Authority, Attention, and the Origins of Monetary Policy Surprises*

David Gauthier

Abstract

At FOMC announcements, markets are surprised by information that was public months earlier, voiced by the members they attend to least. Using Federal Reserve Board members' communications since 1961, I show that attention follows institutional authority rather than informational content. Non-Chair speeches draw little press coverage and market reaction, yet their tone predicts interest-rate forecast errors and the forward-guidance component of monetary policy surprises. The asymmetry attaches to the office, not the person: upon promotion, the same speaker's remarks gain attention and lose their predictive power. A rational-inattention model in which forecasters underweight non-Chairs' influence can explain this puzzle.

JEL D84, E58, G14

Anticipating the path of U.S. monetary policy is among the most consequential forecasting problems in financial markets. Between FOMC meetings, Federal Reserve policymakers communicate almost continuously through speeches, testimony, and public appearances, providing a running public commentary on the economy and the appropriate stance of policy. Yet FOMC announcements still generate sizable monetary policy surprises and substantial asset price movements. Why does monetary policy remain surprising when so much of the relevant information has already been made public before every meeting?

This article argues that part of the answer lies not in what information is available, but in which information markets choose to process. Attention is scarce, and observers of a hierarchical institution allocate it by rank.[1] The gap is directly visible in the press: *Reuters* publishes more than six times as many articles on a Chair or Vice Chair speech as on another governor's. This allocation looks efficient: the Chair sets the agenda and speaks for the Committee. I show it is not. Markets attend to the Chair beyond what the Chair's measured influence warrants, and the communications they discount most carry the information they turn out to have missed. Attention follows perceived authority rather than measured influence.

*Corresponding author: David Gauthier, Banque de France, 31 rue Croix des Petits-Champs, 75001 Paris, France; david.gauthier@banque-france.fr. I thank Michael Bauer, Benjamin Born, Jeremy Fouliard, Valère Fourel, Federico Di Pace, Paul Hubert, Marwan Menaa, Sarah Mouabbi, Giovanni Ricco, Eric Swanson, and Mirko Wiederholt as well as participants at various seminars and conferences for helpful comments and discussions. The views expressed are those of the author and do not necessarily reflect those of the Banque de France or the Eurosystem.

[1] In the Brookings surveys of academic and market observers, the Chair's communications consistently top the usefulness ranking, while a substantial share of respondents report wanting to hear *less* from Board governors other than the Chair (Olson and Wessel, 2016; Wessel and Boocker, 2024; Ahmad and Wessel, 2026).

Answering the question posed above requires separating three things that are usually entangled: members' influence over policy decisions, their institutional rank as Chair, Vice Chair, or governor, and the attention paid to their words. The FOMC is a setting in which all three can be measured separately. Two constructed datasets make the measurement possible. The first is the longest panel of Board-member communication to date: 4,177 speeches and congressional testimonies by every Board member since 1961, scored with a transparent fixed dictionary into a Monetary Stance Signal (MSS) and paired with the FOMC meeting transcripts.[2] The second matches each speech since 1992 to the *Reuters* news articles that cover it, measuring the attention each communication receives—how many articles report a speech, and how faithfully their tone relays the speaker's.

The article establishes four results. First, attention to Federal Reserve communication follows institutional authority, not informational content (Section 2.1). The press relays what non-Chairs say about as faithfully as what the Chair and Vice Chair—henceforth principals—say: article for article, the tone of the coverage tracks the tone of the underlying speech to a similar degree for the two groups. What differs is volume, and expectations follow the volume: a principal's speech draws about six times as many articles as a non-Chair's, the consensus forecast revises roughly four to six times as much per unit of the stance those articles report, and daily short-rate prices move only with the tone of principal coverage.[3]

Second, non-Chair communication carries information about the path of future monetary policy that private agents do not hold (Section 2.2). Jointly with Chair tone, non-Chair speech tone explains 18 percent of the variation in one-year-ahead interest-rate forecast errors, and the predictive content sits entirely in the non-Chair term. The predictability is the same whether expectations are measured from futures prices or from surveys, extends to excess bond returns, most strongly at short maturities, and largely holds out of sample and under conservative inference. Chair and Vice Chair communications carry no comparable predictive content. Predictability of this kind establishes both halves of the argument at once: the discounted speeches carry information the Chair's do not, and that information is not absorbed when it is delivered.

Third, the asymmetry attaches to the office, not the person (Section 2.3). When a governor is promoted to Chair or Vice Chair, press coverage of the same speaker jumps, and the tone of their speeches stops predicting errors in interest-rate forecasts. This reversal is the design that separates authority from identity: the words come from the same individuals before and after promotion, so speaker ability or style cannot account for it. It also ties the first two results together: the predictive content disappears when attention arrives—once markets process a speaker's words on delivery, nothing is left to forecast.

Fourth, the discounted communications account for a measurable part of the monetary policy surprise (Section 3). The Committee drifts from its own policy rule, and forecasters miss the drift: deviations from a Taylor rule fitted to real-time Greenbook nowcasts and projected onto the staff's forecasts correlate with

[2] The dictionary methodology prioritizes interpretability and avoids the look-ahead bias of machine-learning approaches.

[3] Nor is the concentration a supply decision of the press alone: Google searches for policymakers' names show the same concentration (Section 2.1).

their forecast errors. Non-Chair tone predicts these deviations, and so does Chair tone. The difference lies in incorporation: the Chair's words are priced at once, the non-Chairs' with delay or not at all. The delayed part appears at the policy announcements. The tone non-Chair members express between consecutive meetings forecasts the forward-guidance component of the coming surprise, even after conditioning on the full public-news set of Bauer and Swanson (2023b). Monetary policy surprises thus measure more than the Federal Reserve's private information: part of what registers as news is forecastable from communication that had been public for months.

Section 4 measures the attention allocation against a rational benchmark. The benchmark is a rational-inattention model in the tradition of Sims (2003) and Maćkowiak and Wiederholt (2009): forecasters allocate scarce processing capacity across policymakers, and the optimal allocation depends on each member's influence over the Committee's decision and on the variability of their stance. Both inputs are measurable—I recover members' influence over policy decisions from the FOMC meeting transcripts, in the spirit of Chappell et al. (2004)—so the model turns the perceived-authority bias into a quantity. Attention is too concentrated on the Chair. Three separately measured shares summarize the comparison: the Chair and Vice Chair receive more than four-fifths of the press coverage of Board communication; the attention share recovered from forecasters' behavior is 60 percent; the optimal share is 46 percent. Equivalently, forecasters devote to the Chair a level of attention that would be optimal only if the Chair were roughly two-thirds more influential than the average non-Chair; the influence measured from the transcripts runs the other way.

The results are robust to measurement choices: they survive subsamples, rolling windows, smoothing and dictionary perturbations, and externally designed dictionaries (Online Appendices C.6 through C.8). Section 2.4 rules out competing explanations. The asymmetry is not statistical: the non-Chair signal pools several voices, but a single randomly drawn governor predicts no worse than a volume-matched average of several. It is not compensation for risk: premium-free surveys deliver the same non-Chair slopes as futures, and the predictive power declines with maturity, whereas a risk-premium channel would strengthen with it. And it is not a change in the words themselves: across promotions, a speaker's expressed stance, its magnitude, and the density of their policy language are statistically indistinguishable.

The article proceeds as follows. Section 1 describes the speech corpus, the construction of the Monetary Stance Signal, and the market, macroeconomic, and press-coverage data. Sections 2 and 3 establish the attention asymmetry and trace the origin of the surprises. Section 4 measures the attention allocation against the influence weights recovered from FOMC transcripts and quantifies the wedge. Section 5 concludes.

Related Literature. This article speaks to several literatures. A first literature examines how central bank communication shapes financial markets (Kuttner, 2001; Gürkaynak et al., 2005a; Bernanke et al., 2005; Gagnon et al., 2011; Wright, 2012; Cieslak et al., 2019; Gómez-Cram and Grotteria, 2022; Kim et al., 2023; Gorodnichenko et al., 2025), with recent work emphasizing heterogeneity across policymakers: communications by the Chair and Vice Chair dominate market reactions and forecast revisions (Neuhierl and Weber, 2019; Swanson, 2023; Swanson and Jayawickrema, 2023; Ahrens et al., 2025). These studies measure a communication's reach from the market's reaction around the speech itself; I measure attention directly, from the press coverage a speech receives, and show that the market reaction moves with this coverage. I confirm the contemporaneous asymmetry this literature documents and show that it masks the opposite asymmetry in informational content. Closest is Swanson and Jayawickrema (2023), who construct an improved high-frequency policy shock from the Chair's own speeches; the design uses the fact that markets react mainly to the Chair as a measurement device, reading the shock off prices around the communications markets watch most. Here that fact is the object of study: I ask whether the attention markets give the Chair is warranted, and show that the information they turn out to have missed is in the non-Chair communications. Whereas this literature typically infers informational content from contemporaneous price responses, I measure policy stance directly from communication, using a dictionary tailored to central bank language following Apel et al. (2022), in the tradition of Loughran and McDonald (2011).[4]

Second, this article suggests that monetary policy surprises—the high-frequency price movements around FOMC announcements—may reflect less surprise than typically assumed in the literature. One view holds that announcements reveal the Federal Reserve's private information about the economy—the "Fed information effect" (Campbell et al., 2012; Nakamura and Steinsson, 2018; Jarociński and Karadi, 2020; Miranda-Agrippino and Ricco, 2021); another that markets have incomplete knowledge of the Fed's reaction function, so that surprises partly reflect its response to news about the rule followed by the policymaker (Bauer and Swanson, 2023b,a; Bauer et al., 2024). Within the second channel, I identify the source of the underused public information and the friction that keeps it unpriced: part of what appears as news at the announcement had been voiced publicly months earlier, by policymakers markets do not attend to. The predictability is distinct from the correlation of surprises with economic news that Bauer and Swanson (2023b) document: their full news set enters the Section 3.2 regressions as controls. It also gives the imperfect reaction-function knowledge documented in survey beliefs by Bauer et al. (2024) and Sastry (2026) a concrete origin: the underweighted public signals are policymakers' communications, unpriced because of institutional attention. The incorporation operates at the weeks-to-months horizon, complementing the days-scale pre-FOMC drift of Lucca and Moench (2015). And whereas Aruoba and Drechsel (2024) construct the policy shock itself from the Federal Reserve's internal staff documents, the mispricing here concerns the public pre-meeting communication that announcement surprises fail to embed.

[4] Djourelova et al. (2025) examine how coordination in FOMC communication affects its immediate transmission to financial markets, complementary to the question here of why some communication is not reflected in prices until later. Existing work quantifies communication with dictionaries, topic models, hawk–dove classifications, word embeddings, and machine-learning or generative-AI methods (Apel and Grimaldi, 2014; Hansen and McMahon, 2016; Istrefi et al., 2023; Hubert and Labondance, 2021; Handlan, 2022; Cieslak and McMahon, 2023; Cieslak et al., 2023; Gambacorta et al., 2024; Dunn et al., 2024; Bertsch et al., 2024). Generative AI is used here only to assemble the historical corpus and, as a validation exercise, to audit the dictionary's sentence-level classifications (Online Appendix C.8).

A third literature examines limited attention and information rigidities. The model follows Maćkowiak and Wiederholt (2009) in allocating scarce attention, here across policymakers according to perceived institutional authority. The consensus-level underreaction it explains complements evidence on informational rigidities among firms and forecasters (Coibion and Gorodnichenko, 2015; Bouchaud et al., 2019; Born et al., 2025) and on markets' incomplete internalization of the Fed's responsiveness to macroeconomic conditions (Cieslak, 2018; Schmeling et al., 2022). The press-coverage results connect to evidence that coverage steers scarce attention and that neglected assets earn higher subsequent returns (Tetlock, 2007; Barber and Odean, 2008; Da et al., 2011; Engelberg and Parsons, 2011; Fang and Peress, 2009; Mullainathan and Shleifer, 2005): the component of communication that forecasts monetary policy surprises is the one the press leaves uncovered.

Fourth, on the Committee itself, the article connects to research on FOMC deliberations and voting (Hansen et al., 2018; Acosta, 2023; Bordo and Istrefi, 2023; Chappell et al., 1993, 2004), and to Ehrmann et al. (2022), who exploit the statutory rotation of Reserve Bank presidents' voting rights and find that their speeches move markets less in the years they vote. The object studied here is distinct: not the contemporaneous market reaction to a speech, but the predictability of subsequent forecast errors from its tone.

Finally, I add a predictor to the literature on bond-return predictability (Cochrane and Piazzesi, 2011; Ludvigson and Ng, 2009; Joslin et al., 2014; Cieslak and Povala, 2015; Bianchi et al., 2021; Bauer and Chernov, 2024), in which Cieslak (2018) emphasizes short-rate forecast errors. Unlike Cieslak and McMahon (2023), who study communication as a predictor of intermeeting risk-premium movements, I trace the predictability to systematic errors in policy expectations that appear in both monetary policy surprises and bond returns.

1 Data

This section describes the six-decade corpus of Federal Reserve communication and the Monetary Stance Signal (MSS), a fixed-dictionary measure of policy stance; the *Reuters* news data used to measure the coverage each speech receives; and the financial and macroeconomic series used throughout the article. Every explanatory variable is matched to the information set available when forecasts are formed; construction details beyond those needed for the analysis are deferred to Online Appendix A.3.

1.1 Constructing the Monetary Stance Signal

I compile a corpus of public communication by members of the Board of Governors serving on the Federal Open Market Committee (FOMC), obtained from the FRASER archive and the Federal Reserve website. Each document is linked to its date of delivery, its author, and the author's institutional position at the time of communication.[5]

The corpus includes 4,177 speeches and congressional testimonies delivered between 1961 and 2024 by 66 policymakers—1,900 by Chairs and Vice Chairs, and 2,277 by non-Chair governors—all of which enter the baseline signal.[6] The full span is used for the bond-return regressions, whose outcome is the only one that reaches back over it; the press-coverage and promotion designs begin in 1992, when *Reuters* coverage of the Board becomes dense, and the FOMC announcement-surprise design in 1995, when the high-frequency surprise series starts. Publicly available FOMC meeting transcripts cover 479 meetings and conference calls between 1976 and 2020, the endpoint determined by the Federal Reserve's five-year publication lag.

Rather than estimating a statistical language model, I recover policy stance from a fixed dictionary: a predetermined mapping from words to stance with no estimated parameters or look-ahead information, so the measure is exactly replicable. The dictionary builds on the FOMC dictionary of Apel et al. (2022), which pairs keywords describing inflation, economic activity, and labor-market conditions with directional modifiers that classify statements as hawkish or dovish, with negations reversing the classification. I extend that framework with additional vocabulary with which U.S. policymakers discuss the same subjects. The resulting thirteen-term dictionary is reported in Online Appendix A.2. Keyword occurrences are split roughly one-half on prices, one-third on activity, and the remainder on employment. These aggregate shares are close for Chairs and non-Chairs, so the asymmetries documented below are unlikely to reflect a difference in the broad subjects discussed.

The analysis is conducted at the subsentence level, after sentences have been split at commas. The algorithm first searches each subsentence for economically meaningful keywords such as *inflation*, *employment*, or *production*. Conditional on finding a keyword, it then searches within the same subsentence for directional modifiers indicating whether the statement conveys a hawkish or dovish assessment. For example, "inflation is high" is classified as hawkish and "employment is weak" as dovish. A single negation within the subsentence reverses the classification; subsentences containing two or more negations, and the rare subsentences in which both hawkish and dovish modifiers appear, are left unoriented, but still count toward the subsentence total defined below. The negation list (35 terms, common to all keywords) is reported with the dictionary in Online Appendix A.2. The same dictionary and scoring rule are applied to each member's interventions in the FOMC meeting transcripts, yielding the meeting-tone measures $MSS_t^{F,j}$ used in Section 4.[7]

[5] Institutional status is obtained from the Federal Reserve Board and classified as Chair, Vice Chair, or neither. Throughout the paper, Board members serving as neither Chair nor Vice Chair are referred to as *non-Chairs*, and the Chair and Vice Chair as principals.

[6] Bibliographical references, tables, appendices, and other non-substantive material are removed before analysis so that only policymakers' own communications enter the textual analysis.

[7] The stance measure is validated in two ways. First, none of the results hinge on the dictionary's exact composition: the unmodified Apel et al. (2022) dictionary correlates at 0.63 with the baseline and reproduces the predictive results at all but the

The Monetary Stance Signal is constructed separately for each institutional group $j \in \{c, n\}$, where c denotes the Chair and Vice Chair group, and n the non-Chair governors. Each speech i receives the tone score $m_i = (\#HS_i - \#DS_i)/\#TS_i$, where $\#HS_i$, $\#DS_i$, and $\#TS_i$ denote its numbers of hawkish, dovish, and total subsentences. The daily speech signal averages the scores of group j's speeches delivered on day t:

$$\widetilde{MSS}_t^j = \begin{cases} \dfrac{1}{N_t^j} \displaystyle\sum_{i \in \mathscr{S}_t^j} m_i, & \text{if group } j \text{ delivers a speech on day } t, \\ 0, & \text{otherwise,} \end{cases} \tag{1}$$

where $\mathscr{S}_t^j$ is the set of group j's speeches on day t and N_t^j its size; dividing by the subsentence count downweights speeches with little macroeconomic content. Because policymakers' views evolve gradually, I smooth the daily signal using an exponential moving average,

$$MSS_t^j = \omega \widetilde{MSS}_t^j + (1 - \omega) MSS_{t-1}^j, \tag{2}$$

where the weight ω is parameterized in terms of a half-life, and the recursion is implemented as the equivalent normalized exponentially weighted average. Throughout the paper, I use a half-life of sixty days, so that a communication loses half its weight after two months—recent enough to track evolving views while preserving persistence in policymakers' expressed stance.[8]

Figure 1 plots the two Monetary Stance Signals. Both track the postwar policy record from 1961 on: hawkish through the inflation of the 1970s, dovish in the disinflationary decades and at the zero lower bound, and sharply hawkish in the 2022–2023 tightening. The monthly correlation between the two groups' tones is 0.39; the analysis below enters the two signals jointly, so identification comes from their divergences rather than their common movement.

shortest horizons, as does an FOMC-specific external lexicon, while a generic sentiment lexicon does not; the results are also stable under leave-one-out and leave-two-out keyword perturbations. Second, at the sentence level, the dictionary's signed classifications agree with an LLM-classified reference sample and an external annotated benchmark well above chance and comparably to external stance lexicons, and an LLM audit of a stratified sentence sample characterizes the dictionary as a sparse but reliable filter whose classifications, when assigned, agree with the audit rater about four times out of five and with the reference labels about three times in four. Online Appendices C.7 and C.8 report these validation exercises.

[8] The predictive results are unchanged under alternative smoothing parameters, and remain negative and significant at every horizon under a carry-forward rule that holds a member's last observed score until the next speech, and under removing a whole-sample linear trend from the aggregated series (Online Appendix C.7).

FIGURE 1—THE MONETARY STANCE SIGNAL BY INSTITUTIONAL GROUP

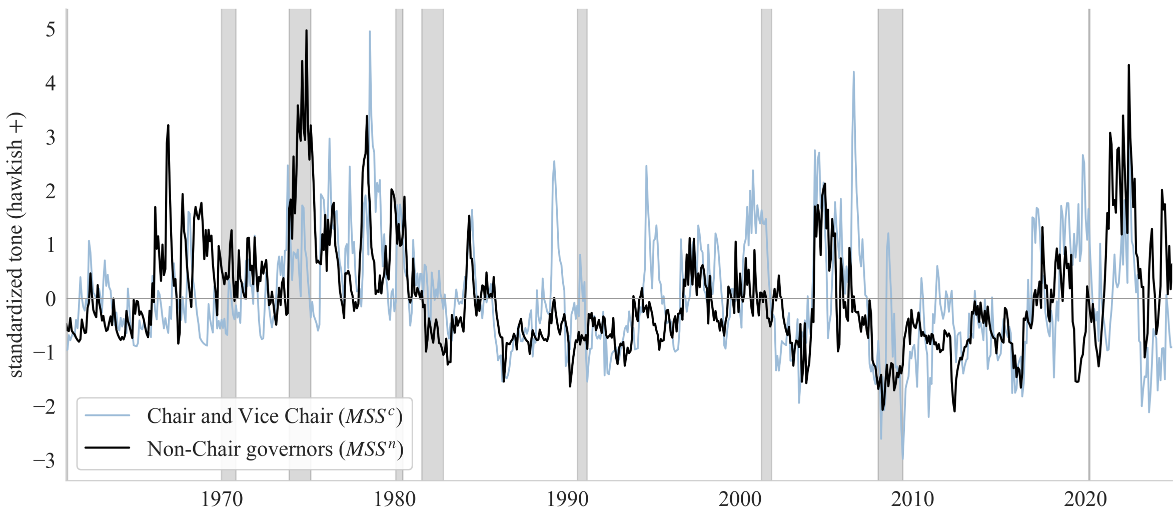

Note: The figure plots month-end values of the smoothed Monetary Stance Signal of Equation (2) for non-Chair governors (MSS^n, black) and the Chair and Vice Chair (MSS^c, light blue), 1961–2024, each demeaned and scaled to unit variance over the sample, so the zero line is the sample mean; higher values are more hawkish. Shaded areas are NBER recessions.

1.2 Press Coverage

To measure the attention devoted to each policymaker's communication, I use the full news archive of *Reuters News* from the Factiva (Dow Jones) database. Coverage of Federal Reserve communication is dense from 1992 onward. Over this period, roughly 100,000 distinct *Reuters* articles mention a Board member by name within a month of one of the member's speeches. I match speeches to the news articles that cover them. A speech's coverage is the number of distinct articles that name the speaker and report the speech within a fixed window around its delivery, with mechanical market wraps and diary round-ups excluded; its coverage tone is the average tone of those articles, scored with the same dictionary as the speeches.The resulting panel assigns coverage to each of the 2,459 Board-member speeches delivered between 1992 and 2024.

1.3 Financial and Macroeconomic Data

I measure expectations of the federal funds rate (FFR) from three complementary sources: federal funds futures at horizons from two months to one year (monthly, sampled on the final trading day), the *Blue Chip Financial Forecasts* (BCFF) survey of about forty-five financial institutions (monthly; sampled quarterly for non-overlapping observations), and, as an additional robustness check, the *Survey of Professional Forecasters* (SPF); for both surveys I use the consensus (mean) forecast. Forecast errors are defined as

$\varepsilon_t^{s,h} = \mathbb{E}_t^s(FFR_{t+h}) - FFR_{t+h}$, where h denotes the forecast horizon and s the source. Positive values correspond to forecasts that were excessively hawkish relative to realized policy. To align speech information with expectations, I measure the MSS three days before each forecast date—the futures sampling date, the SPF collection date, and, for the BCFF, the opening of the survey's response window.[9]

I compute annual holding-period excess returns on n-year zero-coupon Treasuries from the Gürkaynak et al. (2007b) yield curve, along with the first three principal components of yields—the level, slope, and curvature factors (Litterman and Scheinkman, 1991). Yields are available monthly from June 1961 at maturities out to seven years, the panel used in the long-sample regressions, and the full 1–30-year range is complete from 1986.

Macroeconomic controls enter in real time wherever a vintage exists; where the vintage archives do not reach back, the revised series is used. Monthly employment growth and inflation come from the Federal Reserve Bank of Philadelphia's real-time database, and the CFNAI from the Chicago Fed's real-time vintage archive, lagged one month to respect publication delays—two months for the BCFF rounds, which are dated at their first-of-month publication cutoff—and computed from first releases (revised data produce nearly identical results). For specifications estimated at daily, survey-round, or FOMC-meeting frequency, where expectation revisions and price changes may reflect incoming releases rather than communication, I add a panel of real-time ALFRED/FRED releases—nonfarm payrolls, unemployment, headline and core CPI and PCE, industrial production, and real GDP—matched by actual publication date. Online Appendix A.1 lists sources, samples, and mnemonics for the yield, expectations, and macroeconomic reference series; Online Appendix A.3 describes the real-time release panel.

2 Authority, Attention, and Predictive Content

Federal Reserve communication is not processed uniformly. Expected short rates move when the Chair or Vice Chair speaks and barely respond to the speeches of other Board members. Yet the information markets later prove to have missed is concentrated in precisely those overlooked communications. This section documents the tension in three steps. Section 2.1 establishes that attention tracks institutional authority across the full cross-section of speakers. Section 2.2 shows that the predictive content of the same communications runs the other way. Section 2.3 then holds the set of speakers fixed and asks what happens to attention and predictive content when one person's office changes. A closing subsection weighs three alternative readings—a statistical artifact of averaging across speakers, compensation for risk, and a change

[9] This leaves forecasters time to incorporate speeches while ensuring that the signal includes recent communications. No speech after the evaluation date enters the signal. BCFF rounds are dated at their first-of-month publication date, but the survey is answered between the 23rd and 27th of the preceding month, so the signal is evaluated on the 20th of the response month—at least three days before the documented response window opens—and only speeches available to the panelists enter the MSS. Varying the evaluation date between one day and one week before the forecast date for the futures and the SPF and between the 18th and the 22nd of the response month for the BCFF, yields similar results.

in communication style—and finds no support for any of them.

2.1 Attention, Media, and Institutional Authority

Table 1 describes the *Reuters* coverage that Board members' speeches receive, separately for principals and non-Chairs. The table reports two aspects of that coverage: how closely the articles relay the tone of the speech they cover, and how many articles a speech draws. Online Appendix C.1 reports the coverage gap under alternative normalizations and event windows.

TABLE 1—NEWS COVERAGE OF BOARD SPEECHES, BY AUTHORITY

	Principal (Chair, Vice Chair)	Non-Chair (other governors)
Pass-through $\hat{\beta}$ (mean coverage)	0.056***	0.056***
	(0.016)	(0.016)
Correlation ρ	0.173	0.107
Articles per speech	65.0	10.0
N (speeches)	862	1145

Note: The table describes the Reuters News coverage of Board members' speeches, by the speaker's institutional authority (*Principal* = Chair and Vice Chair; *Non-Chair* = other Board governors), over 1991–2024. It takes one observation per matched speech: *Pass-through* $\hat{\beta}$ is the OLS slope of a speech's mean coverage tone on the speech's own tone, with HC1-robust standard errors in parentheses, measuring how faithfully a single article reflects the speech. *Correlation* is the Pearson correlation between speech tone and mean coverage tone. The correlation is not comparable across columns on its own: a group that draws more articles per speech has its coverage-tone mean measured with less noise, which raises the correlation independently of any difference in reporting fidelity; the pass-through slope does not share this bias. *, **, and *** indicate significance at the 10, 5, and 1 percent levels.

Two features stand out. The first is that coverage is faithful for both groups: article for article, the estimated pass-through of speech tone into coverage tone is the same for the two offices. A non-Chair speech that draws coverage at all is relayed as accurately as a principal's.[10]

The second is the gap in how much coverage principal and non-Chair speeches receive. A principal speech draws 65 articles on average against 10 for a non-Chair speech. The pass-through slopes, by contrast, are equal to the third decimal: 0.056 for each office (standard error 0.016), significant at the 1 percent level. The press does not report the two offices differently; it covers one of them far more.

Press coverage is the supply side of attention, and the demand side shows the same concentration. Google Trends searches for policymakers' names jump in the week of a Chair or Vice Chair speech and barely move around a non-Chair speech: pooled across four sample eras, each anchored on the sitting Chair,

[10] Online Appendix C.2 examines fidelity article by article, comparing each covered speech's tone with the mean tone of the articles covering it: the association is positive and significant at every cut.

the speech-week search spike averages $+2.6$ index points for principals against $+0.3$ for non-Chairs, and excluding FOMC-announcement weeks leaves the non-Chair response indistinguishable from zero (Online Appendix B.1). Search volume is not a newsroom's allocation decision, so the concentration is not the media's alone: readers search for the Chair and not for the rest of the Board. Here too the pattern follows the office rather than the individual—searches for "Jerome Powell" were negligible and unmoved by his speeches while he was a governor, and rose sharply around his speeches once he became Chair.

Coverage and searches measure where attention goes. The next results track how this attention shapes the transmission of Federal Reserve communication to forecasts and prices. Table 2 asks whether forecasters revise with the coverage. The outcome is the fixed-target monthly revision of the Blue Chip consensus federal funds rate forecast: the change $r_{t,q} = F_t(\mathit{ffr}_q) - F_{t-1}(\mathit{ffr}_q)$ between consecutive survey rounds in the forecast for the same target quarter q, h quarters ahead. The regressor is the average tone of the *Reuters* articles covering the speeches delivered over the month ending on the 20th of the round's response month, computed separately for coverage of principals and of non-Chairs and entered jointly.[11] Panel A enters the two tone series alone; Panel B adds the changes over the same round in the federal funds rate and in real-time inflation, employment, and activity.

[11] A speech and its covering articles count only through the 20th of the response month, three days before the survey's response window opens, so the coverage is available to the panelists whose revision it is asked to explain.

Table 2—BCFF Forecast Revisions and Speech Coverage

	h = 1Q	h = 2Q	h = 3Q	h = 4Q
A. Average news tone, non-Chair and Chair/Vice				
$\bar{m}_t^n$	0.507**	0.459*	0.573**	0.464**
	(0.238)	(0.261)	(0.269)	(0.236)
$\bar{m}_t^c$	2.334***	2.668***	2.765***	2.623***
	(0.461)	(0.485)	(0.511)	(0.521)
R2adj	0.119	0.121	0.132	0.123
B. Average news tone, ΔFFR_t, ΔCPI_t, $\Delta CFNAI_t$ and ΔEMP_t				
$\bar{m}_t^n$	0.365*	0.301	0.418*	0.342*
	(0.207)	(0.227)	(0.235)	(0.207)
$\bar{m}_t^c$	1.622***	1.920***	1.965***	2.005***
	(0.407)	(0.428)	(0.443)	(0.454)
R2adj	0.240	0.240	0.247	0.219
N (survey rounds)	406	406	406	381

Note: Fixed-target monthly revisions of the Blue Chip Financial Forecasts (BCFF) consensus federal funds rate forecast, $r_{t,q} = F_t(ffr_q) - F_{t-1}(ffr_q)$ in percentage points, for the target quarter q that is h quarters ahead of the survey's own quarter; only adjacent survey rounds are differenced. The sample is the 406 survey rounds from March 1991 to December 2024. The $h = 4$ column is estimated on 25 fewer rounds, all between April 1991 and January 1997, because a round that opens a quarter needs its predecessor to reach five quarters ahead and the survey did not report that far until 1997. The regressors are the average tone of the Reuters articles covering the speeches given in the revision period—scored with the same dictionary as the speeches—computed separately for coverage of non-Chair speakers (superscript n) and of the Chair and Vice Chair (superscript c), and entered jointly; a speech and its articles count only through the 20th of the response month, before the survey's response window opens. Panel B adds the change over the same round in the fed funds rate, inflation, employment and the CFNAI, measured as the change between the real-time vintages available at the two cutoffs—a one-month change observed with roughly a two-month publication lag. The round-level standard deviations of the two regressors are 0.0427 (non-Chair) and 0.0246 (principal); the per-standard-deviation magnitudes quoted in the text scale the coefficients by these values. The Wald p-values for equality of the two coefficients at $h = 1$–4 are 0.00/0.00/0.00/0.00 in Panel A and 0.01/0.00/0.00/0.00 in Panel B. Newey and West (1987) standard errors (6 monthly lags) in parentheses. *, **, and *** indicate significance at the 10, 5, and 1 percent levels, respectively.

Forecasters revise with the tone of the coverage, and they revise much more with the tone the press attributes to a principal. In Panel B the principal coefficient is significant at the 1 percent level at every horizon and is worth 4.0 to 4.9 basis points of revision per standard deviation of coverage tone; the non-Chair coefficient is worth 1.3 to 1.8 basis points and is significant only at the 10 percent level, and not at every horizon. The test of equality rejects at every horizon. Attention therefore enters twice: the press relays a principal's stance in about six times as many articles (Table 1), and forecasters revise roughly four to six times as much per unit of the stance those articles report—a ratio of about three per standard deviation, since non-Chair coverage tone is the more volatile series. Part of that second gap is measurement: a non-Chair speech's coverage tone is averaged over 10 articles against 65 for a principal, so it is the noisier regressor, and measurement error attenuates its coefficient toward zero. Prices track the same allocation: each trading day's change in federal funds futures and in one- and two-year Treasury yields moves with the tone the press attributes to principals and not with the tone it attributes to non-Chairs; the non-Chair coefficients are never distinguishable from zero (Online Appendix B.2).

Communication reaches expectations and prices through the coverage it receives, so the reaction to that coverage is part of the effect of communication itself. If principals' words are processed on delivery while non-Chair words go largely unprocessed, then whatever non-Chair speeches carry and principal speeches do not should not be priced when it is spoken; it should appear later, as predictable errors. The next subsection tests that prediction.

2.2 Predictive Content by Speaker Status

The mirror image of this attention gap is that the overlooked speeches contain information left unused by markets and forecasters. The predictive results in this section come from a single specification. For an outcome y_t—either an interest-rate forecast error or an excess bond return—I estimate

$$y_t = \gamma_0 + \gamma_1 MSS_t^j + \gamma_2 X_t + \varepsilon_t, \tag{3}$$

where MSS_t^j is the Monetary Stance Signal built from the speeches of group $j \in \{c, n\}$ (principals or non-Chairs) and X_t collects the real-time controls of Section 1. Outcomes are oriented so that higher values denote excessively hawkish expectations; a negative γ_1 therefore indicates that a hawkish speech tone today anticipates monetary policy tightening not yet embedded in forecasts or prices.

The tone of non-Chair speeches predicts short-rate forecast errors; the tone of Chair speeches carries no comparable content. The finding holds across expectations measures and inference methods. Table 3 reports the regressions for Blue Chip Financial Forecasts errors. Panel A includes the two tone series alone, and Panel B adds the current federal funds rate, one-year employment and inflation growth, and the CFNAI, each control entering as the most recent release available at the round's first-of-month publication date, with Newey and West (1987) standard errors.

TABLE 3—PREDICTIVE REGRESSIONS OF BCFF SHORT-TERM RATE FORECAST ERRORS

	h = 1Q	h = 2Q	h = 3Q	h = 4Q
A. MSS_t^n and MSS_t^c				
MSS_t^n	-1.374***	-3.773***	-6.824***	-10.216***
	(0.519)	(1.040)	(1.575)	(2.295)
MSS_t^c	0.681	1.149	1.564	2.452
	(0.539)	(1.209)	(1.909)	(2.401)
R2adj	0.042	0.096	0.146	0.184
B. MSS_t, FFR_t, ΔCPI_t, $CFNAI$ and ΔEMP_t				
MSS_t^n	-1.542**	-3.510***	-6.078***	-8.769***
	(0.623)	(1.204)	(1.599)	(2.029)
MSS_t^c	0.489	0.790	1.007	1.760
	(0.535)	(1.211)	(1.886)	(2.315)
R2adj	0.068	0.138	0.211	0.264
N (quarters)	136	136	136	136

Note: The table presents predictive regressions of Blue Chip Financial Forecasts (BCFF) federal funds rate forecast errors at horizon h, with h ranging from 1 to 4 quarters ahead; forecast errors are measured in percentage points of the federal funds rate. These regressions estimate Equation (3), where X_t denotes the Panel B controls and is empty in Panel A. Newey and West (1987) standard errors (6 quarterly lags) in parentheses. The full set of controls in Panel B includes the current quarter FFR_t, annual employment growth ΔEMP_t, year-on-year inflation ΔCPI_t, and the $CFNAI$; each control enters as the most recent release available at the round's first-of-month publication date. All regressions include a constant. *, **, and *** indicate significance at the 10, 5, and 1 percent levels, respectively. The data cover the period from 1991Q1 to 2024Q4 for a total of 136 observations.

The table shows three things. The predictive content belongs to non-Chairs: Chair coefficients are never correctly signed and are less than a third the size of the non-Chair coefficients at the horizons where the effect is strongest. The effect strengthens with the forecast horizon—from -1.37 at one quarter to -10.22 at four in Panel A—consistent with tone signaling a shift in the policy stance that builds over subsequent quarters, so that a forecast that fails to embed it accumulates errors as the target recedes. And it is little changed by the macroeconomic controls, consistent with the MSS carrying information beyond observed economic conditions.

The same regression run on federal funds futures gives the same answer at monthly horizons (Table 4). The survey and the futures market differ in frequency, coverage, and institutional setting, so a mechanism specific to futures pricing is unlikely to generate both results. Section 2.4 returns to this agreement as evidence against a risk-premium reading of the predictability.[12]

[12]The predictive pattern also holds in the SPF, which forecasts the three-month Treasury-bill rate at quarterly frequency, and the predictive power declines with bond maturity beyond the one-year note, consistent with the MSS forecasting the expected path of short rates.

Nor is the result driven by the zero-lower-bound decade: dropping 2008–2015 strengthens the one-year BCFF slope (-11.19 against -10.22 in the full sample). In the pre-2008 sample, which predates the zero lower bound entirely, the survey slopes are significant at the 5 percent level or better at every horizon, and the futures slopes remain significant at the 5 percent level at the four- and twelve-month horizons, and at the 10 percent level at eight months. In rolling ten-year windows the relationship sharpens as Federal Reserve communication expands: it first attains sustained significance in windows ending in 2004 and is significant in most windows thereafter. Online Appendices C.3, C.4, and C.6 report these results.

The Chair series again carries no correctly signed content at any horizon.[13] Its only significant cells sit at the shortest futures horizons with the opposite sign to the underreaction prediction. Taken at face value, these suggest overreaction to attended communication in the spirit of Bordalo et al. (2020a): markets move too far on the Chair's words and the front of the futures curve subsequently unwinds.

TABLE 4—PREDICTIVE REGRESSIONS OF FFR-FUTURES FORECAST ERRORS

	h = 2M	h = 4M	h = 6M	h = 8M	h = 10M	h = 12M
A. MSS^n_t and MSS^c_t						
MSS^n_t	-0.553**	-1.893***	-3.560***	-5.159***	-7.349***	-9.527***
	(0.221)	(0.573)	(0.926)	(1.301)	(1.636)	(1.930)
MSS^c_t	0.303*	1.000**	1.333	0.934	0.400	0.663
	(0.158)	(0.503)	(0.933)	(1.223)	(1.197)	(1.289)
R2adj	0.024	0.079	0.115	0.158	0.343	0.376
B. MSS_t, FFR_t, ΔCPI_t, $CFNAI_{t-1}$ and ΔEMP_t						
MSS^n_t	-0.593**	-1.930***	-3.366***	-4.596***	-6.466***	-8.345***
	(0.239)	(0.606)	(0.947)	(1.342)	(1.955)	(2.363)
MSS^c_t	0.298**	1.018**	1.463*	1.192	0.605	0.955
	(0.140)	(0.483)	(0.885)	(1.180)	(1.156)	(1.288)
R2adj	0.065	0.136	0.185	0.204	0.364	0.392
N (months)	433	432	413	368	252	252

Note: The table presents predictive regressions of FFR futures forecast errors at horizon h, with h ranging from 2 to 12 months ahead; forecast errors are measured in percentage points of the federal funds rate. These regressions estimate Equation (3), where X_t denotes the Panel B controls and is empty in Panel A. Newey and West (1987) standard errors (18 monthly lags) in parentheses. The full set of controls in Panel B includes the current month FFR_t, annual employment growth ΔEMP_t, year-on-year inflation ΔCPI_t, and $CFNAI_{t-1}$; the $t-1$ dating is the latest release available at the month-end sampling date, the convention of every month-end-dated design (Online Appendix A.3). Each column's sample spans its contract's availability—longer-horizon contracts are listed later, the one-year contracts from January 2004—so N falls with the horizon. All regressions include a constant. *, **, and *** indicate significance at the 10, 5, and 1 percent levels, respectively. Details for the dates covered in this table are provided in Table A.2 in the Online Appendix.

[13] Nor is the null a signal-to-noise artifact: the Chair series is the more reliable of the two by a split-half test (Online Appendix C.9).

Figure 2 plots the raw series behind these regressions. The co-movement between the futures errors and the inverted non-Chair MSS builds with the forecast horizon, and the figure is the visual counterpart of Table 4: the shortest-dated errors sit near zero and barely track the tone, while the errors nine to twelve months ahead follow it closely. The errors cluster at turns in the policy stance, where a tone that signals a building stance shift should find forecasts wrong: around the easings of 2007–2009 and 2019–2020 and the tightening of 2022–2023, the year-ahead contracts missed most, and the tone was already moving while they were being priced.

FIGURE 2—FFR-FUTURES FORECAST ERRORS AND THE (INVERTED) MSS

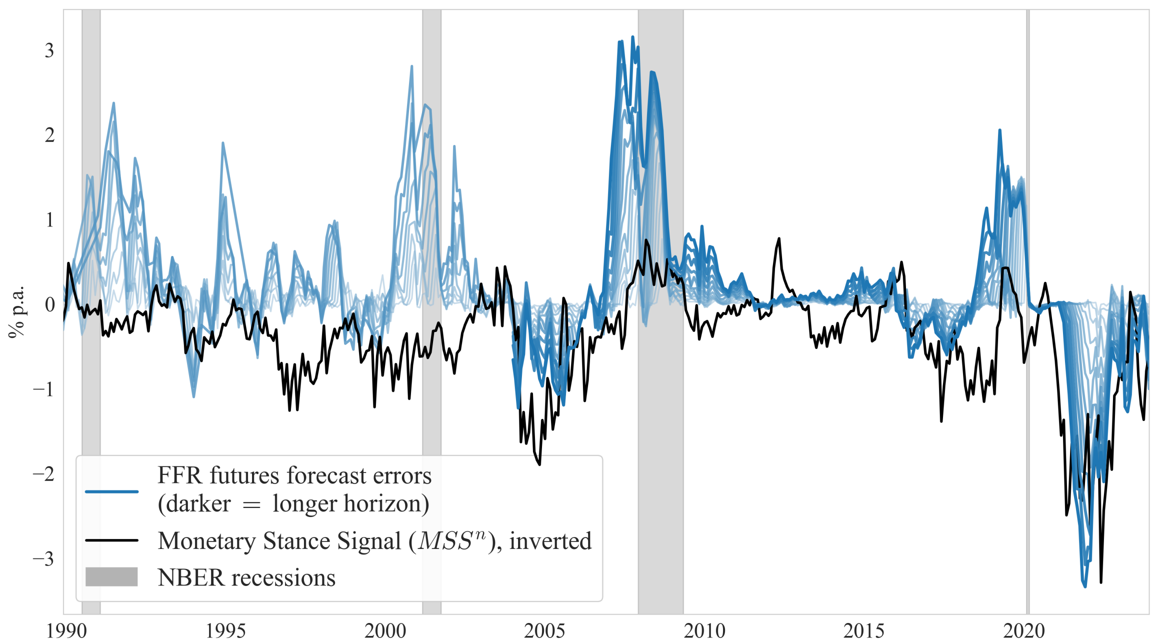

Note: The figure plots federal funds futures forecast errors, dated at the time of prediction, at every monthly horizon from two to twelve months ahead (blue; longer horizons are darker and drawn heavier); Table 4 reports the regressions at the even horizons. The black line is the inverted non-Chair MSS, rescaled to the axis of the errors. The figure shows January 1990 through December 2023. Shaded areas correspond to NBER recession dates.

The same mispricing arises for Treasury returns, over a far longer sample. Short-rate forecasts are available only over short samples (futures) or at quarterly frequency (surveys), which limits the number of observations. Rate-based outcomes also miss policies that act through channels other than direct changes in the federal funds rate. Because the excess bond return (EBR) largely embeds short-rate forecast errors (Cieslak, 2018), I use it as a second outcome: it is available monthly back to June 1961, spanning the entire speech sample.[14]

Table 5 reports the results. Panels A and B present specifications that include the first three principal components of yields together with the non-Chair MSS, estimated over two different sample periods. Panels C and D extend the regressions with auxiliary predictors. The table reports Newey and West (1987) p-values with 18 lags and the more conservative reverse-regression p-values of Hodrick (1992), in the delta-method implementation of Wei and Wright (2013), each also under the bias-corrected bootstrap of Bauer and Hamilton (2018)—four inference methods in all. Columns 1–7 present results for bonds with maturities of two, three, five, seven, ten, fifteen, and twenty-five years. To facilitate comparison across maturities, excess returns are standardized by bond duration, $rx_t^{(n)}/n$.

[14]The determinants of variation in excess bond returns have been extensively studied since Campbell and Shiller (1991). Following the standard approach in the bond-return predictability literature, I estimate Equation (3) with the dependent variable y_t defined as the annual holding-period excess return on an n-year zero-coupon bond, $rx_t^{(n)}$, and with X_t including the macro variables and yield-curve components described earlier.

TABLE 5—PREDICTIVE REGRESSIONS OF EXCESS BOND RETURNS

	n=2Y	n=3Y	n=5Y	n=7Y	n=10Y	n=15Y	n=25Y
A. MSS and PCs (Nobs: 470 months)							
MSS_t^n	-37.627	-43.607	-37.217	-27.936	-18.143	-11.096	-7.37
HAC p-value	0.0	0.001	0.013	0.064	0.234	0.474	0.63
RR p-value	0.007	0.012	0.038	0.093	0.205	0.327	0.476
Bootstrap HAC p-value	0.003	0.015	0.053	0.161	0.331	0.558	0.682
Bootstrap RR p-value	0.022	0.03	0.063	0.142	0.248	0.396	0.522
R2adj	0.33	0.325	0.324	0.316	0.292	0.265	0.246
B. MSS and PCs (Nobs: 763 months)							
MSS_t^n	-30.774	-36.702	-34.265	-28.911			
HAC p-value	0.013	0.008	0.008	0.014			
RR p-value	0.029	0.028	0.034	0.044			
Bootstrap HAC p-value	0.042	0.027	0.034	0.062			
Bootstrap RR p-value	0.075	0.064	0.08	0.11			
R2adj	0.184	0.179	0.194	0.211			
C. MSS, PCs, τ_t^{CPI}, $CFNAI_{t-1}$ and ΔEMP_t (Nobs: 470 months)							
MSS_t^n	-35.045	-41.674	-37.928	-30.974	-23.516	-18.358	-15.79
HAC p-value	0.0	0.0	0.003	0.011	0.04	0.096	0.13
RR p-value	0.011	0.016	0.033	0.059	0.102	0.136	0.187
Bootstrap HAC p-value	0.001	0.006	0.022	0.054	0.115	0.186	0.221
Bootstrap RR p-value	0.023	0.035	0.049	0.091	0.144	0.165	0.254
R2adj	0.368	0.367	0.366	0.358	0.335	0.311	0.291
D. MSS, PCs, τ_t^{CPI} and ΔEMP_t (Nobs: 763 months)							
MSS_t^n	-24.288	-29.431	-27.508	-22.775			
HAC p-value	0.017	0.011	0.014	0.028			
RR p-value	0.063	0.062	0.077	0.099			
Bootstrap HAC p-value	0.063	0.039	0.054	0.083			
Bootstrap RR p-value	0.121	0.137	0.151	0.182			
R2adj	0.214	0.212	0.227	0.242			

Note: The table presents predictive regressions of annual bond excess returns on MSS_t, three yield principal components (PC_t), and additional control variables; the dependent variable is the excess return for a one-year holding period on an n-year bond, standardized by bond duration ($rx_t^{(n)}/n$). These regressions estimate Equation (3), with y_t the n-year excess bond return and X_t the yield principal components and additional controls. Regressions are estimated at a monthly frequency; each month, excess returns earned over the subsequent twelve months are forecast. The additional controls are the inflation trend, the employment growth rate, and the CFNAI. The inflation trend is the discounted moving average τ_t^{CPI} of Cieslak and Povala (2015) ($\rho = 0.987$, window $N = 120$ months). Panels A and C cover the period from 1985M11 to 2024M12; Panels B and D cover the period from 1961M6 to 2024M12. Purchase months run to the end of the sample, so the final annual return is realized in December 2025. Panel D omits the CFNAI, whose series begins in 1967, after the start of the six-decade sample. Maturities beyond seven years are unavailable in Panels B and D because the long-maturity yield series begin in 1971. p-values are reported for three alternative standard error corrections: (1) Newey and West (1987) (HAC) standard errors with 18 lags, (2) the reverse-regression correction (RR) of Hodrick (1992), and (3) bootstrap-based inference following Bauer and Hamilton (2018). The reverse regression relies on predicting monthly returns using annual averages of predictors, rather than predicting annual overlapping returns using the month-t value of a predictor. The delta-method implementation of reverse regressions follows Wei and Wright (2013). Bootstrap regressions are computed using the bias-corrected bootstrap specification of Bauer and Hamilton (2018).

The results parallel those for short-rate forecast errors. The non-Chair MSS predicts excess returns in both samples, and the predictability is strongest at short maturities, while Chair tone has no predictive power for the EBR (not reported, to save space). The specification asks more of the tone than the standard predictability regressions ask of their predictors: every panel conditions on the level, slope, and curvature of the yield curve, so the tone forecasts returns beyond the information in current yields, and the Bauer and Hamilton (2018) bootstrap reported throughout was designed to test precisely such claims of predictability from outside the curve. Panels C and D add the Cieslak and Povala (2015) trend-inflation predictor, which carries strong predictive power for longer-maturity returns, together with real-time employment growth; Panel C adds the activity index as well. The predictability survives: in the post-1985 sample, significance holds through five-year maturities under all four inference methods at or near the 5 percent level, with or without the additional predictors (Panels A and C); in the six-decade sample it holds at the 10 percent level or better under every method through five-year maturities without the macroeconomic controls (Panel B), and only the most conservative combination—the bootstrapped reverse regression with the controls added (Panel D)—no longer rejects at conventional levels (p between 0.12 and 0.18).

A one-standard-deviation more dovish non-Chair tone predicts roughly 25 basis points of additional annual excess return per year of bond duration—about 50 basis points on a two-year note (23 and 47 basis points on the six-decade sample).[15] The two outcomes are one result: because the EBR embeds the short-rate forecast errors, the communications that forecast the errors should forecast returns as well, and they do, back to 1961, decades before the survey and futures data begin. The maturity profile agrees: the predictability sits at short maturities, where a mispriced short-rate path—rather than a term premium—would put it. The result also holds out of sample: by the Clark and West (2007) test, adding the non-Chair MSS to the predictors significantly lowers the mean squared prediction error for bond returns at every reported maturity, for the survey errors under every benchmark at every horizon, and for the futures errors in every cell but one, the constant-only benchmark at the shortest horizon (Online Appendix C.5).

2.3 Evidence from Promotions to Chair and Vice Chair

The two preceding subsections compare different speakers with different offices. They leave open whether markets discount non-Chair communication because of the individuals delivering it or because of the institutional position they hold. This subsection keeps the composition of speakers fixed and varies the office: the policymakers who transition from governor to Chair or Vice Chair are observed in both roles, so speaker

[15] The predictable component is statistically significant but economically modest once translated into a trading rule: the gross Sharpe improvement is small, and realistic transaction costs absorb it. A simple implementable rule—a position in the two-year note proportional to minus the standardized non-Chair MSS, capped at unit exposure—raises the Sharpe ratio of the 1961–2024 two-year annual-return position from 0.23 to 0.29 before costs. The rule turns over about 0.22 units per month, so with transaction costs between 2 and 10 basis points per unit of turnover, the improvement shrinks to roughly 0.01 or reverses outright. The exercise is in-sample and therefore an upper bound. This configuration is what a limited-attention account implies: the mispricing is too small, relative to the cost of monitoring every governor's speech, to attract the capital that would eliminate it, so it persists.

ability and style are differenced out, and what remains is what changes with the office.[16] This is the cleanest separation of authority from identity the setting allows: attention and predictive content move on the same speakers, in opposite directions.

Attention moves first and visibly. Figure 3 tracks press coverage per speech in event time around promotion for the switchers observed on both sides of the transition within the media window: coverage is flat in the years before nomination, rises in the nomination year itself as the appointment becomes news, and then runs between three and six times the pre-nomination level.

FIGURE 3—PRESS COVERAGE AROUND PROMOTION TO CHAIR OR VICE CHAIR

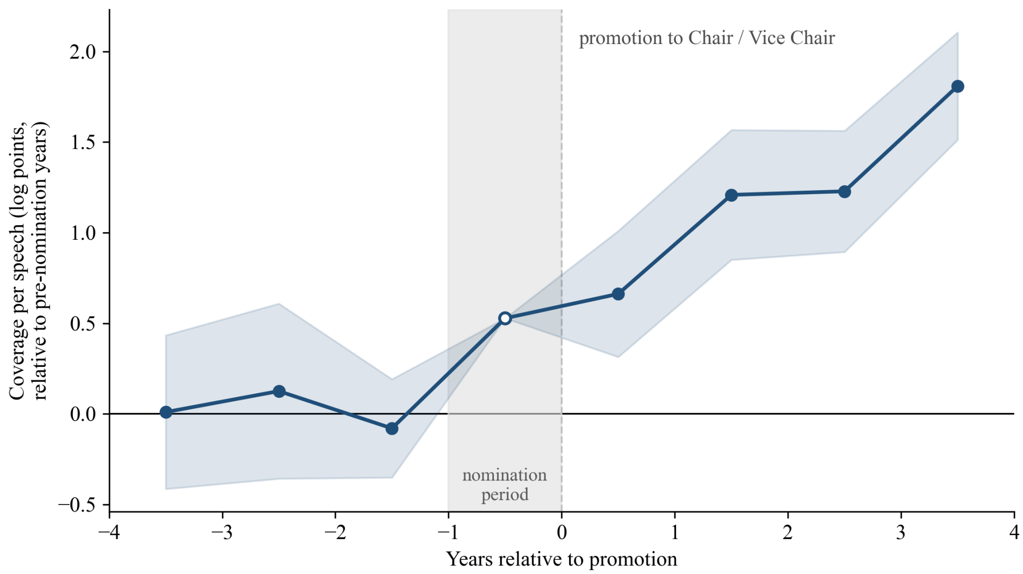

Note: The figure plots coefficients on event-year bins from a regression of $\log(1+\text{Mentions})$—distinct *Reuters* articles in the $(-1,5)$-day window around each speech—on years relative to the speaker's promotion to Chair or Vice Chair, with speaker and year fixed effects and three speech-level controls (the log word count of the speech, the days between the speech and the nearest FOMC meeting, and the log count of Federal Reserve–related *Reuters* articles in the days around the speech), for the seven Board members observed both before and after promotion (865 speeches, 1992–2024). Event time is measured from the speaker's first speech in principal office. Coefficients are normalized to the count-weighted mean coverage of the pre-nomination years (event years ≤ -2), and the terminal bins pool all earlier and later years. The grey band and open marker indicate the nomination year, whose coverage already reflects the announced promotion. The band around the plotted coefficients is the 95 percent confidence band, with standard errors clustered by speaker.

Attention and predictive content are then measured in one table, on one set of speakers and one forecast

[16]The corpus contains thirteen members whose careers run from governor to Chair or Vice Chair. Seven of them—in order of promotion, Roger Ferguson, Ben Bernanke, Donald Kohn, Janet Yellen, Jerome Powell, Lael Brainard, and Philip Jefferson—are covered by the news file on both sides of the transition and form the sample of Figure 3 and Table 6; the bond-return design of Online Appendix B.3 uses all thirteen.

object. Table 6 follows the seven members the news file covers on both sides of promotion through the coverage their speeches draw, the forecast revisions that coverage moves, and the forecast errors their own words predict.

TABLE 6—ATTENTION, PRICING, AND INFORMATION ACROSS PROMOTION TO CHAIR OR VICE CHAIR

	Before (governor)	After (Chair/Vice)	Wald (p)
Panel A. Coverage of the speech (per speech)			
Articles per speech	14.2	62.7	
Pass-through $\hat{\beta}$ (mean coverage)	0.021	0.062***	0.216
	(0.026)	(0.021)	
Correlation ρ	0.050	0.204	
N (speeches)	249	446	
Panel B. BCFF short-rate forecast revisions on the tone of that coverage			
$h = 1Q$ ($N = 350$)	0.422***	1.285***	0.025
	(0.150)	(0.332)	
$h = 2Q$ ($N = 350$)	0.417**	1.664***	0.004
	(0.172)	(0.380)	
$h = 3Q$ ($N = 350$)	0.441***	1.907***	0.001
	(0.168)	(0.403)	
$h = 4Q$ ($N = 345$)	0.409***	2.001***	0.001
	(0.153)	(0.420)	
Panel C. BCFF short-rate forecast errors on the speakers' own tone			
$h = 1Q$	-0.054***	0.022	0.027
	(0.018)	(0.022)	
$h = 2Q$	-0.135***	0.007	0.048
	(0.039)	(0.049)	
$h = 3Q$	-0.254***	-0.005	0.015
	(0.067)	(0.076)	
$h = 4Q$	-0.408***	0.007	0.001
	(0.116)	(0.094)	
N (quarters)	116		

Note: The promotion design on one group of speakers: the 7 Board members whose career runs Governor → Chair or Vice Chair and whom the *Reuters News* file covers on both sides of their own promotion (Ferguson, Bernanke, Kohn, Yellen, Powell, Brainard, Jefferson), with 281 pre-promotion and 583 post-promotion speeches over 1992–2024. *Before* and *After* are the same people either side of their own promotion, so the design holds the speakers fixed and varies the office. *Panel A* is Table 1 on this group, one observation per matched speech, with HC1 standard errors; its N counts the matched speeches only, while the speech totals above count every speech either side of promotion. *Panel B* is the Blue Chip Financial Forecasts (BCFF) revision regression of Table 2 with its macroeconomic controls and the pre- and post-promotion coverage entered jointly, over up to 350 rounds from November 1995 to December 2024 (endpoint rounds lose a target quarter at $h = 4Q$). The round-level average coverage tone has standard deviation 0.0428 before and 0.0245 after promotion; the per-standard-deviation magnitudes quoted in the text scale the Panel B coefficients by these values. *Panel C* is the joint specification of Table 3 with the pre- and post-promotion tones in place of the two offices, over 1995–2024, both scaled by the standard deviation of the pre-promotion tone. The *Wald* column tests the after-promotion interaction in Panel A and equality of the raw coefficients in Panels B and C, in Panel B per unit of tone rather than per standard deviation. Newey and West (1987) standard errors (6 monthly lags in Panel B, 6 quarterly lags in Panel C) in parentheses. *, **, and *** indicate significance at the 10, 5, and 1 percent levels, respectively.

Panel A measures the coverage itself. The same speaker draws 14.2 articles per covered speech as a governor and 62.7 after promotion, and article for article the stance passes into that coverage at 0.021 before promotion against 0.062 after. The article counts are close to the cross-section values of Table 1, where a non-Chair speech draws 10 articles against 65 for a principal; promotion moves a speaker from one column of that table to the other. The two per-article slopes are not statistically distinguishable (Wald $p = 0.22$), and the pre-promotion slope is not distinguishable from zero either: it is estimated on 249 matched speeches with a standard error of 0.026, so this design does not settle whether per-article fidelity changes at promotion. What promotion multiplies is the volume of coverage. The correlation between a speech's tone and its mean coverage tone is 0.05 before promotion and 0.20 after, but a post-promotion speech's coverage tone is averaged over 63 articles against 14 before, so the gap in correlations partly reflects the precision with which each is measured.

Panel B carries that coverage into the published consensus. It re-estimates the Blue Chip revision regression of Table 2 on these seven members alone, entering the coverage of the same individuals' pre-promotion and post-promotion speeches jointly. Both coefficients are significant at the 5 percent level or better at every horizon, and the Wald test rejects equality of the two at every horizon, with p-values between 0.001 and 0.025. Per standard deviation the tone of post-promotion coverage is the larger channel, between 3.2 and 4.9 basis points against 1.8 to 1.9 before promotion. The pattern of Section 2.1 thus repeats within the person: the same speaker's reported stance moves the consensus three to five times as much per unit of tone once the office changes, and it reaches the consensus through four times as many articles.

The predictive content, in Panel C, moves the opposite way. The regression is that of Section 2.2—the Blue Chip consensus forecast error on speech tone—run on the switchers alone, entering the same individuals' pre-promotion tone and post-promotion tone jointly. Both tones are scaled by the standard deviation of the pre-promotion tone, so the coefficients read per standard deviation rather than in the raw units of Table 3. Before promotion their tone predicts the error at every horizon, from -0.054 one quarter out to -0.408 four quarters out, each at the 1 percent level. After promotion the coefficients are indistinguishable from zero, and a Wald test rejects equality of the two at every horizon. The same reversal appears in excess bond returns, on the wider sample of thirteen switchers that does not require the news file to cover both sides of the promotion (Online Appendix B.3).

The predictive content of a given person's communication thus falls just as markets begin to attend to it. What changes across the promotion is the attention the speaker commands, not the content of what they say: a switcher's stance is statistically indistinguishable across the transition, as the next subsection shows.

2.4 Alternative Explanations

Taken together, the full cross-section, six decades of bond returns, and the status-change reversals place the asymmetry at odds with full-information rational expectations (Coibion and Gorodnichenko, 2015; Bouchaud et al., 2019; Born et al., 2025); the last, on a sample that holds the identity of speakers fixed, makes the case most cleanly. I investigate three competing explanations for these facts: an artifact of averaging across speakers, compensation for risk, and a change in communication style.

The first is an artifact of averaging: the non-Chair signal pools the speeches of several governors while the Chair signal tracks one or two, so a cross-sectional average of multiple voices could be the better statistical signal for reasons that have nothing to do with attention. The scope for this is small from the outset: on average only 3.3 distinct governors speak within a rolling quarter before each survey round. I use a Monte Carlo test to address the question directly (Online Appendix B.4). Across 150 random draws for the forecast errors and 100 for bond returns, the signal is rebuilt from a restricted speech set and the headline regressions re-estimated: *single-voice* draws keep one randomly chosen governor's speeches per quarter, so nothing is averaged across speakers, while *volume-matched placebo* draws keep the same number of speeches drawn at random across governors. If averaging drove the predictive power, the single-voice estimates would be weaker than the placebo's. They are not: the single-voice and placebo distributions are not distinguishable in either design (Mann–Whitney $p = 0.33$ and $p = 0.51$), and the estimated share of the signal attributable to averaging is 0.00 in both. A lone governor predicts about as well as the average of several, so the discount does not run through diversifiable speaker noise.

The second is that the predictive power of speech tone reflects a time-varying risk premium (Miranda-Agrippino, 2023). Two features of the evidence cut against it. First, as Section 2.2 notes, the non-Chair slopes are nearly identical whether expectations are measured from futures or from surveys: futures may embed premia at longer horizons (Krueger and Kuttner, 1996; Gürkaynak et al., 2007a; Piazzesi and Swanson, 2008; Hamilton, 2009), whereas the BCFF and SPF surveys are free of them by construction. And since the EBR largely embeds short-rate forecast errors (Cieslak, 2018), the predictability of those errors already implies bond-return predictability. Second, the predictive power of the non-Chair MSS declines with bond maturity (Table 5), the opposite of a premium channel, which strengthens at the long end; and the sign of the coefficient would require short-maturity premia to be countercyclical—investors demanding unusually high compensation to hold short-term Treasuries precisely during recessions, at odds with these securities' safe-asset role. Premia in short-rate futures and policy surprises are in any case economically small (Piazzesi and Swanson, 2008; Schmeling et al., 2022).[17]

The third is a change in communication style. The switcher design compares the same speakers across offices, but a promoted governor might simply speak differently upon promotion—more guardedly, on different topics—so that the change in predictive content reflects the words rather than the attention paid to them. This matters only if the change touches the expressed stance, and it does not detectably. In regres-

[17] For the EBR regressions, the macro-finance literature also offers spanning-based accounts of return predictability beyond the yield curve: Duffee (2011) and Joslin et al. (2014) emphasize hidden or unspanned factors, whose effects on short-rate expectations and term premia offset across maturities, while Cochrane and Piazzesi (2011) and Duffee (2011) note that measurement noise may obscure the risk-premium information embedded in current yields. These accounts describe how a predictor can forecast returns without appearing in yields; they do not speak to the forecast-error evidence, which is measured directly from expectations data rather than inferred from prices.

sions of four features of a speech—its stance m, the absolute stance $|m|$, its dictionary-keyword density, and its length—on the status indicator, with speaker fixed effects and standard errors clustered by speaker, the stance measures do not move detectably across promotion (Online Appendix B.5). The one feature that does move is speech length, which falls by roughly 620 words. Length does not enter the signal, which scores the direction of stance among policy sentences rather than the number of words, so the shift leaves the measured stance untouched.

The next section asks where the mispriced information comes from.

3 Rule Deviations and Monetary Policy Surprises

The previous section establishes that non-Chair speeches carry information markets leave unpriced. This section identifies what that information is about: deviations of policy from the Committee's own rule, which appear at the FOMC announcement as policy surprises. I establish three results. First, forecasters' errors correlate with policymakers' deviations from the monetary policy rule. Second, non-Chair speech tone predicts these deviations. Third, the tone non-Chair members express between consecutive meetings forecasts the forward-guidance component of the next announcement's surprise, even after conditioning on the full public-news set.[18]

3.1 Speech Tone and Deviations from the Policy Rule

To measure how deviations from the interest-rate rule relate to forecasters' prediction errors, I estimate the FOMC's reaction function on the projections available to it in real time. The policy rate at each Greenbook date s is regressed on the staff's *nowcast* of unemployment and inflation,

$$i_s = \alpha + \beta\, \mathbb{E}_s^{GB}[u_s] + \gamma \mathbb{E}_s^{GB}[\pi_s] + \varepsilon_s, \tag{4}$$

where $\mathbb{E}_s^{GB}[\cdot]$ denotes the Federal Reserve Board staff's Greenbook/Tealbook nowcast formed at s, and i_s is the policy rate in the first month of quarter s. The Taylor rule is estimated recursively by exponentially weighted least squares with a forgetting factor $\delta = 0.99$—a half-life of roughly seventeen years, so the coefficients drift slowly. The vintage-t estimates $(\hat{\alpha}_t, \hat{\beta}_t, \hat{\gamma}_t)$ use Greenbook dates through the one preceding t only.[19] The *expected deviation* at horizon h is obtained by projecting this rule onto the staff's h-quarter-ahead

[18] The link between central banks' deviations from policy rules and their communications is emphasized in Yellen (2013), Yellen (2017), Tietz (2018), and Schmeling et al. (2022).

[19] $\delta = 0.99$ matches the slow drift documented for estimates of the Federal Reserve's reaction function (Primiceri, 2005) and is equivalent to a constant-gain learning rate of $1-\delta = 0.01$ per quarter, at the persistent end of the adaptive-learning literature on monetary policy (Eusepi and Preston, 2018). Across half-lives between roughly three and thirty-five years ($\delta \in [0.95, 0.995]$), the Panel A correlations below rise with the horizon, and the non-Chair/Chair adjusted-R^2 profiles are unchanged. The Greenbook's five-year confidentiality lag caps the sample about five years before the present, and the initial estimation window spans ten years

forecast, $\hat{\imath}_{t+h|t} = \hat{\alpha}_t + \hat{\beta}_t\, \mathbb{E}_t^{GB}[u_{t+h}] + \hat{\gamma}_t\, \mathbb{E}_t^{GB}[\pi_{t+h}]$, and comparing it to the realized rate: $\varepsilon_{t+h}^{Taylor} = \hat{\imath}_{t+h|t} - i_{t+h}$, positive when policy turns out more accommodative than the forward-looking rule implied.[20]

Deviations from the rule are what forecasters miss. Panel A of Table 7 reports correlations between BCFF forecast errors and the deviations for each prediction quarter: near zero at the one-quarter horizon, rising to 0.46 at one year, and statistically significant beyond a single quarter. The forecast error and the deviation both subtract the realized rate, so part of the first-row correlation is mechanical. The second row instead correlates the two *predictors*—the rule's forecast $\hat{\imath}_{t+h|t}$ and the survey forecast $\mathbb{E}_t^{BCFF}\, i_{t+h}$—neither of which contains the realized rate. The two correlate at 0.44 to 0.56, significant at every horizon: forecasters largely follow where the real-time rule points.

TABLE 7—TAYLOR-RULE DEVIATIONS: FORECAST ERRORS AND SPEECH TONE

	$h=1Q$	$h=2Q$	$h=3Q$	$h=4Q$
A. Correlation with BCFF forecast errors				
$\rho(\varepsilon_{t+h}^{Taylor},\ \varepsilon_t^{BCFF,h})$	-0.01	0.20**	0.34***	0.46***
$\rho(\hat{\imath}_{t+h|t},\ \mathbb{E}_t^{BCFF}\, i_{t+h})$	0.45***	0.44***	0.50***	0.56***
N (quarters)	120	120	120	120
B. Predictive regression: both tone series				
MSS_t^n	-6.048	-15.162**	-17.260***	-18.155***
	(6.281)	(6.121)	(6.048)	(5.277)
MSS_t^c	-15.587***	-10.780***	-9.121***	-7.482*
	(2.919)	(3.009)	(3.436)	(3.830)
R2adj	0.323	0.321	0.335	0.324
N (quarters)	114	114	114	114

Note: The Taylor rule is estimated in real time by regressing the policy rate on the Federal Reserve Board staff's (Greenbook/Tealbook) *nowcast* of headline CPI inflation and unemployment—one exponentially weighted reaction function ($\delta = 0.99$, ten-year initial window, estimated through the preceding Greenbook only)—and then projecting the fitted rule onto the staff's h-quarter-ahead forecast to obtain $\hat{\imath}_{t+h|t}$, a real-time forecast of the target-date rate. The expected deviation is $\varepsilon_{t+h}^{Taylor} = \hat{\imath}_{t+h|t} - i_{t+h}$, in percentage points of the federal funds rate. Greenbook forecasts are released with a five-year lag. The sample covers 1991:Q2–2021:Q4 (120 quarters at each horizon in Panel A, 114 in Panel B; the stated span is the union across horizons). Panel A correlates the deviations with h-quarter-ahead Blue Chip Financial Forecasts (BCFF) short-rate forecast errors, averaged within each prediction quarter. Its second row correlates the two predictors directly—$\hat{\imath}_{t+h|t}$ and the BCFF forecast $\mathbb{E}_t^{BCFF}\, i_{t+h}$. Stars in Panel A use the Pearson correlation p-values, unadjusted for the overlap of the h-quarter-ahead objects. Panel B regresses the deviations on non-Chair and Chair speech tone, MSS_t^n and MSS_t^c, each the last observation available as of the BCFF survey date and both entered jointly in a single regression, with Newey and West (1987) standard errors (6 quarterly lags) in parentheses. All regressions include a constant. *, **, and *** indicate significance at the 10, 5, and 1 percent levels, respectively.

If forecasters miss the rule deviations, then any public signal of the deviations is information they fail

(40 quarterly observations).

[20] The deviation is the rule's forecast error for the target-date rate. It can reflect drift in the coefficients, the staff's forecast errors about future conditions passed through the rule, misspecification of the rule's inputs or functional form, and monetary policy shocks.

to incorporate. Panel B of Table 7 asks whether speech tone is such a signal: it regresses the deviation $\varepsilon_{t+h}^{Taylor}$ on both tone series at once, one regression per horizon. A hawkish (dovish) non-Chair tone predicts, from two quarters out, a tighter (looser) policy stance than the rule implies, and the relationship strengthens with the horizon and is significant at the 5 percent level or better at every horizon from two quarters out. Chair tone shows the reverse profile: it is strongest one quarter out and weakens steadily through one year, though it remains significant at every horizon, at the 10 percent level by one year. Between one quarter and four the two offices reverse in magnitude—the non-Chair coefficient roughly triples while the Chair's halves; the reversal rests on the point estimates, without a test of the difference. The asymmetry this article documents therefore lies in incorporation, not content. From two quarters out—the horizons at which Panel A shows the deviation to be what forecasters miss—both offices' tones are informative about the deviation. Yet only non-Chair tone forecasts the errors (Tables 3 and 4).[21]

3.2 Speech Tone and Monetary Policy Surprises

The rule-deviation evidence shows that non-Chair speeches carry the stance information that forecasters miss over the following quarters. If markets underweight those speeches, part of what they say should enter prices only when the decision forces it there: at the announcement itself. For each scheduled meeting, I average each group's speech tone over the most recent intermeeting window and ask whether that tone forecasts the surprise at the announcement.[22] The windows follow the policy calendar, and every speech in them is public before the announcement, so under full incorporation the tone should forecast nothing. The surprises are the Gürkaynak et al. (2005b) target and path factors and the overall Nakamura and Steinsson (2018) policy-news shock.[23] Each surprise is regressed on both groups' tone, without controls and then with the full Bauer and Swanson (2023b) public-news set added as controls. The controlled regressions ask whether the words carry information beyond what public macroeconomic and financial data already predict.

Table 8 reports the result. Non-Chair tone expressed over the intermeeting period before a decision forecasts the *path* factor of the upcoming announcement—a hawkish tone forecasting a hawkish path surprise, about 0.17 standard deviations per standard deviation of tone, significant at the 1 percent level, net of the public-news controls. The target factor carries no loading, while the overall policy-news shock loads more weakly (about 0.08, at the 10 percent level). Forward guidance—the statement's language about the future path of policy—is where a stance that binds over quarters should appear. The Chair-group tone

[21] Lucca and Trebbi (2011) also find that the semantic orientation of FOMC statements explains a substantial share of Taylor-rule deviations. Their signal is the Committee's official and most closely watched communication; the results here show that the same rule information travels through the communications markets watch least.

[22] A single intermeeting window does not always contain speeches from both groups; a group with no speech in the last window carries its average over the preceding window instead, which keeps the meeting in the sample. Results are similar when the tone is averaged over the two most recent windows.

[23] I use the published high-frequency series of Acosta et al. (2024), which extends these measures through the recent tightening cycle by splicing SOFR futures onto the Eurodollar series; scheduled meetings only.

forecasts none of the three surprise measures: the words markets watch most closely have no measurable forecasting power for the announcement surprise. This is what full incorporation would imply, though the null is also consistent with the Chair's intermeeting tone carrying no signal; Section 4 separates the two.

TABLE 8—MONETARY POLICY SURPRISES AND SPEECH TONE

	NS	Target	Path
A. Without Bauer–Swanson news controls			
$SpeechTone^{n}_{(t-1,t)}$	0.133**	-0.046	0.229***
	(0.067)	(0.062)	(0.057)
$SpeechTone^{c}_{(t-1,t)}$	0.011	0.011	0.004
	(0.048)	(0.070)	(0.035)
R^2	0.02	0.00	0.05
B. With Bauer–Swanson news as controls			
$SpeechTone^{n}_{(t-1,t)}$	0.081*	-0.061	0.172***
	(0.048)	(0.054)	(0.039)
$SpeechTone^{c}_{(t-1,t)}$	-0.027	0.002	-0.039
	(0.062)	(0.078)	(0.032)
R^2	0.13	0.02	0.17
N	228	228	228

Note: Predictive regressions of announcement surprises on intermeeting speech tone, estimated jointly: each surprise is regressed on both groups' intermeeting tone. $SpeechTone^{j}_{(t-1,t)}$ is the average MSS tone of group $j \in \{n, c\}$ (non-Chair governors; Chair and Vice Chair) over the days strictly between scheduled meetings $t-1$ and t; a group with no speech in that window carries its average over the preceding window instead. Dependent variables: the Acosta et al. (2024) policy-news shock (*NS*) and the target and path factors. Panel A regresses the raw surprises; Panel B adds the full Bauer and Swanson (2023b) public-news set as controls, so the tone coefficients are net of public macro and financial data. Sample: scheduled FOMC meetings, May 1995–December 2023; both panels share it (meetings with the full control set). Tones and surprises standardized on the regression sample. Newey and West (1987) standard errors (18 lags at the meeting frequency) in parentheses. *, **, and *** indicate significance at the 10, 5, and 1 percent levels, respectively.

An Illustrative Episode: February to September 2009. On February 11, 2009, in a public speech that drew the modest coverage typical of a governor's, Governor Elizabeth Duke described "high rates of job loss and weak income growth" restraining housing demand and housing-sector weakness that "remains a significant drag on the macroeconomy." Her tone sat in the far dovish tail of the non-Chair signal. The private record of the September 2009 FOMC meeting shows the same view: Duke found the statement's opening paragraph "awfully rosy" and said she could "envision a good case for actually extending our purchases" into the second quarter of 2010. The Committee extended its purchases and reaffirmed its commitment to continued accommodation, and the forward-guidance component of the announcement registered one of the sample's largest dovish path surprises, 2.3 standard deviations on the 1995–2020 surprise sample—the transcript-capped window of Section 4. The dovish tilt her public words had signaled months earlier was still moving prices at the announcement. The rule-deviation regressions above (Table 7) show the episode is representative rather than exceptional: tone voiced quarters ahead forecasts the deviation from the rule that forecasters miss. The tone of the final intermeeting window before the September meeting was itself close to neutral, so the surprise reads as the delayed pricing of a stance voiced months earlier, at the horizon of Table 7, rather than of the last window's speeches.

Together, the two subsections show that forecasters miss the deviations from the rule that build over quarters (Panel A), non-Chairs voice the direction of that drift in advance (Panel B), and, consistent with markets underweighting them, the drift reappears as the predictable forward-guidance component of the announcement (Table 8). Non-Chair tone thus forecasts the path of policy, the Committee's evolving stance relative to the rule, rather than a contemporaneous level surprise. Neither the regressions nor the episode explains why public words from a sitting governor are processed only at the announcement, if at all. The next section measures the underweighting directly, against the benchmark of an optimal attention allocation.

4 Attention Versus Measured Influence

The preceding sections show markets attending to the Chair while the information sits with non-Chairs. This section asks whether that allocation is justified—whether non-Chairs simply matter less—or whether it gives non-Chairs less attention than their measured influence warrants. The benchmark is a rational-inattention model in the tradition of Sims (2003) and Maćkowiak and Wiederholt (2009).

4.1 A Rational-Inattention Benchmark

The model has a single period. Policymakers—classified as Chair or non-Chair—communicate their views and then set the interest rate based on their private assessment of the economy; a representative forecaster predicts that rate, allocating costly attention across policymakers' noisy speech signals depending on the speaker's institutional status.[24] The interest-rate decision is linear in policymakers' stances,

$$i_t = \nu_c\, \alpha_t^c + \nu_n\, \alpha_t^n + g(X_t) + \varepsilon_t^i, \tag{5}$$

where ν_c and ν_n are the fixed weights of the Chair and non-Chair members in the decision, $g(\cdot)$ is the systematic component of the rule—known to the forecaster, who also observes the state of the economy X_t—and ε_t^i is a monetary policy shock. The stances α_t^c and α_t^n are communicated in speeches, but the

[24] Throughout, the terms *views* and *stances* are used interchangeably. Superscripts c,n index the two groups throughout: group c pools the Chair and Vice Chair, group n the remaining members.

forecaster has limited capacity to process them: she encodes each group's communication into a private signal $S_t^j = \alpha_t^j + \eta_t^j$, $j \in \{c, n\}$, whose independent Gaussian noise η_t^j is *processing* error, with precision set by the attention κ_c, κ_n (in bits) allocated subject to a capacity constraint $\kappa_c + \kappa_n \leq \kappa$. The noise is the forecaster's, not the speaker's: the communication itself reveals the stance, so an econometrician reading the full text recovers α_t^j, while the capacity-constrained forecaster carries away only the noisy summary S_t^j.

Online Appendix D collects the model's derivations, including the proof of the proposition below; the optimal allocation follows Sims (2003) and Maćkowiak and Wiederholt (2009). Attention tilts toward the group whose stance is more influential for the decision or more variable: optimal attention to the Chair is increasing in the single index

$$x = \frac{\sigma_c^2}{\sigma_n^2}\left(\frac{\nu_c}{\nu_n}\right)^2, \tag{6}$$

where σ_j^2 is the variance of group j's stance. Concentrating attention on the Chair can therefore be optimal in principle—when the Chair's stance is more important or more volatile than the non-Chairs'—and the forecaster's predicted rate then responds strongly to Chair speeches and little to non-Chair speeches; whether it is optimal in fact is a measurement question.

Endogeneity of Monetary Policy Surprises. The model's central implication is that monetary policy surprises are a by-product of the attention allocation. After the decision i_t is announced, the monetary policy surprise $Surprise_t \equiv i_t - \mathbb{E}(i_t \mid \mathscr{H}_t)$ loads on each group's latent stance through the attention paid to its communication, where $\mathscr{H}_t = \{S_t^c, S_t^n, X_t\}$ is the forecaster's information set.

Proposition 1 (Endogeneity of monetary policy surprises)**.** *In equilibrium, the monetary policy surprise loads on group j's latent stance α_t^j with coefficient*

$$\beta_j \;\equiv\; \nu_j(1-\lambda_j) \;=\; \nu_j 2^{-2\kappa_j^*}, \qquad j \in \{c, n\}, \tag{7}$$

where $\lambda_j = 1 - 2^{-2\kappa_j^*}$ is the share of group j's stance priced before the announcement. For any group with non-zero influence ($\nu_j \neq 0$), $|\beta_j|$ is strictly decreasing in the attention κ_j^* allocated to that group, with $\beta_j = \nu_j$ at $\kappa_j^* = 0$ and $\beta_j \to 0$ as $\kappa_j^* \to \infty$.

A heavily attended group contributes little to the surprise even when its influence is large; an under-attended group generates predictable surprises in proportion to its influence—exactly the non-Chair pattern of Section 3.[25]

[25] The model's parameter space is one-sided: β_j lies in $(0, \nu_j]$ for positive influence, so attention drives a loading to zero, never below it. Overreaction lies outside the framework and marks its boundary (Bordalo et al., 2020a); the exercise below conditions on the underreaction sign, and Online Appendix D.1 discusses the boundary and the sign condition, together with the short-horizon

4.2 Taking the Model to the Data

The model reduces the measurement to two regressions on one pair of stance proxies. The surprise regression identifies the loadings $\beta_j = \nu_j(1-\lambda_j)$ of Proposition 1—the product of influence and inattention—while the decision regression identifies the influence weights ν_j on their own; dividing the first by the second isolates the inattention, which the proposition's mapping converts into bits of attention, and the recovered allocation can then be set against the optimal one implied by the index x. This subsection carries out the first step; Section 4.3 completes the chain.

As proxies for α_t^c and α_t^n, I use the meeting intervention tones $MSS_t^{F,c}$ and $MSS_t^{F,n}$: the per-member average intervention tone of each group inside the FOMC meeting, computed from the transcripts with the dictionary-based approach of Section 1.1, with the non-Chair series orthogonalized to the Chair's.[26] Two facts support the proxies: the interest-rate decision loads significantly, with the correct sign, on both groups' meeting tones—the regression underlying the influence estimates of Section 4.3—and public speeches foreshadow members' positions in subsequent deliberations at the individual level (Online Appendix D.4).

With these proxies, I estimate the surprise loadings $\beta_j = \nu_j(1-\lambda_j)$ directly, adding the real-time data-news controls of Table 9 and using the fact that the noise terms are orthogonal to the stances:

$$Surprise_t = \underset{(0.091)}{0.127}\ \alpha_t^c + \underset{(0.091)}{0.509}\ \alpha_t^n + \varepsilon_t, \tag{8}$$

where the standard errors in parentheses are Newey–West (18 lags) and the controls are not shown. The monetary policy surprise $Surprise_t$ is the Nakamura and Steinsson (2018) policy-news shock (*NS*) in the Acosta et al. (2024) extension used throughout Section 3.[27] The Chair loading is statistically indistinguishable from zero, consistent with Chair stances being largely priced before the announcement; the non-Chair loading is significantly positive, so non-Chair meeting stances help explain subsequent surprises.

These loadings alone do not establish inattention. Each is the product of influence and inattention, so the near-zero Chair loading is consistent with full attention ($\lambda_c \to 1$) or with negligible Chair influence ($\nu_c \to 0$), and the large non-Chair loading with under-attention ($\lambda_n < 1$) or with large non-Chair influence. The rational reading—non-Chairs matter less, so they are watched less—survives at this point; the next subsection confronts it with the influence measured from the transcripts.

Chair cells of Table 4, and collects the model's remaining maintained assumptions with the direction in which each pushes the recovered quantities.

[26]Assigning the Committee's common tone component to the Chair reflects the institution: the Chair sets the meeting's agenda, speaks to a policy recommendation the staff has prepared under the Chair's direction, and anchors the consensus other members react to (Chappell et al., 2004), so the common component is most plausibly the Chair's. The choice is also the conservative one for everything that follows: it credits the Chair with the largest defensible influence weight.

[27]The sample runs from February 1995, when the series begins, to December 2020, capped by the five-year transcript lag (207 scheduled meetings).

4.3 Separating Influence from Inattention

The decision regression supplies the missing piece. Projecting the realized rate decision on the same meeting-tone series identifies the influence weights ν_c and ν_n directly (Table 9); the regression is augmented, like the surprise regression, with real-time data-news controls so that neither loading absorbs the response to incoming releases; the approach follows Chappell et al. (2004, 2005), who recover members' influence over adopted policy from the historical record of their meeting stances. The point estimate contradicts the natural prior: the orthogonalized non-Chair weight, $\nu_n = 4.26$, exceeds the Chair weight, $\nu_c = 3.23$. The ranking holds in the uncontrolled decision regression and strengthens under every alternative allocation of the Committee's common tone component, because the baseline credits that component entirely to the Chair. An individual-stance estimator in the manner of Chappell et al. (2004), which projects the decision on members' own meeting tones, corroborates the ranking: it places the Chair's share of influence over the rate decision—the object the surprise is defined against—at 31 to 35 percent for the Chair alone, and at 22 to 29 percent for the Chair-and-Vice group under a symmetric partition (Online Appendix Table D.2), below every attention benchmark that follows. What the argument requires is weaker still: that non-Chair influence not be negligible relative to the Chair's.

The surprise and decision regressions carry identical regressors (see the table note), so dividing the surprise loadings of Equation (8)—which Table 9 reports—by the corresponding influence weights isolates the inattention,

$$1-\lambda_c = \frac{\nu_c(1-\lambda_c)}{\nu_c} = \frac{0.127}{3.23} \approx 0.039, \qquad 1-\lambda_n = \frac{\nu_n(1-\lambda_n)}{\nu_n} = \frac{0.509}{4.26} \approx 0.119, \tag{9}$$

and the attention allocation follows in bits from $\kappa_j = -\frac{1}{2}\log_2(1-\lambda_j)$, inverting the mapping $\lambda = 1-2^{-2\kappa}$ of Proposition 1, now evaluated at the realized rather than the optimal allocation.[28] This casts the asymmetry in interpretable units: forecasters extract $\kappa_c \approx 2.33$ bits about the Chair stance but only $\kappa_n \approx 1.53$ bits about the non-Chair stance (Table 9).

[28] Derived quantities here and below are computed from unrounded estimates; ratios of the rounded table entries can differ in the last digit. Information flow is measured in bits (see Sims, 2003). Allocating 1, 2, and 3 bits of information flow to the problem of tracking a Gaussian white noise process yields a ratio of posterior variance to prior variance of 1/4, 1/16, and 1/64, respectively. The bit *levels* inherit the scale of the surprise series; the realized-versus-optimal gap and the influence ratio below do not.

TABLE 9—ATTENTION ALLOCATION: RECOVERED VERSUS OPTIMAL

		Chair & Vice (c)	Non-Chair (n)
Influence ν_j	(decision reg., macro controls)	3.23	4.26
Surprise loading $\nu_j(1-\lambda_j)$	(surprise reg., macro controls)	0.13	0.51
Inattention $1-\lambda_j$		0.039	0.119
Attention $\kappa_j = -\frac{1}{2}\log_2(1-\lambda_j)$ (bits)		2.33	1.53
Optimal attention κ_j^* (bits)		1.76	2.10
Over-attention gap $\kappa_c - \kappa_c^*$ (bits)		0.57	

Note: The table combines a decision regression, which identifies the influence weights ν_c and ν_n, with a monetary-policy-surprise regression, whose coefficients identify $\nu_c(1-\lambda_c)$ and $\nu_n(1-\lambda_n)$. The surprise is the Nakamura and Steinsson (2018) policy-news shock in the Acosta et al. (2024) extension, as in Equation (8); the decision is the change in the federal funds target (the upper bound of the target range after 2008), built from the official daily target series so that no-change meetings enter the sample. The decision regression covers the 224 scheduled meetings of 1988–2020, the surprise regression the 207 scheduled meetings of 1995–2020; both add the real-time data-news controls (the first releases of CPI, industrial production, nonfarm payrolls, real GDP, and unemployment), and differ from the regressions of Table D.1 in both controls and sample. Both use the same meeting-tone series $MSS_t^{F,c}$ and (Chair-orthogonalized) $MSS_t^{F,n}$, so the ratios $1-\lambda_j$ are unit-free; $\kappa_j = -\frac{1}{2}\log_2(1-\lambda_j)$ bits. The optimal allocation is $\kappa_c^* = \frac{1}{2}\kappa + \frac{1}{4}\log_2 x$, with $\kappa = 3.87$, $\kappa_n^* = \kappa - \kappa_c^*$, and $x = (\sigma_c^2/\sigma_n^2)(\nu_c/\nu_n)^2 = 0.62$ (variance ratio 1.09, decision-regression sample). Both regressions use Newey and West (1987) standard errors (18 lags). All entries are point estimates computed from unrounded regression coefficients; no standard errors are reported for the derived rows.

4.4 The Attention Wedge

The model delivers a test of the allocation. Optimal Chair attention rises with the index $x = (\sigma_c^2/\sigma_n^2)(\nu_c/\nu_n)^2$ of Equation (6), and every input is measurable: the variance ratio $\sigma_c^2/\sigma_n^2 \approx 1.09$ from the meeting-tone series (Table 9, notes) and the influence ratio $\nu_c/\nu_n \approx 0.76$ from the decision regression together imply $x \approx 0.62$. The orthogonalization behind σ_n^2 raises x, so the input is conservative for what follows. Evaluating the interior optimum, $\kappa_c^* = \frac{1}{2}\kappa + \frac{1}{4}\log_2(x)$, at the realized capacity $\kappa = \kappa_c + \kappa_n \approx 3.9$ bits gives $\kappa_c^* \approx 1.8$ bits and $\kappa_n^* \approx 2.1$ bits: because non-Chairs carry the larger influence ($\nu_n > \nu_c$) while the Chair's stance is only somewhat more variable, the optimum tilts toward non-Chairs.

The realized allocation tilts the other way: non-Chairs carry the larger decision weight, yet receive the smaller share of attention. The shortfall is 0.57 bits—non-Chairs are given 40 percent of capacity against an optimal 54. In ratio form the same wedge says the split would be optimal only if the average non-Chair were perceived as 0.60 times as influential as the Chair, against the measured 1.32.

Two measures outside the model corroborate the direction. Non-Chair speeches attract only about 17 percent of the *Reuters* coverage of Board communication (computed from the per-speech article counts and speech counts of Table 1), well below both their realized 40 percent attention share and the optimal 54 percent; article counts and bits are not the same units, so the comparison is directional rather than a bound. And Chappell et al. (2004) put the Chair at 40 to 50 percent of the Committee's voting weight over the Burns years, leaving the other eleven voters with at least half; Burns chaired an unusually centralized committee, so that half is a lower bound. Chappell's remainder includes the Vice Chair and the Reserve Bank presidents, so this comparison too is directional—and it still sits far above the attention non-Chair governors are given.

This is the content of *perceived authority*: at the point estimates, forecasters allocate capacity as though non-Chairs carried less weight in the policy process than the decision regression reveals, discounting them in a way that tracks institutional standing rather than measured influence—the mirror of attention drawn to salient, high-authority sources (Bordalo et al., 2020b). A natural alternative is that discounting non-Chairs is rational because the Chair aggregates the Committee's views; the orthogonalization already credits the common component of meeting tone to the Chair, so v_c embeds any aggregator role and non-Chairs remain under-attended even against that generous measure. Both inputs draw on FOMC transcripts released with a five-year lag, so the exercise establishes that the allocation departed from the optimum ex post, not that the departure was exploitable in real time.

5 Conclusion

This article documents an institutional attention bias in the processing of Federal Reserve communication. Financial markets, private forecasters, and the press concentrate on the Chair and Vice Chair, whose speeches move prices on the day they are delivered, yet the discounted communications of the remaining members are the ones that carry unpriced information—predicting interest-rate forecast errors, excess bond returns, and the forward-guidance component of subsequent policy surprises. The asymmetry is a property of the office, not the person: promotion moves attention with institutional status, while the expressed stance is unchanged across the transition. Measured against the influence and tone variance recovered from FOMC transcripts, the attention allocation gives non-Chairs less capacity than a rational-inattention benchmark warrants. Attention follows perceived authority rather than measured influence.

The finding bears directly on how monetary policy shocks are identified. High-frequency announcement surprises are treated as exogenous news because they are assumed unforecastable from prior public information; the evidence recasts that exogeneity as a property of the market's attention allocation rather than of the information environment. A measurable part of the surprise—concentrated in its forward-guidance dimension—is forecastable from speeches delivered months earlier, even after conditioning on the news set of Bauer and Swanson (2023b).

For central banks, the reach of communication depends on who is heard as well as on what is said. The FOMC's decentralized structure supplies a stream of individually informative signals, but the price system absorbs mainly those carrying the right institutional label. Channel choice and the assignment of speaking roles are therefore policy levers, and more communication from under-attended members is not, by itself, more transparency: decentralized communication improves transmission only if markets attend to the information it carries, and the measured allocation implies they do so incompletely.

Part of the announcement surprise is therefore not new information: it is the delayed pricing of communication that had been public for months.

Online Appendix

An Online Appendix collects the data-construction details, robustness tests, and additional results referred to throughout. It is available here.

References

Acosta, Miguel. 2023. "A New Measure of Central Bank Transparency and Implications for the Effectiveness of Monetary Policy." *International Journal of Central Banking* 19 (3): 49–97.

Acosta, Miguel, Connor M. Brennan, and Margaret M. Jacobson. 2024. "Constructing high-frequency monetary policy surprises from SOFR futures." *Finance and Economics Discussion Series* (2024-034): 1–9. 10.17016/feds.2024.034.

Ahmad, Sarah, and David Wessel. 2026. "Grading Fed communications: A 2026 survey of Fed watchers." survey, The Brookings Institution, Hutchins Center on Fiscal and Monetary Policy, https://www.brookings.edu/articles/grading-fed-communications-a-2026-survey-of-fed-watchers/.

Ahrens, Maximilian, Deniz Erdemlioglu, Michael McMahon, Christopher J. Neely, and Xiye Yang. 2025. "Mind your language: Market responses to central bank speeches." *Journal of Econometrics* 249 105921. 10.1016/j.jeconom.2024.105921.

Apel, Mikael, Marianna Blix Grimaldi, and Isaiah Hull. 2022. "How Much Information Do Monetary Policy Committees Disclose? Evidence from the FOMC's Minutes and Transcripts." *Journal of Money, Credit and Banking* 54 (5): 1459–1490. 10.1111/jmcb.12885.

Apel, Mikael, and Marianna Blix Grimaldi. 2014. "How Informative Are Central Bank Minutes?." *Review of Economics* 65 (1): 53–76. 10.1515/roe-2014-0104.

Aruoba, S. Borağan, and Thomas Drechsel. 2024. "Identifying Monetary Policy Shocks: A Natural Language Approach." Working Paper32417, National Bureau of Economic Research. 10.3386/w32417.

Barber, Brad M., and Terrance Odean. 2008. "All That Glitters: The Effect of Attention and News on the Buying Behavior of Individual and Institutional Investors." *The Review of Financial Studies* 21 (2): 785–818.

Bauer, Michael, and Mikhail Chernov. 2024. "Interest Rate Skewness and Biased Beliefs." *Journal of Finance* 79 (1): 173–217. 10.1111/jofi.13276.

Bauer, Michael D., and James D. Hamilton. 2018. "Robust bond risk premia." *Review of Financial Studies* 31 (2): 399–448. 10.1093/rfs/hhx096.

Bauer, Michael D., Carolin E. Pflueger, and Adi Sunderam. 2024. "Perceptions About Monetary Policy." *Quarterly Journal of Economics* 139 (4): 2227–2278. 10.1093/qje/qjae021.

Bauer, Michael D., and Eric T. Swanson. 2023a. "A Reassessment of Monetary Policy Surprises and High-Frequency Identification." *NBER Macroeconomics Annual* 37 (1): 87–155. 10.1086/723574.

Bauer, Michael D., and Eric T. Swanson. 2023b. "An Alternative Explanation for the 'Fed Information Effect'." *American Economic Review* 113 (3): 664–700. 10.1257/aer.20201220.

Bernanke, Ben S., Jean Boivin, and Piotr Eliasz. 2005. "Measuring the effects of monetary policy: A factor-augmented vector autoregressive (FAVAR) approach." *Quarterly Journal of Economics* 120 (1): 387–422. 10.1162/0033553053327452.

Bertsch, Christoph, Isaiah Hull, Robin L. Lumsdaine, and Xin Zhang. 2024. "Four Facts about International Central Bank Communication." technical report, Sveriges Riksbank, Stockholm. 10.2139/ssrn.4773333.

Bianchi, Daniele, Matthias Büchner, and Andrea Tamoni. 2021. "Bond Risk Premiums with Machine Learning." *Review of Financial Studies* 34 (2): 1046–1089. 10.1093/rfs/hhaa062.

Bordalo, Pedro, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer. 2020a. "Overreaction in Macroeconomic Expectations." *American Economic Review* 110 (9): 2748–2782. 10.1257/aer.20181219.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2020b. "Memory, Attention, and Choice." *The Quarterly Journal of Economics* 135 (3): 1399–1442.

Bordo, Michael, and Klodiana Istrefi. 2023. "Perceived FOMC: The making of hawks, doves and swingers." *Journal of Monetary Economics* 136 125–143. 10.1016/j.jmoneco.2023.03.001.

Born, Benjamin, Zeno Enders, and Gernot J. Müller. 2025. "On FIRE, News, and Expectations." *The Routledge Handbook of Economic Expectations in Historical Perspective* 178–194. 10.4324/9781003379010-11.

Bouchaud, Jean Philippe, Philipp Krüger, Augustin Landier, and David Thesmar. 2019. "Sticky Expectations and the Profitability Anomaly." *Journal of Finance* 74 (2): 639–674. 10.1111/jofi.12734.

Campbell, Jeffrey R., Charles L. Evans, Jonas D.M. Fisher, and Alejandro Justiniano. 2012. "Macroeconomic effects of federal reserve forward guidance." *Brookings Papers on Economic Activity* (1): 1–54. 10.1353/eca.2012.0004.

Campbell, John Y., and Robert J. Shiller. 1991. "Yield spreads and interest rate movements: A bird's eye view." *Review of Economic Studies* 58 (3): 495–514. 10.2307/2298008.

Chappell, Henry W., Thomas M. Havrilesky, and Rob Roy McGregor. 1993. "Partisan Monetary Policies: Presidential Influence through the Power of Appointment." *The Quarterly Journal of Economics* 108 (1): 185–218. 10.2307/2118499.

Chappell, Henry W., Rob Roy McGregor, and Todd Vermilyea. 2004. "Majority Rule, Consensus Building, and the Power of the Chairman: Arthur Burns and the FOMC." *Journal of Money, Credit and Banking* 36 (3): 407–422. 10.1353/mcb.2004.0053.

Chappell, Henry W., Rob Roy McGregor, and Todd Vermilyea. 2005. *Committee Decisions on Monetary Policy: Evidence from Historical Records of the Federal Open Market Committee*. Cambridge, MA: MIT Press.

Cieslak, Anna. 2018. “Short-rate expectations and unexpected returns in treasury bonds.” *Review of Financial Studies* 31 (9): 3265–3306. 10.1093/rfs/hhy051.

Cieslak, Anna, Stephen Hansen, Michael McMahon, and Song Xiao. 2023. “Policymakers’ Uncertainty.” technical report, Social Science Research Network, Rochester, NY. 10.2139/ssrn.4631013.

Cieslak, Anna, and Michael McMahon. 2023. “Tough Talk: The Fed and the Risk Premium.” technical report, Social Science Research Network, Rochester, NY. 10.2139/ssrn.4560220.

Cieslak, Anna, Adair Morse, and Annette Vissing-Jorgensen. 2019. “Stock Returns over the FOMC Cycle.” *Journal of Finance* 74 (5): 2201–2248. 10.1111/jofi.12818.

Cieslak, Anna, and Pavol Povala. 2015. “Expected Returns in Treasury Bonds.” *Review of Financial Studies* 28 (10): 2859–2901. 10.1093/rfs/hhv032.

Clark, Todd E., and Kenneth D. West. 2007. “Approximately normal tests for equal predictive accuracy in nested models.” *Journal of Econometrics* 138 (1): 291–311. 10.1016/J.JECONOM.2006.05.023.

Cochrane, John H., and Monika Piazzesi. 2011. “Decomposing the Yield Curve.” *SSRN Electronic Journal*. 10.2139/ssrn.1333274.

Coibion, Olivier, and Yuriy Gorodnichenko. 2015. “Information rigidity and the expectations formation process: A simple framework and new facts.” *American Economic Review* 105 (8): 2644–2678. 10.1257/aer.20110306.

Da, Zhi, Joseph Engelberg, and Pengjie Gao. 2011. “In Search of Attention.” *Journal of Finance* 66 (5): 1461–1499. 10.1111/j.1540-6261.2011.01679.x.

Djourelova, Milena, Filippo Ferroni, Leonardo Melosi, and Alessandro Villa. 2025. “One Fed, Many Voices: Coordinated Communication vs. Transparent Debate.” *SSRN Electronic Journal*. 10.2139/ssrn.5778622.

Duffee, Gregory R. 2011. “Information in (and not in) the term structure.” *Review of Financial Studies* 24 (9): 2895–2934. 10.1093/rfs/hhr033.

Dunn, Wendy, Ellen E. Meade, Nitish Ranjan Sinha, and Raakin Kabir. 2024. “Using Generative AI Models to Understand FOMC Monetary Policy Discussions.” 10.17016/2380-7172.3678.

Ehrmann, Michael, Robin Tietz, and Bauke Visser. 2022. “Voting Right Rotation, Behavior of Committee Members and Financial Market Reactions: Evidence from the U.S. Federal Open Market Committee.” *Journal of International Money and Finance* 126 102765.

Engelberg, Joseph E., and Christopher A. Parsons. 2011. “The Causal Impact of Media in Financial Markets.” *The Journal of Finance* 66 (1): 67–97.

Eusepi, Stefano, and Bruce Preston. 2018. “The science of monetary policy: An imperfect knowledge perspective.” *Journal of Economic Literature* 56 (1): 3–59. 10.1257/jel.20160889.

Fang, Lily, and Joel Peress. 2009. “Media coverage and the cross-section of stock returns.” *Journal of Finance* 64 (5): 2023–2052. 10.1111/j.1540-6261.2009.01493.x.

Gagnon, Joseph, Matthew Raskin, Julie Remache, and Brian Sack. 2011. “The financial market effects of the federal reserve’s large-scale asset purchases.” *International Journal of Central Banking* 7 (1): 3–43.

Gambacorta, Leonardo, Byeungchun Kwon, Taejin Park, Pietro Patelli, and Sonya Zhu. 2024. "CB-LMs: language models for central banking." BIS Working Paper1215, Bank for International Settlements, https://www.bis.org/publ/work1215.htm.

Gómez-Cram, Roberto, and Marco Grotteria. 2022. "Real-time price discovery via verbal communication: Method and application to Fedspeak." *Journal of Financial Economics* 143 (3): 993–1025. 10.1016/j.jfineco.2021.12.004.

Gorodnichenko, Yuriy, Tho Pham, and Oleksandr Talavera. 2025. "Central bank communication on social media: What, to whom, and how?." *Journal of Econometrics* 249 105869. 10.1016/j.jeconom.2024.105869.

Gürkaynak, Refet S., Brian P. Sack, and Eric T. Swanson. 2007a. "Market-based measures of monetary policy expectations." *Journal of Business and Economic Statistics* 25 (2): 201–212. 10.1198/073500106000000387.

Gürkaynak, Refet S., Brian Sack, and Eric Swanson. 2005a. "The sensitivity of long-term interest rates to economic news: Evidence and implications for macroeconomic models." *American Economic Review* 95 (1): 425–436. 10.1257/0002828053828446.

Gürkaynak, Refet S., Brian Sack, and Eric T. Swanson. 2005b. "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, https://www.ijcb.org/journal/ijcb05q2a2.htm.

Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright. 2007b. "The U.S. Treasury yield curve: 1961 to the present." *Journal of Monetary Economics* 54 (8): 2291–2304. 10.1016/j.jmoneco.2007.06.029.

Hamilton, James D. 2009. "Daily changes in fed funds futures prices." *Journal of Money, Credit and Banking* 41 (4): 567–582. 10.1111/j.1538-4616.2009.00223.x.

Handlan, Amy. 2022. "Text Shocks and Monetary Surprises: Text Analysis of FOMC Statements with Machine Learning." Working paper, Brown University.

Hansen, Stephen, and Michael McMahon. 2016. "Shocking language: Understanding the macroeconomic effects of central bank communication." *Journal of International Economics* 99 S114–S133. 10.1016/j.jinteco.2015.12.008.

Hansen, Stephen, Michael McMahon, and Andrea Prat. 2018. "Transparency and deliberation within the FOMC: A computational linguistics approach." *Quarterly Journal of Economics* 133 (2): 801–870. 10.1093/qje/qjx045.

Hodrick, Robert J. 1992. "Dividendyields." *The Review of Financial Studies* 5 (3): 357–386.

Hubert, Paul, and Fabien Labondance. 2021. "The signaling effects of central bank tone." *European Economic Review* 133. 10.1016/j.euroecorev.2021.103684.

Istrefi, Klodiana, Florens Odendahl, and Giulia Sestieri. 2023. "Fed communication on financial stability concerns and monetary policy decisions: Revelations from speeches." *Journal of Banking and Finance* 151. 10.1016/j.jbankfin.2023.106820.

Jarociński, Marek, and Peter Karadi. 2020. "Deconstructing monetary policy surprises-The role of information shocks." *American Economic Journal: Macroeconomics* 12 (2): 1–43. 10.1257/mac.20180090.

Joslin, Scott, Marcel Priebsch, and Kenneth J. Singleton. 2014. "Risk premiums in dynamic term structure models with unspanned macro risks." *Journal of Finance* 69 (3): 1197–1233. 10.1111/jofi.12131.

Kim, Alex G., Maximilian Muhn, and Valeri V. Nikolaev. 2023. "From Transcripts to Insights: Uncovering Corporate Risks Using Generative AI." technical report, Social Science Research Network, Rochester, NY. 10.2139/ssrn.4593660.

Krueger, Joel T., and Kenneth N. Kuttner. 1996. "The fed funds futures rate as a predictor of federal reserve policy." *Journal of Futures Markets* 16 (8): 865–879. 10.1002/(SICI)1096-9934(199612)16:8<865::AID-FUT2>3.0.CO;2-K.

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.

Litterman, Robert, and José Scheinkman. 1991. "Common factors affecting bond returns." *The Journal of Fixed Income* 1 (1): 54–61.

Loughran, Tim, and Bill McDonald. 2011. "When Is a Liability Not a Liability? Textual Analysis, Dictionaries, and 10-Ks." *Journal of Finance* 66 (1): 35–65. 10.1111/j.1540-6261.2010.01625.x.

Lucca, David O., and Emanuel Moench. 2015. "The Pre-FOMC announcement drift." *Journal of Finance* 70 (1): 329–371. 10.1111/jofi.12196.

Lucca, David O., and Francesco Trebbi. 2011. "Measuring Central Bank Communication: An Automated Approach with Application to FOMC Statements." technical report, Social Science Research Network, Rochester, NY. 10.2139/ssrn.1470443.

Ludvigson, Sydney C., and Serena Ng. 2009. "Macro factors in bond risk premia." *Review of Financial Studies* 22 (12): 5027–5067. 10.1093/rfs/hhp081.

Maćkowiak, Bartosz, and Mirko Wiederholt. 2009. "Optimal sticky prices under rational inattention." *American Economic Review* 99 (3): 769–803. 10.1257/aer.99.3.769.

Miranda-Agrippino, Silvia. 2023. "Unsurprising Shocks: Information, Premia, and the Monetary Transmission." Staff Working Paper626, Bank of England, https://www.bankofengland.co.uk/working-paper/2016/unsurprising-shocks-information-premia-and-the-monetary-transmission.

Miranda-Agrippino, Silvia, and Giovanni Ricco. 2021. "The Transmission of Monetary Policy Shocks." *American Economic Journal: Macroeconomics* 13 (3): 74–107. 10.1257/mac.20180124.

Mullainathan, Sendhil, and Andrei Shleifer. 2005. "The Market for News." *American Economic Review* 95 (4): 1031–1053.

Nakamura, Emi, and Jón Steinsson. 2018. "High-frequency identification of monetary non-neutrality: The information effect." *Quarterly Journal of Economics* 133 (3): 1283–1330. 10.1093/QJE/QJY004.

Neuhierl, Andreas, and Michael Weber. 2019. "Monetary policy communication, policy slope, and the stock market." *Journal of Monetary Economics* 108 140–155. 10.1016/j.jmoneco.2019.08.005.

Newey, Whitney K., and Kenneth D. West. 1987. "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica* 55 (3): 703–708.

Olson, Peter, and David Wessel. 2016. "Federal Reserve communications: Survey results." survey, The Brookings Institution, Hutchins Center on Fiscal and Monetary Policy, https://www.brookings.edu/wp-content/uploads/2016/11/fed-communications-survey-results.pdf.

Piazzesi, Monika, and Eric T. Swanson. 2008. "Futures prices as risk-adjusted forecasts of monetary policy." *Journal of Monetary Economics* 55 (4): 677–691. 10.1016/j.jmoneco.2008.04.003.

Primiceri, Giorgio E. 2005. "Time varying structural vector autoregressions and monetary policy." *Review of Economic Studies* 72 (3): 821–852. 10.1111/j.1467-937X.2005.00353.x.

Sastry, Karthik A. 2026. "Disagreement About Monetary Policy." *American Economic Journal: Macroeconomics* 18 (2): 486–522. 10.1257/mac.20220266.

Schmeling, Maik, Andreas Schrimpf, and Sigurd A.M. Steffensen. 2022. "Monetary policy expectation errors." *Journal of Financial Economics* 146 (3): 841–858. 10.1016/j.jfineco.2022.09.005.

Sims, Christopher A. 2003. "Implications of rational inattention." *Journal of Monetary Economics* 50 (3): 665–690. 10.1016/S0304-3932(03)00029-1.

Swanson, Eric T. 2023. "The Importance of Fed Chair Speeches as a Monetary Policy Tool." *AEA Papers and Proceedings* 113 394–400. 10.1257/pandp.20231073.

Swanson, Eric T, and Vishuddhi Jayawickrema. 2023. "Speeches by the Fed Chair Are More Important Than FOMC Announcements: An Improved High-Frequency Measure of U.S. Monetary Policy Shocks." working paper, University of California, Irvine.

Tetlock, Paul C. 2007. "Giving content to investor sentiment: The role of media in the stock market." *Journal of Finance* 62 (3): 1139–1168. 10.1111/j.1540-6261.2007.01232.x.

Tietz, Robin. 2018. "How Does the Fed Manage Interest Rate Expectations?." technical report, Social Science Research Network, Rochester, NY. 10.2139/ssrn.3254094.

Wei, Min, and Jonathan H. Wright. 2013. "Reverse regressions and long-horizon forecasting." *Journal of Applied Econometrics* 28 (3): 353–371. 10.1002/jae.1274.

Wessel, David, and Sam Boocker. 2024. "Grading Fed communications: A 2024 survey of Fed watchers." survey, The Brookings Institution, Hutchins Center on Fiscal and Monetary Policy, https://www.brookings.edu/articles/grading-fed-communications-a-2024-survey-of-fed-watchers/.

Wright, Jonathan H. 2012. "What does monetary policy do to long-term interest rates at the zero lower bound?." *Economic Journal* 122 (564): F447–F466. 10.1111/j.1468-0297.2012.02556.x.

Yellen, Janet L. 2013. "Communication in Monetary Policy : a speech at the Society of American Business Editors and Writers 50th Anniversary Conference, Washington, D.C., April 4, 2013." apr, https://ideas.repec.org/p/fip/fedgsq/625.html.

Yellen, Janet L. 2017. "Panel: At the heart of policy: challenges and opportunities of central bank communication." https://www.youtube.com/watch?v=DI7p-g51O8g.