23>

FEDERAL RESERVE BANK *of* CHICAGO

*Working papers are not edited, and all opinions are the responsibility of the author(s). The views expressed do not necessarily reflect the views of the Federal Reserve Bank of Chicago or the Federal Reserve System.



One Fed, Many Voices: Coordinated Communication vs. Transparent Debate*

Milena Djourelova[†]

USC

Filippo Ferroni[‡]

University of Bologna

Leonardo Melosi[§]

EUI & CEPR

Alessandro T. Villa[¶]

FRB Chicago

March 27, 2026

Abstract

We study how observed coordination in central bank communication affects monetary policy transmission. We assess the alignment of 481 speeches by FOMC members since 2007 with post-meeting communication using a text-similarity measure (cosine similarity). To identify policy shocks, we use high-frequency changes in eurodollar futures around each speech. On average, these monetary policy speech surprises have no significant effects on inflation expectations or equity prices. However, speech surprises more closely aligned with prior FOMC communications strengthen monetary policy transmission, while those reflecting visible disagreement significantly weaken it. A general equilibrium model with incomplete information rationalizes these findings.

JEL codes: E52; E58; G14; C55; D83.

Keywords: Monetary policy transmission, text analysis, Cosine similarity, Large language models, FOMC speeches, High-frequency identification

*We thank Spencer Krane for his useful comments. We are also thankful to Karen Petrosyan and Sankalp Yadav for their excellent research assistance. The views in this paper are solely those of the authors and should not be interpreted as reflecting the views of the Federal Reserve Bank of Chicago, or any other person associated with the Federal Reserve System.

[†]University of Southern California. 3620 South Vermont Avenue, LA, USA. E-mail: mdjourel@gmail.com.

[‡]University of Bologna. Piazza Antonio Scaravilli 2, Bologna, Italy. E-mail: filippo.ferroni2@unibo.it.

[§]European University Institute. Via delle Fontanelle 18, Fiesole, Italy. E-mail: leonardo.melosi@eui.eu.

[¶]Federal Reserve Bank of Chicago. 230 S. LaSalle St., Chicago, USA. E-mail: forket86@gmail.com.

1 INTRODUCTION

In the United States—as in many other countries—monetary policy is set by a committee, the Federal Open Market Committee (FOMC), which brings together the Board of Governors and the presidents of the 12 Federal Reserve Banks. Both governors and Reserve Bank presidents regularly deliver public speeches conveying their views on the economic outlook and the appropriate stance of policy. At times, these remarks closely reflect the Committee’s consensus; at others, they reveal a broader range of individual perspectives. We leverage this variation in speech alignment to study whether coordination and transparent debate affect monetary policy transmission.

We assemble a dataset of 481 speeches delivered by FOMC participants since 2007. Each speech is time-stamped to the minute, allowing us to match it with five-minute eurodollar futures prices and extract the associated high-frequency “monetary policy speech” (MPS) surprise. We measure the verbal alignment of each speech with the most recent FOMC statement using *cosine similarity*, a standard text-based metric in natural language processing. This measure of alignment provides a proxy for coordination in policy communication because greater alignment with the FOMC’s statement—whether deliberate or incidental—captures the extent to which individual speeches reinforce or dilute a unified policy signal from the private sector’s perspective. We therefore interpret measured alignment as the degree of coordinated signaling, regardless of policymakers’ underlying intent.

Three key findings emerge from our analysis. First, on average—before accounting for differences in textual alignment with the preceding Chair’s press conference—monetary policy surprises have no statistically significant effect on inflation expectations or equity prices. Second, greater alignment of member speeches with the Chair’s post-meeting communication amplifies policy signals: monetary policy tightening (easing) raises (lowers) short-term yields and lowers (raises) inflation expectations and equity prices. Third, less similar speeches attenuate prior policy effects, with market reactions muted or even reversed.

The cosine similarity measure represents each text as a weighted list of words, where *term frequency-inverse document frequency* (TF-IDF) assigns greater weight to words that are distinctive within the corpus. The cosine similarity then measures the angle between two such representations: a value closer to one indicates greater alignment with the Chair’s message, while a value closer to zero indicates that the messages are orthogonal. This provides a useful and convenient measure of the extent to which a speech reinforces or departs from the Chair’s post-meeting communication.

Our empirical analysis proceeds in two steps. First, we estimate daily local projections of Treasury yields, breakeven inflation, and equity prices on monetary policy speech (MPS)

surprises. Unconditionally, the effects are modest: a one-basis-point MPS tightening raises medium-term yields by about one basis point, with little persistence, and has no discernible impact on market-implied inflation expectations. However, conditioning on similarity reveals a striking pattern. Speeches that closely align with the Chair’s message produce a front-loaded tightening: short-term yields rise and breakeven inflation expectations fall, mirroring the transmission of standard monetary policy surprises.

In contrast, speeches that differ from the Chair’s message lift the entire yield curve, raise inflation expectations, and are accompanied by equity gains, suggesting that markets interpret them as signals about future fundamentals rather than about the policy stance itself. In this sense, speeches that differ from the Chair’s press conference may reflect more persistent shifts in the policy framework, such as a focus on long-term objectives or structural considerations, which can give rise to neo-Fisherian dynamics. Alternatively, these deviations may convey new information about the macroeconomic outlook, consistent with the Delphic (or informational) effect emphasized in the literature.

Because bag-of-words similarity can miss semantic proximity created by paraphrasing or differences in rhetorical style, we complement our baseline TF-IDF cosine measure with an LLM-based semantic similarity score. This alternative measure evaluates whether a speech and the Chair’s press conference convey the same underlying message, beyond shared wording. Using the LLM-based index in our interaction regressions delivers the same qualitative conclusions: speeches that are more aligned with the Chair amplify the effects of policy surprises, while less aligned speeches are associated with markedly different market responses. These patterns are robust to alternative event windows, and additional checks—excluding Chair speeches and filtering similarity through contemporaneous FOMC deliberations—confirm the results.

To rationalize these findings, we propose a simple imperfect information model embedding a standard New Keynesian core in which private agents form expectations not only from current macroeconomic conditions, such as inflation, output, and interest rates, but also from forward-looking central bank speeches. These speeches serve as noisy signals of next-period monetary policy and natural rate innovations. We show that this simple framework can replicate the evidence: when speeches are aligned with each other, they strengthen policy effectiveness, while low-similarity speeches cause inflation to increase following anticipated contractionary monetary shocks.

Our findings offer new evidence that monetary policy communication is not a “one-man show.” Markets extract meaningful information from a broader set of officials, and transparency can emerge through these decentralized channels.

Literature Review. This paper contributes to the growing literature on central bank communication, which has traditionally focused on formal and centralized communication tools such as policy statements, transcripts, and press conferences. Early contributions include [Hansen et al. \(2017\)](#) and [Cieslak and Schrimpf \(2019\)](#), who study the effects of statements and transcripts released after FOMC meetings. [Coibion et al. \(2022\)](#) examine the market impact of Chair press conferences, while [Aruoba and Drechsel \(2024\)](#) analyze documents prepared by Fed staff prior to policy decisions.

In contrast, a newer strand of research focuses on speeches by FOMC members, particularly regional Fed presidents and Governors. Although this literature is still developing, it is expanding rapidly. [Neuhierl and Weber \(2019\)](#) show that speeches by the Fed Chair or Vice Chair help predict the slope of the yield curve. [Ehrmann et al. \(2021\)](#) demonstrate that voting rights influence the number and tone of speeches delivered by Fed presidents, with more frequent and impactful speeches given when members are not voting.

In the context of high-frequency identification, [Bauer and Swanson \(2023a\)](#) reassess the interpretation of monetary policy surprises, arguing for a more refined decomposition of policy shocks that distinguishes between pure monetary policy actions and information effects. Similarly, [Swanson \(2023\)](#) and [Swanson and Jayawickrema \(2024\)](#) extend high-frequency approaches to unscheduled communications, including speeches and testimonies. [Swanson \(2023\)](#) and [Swanson and Jayawickrema \(2024\)](#) highlight that speeches by the Fed Chair can be more influential than official FOMC statements in shaping market expectations. Relatedly, [Malmendier et al. \(2021\)](#) study whether policymakers’ views on monetary policy can be inferred from their speech content, including their choice of language and tone. [Istrefi et al. \(2023\)](#) explore whether FOMC members’ expressed concerns about financial stability—captured through a constructed index based on their speeches—help explain policy actions.

Another challenge lies in the fact that the relative importance of these effects likely depends on the language used by policymakers in their communications ([Lunsford, 2020](#)). Language and tone can act as signals themselves, shaping market expectations independently of the policy stance conveyed.

Building on this literature, our paper focuses on the market effects of speeches by all FOMC members—not just the Chair—and introduces a novel textual similarity measure that captures how aligned a speech is with the Chair’s post-meeting press conference. It represents the first attempt to combine the intraday market reactions to monetary policy communication with the textual content of such communication. We show that speeches closely aligned with the Chair’s message produce effects similar to conventional monetary shocks, while those emphasizing different themes affect markets through distinct channels—highlighting

that decentralized communication can enhance transparency, yet also introduce interpretive complexity.

The remainder of the paper is organized as follows. Section 2 describes the construction of the textual datasets. Section 3 introduces the high-frequency database, examines financial market reactions to speeches in isolation, and extends the analysis by interacting the similarity metrics with the high-frequency identification measure. Section 4 presents a set of robustness checks. Section 5 presents a simple model that rationalizes the empirical findings. Section 6 concludes.

2 DATA SOURCES

Our analysis draws on the following three textual sources related to Federal Reserve communications, which we transform into structured datasets capturing the timing, authorship, and content of monetary policy deliberations and statements.

FOMC Press Conferences. Beginning in 2011, the Fed Chair has held press conferences following FOMC meetings. We obtain the full text of these conferences from the same Federal Reserve archive. We parse each document to isolate individual questions and responses, generating a dataset that includes the date, the text of each journalist’s question, and the corresponding reply from the Chair. This allows us to quantify the tone, content, and structure of real-time public communication by the Fed.

Public Speeches by FOMC Participants. We also compile a corpus of public speeches by individual FOMC participants—including Board members and Presidents of the twelve regional Federal Reserve Banks—delivered between 2007 and 2018. The majority of these speeches are drawn from the FRASER archive ([link](#)), with additional documents obtained directly from Reserve Bank websites (e.g., [Minneapolis Fed](#)). Where available, we record the exact start and end times of speeches. In cases where end times are not reported, we impute them based on speech length and an estimated words-per-minute rate to ensure comparability across events.

FOMC Meeting Transcripts. Finally, we collect a corpus covering internal FOMC deliberations. Since 2007, the Federal Reserve has published verbatim transcripts of FOMC meetings with a five-year lag. We collect these transcripts from the Federal Reserve Board’s official archive and convert them from PDF into a structured format.¹ For each meeting,

¹The original transcripts are available at this [link](#).

we extract the date, identify individual speakers using a harmonized naming convention, and record the full text of their interventions. This dataset enables speaker-level analysis of deliberation dynamics within the FOMC.

2.1 TEXT SIMILARITY BETWEEN INDIVIDUAL AND CENTRALIZED COMMUNICATION

To quantify the alignment between individual and centralized communications of the Federal Reserve, we construct a measure of similarity between the speeches given by individual FOMC members and those delivered by the Fed Chair in the most recent FOMC press conference.

We begin by standardizing the raw textual data: (i) converting all text to lowercase and remove punctuation and non-alphanumeric characters; (ii) lemmatizing words to their base forms to account for inflectional variation; (iii) removing common stopwords using the standard English dictionary from the Natural Language Toolkit (NLTK); and (iv) manually harmonizing the names of FOMC participants across datasets to ensure consistent speaker identification.

We then represent each text in the corpus numerically using the Term Frequency–Inverse Document Frequency (TF-IDF) method. This assigns a weight to each word in a document based on two components: term frequency, which reflects how often a word appears in a given document, and inverse document frequency, which downweights words that are common across many documents and thus less informative. The resulting TF-IDF vectors provide a high-dimensional representation of each statement that captures its semantic content.

Once statements are vectorized, we compute pairwise cosine similarity scores to quantify the textual alignment between individual FOMC participants and the Fed Chair. Cosine similarity captures the angular distance between two TF-IDF vectors, with values closer to one indicating greater semantic similarity. For each meeting, we compare a participant’s full set of statements to the complete set of the Chair’s remarks—including both prepared comments and responses to journalists—thereby capturing alignment with the totality of the Chair’s public communication. This procedure yields a participant-level similarity score for each meeting. We apply the same methodology to FOMC member speeches, comparing their content to the Chair’s full press conference remarks to assess how closely individual public communications echo the Chair’s messaging.

2.2 ILLUSTRATIVE EXAMPLES AND DISPERSION ACROSS FOMC-MEMBERS

To provide intuition for the cosine similarity scores computed in our analysis, we present a set of examples with particularly high values. Table 1 reports excerpts from selected

FOMC speeches and the corresponding Chair’s press conference remarks with which they exhibit the greatest similarity. These examples help validate that high similarity scores reflect substantive alignment in both language and content.

In the top-ranked case, Stanley Fischer’s February 2016 speech closely mirrors Janet Yellen’s December 2015 press conference, with a cosine similarity score of 0.81. Both emphasize the decision to raise the federal funds rate, the continued accommodative stance of monetary policy, and confidence in inflation returning to target. The alignment is not only thematic but also lexical, with near-identical phrasing across several statements.

Other high-similarity examples demonstrate similar patterns. Powell’s 2013 speech reiterates Bernanke’s emphasis on subdued inflation and the persistence of accommodative financial conditions. Likewise, speeches by Fischer (2015) and Brainard (2017) echo the Chair’s outlook on GDP growth, labor market strength, and external risks such as net exports and global uncertainty. Across these examples, high cosine similarity arises from both shared vocabulary and convergence in policy framing.

Together, these examples illustrate that our similarity measure captures meaningful variation in the degree to which individual FOMC participants’ public statements reflect or anticipate the policy narrative presented by the Chair.

Table 1: High Cosine Similarity Examples.

Information	Speech	Press Conference
Fischer (2016/02/01) Yellen (2015/12/16) <i>Score: 0.8133</i>	“...decided to raise the target range for the federal funds rate by $\frac{1}{4}$ percentage point, to $\frac{1}{4}$ to $\frac{1}{2}$ percent.” “...confidence that inflation would return to our 2 percent goal over the medium term.” “...monetary policy remains accommodative after the small increase in the federal funds.”	“...raise the target range for the federal funds rate by $\frac{1}{4}$ percentage point, bringing it to $\frac{1}{4}$ to $\frac{1}{2}$ percent.” “...reasonably confident that inflation would move back to its 2 percent objective over the medium term.” “...after today’s increase, the stance of monetary policy remains accommodative.”
Powell (2013/06/27) Bernanke (2013/06/19) <i>Score: 0.7945</i>	“Inflation is currently running below the FOMC’s 2 percent long-term objective.” “...purchased assets will remain on the Fed’s balance sheet for some time and continue to put downward pressure on rates.”	“Inflation has been running below the Committee’s longer-run objective of 2 percent.” “...adjusting the pace of asset purchases and its forward guidance regarding the target for the federal funds rate.”

Continued on next page

Table 1: High Cosine Similarity Examples. (Continued)

Information	Speech	Press Conference
Fischer (2015/10/11) Yellen (2015/09/17) <i>Score: 0.76</i>	• “Gross domestic product (GDP) growth in the first half of 2015 is now estimated to have been at an annual rate of 2-1/4 percent, and private forecasters are projecting GDP to continue to rise.” • “Moreover, net exports have served as a significant drag on growth over the past year and recent global economic and financial developments highlight the risk that a slowdown in foreign growth might restrain U.S. economic activity.” • “However, that is an expectation, not a commitment. Both the timing of the first rate increase and any subsequent adjustments to the federal funds rate target will depend critically on future developments in the economy.”	• “U.S. real gross domestic product is estimated to have expanded at a 2 $\frac{1}{4}$ percent pace in the first half of the year, a notably stronger outcome than expected in June.” • “Moreover, net exports were a substantial drag on GDP growth during the first half of the year, reflecting the earlier appreciation of the dollar and weaker foreign demand.” • “...that the timing of the initial increase in the federal funds rate will depend on the Committee’s assessment of the implications of incoming information for the economic outlook. To be clear, our decision will not hinge on any particular data release or on day-to-day movements in financial markets.”
Brainard (2017/05/30) Yellen (2017/03/15) <i>Score: 0.74</i>	• “The labor market has continued to strengthen. Payroll growth has averaged 175,000 over the past three months, more than sufficient to absorb new entrants into the labor market.” • “...the underlying momentum of the domestic expansion remains solid. While U.S. consumption was weak in the first quarter of 2017, the data so far are consistent with a rebound in the current quarter. Moreover, financial conditions remain supportive of continued economic expansion despite some recent volatility.”	• “Job gains averaged about 200,000 per month over the past three months, maintaining the solid pace we have seen over the past year.” • “...a reassessment of the economic outlook or of the appropriate course for monetary policy. I’ll have more to say about monetary policy shortly, but first I’ll review recent economic developments and the outlook. The economy continues to expand at a moderate pace. Solid income gains and relatively high levels of consumer sentiment and wealth have supported household spending growth.”

Continued on next page

Table 1: High Cosine Similarity Examples. (Continued)

Information	Speech	Press Conference
Brainard (2016/02/26) Yellen (2015/12/16) <i>Score: 0.74</i>	• “The U.S. unemployment rate is now under 5 percent, compared with 10 percent at its recent peak. Even so, there is evidence that some labor market slack still remains. The United Kingdom has experienced a rapid drop in unemployment to 5.1 percent, as low as pre-crisis levels, and labor force participation has remained relatively strong.” • “...many economies face common negative shocks or where negative shocks in one country are quickly transmitted across borders, it is natural to consider whether coordination can improve outcomes. Under certain conditions—such as flexible exchange rates, deep and well-regulated financial markets, and flexible product and labor markets—policies designed for the domestic economy can readily offset any spillovers from economic conditions abroad, and policies designed to address domestic conditions can achieve desirable outcomes both within the national economy and more broadly.”	• “...So far this year, a total of 2.3 million jobs have been added to the economy, and over the most recent three months, job gains have averaged an estimated 218,000 per month, similar to the average pace since the beginning of the year. The unemployment rate, at 5 percent in November, is down 0.6 percentage point from the end of last year and is close to the median of FOMC participants’ estimates of its longer-run normal level.” • “We have considered the risks to the outlook and worried about the fact that with interest rates at zero, we have less scope to respond to negative shocks than to positive shocks that would call for a tightening of policy. That is a factor that has induced us to hold rates at zero for this long. But we recognize that policy is accommodative, and if we do not begin to slightly reduce the amount of accommodation, the odds are good that the economy would end up overshooting both our employment and inflation objectives.”

Continued on next page


Table 1: High Cosine Similarity Examples. (Continued)

Information	Speech	Press Conference
Williams (2011/05/04) Bernanke (2011/04/27) <i>Score: 0.72</i>	“But, that doesn’t mean we’ve regained all the ground lost during the recession. Not by a long shot. In fact, despite adding about $1\frac{1}{2}$ million jobs over the past 13 months, there are still over 7 million fewer jobs in the United States than we had before the downturn. The recovery has sputtered at times and our forward progress has been disappointingly slow.” “...Commerce Department estimated that real gross domestic product—the broadest measure of U.S. economic activity—rose just 1.8 percent on an annual basis during the first three months of this year. This is quite a bit slower than the 2.8 percent growth we saw in 2010. Unusually severe winter weather earlier this year and a few other transitory factors held down first-quarter growth.”	“...it is encouraging to see the improvement that we’ve seen in recent months. That being said, the pace of improvement is still quite slow, and we are digging ourselves out of a very, very deep hole. We are still something like 7 million-plus jobs below where we were before the crisis. And so, clearly, the fact that we’re moving in the right direction, even though that’s encouraging, doesn’t mean that the labor market is in good shape.” “We haven’t seen the GDP number yet. But we, like most private-sector forecasters, are expecting a relatively weak number for the first quarter, maybe something a little under 2 percent. Most of the factors that account for the slower growth in the first quarter appear to us to be transitory. They include things like, for example, lower defense spending than was anticipated, which will presumably be made up in a later quarter.”

Notes: This table reports illustrative speech–press-conference pairs with high TF–IDF cosine similarity. Entries show short excerpts (ellipses denote omitted text). Scores are computed using TF–IDF representations and cosine similarity.

Figure 1 illustrates the cross-member dispersion in cosine similarity. The similarity distribution spans nearly the full possible range (0.10–0.90), indicating that individual speeches exhibit substantial heterogeneity in how closely they echo the Chair’s press-conference narrative. Individual members sometimes reproduce the Chair’s language almost verbatim but at other times diverge markedly in emphasis or tone. These patterns corroborate a “not a one-man show, but not egalitarian either” view of FOMC communications: Governors (and particularly the Chair) tend to cluster at higher similarity scores, reflecting coordinated messaging, whereas regional Presidents display greater dispersion, reflecting idiosyncratic emphasis or policy nuance.

Figure 1: Dispersion in Speech–Press–Conference Similarity across FOMC Members.



Notes: This figure plots, for each FOMC member, the distribution of cosine-similarity scores between that member’s public speeches and the Chair’s most recent post-meeting press conference. Each dot represents one speech–press-conference pair; dot size is proportional to the relative frequency of the score within the member’s distribution. Higher values indicate greater textual alignment with the Chair’s press-conference communication.

2.3 ALTERNATIVE SIMILARITY MEASURES

To validate the cosine-similarity measure and to probe the robustness of our results, we consider two alternative methods for measuring the alignment between individual speeches and centralized messaging in FOMC press conferences.

2.3.1 LLM-based method

First, while our baseline cosine-similarity approach is transparent and widely used, it relies on a bag-of-words representation that abstracts from higher-order linguistic structure and contextual meaning and may therefore miss important variation arising from more nuanced semantic differences. To assess whether this limitation affects our results, we complement it with a more precise measure derived from a large language model (LLM)–based annotation.

Specifically, for each speech–press-conference pair, we query ChatGPT 5.2 (Thinking) using the following prompt: “You are scoring semantic similarity between two texts. Return a continuous score between 0 and 1. Focus on meaning (topics/claims/stance), not style.” We also extract a verbal explanation for the score, which we use for manual validation. This

procedure yields an alternative measure of textual alignment that accounts for more complex semantic relationships, paraphrasing, and contextual cues that may not be fully captured by word-frequency methods.

Appendix Figure A2 compares the TF-IDF cosine similarity scores to the corresponding LLM-based similarity measures across all speech–press-conference pairs in the sample. The two measures are strongly correlated, with a sample correlation of 0.77. At the same time, the dispersion around the least-square fitting line suggests that the LLM-based measure incorporates additional dimensions of meaning beyond pure word co-occurrence.

Second, an additional advantage of the LLM-based method is that it helps interpret the source of low similarity between individual speeches and press conferences. In particular, the same framework allows us to distinguish whether low similarity reflects discussion of different topics, disagreement on a common topic, or—conditional on disagreement—differences in monetary policy stance versus assessments of economic conditions. Specifically, we use a follow-up prompts to classify (i) whether the two texts talk about the same topic; (ii) conditional on talking about the same topic, whether they express agreement or disagreement; (iii) whether the disagreement stems from different (similar) monetary policy stances or from different assessments of economic conditions. Figure 2 summarizes this classification scheme. We find that 45% of texts talk about the same topic. Conditional on topic overlap, 78% agree, whereas 22% disagree. Conditional on disagreement, the source is split roughly evenly between monetary policy stance and economic assessment. Therefore, low similarity primarily reflects topic divergence rather than disagreement conditional on discussing the same topic. In the robustness analysis Subsection 4, we focus on the “Yes–Yes” cases—pairs that discuss the same topic and agree—and show that conditioning on these subcells yields results very similar to those based on our high-similarity index.

2.3.2 Filtering out non-FOMC-relevant topics

The second potential weakness of our baseline measure is that it does not distinguish between topics strictly related to Fed policy versus other unrelated topics.

To ensure that measured similarity is topically relevant to the FOMC deliberations, we construct an alternative measure defined as the product of a speech’s similarity to the Chair’s press conference, and its similarity to the contemporaneous FOMC deliberation (using the text of FOMC meeting transcripts). Because full meeting transcripts are available only through 2018, the resulting series terminates that year.

Appendix Figure A3 shows that the resulting measure is highly correlated with our simpler baseline measure, suggesting that topics other than Fed policy do not drive this variation.




Figure 2: LLM-based decomposition of similarity across text pairs.

3 SPEECHES, CONTENT, AND MARKET REACTIONS

This section examines how FOMC members’ speeches affect financial markets and how those effects vary with each speech’s textual similarity to the Chair’s post-meeting press conference. In Subsection 3.1, we construct a high-frequency measure of monetary policy surprises around speeches and estimate their impact on interest rates and inflation expectations via local projections. In Subsection 3.2, we introduce cosine-similarity-based measures of textual proximity to the Chair’s post-meeting press conference and show that market responses vary systematically with the degree of similarity.

3.1 HIGH-FREQUENCY SPEECH SURPRISES AND MARKET REACTIONS

In this section, we document how financial markets respond to speeches by FOMC members. We first describe the construction of a high-frequency dataset capturing intraday interest-rate movements around each speech. We then use daily local projections to estimate both the statistical and economic significance of speech-induced surprises and the persistence of these effects over subsequent days.

Intraday data and event window. Our surprise measure is based on five-minute averages of eurodollar (ED) futures at one- to four-quarter maturities over the period 2011–2022. We retain only speeches delivered during U.S. equity-market hours on business days, and we exclude any events when two or more FOMC members speak simultaneously. For each qualifying speech, we define an event window from ten minutes before delivery to twenty minutes after conclusion. Speech end times are imputed by multiplying each transcript’s

word count by a standard words-per-minute rate.

Monetary policy speech (MPS) surprise. Within each 30-minute window, we compute intraday changes in all four maturities and extract their first principal component (PC1). We then rescale PC1 to have unit variance in the ED4 contract; this rescaled series is our *monetary policy speech* (MPS) surprise. The top panel of Appendix Figure A1 plots the MPS time series, with red markers indicating speeches by Board members, while the bottom panel displays its autocorrelation function, which confirms the absence of serial correlation. In total, our sample comprises 481 speeches, of which 403 are delivered by Board members and 78 by Reserve Bank Presidents.

Market Response. We assess how markets react to the FOMC member speech by running the following regression:

$$y_{t+h} - y_{t-1} = \alpha_h + \beta_h MPS_t + u_{h,t} \quad (1)$$

where t represents the day when the speech occurs and $h = 1, \dots, H$ is the number of days after the speech, MPS_t is the factor describing the intraday reaction of the short- to medium-term yield curve to monetary policy communication described in the previous section, and u_t is an error term. As previous work does, we investigate the reaction of the yield curve and market-based measures of expected inflation. More precisely, we consider the following set of financial data: the 3- and 6-month Treasury Bill (TBill), the yield on U.S. Treasury Securities at 3-Month, 1-, 5- and 10-Year Constant Maturity (GBY), Breakeven (BE) Inflation at five and ten year horizons.

Table 2 reports the estimated values for β_h for $h = 1, 3, 5$ when all speeches are included. These coefficients represent the impact that the speech has on financial instruments on the day of the speech, one day after the speech and five days after the speech along with the portion of variance explained (R^2). A speech that increases the intraday ED4 does not have a strong impact on short maturity but does have an effect on interest rates with the maturity between one and five years. Although the effect seems to be persistent and does not vanish after 5 days, speeches explain a tiny and diminishing portion of interest rate fluctuations over time. While not statistically significant, the point estimates of the effect on BE inflation have the expected sign (i.e. negative). The effects on the stock market, volatility and spread are small and non significant (not shown here).

Magnitudes are small. Daily and intraday variations are expressed in percent and the range values of the latter are bounded between ± 0.03 , i.e. 3 basis points. So a 1 bps increase in the MPS increases the 5y GBY by 1.12 basis points at one year maturity the day after

the speech.

Table 2: Market Responses to Monetary Policy Speech (MPS) Surprises.

	$h = 1$		$h = 3$		$h = 5$	
	β_h	Adj. R^2	β_h	Adj. R^2	β_h	Adj. R^2
3M TBill	0.44***	0.07	0.28*	0.01	0.34*	0.00
6M TBill	0.49***	0.09	0.53***	0.03	0.62**	0.02
3M GBY	0.42***	0.07	0.32**	0.01	0.31*	0.00
1Y GBY	0.94**	0.19	0.95***	0.07	1.08**	0.05
5Y GBY	1.07***	0.07	1.03***	0.04	1.33***	0.04
10Y GBY	0.95***	0.05	0.95***	0.03	1.20***	0.03
10Y BE Inflation	-0.02	0.00	-0.13	0.00	0.07	0.00
5Y BE Inflation	0.00	0.00	-0.16	0.00	0.05	0.00
SP500	0.08	0.01	-0.01	0.00	-0.03	0.00

Notes: Each β_h is from a separate regression of $y_{t+h} - y_{t-1}$ on the monetary policy speech (MPS) surprise on speech days. The MPS surprise is the first principal component of intraday changes in ED1–ED4 eurodollar futures in a narrow window around each speech, scaled to have a unit impact on ED4. Adj. R^2 is reported for each regression. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

3.2 HIGH-FREQUENCY SPEECH SURPRISES, SPEECH CONTENT, AND MARKET REACTIONS

As documented in Subsection 3.1, FOMC member speeches can induce significant shifts in the yield curve but generate ambiguous responses in inflation expectations and equity prices. This pattern implies two key points. First, market participants systematically update their outlooks in response to speech content. Second, the direction of these updates—whether the speech is interpreted as positive or negative news—cannot be inferred from yield-curve shifts alone. Rather than reflecting a lack of informational value, this interpretative ambiguity is likely a natural consequence of the heterogeneous nature of speeches. Some speeches emphasize projections for real economic activity, while others convey information about the stance and future path of monetary policy. Empirical work has shown that these different information types can produce similar yield-curve responses but divergent effects on inflation expectations and asset prices (e.g., [Andrade and Ferroni, 2021](#)). Moreover, variation in focus on short-run versus long-run dynamics, as well as differences in individual FOMC participants’ perspectives, further contribute to the multiplicity of market interpretations.

Speech content is inherently multidimensional, making it difficult to construct concise summary metrics. We therefore focus on one interpretable margin: the degree of textual similarity between an individual speech and the Chair’s post-meeting press conference statement. This focus is motivated by two considerations. First, it allows us to assess whether financial markets respond differently when a speech reiterates guidance already conveyed in the press conference versus when it emphasizes distinct information, thereby evaluating the marginal value of message repetition. Second, anchoring our analysis to the press conference—a benchmark widely studied in the literature on monetary policy communication and its effects on financial markets and macroeconomic outcomes (e.g., [Bauer and Swanson, 2023a,b](#))—facilitates direct comparison with existing work.

We organize the analysis in three steps. First, we introduce the text-based measure of each speech’s similarity to the Chair’s post-meeting press conference. Second, we examine how the correlation between monetary policy speech surprises and daily changes in key financial variables varies across different levels of similarity. Finally, we estimate dynamic projections in which speech surprises are interacted with similarity measures to quantify how the content of a speech conditions its effect on yields, inflation expectations, and other asset prices over time.

In the subsequent analysis, we classify each speech into one of three categories based on their similarity score, ℓ_t : (i) *Far*: $\ell_t < \mu(\ell) - \sigma(\ell)$, (ii) *Neutral*: $\mu(\ell) - \sigma(\ell) \leq \ell_t \leq \mu(\ell) + \sigma(\ell)$, and (iii) *Close*: $\ell_t > \mu(\ell) + \sigma(\ell)$, where $\mu(\ell)$ and $\sigma(\ell)$ denote the sample mean and standard

deviation of $\{\ell_t\}$, respectively.

We then compute the correlation between the monetary policy surprise of speeches in each of the three categories with the daily changes in the 3-month Treasury Bill (TBill), the yield on U.S. Treasury Securities at 1-, 5- and 10-Year Constant Maturity (GBY), and Breakeven (BE) Inflation at five and ten year horizons. Table 3 reports the correlations between financial variables and monetary policy speech surprise conditional on its similarity to the press conference.

Consistent with the findings of Subsection 3.1, monetary policy speech surprises are typically correlated with daily shifts in the yield curve but exhibit little association with market-based inflation expectations (see the “All” column in Table 3). When we partition speeches based on their similarity score, distinct patterns emerge. Speeches that exhibit a low cosine similarity generate positive co-movement between nominal yields and breakeven inflation—especially at longer maturities—whereas speeches that exhibit a high cosine similarity produce yield increases that coincide with declines in inflation expectations. The one-year yield response is particularly pronounced for high cosine similarity speeches. High cosine similarity speeches consistently yield negative correlations between nominal interest rates and inflation expectations.

Table 3: Correlations between Financial Variables and MPS Surprises by Similarity Bin.

	$h = 1$ (one day after)				$h = 3$ (three days after)				$h = 5$ (five days after)			
	All	Far	Neutral	Close	All	Far	Neutral	Close	All	Far	Neutral	Close
3M TBill	0.28	0.26	0.03	0.53	0.11	0.30	-0.06	0.08	0.11	0.30	-0.06	0.08
1Y GBY	0.44	0.30	0.15	0.82	0.27	0.37	0.03	0.66	0.27	0.37	0.03	0.66
5Y GBY	0.28	0.30	0.17	0.49	0.21	0.41	0.08	0.34	0.21	0.41	0.08	0.34
10Y GBY	0.24	0.22	0.13	0.43	0.18	0.30	0.08	0.31	0.18	0.30	0.08	0.31
10Y BE Inflation	-0.01	0.20	-0.04	-0.07	-0.04	0.19	-0.09	-0.12	-0.04	0.19	-0.09	-0.12
5Y BE Inflation	0.00	0.24	0.00	-0.14	-0.04	0.25	-0.05	-0.24	-0.04	0.25	-0.05	-0.24

Notes: Entries are sample correlations between the monetary policy speech (MPS) surprise on day t and daily changes in each financial variable measured $h \in \{1, 3, 5\}$ days after the speech. “All” uses all speeches; “Far,” “Neutral,” and “Close” split speeches into terciles based on the speech–press-conference similarity index ℓ_t : Far if $\ell_t < \mu(\ell) - \sigma(\ell)$, Close if $\ell_t > \mu(\ell) + \sigma(\ell)$, and Neutral otherwise.

Market Response. To assess how variation in speech content relative to the Chair’s message influences the market impact of speech-induced surprises, we estimate the following interaction specification.

$$y_{t+h} - y_{t-1} = \alpha_h + \beta_h \text{MPS}_t + \gamma_h \ell_t + \delta_h (\text{MPS}_t \times \ell_t) + u_{j,h,t}, \quad (2)$$

for horizons $h = 0, 1, \dots, H$, where ℓ_t^1 is our baseline speech–press-conference cosine similarity. By allowing ℓ_t to vary continuously in $[0, 1]$, we trace how the marginal effect of the speech surprise,

$$\frac{\partial}{\partial \text{MPS}_t}(y_{t+h} - y_{t-1}) = \beta_h + \delta_h \ell_t,$$

evolves with similarity to the Chair’s message. Confidence bands for this varying impact are constructed so that


$$\text{Var}(\beta_h + \delta_h \ell_t) = \text{Var}(\beta_h) + \ell_t^2 \text{Var}(\delta_h) + 2 \ell_t \text{Cov}(\beta_h, \delta_h).$$

Figure 3 displays the financial market responses one day after a monetary policy speech, using our baseline measure of similarity. Each panel plots the effect of a tightening monetary policy surprise as a function of similarity between the speech and the Chair’s post-meeting message: values on the left correspond to speeches that exhibit a low cosine similarity, while those on the right are closer to the press conference. The black circle denotes the unconditional (average) effect across all observations.

The results yield two main insights. First, when the speech content exhibits a high cosine similarity (toward the right of the plots), short-to-intermediate interest rates (six months to one year) exhibit a more pronounced response, while inflation expectations decline—unlike in other regions where they are typically insignificant or mildly positive. In contrast, the long end of the yield curve shows weaker and statistically insignificant movements. This pattern is broadly consistent with a temporary monetary policy tightening, where markets expect near-term adjustments without implications for the long-run stance.

Second, when the speech content exhibits a low cosine similarity (toward the left), the response of the term structure steepens: interest rates at longer maturities react more strongly than short-term rates. This is accompanied by rising inflation expectations and a positive stock market response. These patterns can be interpreted in two complementary ways. One possibility is that the speech reveals new information about the macroeconomic outlook, consistent with a Delphic or information effect as discussed in the literature. Alternatively, speeches that differ substantially in content from the Chair’s message may emphasize long-term objectives or structural considerations, potentially giving rise to neo-Fisherian dynamics.

Figure 3: Financial-Market Responses to MPS Surprises across Speech–Press–Conference Similarity.



Notes: Each panel reports the estimated marginal effect of a tightening monetary policy speech (MPS) surprise on the indicated financial variable, as a function of the speech–press–conference similarity index $\ell_t \in [0, 1]$, measured one day after the speech. The solid blue line plots $\beta_n + \delta_n \ell_t$ from equation (2). Shaded areas and dotted lines denote pointwise confidence bands. The black circle indicates the unconditional effect β_n estimated without the interaction term, evaluated at the sample mean of ℓ_t ; vertical bars report its confidence interval.

Daily Local Projections. To isolate financial market responses to monetary policy speeches that closely resemble the Chair’s message, we restrict the sample to episodes in which the speech content exhibits a high degree of textual similarity to the post-meeting press conference. Specifically, we retain only those speeches for which the cosine similarity index exceeds one standard deviation above the mean. For this restricted sample, we estimate the following regression:

$$y_{t+h} - y_{t-1} = \alpha_h + \beta_h MPS_t^* + u_{h,t},$$

for $h = 0, 1, \dots$, where MPS_t^* captures the high-frequency variation in interest rates associated with speeches classified as highly similar to the Chair’s press conference.

Figure 4 plots the responses of financial variables over time. The black line shows the baseline response to all monetary policy speeches, while the blue line isolates the response to speeches with high textual similarity. The results are consistent with the baseline findings: tightening speeches that resemble the Chair’s message lead to a significant increase in interest rates at maturities between six months and one year, a persistent decline in inflation expectations, and little to no effect at the long end of the yield curve.


4 ROBUSTNESS

We assess robustness along six dimensions: (i) varying the event-study horizon, (ii) excluding speeches delivered by the FOMC Chair, (iii) replacing TF-IDF cosine similarity with an LLM-based semantic similarity score, (iv) restricting attention to text pairs classified by the LLM as discussing the same topic and agreeing, (v) using an alternative similarity index that filters out non-FOMC-relevant topics by conditioning similarity on contemporaneous FOMC deliberations, and (vi) verifying that cosine similarity does not simply proxy for the time elapsed since the preceding Chair press conference.

4.1 ALTERNATIVE TIME-HORIZONS

Our baseline results evaluate the interaction specification at horizon $h=1$, i.e., using daily changes measured one day after each speech. Appendix Figures A4 and A5 replicate Figure 3 for $h=0$ (same-day changes) and $h=2$ (two-day changes). The qualitative pattern is unchanged. When a speech is more closely aligned with the Chair’s post-meeting press conference, a tightening MPS surprise primarily affects short-to-intermediate maturities and is associated with lower breakeven inflation. When similarity is low, the response tilts toward longer maturities and inflation expectations do not decline. Overall, the similarity-

Figure 4: Dynamic Responses to MPS Surprises: High-Similarity Speeches vs. Full Sample.



Notes: This figure reports daily local-projection impulse responses of the indicated financial variables to a one-unit tightening monetary policy speech (MPS) surprise. The black line shows estimates using all speeches in the sample. The blue line restricts the sample to speeches with high textual similarity to the Chair's most recent post-meeting press conference (similarity exceeding one standard deviation above the mean). Shaded areas denote pointwise confidence bands. The horizontal axis reports days after the speech.

conditioned heterogeneity in market responses is not sensitive to the choice of horizon.

4.2 EXCLUDING FOMC CHAIR SPEECHES

To probe the robustness of our baseline results, we first verify that they are not driven by the influence of the FOMC Chair exerted through non-FOMC speeches. Figures A6 and A7 report the resulting estimates and are the direct analogues of Figures 3 and 4 in the main text, respectively. As demonstrated by these figures, our results are hardly affected once we remove Chair speeches.

This suggests that markets extract clues about monetary policy not only from the Chair speeches but also from the plurality of the committee voices.

4.3 ALTERNATIVE SIMILARITY MEASURE: LLM-BASED SIMILARITY

While transparent, our simple bag-of-words approach of measuring similarity between individual speeches and press conferences may miss semantic proximity that arises through paraphrasing, contextual framing, or differences in wording. For robustness, we repeat the analysis in Subsection 3.2 using an LLM-based semantic similarity score in place of cosine similarity.

Specifically, we re-estimate equation (2) replacing the cosine-similarity index ℓ_t with its LLM-based counterpart $\tilde{\ell}_t \in [0, 1]$, and trace the implied marginal effect of a monetary policy speech surprise, $\beta_h + \delta_h \tilde{\ell}_t$, across values of semantic similarity.

Figure A8 (LLM analogue of Figure 3) delivers the same qualitative message as the cosine-similarity results. When speeches are semantically close to the Chair’s message, the tightening surprise produces larger increases at short-to-intermediate maturities (roughly six months to one year) and a decline in market-based inflation expectations, with comparatively muted movements at the long end. When speeches are semantically distant, the term-structure response tilts toward longer maturities and inflation expectations do not fall, consistent with a more “information/Delphic” interpretation of those communications.

We also replicate the dynamic exercise underlying Figure 4. Mirroring the baseline cutoff, we classify a speech as “high similarity” when its LLM similarity exceeds one standard deviation above the mean and estimate daily local projections on this restricted sample. Figure A9 (LLM analogue of Figure 4) again closely matches the baseline: semantically aligned tightening speeches generate a more pronounced and persistent response in short-to-intermediate rates and real rates, alongside a sustained decline in breakeven inflation, while longer-maturity yields move little. Overall, the near-coincidence between the cosine-

and LLM-based results indicates that our findings are robust to an alternative notion of similarity that directly targets semantic content.

4.4 LLM-BASED TOPIC AGREEMENT

We also examine whether our results continue to hold when we restrict attention to the subset of speech–press-conference pairs that, according to the LLM classification summarized in Figure 2, discuss the same topic and agree. This restriction isolates the “Yes–Yes” subcells.² Appendix Figure A12 shows that conditioning on these subcells yields dynamic responses very similar to those obtained using our high-similarity index. This finding further supports the interpretation that our baseline results capture meaningful alignment in policy communication.

4.5 ALTERNATIVE SIMILARITY MEASURE: FILTERING OUT NON-FOMC RELEVANT TOPICS

We verify that our results capture similarity in FOMC-relevant content rather than unrelated topics. To do so, we replicate the baseline analysis using an alternative similarity measure. Instead of the cosine similarity between each speech and each press conference, we compute (i) the cosine similarity between the speech and FOMC deliberation text and (ii) the cosine similarity between the press conference and the same deliberation text, and use the product of these two similarities as our measure. The results (Appendix Figures A10 and A11) are consistent with those obtained using the baseline measure.³

4.6 ELAPSED TIME AND COSINE SIMILARITY

Finally, we assess whether cosine similarity merely proxies for the time elapsed since the preceding Chair press conference. In a regression of cosine similarity on the number of days between the press conference and the speech, the estimated coefficient is small and statistically indistinguishable from zero, suggesting that our similarity measure is not driven by a mechanical decay in overlap over time.

²Rather than a continuous measure, this variable is an indicator equal to one for text pairs in that category.

³The difference is that these results exhibit wider confidence intervals, which likely reflects the much shorter sample period available for that specification.

5 A SIMPLE MODEL WITH SPEECH-DRIVEN EXPECTATIONS

We embed a standard New Keynesian (NK) core in which private agents form expectations not only based on current inflation, output, and interest rates, but also on forward-looking central-bank speeches. Speeches serve as noisy signals of next-period fundamental innovations.

5.1 MODEL

We retain the familiar NK core:

$$\begin{aligned} \text{IS curve:} \quad & x_t = \widehat{\mathbb{E}}_t[x_{t+1}] - \frac{1}{\sigma}(i_t - \widehat{\mathbb{E}}_t[\pi_{t+1}] - r_t^n), \\ \text{NKPC:} \quad & \pi_t = \beta \widehat{\mathbb{E}}_t[\pi_{t+1}] + \kappa x_t + u_t, \\ \text{Taylor rule:} \quad & i_t = \psi_\pi \pi_t + \psi_x x_t + \theta_t, \end{aligned}$$

where x_t indicates output gap, π_t inflation, i_t the nominal rate, r_t^n the natural real rate, u_t a cost-push shock, and θ_t a monetary-policy shock. The unobserved fundamentals (θ_t, r_t^n, u_t) follow known AR(1):

$$\begin{aligned} \theta_{t+1} &= \rho_\theta \theta_t + \varepsilon_{t+1}^\theta, \\ r_{t+1}^n &= \rho_n r_t^n + \varepsilon_{t+1}^n, \\ u_{t+1} &= \rho_u u_t + \varepsilon_{t+1}^u, \end{aligned}$$

with $\varepsilon_t \sim \mathcal{N}(0, \Sigma)$. Unlike the standard NK model, agents at each date t observe a scalar speech signal s_t that provides noisy information about the next-period fundamental monetary policy shock:

$$s_t = \zeta_t \varepsilon_{t+1}^\theta + \eta_t, \quad \eta_t \sim \mathcal{N}(0, \sigma_s^2),$$

where $\zeta_t \sim \text{Beta}(\alpha, \beta)$, so that $\zeta_t \in [0, 1]$, captures the speech's similarity to the Chair's press conference.

A higher ζ_t indicates that the speech is more informative about the upcoming policy shock, in the sense that it reinforces the message of the press conference, reduces uncertainty, and increases investor confidence in the expected policy path. Intuitively, ζ_t captures the “policy-focus” of central-bank communication: a high ζ_t means speeches closely echo—and thus reinforce—the Fed's imminent rate decision.

In this context, agents form expectations given four linear forecasting rules:

$$\begin{aligned}\widehat{\mathbb{E}}_t[\pi_{t+1}] &= \alpha_0^\pi + \alpha_1^\pi \pi_t + \alpha_2^\pi x_t + \alpha_3^\pi i_t + \alpha_4^\pi s_t + \alpha_5^\pi \widehat{\zeta}_t, \\ \widehat{\mathbb{E}}_t[x_{t+1}] &= \alpha_0^x + \alpha_1^x \pi_t + \alpha_2^x x_t + \alpha_3^x i_t + \alpha_4^x s_t + \alpha_5^x \widehat{\zeta}_t, \\ \widehat{\mathbb{E}}_t[\theta_{t+1}] &= \alpha_0^\theta + \alpha_1^\theta \theta_t + \alpha_2^\theta s_t + \alpha_3^\theta \widehat{\zeta}_t, \\ \widehat{\mathbb{E}}_t[r_{t+1}^n] &= \alpha_0^n + \alpha_1^n r_t^n + \alpha_2^n s_t + \alpha_3^n \widehat{\zeta}_t,\end{aligned}$$

which depend on the speech signal s_t and the speech similarity $\widehat{\zeta}_t = \zeta_t - \bar{\zeta}$ expressed in linear deviation from a steady-state value $\bar{\zeta}$. These coefficients (α) are pinned down by a fixed-point algorithm, ensuring that when agents simulate the model using their own forecast rules, OLS on simulated paths reproduces exactly those same rules.

Plugging the four linear forecasting rules into the NK equilibrium yields a linear system $\Gamma_0(\alpha) z_t + \Gamma_1(\alpha) \widehat{\mathbb{E}}_t[z_{t+1}] + C(\alpha) = 0$, which we solve to get the perceived law

$$z_{t+1} = G(\alpha) z_t + M(\alpha) \varepsilon_{t+1}.$$

The fixed-point over (α) requires that when one simulates this system and regresses each LHS variable on the RHS regressors (including s_t), one recovers the same α .

5.2 CALIBRATION

Table 4 summarizes the model's parameterization. Standard values are assigned to preference and policy rule parameters, including an intertemporal elasticity of substitution $\sigma = 2$, Taylor rule coefficients $\psi_\pi = 1.5$ and $\psi_x = 0.3$, and a quarterly discount factor $\beta = 0.99$, consistent with a 4% annual steady-state interest rate. The slope of the Phillips curve is calibrated to $\kappa = 0.0429$, implying an average price duration of four quarters under Calvo pricing. Monetary policy shocks are assumed to be transitory, with persistence $\rho_\theta = 0.10$ and standard deviation $\sigma_\theta = 0.0025$, corresponding to a 25-basis-point quarterly innovation. The natural-rate shock is more persistent ($\rho_{r^n} = 0.9$), reflecting the slow-moving nature of underlying fundamentals, and features a 20-basis-point standard deviation. Following [Smets and Wouters \(2007\)](#), the cost-push shock exhibits high persistence ($\rho_u = 0.89$) and modest volatility ($\sigma_u = 0.0014$). The signal noise standard deviation is set to 0.01 to capture a relatively noisy information environment. Finally, the parameters of the Beta distribution governing the signal weight ζ_t are calibrated to match a mean of 0.369 and a standard deviation of 0.124, which are calculated from data as the sample mean and standard deviation of cosine similarity between the Chair's speech and the corresponding press conference. This implies shape parameters $a = 5.176$ and $b = 8.859$.

Table 4: Calibration of Model Parameters.

Parameter	Value	Description	Source/Target
σ	2.00	Intertemporal elasticity of substitution	Standard
ψ_π	1.50	Taylor-rule coefficient on inflation	Taylor principle
ψ_x	0.30	Taylor-rule coefficient on output gap	Standard
κ	0.0429	NK Phillips-curve slope	Avg. price duration of 4 quarters
β	0.99	Quarterly discount factor	4% annual steady-state return
ρ_θ	0.10	Persistence of monetary-policy shock	Transitory policy shocks
σ_θ	0.0025	Std. dev. of monetary-policy shock	25 bps quarterly innovation
ρ_{r^n}	0.90	Persistence of natural-rate shock	Slow-moving fundamentals
σ_{r^n}	0.0020	Std. dev. of natural-rate shock	20 bps quarterly innovation
ρ_u	0.89	Persistence of cost-push shock	Smets and Wouters (2007)
σ_u	0.0014	Std. dev. of cost-push shock	Smets and Wouters (2007)
σ_s	0.01	Std. dev. of signal noise	100 bps (high-noise benchmark)
a	5.176	Beta shape parameter for ζ_t	Matches $\mathbb{E}[\zeta_t] = 0.369$
b	8.859	Beta shape parameter for ζ_t	Matches $\text{Std}(\zeta_t) = 0.124$

Notes: The model is calibrated at a quarterly frequency. The speech-similarity weight ζ_t is modeled as $\text{Beta}(a, b)$; a and b are calibrated to match the sample mean and standard deviation of the similarity measure in the data.

5.3 EFFECT OF SPEECHES-AUGMENTED MONETARY POLICY

We impose a one-time tightening, in the form of a 25 bps shock to θ_t at $t = 1$ and trace four quarters of responses under three speech-information experiments. In all cases ζ_t is perturbed only at $t = 1$ and then reverts to its steady-state $\bar{\zeta}$ immediately after.

5.3.1 Experiment 1: Coordinated Communication


This experiment compares the dynamic effects of a monetary policy tightening across three levels of initial speech similarity ($\zeta_1 = 0, 0.5, 1$), relative to a no-speech counterfactual in which $\zeta_t \equiv 0$ for all t . Figure 5 reports the impulse responses of the nominal interest rate i_t , expected short rate $\mathbb{E}_t[i_{t+1}]$, inflation π_t , and expected inflation $\mathbb{E}_t[\pi_{t+1}]$ over four quarters.

Two regimes emerge around the steady-state similarity level $\bar{\zeta} = 0.4$. When the initial similarity ζ_1 exceeds $\bar{\zeta}$, the monetary policy shock is more powerful: while the path of the nominal rate is largely unaffected, agents revise down their expectations of future short rates more sharply, realized inflation falls more steeply, and expected inflation declines more

rapidly and persistently. This amplification arises when more similar speeches reinforce the Chair’s message, allowing agents to better anticipate the future trajectory of monetary policy shocks. This mechanism mirrors the empirical evidence in Figure 3, which shows that inflation expectations decline more when policy speeches are closely aligned with the Chair.

Conversely, when $\zeta_1 < \bar{\zeta}$, speeches are less informative than usual, and the same shock elicits a weaker response: expectations adjust less, and inflation outcomes are more muted. Again, this pattern is consistent with Figure 3, where inflation expectations rise following a tightening surprise accompanied by an uncoordinated speech.

Figure 5: Model Impulse Responses to a Monetary Policy Shock under Varying Speech Informativeness.



Notes: This figure reports model impulse responses to a one-time 25 basis point tightening shock to the monetary policy disturbance. Responses are shown under different values of the speech-informativeness (similarity) parameter $\zeta_1 \in \{0, 0.5, 1\}$, which governs the weight of the speech signal on the upcoming policy innovation, and under a no-speech counterfactual ($\zeta_t \equiv 0$). The panels plot the short nominal rate i_t , the expected short nominal rate $E_t[i_{t+1}]$, inflation π_t , and expected inflation $E_t[\pi_{t+1}]$ over quarters.

5.3.2 Experiment 2: Delphic Signals and Natural-Rate News

We next allow speeches to convey information not only about future monetary policy shocks, but also about forthcoming movements in the natural rate. Specifically, the speech signal is modified as:


$$s_t = \zeta_t \varepsilon_{t+1}^\theta + (1 - \zeta_t) \varepsilon_{t+1}^n + \eta_t.$$

This formulation captures the idea that less coordinated speeches may reflect persistent shifts in the policy framework or the macroeconomic outlook, consistent with Delphic signaling. While the composition of the signal changes, the fundamental mechanism remains: higher similarity concentrates information on the policy shock, enhancing its effects; lower similarity shifts attention to broader fundamentals.

We impose a 25-basis-point monetary policy tightening alongside a simultaneous natural-rate shock calibrated so that expected inflation remains unchanged on impact in the no-speech baseline. Figure 6 plots the path of expected inflation under three values of ζ_1 , holding all else constant.

The qualitative results mirror those in Experiment 5.3.1: higher similarity amplifies the response, while lower similarity induces a reversal consistent with Delphic interpretation. The key difference is that all responses are shifted upward due to the direct effect of the natural-rate shock on inflation expectations. This upward shift reflects agents' improved inference about structural fundamentals when speeches are less aligned with the Chair.

Figure 6: Model Impulse Responses with Delphic Signals: Joint Monetary Policy and Natural-Rate Shocks.



Notes: This figure reports model impulse responses to a one-time 25 basis point tightening shock to the monetary policy disturbance combined with a 55 basis point innovation to the natural rate. The speech signal loads on both shocks according to $s_t = \zeta_t \varepsilon_{t+1}^\theta + (1 - \zeta_t) \varepsilon_{t+1}^n + \eta_t$, so lower similarity (smaller ζ_t) shifts the informational content of speeches toward natural-rate news. Responses are shown for different values of $\zeta_1 \in \{0, 0.5, 1\}$, with ζ_t reverting to its steady-state value after impact.

6 CONCLUSIONS

We analyze the impact of public speeches delivered by FOMC members other than the Chair. Using a novel dataset of 481 time-stamped speeches since 2007, we extract high-frequency monetary policy surprises and measure each speech's textual similarity to the Chair's post-meeting press conference. This framework allows us to assess not only whether committee-based communications shape market expectations and the transmission of monetary policy, but also how these effects depend on the degree of alignment with the Chair's message.

On average, market responses to individual speeches are modest. Conditioning on textual similarity, however, reveals a sharp asymmetry. Speeches that closely echo the Chair's

press conference raise short-term interest rates, lower breakeven inflation expectations, and transmit to real activity in ways consistent with conventional monetary policy shocks. By contrast, speeches that diverge in content lift the entire yield curve, increase inflation expectations, and are accompanied by stock market gains—suggesting that markets interpret them as signals about future fundamentals rather than immediate policy changes. A simple New Keynesian model with noisy speech signals accounts for these findings: alignment with the Chair strengthens policy effectiveness, while divergence attenuates the impact of monetary disturbances.

These results demonstrate that monetary policy communication is a central channel for the transmission of policy impulses. Importantly, they show that markets listen not only to the Fed Chair but also to the broader committee, and that the decentralized nature of FOMC communication has material consequences for financial markets and the economy. Greater coordination among members can enhance the effectiveness of monetary policy, while lack of coordination can dilute it.

Future work should examine whether markets assign different weights to individual voices within the Committee—responding more strongly to some members’ speeches than to others—or whether it is primarily the topic or other characteristics of a speech that drive market reactions to FOMC communication. More broadly, our findings highlight the value of integrating textual analysis into high-frequency identification strategies and open new avenues for studying how decentralized communication shapes the transmission of monetary policy.

REFERENCES


- Andrade, P. and F. Ferroni (2021). Delphic and odyssean monetary policy shocks: Evidence from the euro area. *Journal of Monetary Economics* 117, 816–832.
- Aruoba, S. B. and T. Drechsel (2024, May). Identifying monetary policy shocks: A natural language approach.
- Bauer, M. D. and E. T. Swanson (2023a, March). An alternative explanation for the “fed information effect”. *American Economic Review* 113(3), 664–700.
- Bauer, M. D. and E. T. Swanson (2023b). A reassessment of monetary policy surprises and high-frequency identification. *NBER Macroeconomics Annual* 37, 87–155.
- Cieslak, A. and A. Schrimpf (2019). Non-monetary news in central bank communication. *Journal of International Economics* 118, 293–315.
- Coibion, O., Y. Gorodnichenko, and M. Weber (2022). Monetary policy communications and their effects on household inflation expectations. *Journal of Political Economy* 130(6),

1537–1584.

- Ehrmann, M., R. Tietz, and B. Visser (2021, June). Voting right rotation, behavior of committee members and financial market reactions: evidence from the U.S. Federal Open Market Committee. Working Paper Series 2569, European Central Bank.
- Hansen, S., M. McMahon, and A. Prat (2017, 10). Transparency and deliberation within the fomc: A computational linguistics approach*. *The Quarterly Journal of Economics* 133(2), 801–870.
- Istrefi, K., F. Odendahl, and G. Sestieri (2023). Fed communication on financial stability concerns and monetary policy decisions: Revelations from speeches. *Journal of Banking Finance* 151, 106820.
- Lunsford, K. G. (2020, September). Policy language and information effects in the early days of federal reserve forward guidance. *American Economic Review* 110(9), 2899–2934.
- Malmendier, U., S. Nagel, and Z. Yan (2021). The making of hawks and doves. *Journal of Monetary Economics* 117, 19–42.
- Neuhierl, A. and M. Weber (2019). Monetary policy communication, policy slope, and the stock market. *Journal of Monetary Economics* 108, 140–155. “Central Bank Communications: From Mystery to Transparency” May 23–24, 2019 Annual Research Conference of the National Bank of Ukraine Organized in cooperation with Narodowy Bank Polski.
- Smets, F. and R. Wouters (2007, June). Shocks and frictions in us business cycles: A bayesian dsge approach. *American Economic Review* 97(3), 586–606.
- Swanson, E. T. (2023, May). The importance of fed chair speeches as a monetary policy tool. *AEA Papers and Proceedings* 113, 394–400.
- Swanson, E. T. and V. Jayawickrema (2024). Speeches by the fed chair are more important than fomc announcements: An improved high-frequency measure of us monetary policy shocks. Working Paper.


A APPENDIX

Figure A1: Monetary Policy Speech (MPS) Surprise.




Notes: The top panel plots the monetary policy speech (MPS) surprise constructed from five-minute changes in eurodollar futures around each speech. The blue line shows the MPS series for all speeches; red markers indicate speeches delivered by Board of Governors (BoG) members. The bottom panel reports the sample autocorrelation function of the MPS series; horizontal lines denote 95% confidence bands under the null of zero autocorrelation.

Figure A2: Similarity between individual speeches and press conferences: Baseline measure vs. an LLM-derived one.



Notes: This figure plots TF-IDF cosine similarity between FOMC member speeches and the Chair's post-meeting press conference (x-axis) against an LLM-based semantic similarity score (y-axis). Each point corresponds to a speech-press-conference pair. The solid black line is the 45-degree line. The in-panel value reports the sample correlation between the two measures ($\rho = 0.77$). The horizontal banding reflects discretization of the LLM scores (e.g., rounding or use of focal values in the model's numeric outputs).


Figure A3: Similarity between individual speeches and press conferences: Baseline measure vs. filtering out non-FOMC-relevant topics.



Notes: The top panel plots two measures of textual alignment with the Chair’s post-meeting press conference. The solid blue line is the baseline cosine similarity between each speech and the preceding Chair press conference, $coSim_{S,PC}$. The dash-dot black line is an alternative measure that down-weights speeches less related to contemporaneous FOMC deliberations, defined as $coSim_{T,PC} \times coSim_{T,S}$, where T denotes the transcript of the corresponding FOMC meeting; because transcripts are available only through 2018, this series ends in 2018. The bottom panel plots $coSim_{S,PC}$ against $coSim_{T,PC} \times coSim_{T,S}$ for speeches with available transcripts; the in-panel value reports their sample correlation.


A.1 ROBUSTNESS RESULTS: CHANGE THE TIME-HORIZON

Figure A4: Financial-Market Responses to MPS Surprises across Speech–Press-Conference Similarity ($h = 0$).



Notes: Each panel reports the estimated marginal effect of a tightening monetary policy speech (MPS) surprise on the indicated financial variable, as a function of the speech–press-conference similarity index $\ell_t \in [0, 1]$, measured one day after the speech. The solid blue line plots $\beta_n + \delta_n \ell_t$ from equation (2). Shaded areas and dotted lines denote pointwise confidence bands. The black circle indicates the unconditional effect β_n estimated without the interaction term, evaluated at the sample mean of ℓ_t ; vertical bars report its confidence interval.


Figure A5: Financial-Market Responses to MPS Surprises across Speech–Press–Conference Similarity ($h = 2$).



Notes: Each panel reports the estimated marginal effect of a tightening monetary policy speech (MPS) surprise on the indicated financial variable, as a function of the speech–press–conference similarity index $\ell_t \in [0, 1]$, measured one day after the speech. The solid blue line plots $\beta_n + \delta_n \ell_t$ from equation (2). Shaded areas and dotted lines denote pointwise confidence bands. The black circle indicates the unconditional effect β_n estimated without the interaction term, evaluated at the sample mean of ℓ_t ; vertical bars report its confidence interval.


A.2 ROBUSTNESS RESULTS: EXCLUDING CHAIR SPEECHES

Figure A6: Responses to MPS Surprises by Speech–Press-Conference Similarity (Excluding Chair Speeches).



Notes: This figure shows the response of financial variables to a tightening monetary policy speech (MPS) surprise across different levels of textual similarity between the speech and the Chair's most recent post-meeting press conference, measured one day after the speech. The solid blue line reports the estimated marginal effect as a function of similarity; dotted blue lines denote pointwise confidence bands. Black circles indicate the unconditional (average) effect estimated without interacting the MPS surprise with the similarity measure. The sample excludes speeches delivered by the FOMC Chair.


Figure A7: Dynamic Responses to MPS Surprises (High Similarity Speeches, Excluding Chair Speeches).



Notes: This figure reports daily local-projection impulse responses of financial variables to a tightening monetary policy speech (MPS) surprise. The blue line shows responses estimated using only speeches whose speech–press-conference cosine similarity exceeds one standard deviation above its sample mean; the black line shows the baseline response using all speeches. Dotted lines denote pointwise confidence bands around each response. The sample excludes speeches delivered by the FOMC Chair.


A.3 ROBUSTNESS RESULTS: LLM-BASED SIMILARITY MEASURE

Figure A8: LLM-Based Similarity and Financial-Market Responses to Monetary Policy Speech Surprises ($h = 1$).



Notes: Each panel plots the estimated effect of a tightening monetary policy speech (MPS) surprise on the indicated financial variable as a function of LLM-based semantic similarity between the speech and the Chair's most recent post-meeting press conference (horizontal axis). The solid blue line shows the fitted marginal effect $\beta_h + \delta_h \tilde{\ell}_t$ evaluated at $h = 1$ day after the speech; shaded areas (and dotted boundaries) report the corresponding confidence band. Black circles denote the unconditional effect estimated without interacting the MPS surprise with similarity. The legend distinguishes the unconditional (HF) effect from the similarity-conditioned (HF & LLM similarity) effect.


Figure A9: Dynamic Responses to Speech Surprises: High LLM-Based Similarity Sample.



Notes: This figure reports daily local-projection responses of financial variables to a tightening monetary policy speech (MPS) surprise over horizons $h = 0, \dots, 10$. The blue line isolates responses using only speeches with high LLM-based semantic similarity to the Chair's most recent post-meeting press conference (similarity above one standard deviation from the sample mean); the gray/black line reports the corresponding response using all speeches. Shaded areas and dotted lines denote confidence bands around the respective responses. The bottom-right tile is reserved for the legend.

A.4 ROBUSTNESS RESULTS: FILTERING OUT NON-FOMC-RELEVANT TOPICS


Figure A10: Financial-Market Responses by Similarity to the Chair and the FOMC Discussion ($h = 1$).



Notes: Each panel plots the estimated effect of a tightening monetary policy speech (MPS) surprise on the indicated financial variable as a function of the combined similarity between the speech, the transcript the FOMC discussion and the Chair’s most recent post-meeting press conference (horizontal axis). The solid blue line shows the fitted marginal effect $\beta_h + \delta_h \tilde{\ell}_t$ evaluated at $h = 1$ day after the speech; shaded areas (and dotted boundaries) report the corresponding confidence band. Black circles denote the unconditional effect estimated without interacting the MPS surprise with similarity. The legend distinguishes the unconditional (HF) effect from the similarity-conditioned (HF & LLM similarity) effect.


A.5 ROBUSTNESS RESULTS: “YES-YES” SUBCELLS OF FIGURE 2

Figure A11: Dynamic Responses to Speech Surprises: High similarity to the Chair and the FOMC Discussion.



Notes: This figure reports daily local-projection responses of financial variables to a tightening monetary policy speech (MPS) surprise over horizons $h = 0, \dots, 10$. The blue line isolates responses using only speeches with high combined similarity between the speech, the transcript of the FOMC discussion and the Chair's most recent post-meeting press conference (similarity above one standard deviation from the sample mean); the gray/black line reports the corresponding response using all speeches. Shaded areas and dotted lines denote confidence bands around the respective responses. The bottom-right tile is reserved for the legend.

Figure A12: Dynamic Responses to Speech Surprises: “Yes-Yes” subcells of 2.



Notes: This figure reports daily local-projection responses of financial variables to a tightening monetary policy speech (MPS) surprise over horizons $h = 0, \dots, 10$. The blue line isolates responses using only speeches that talk about the same topic and where the speakers agree with each others. The classification scheme follows the LLM procedure summarized by Figure 2. The gray/black line reports the corresponding response using all speeches. Shaded areas and dotted lines denote confidence bands around the respective responses. The bottom-right tile is reserved for the legend.