Supplemental Appendix

Authority, Attention, and the Origins of Monetary Policy Surprises

David Gauthier*

A Data

This appendix records the sources of every series used in the article (Appendix A.1), the dictionary behind the Monetary Stance Signal (Appendix A.2), and the construction details for the expectations, return, and macroeconomic variables (Appendix A.3).

A.1 Data Sources

Tables A.1, A.2, and A.3 list the source, sample, and series identifiers for the yields, the interest-rate expectations, and the macroeconomic variables.

*Banque de France, Email: david.gauthier@banque-france.fr.

TABLE A.1—SOURCES FOR YIELDS

Variable	First obs	Last obs	Nobs	Mnemonics	Source
Y01	1961-06-30	2024-12-31	763	SVENY01	GSW (2007)
Y02	1961-06-30	2024-12-31	763	SVENY02	GSW (2007)
Y03	1961-06-30	2024-12-31	763	SVENY03	GSW (2007)
Y04	1961-06-30	2024-12-31	763	SVENY04	GSW (2007)
Y05	1961-06-30	2024-12-31	763	SVENY05	GSW (2007)
Y06	1961-06-30	2024-12-31	763	SVENY06	GSW (2007)
Y07	1961-06-30	2024-12-31	763	SVENY07	GSW (2007)
Y08	1971-08-31	2024-12-31	641	SVENY08	GSW (2007)
Y09	1971-08-31	2024-12-31	641	SVENY09	GSW (2007)
Y10	1971-08-31	2024-12-31	641	SVENY10	GSW (2007)
Y11	1971-11-30	2024-12-31	638	SVENY11	GSW (2007)
Y12	1971-11-30	2024-12-31	638	SVENY12	GSW (2007)
Y13	1971-11-30	2024-12-31	638	SVENY13	GSW (2007)
Y14	1971-11-30	2024-12-31	638	SVENY14	GSW (2007)
Y15	1971-11-30	2024-12-31	638	SVENY15	GSW (2007)
Y16	1981-07-31	2024-12-31	522	SVENY16	GSW (2007)
Y17	1981-07-31	2024-12-31	522	SVENY17	GSW (2007)
Y18	1981-07-31	2024-12-31	522	SVENY18	GSW (2007)
Y19	1981-07-31	2024-12-31	522	SVENY19	GSW (2007)
Y20	1981-07-31	2024-12-31	522	SVENY20	GSW (2007)
Y21	1985-11-30	2024-12-31	470	SVENY21	GSW (2007)
Y22	1985-11-30	2024-12-31	470	SVENY22	GSW (2007)
Y23	1985-11-30	2024-12-31	470	SVENY23	GSW (2007)
Y24	1985-11-30	2024-12-31	470	SVENY24	GSW (2007)
Y25	1985-11-30	2024-12-31	470	SVENY25	GSW (2007)
Y26	1985-11-30	2024-12-31	470	SVENY26	GSW (2007)
Y27	1985-11-30	2024-12-31	470	SVENY27	GSW (2007)
Y28	1985-11-30	2024-12-31	470	SVENY28	GSW (2007)
Y29	1985-11-30	2024-12-31	470	SVENY29	GSW (2007)
Y30	1985-11-30	2024-12-31	470	SVENY30	GSW (2007)

Note: This table contains information on the bond yields used to compute the excess bond returns. Maturities range from 1 to 30 years. Each row reports the availability of that maturity's own series: the 1–7-year maturities used in the long-sample regressions begin in June 1961, while the longest maturities enter from late 1985, so the balanced 1–30-year panel is complete only from November 1985 onward (see Appendix A.3). All yields are sourced from Gürkaynak et al. (2007).

TABLE A.2—SOURCES FOR INTEREST-RATE EXPECTATIONS

Variable	First obs	Last obs	Nobs	Mnemonics	Horizons	Source
30 Day FFR Futures	1988-12-31	2024-12-31	433	ff2 Comdty	1	Bloomberg
30 Day FFR Futures	1988-12-31	2024-12-31	433	ff3 Comdty	2	Bloomberg
30 Day FFR Futures	1988-12-31	2024-12-31	433	ff4 Comdty	3	Bloomberg
30 Day FFR Futures	1988-12-31	2024-12-31	432	ff5 Comdty	4	Bloomberg
30 Day FFR Futures	1988-12-31	2024-12-31	429	ff6 Comdty	5	Bloomberg
30 Day FFR Futures	1988-12-31	2024-12-31	413	ff7 Comdty	6	Bloomberg
30 Day FFR Futures	1989-02-28	2024-12-31	392	ff8 Comdty	7	Bloomberg
30 Day FFR Futures	1989-01-31	2024-12-31	368	ff9 Comdty	8	Bloomberg
30 Day FFR Futures	2004-01-31	2024-12-31	252	ff10 Comdty	9	Bloomberg
30 Day FFR Futures	2004-01-31	2024-12-31	252	ff11 Comdty	10	Bloomberg
30 Day FFR Futures	2004-01-31	2024-12-31	252	ff12 Comdty	11	Bloomberg
30 Day FFR Futures	2004-01-31	2024-12-31	252	ff13 Comdty	12	Bloomberg
FFR Forecast	1991-03-01	2024-12-01	136		1	BCFF
FFR Forecast	1991-03-01	2024-12-01	136		2	BCFF
FFR Forecast	1991-03-01	2024-12-01	136		3	BCFF
FFR Forecast	1991-03-01	2024-12-01	136		4	BCFF
3-months T-bill	1981-08-16	2024-11-11	174		1	SPF
3-months T-bill	1981-08-16	2024-11-11	174		2	SPF
3-months T-bill	1981-08-16	2024-11-11	174		3	SPF
3-months T-bill	1981-08-16	2024-11-11	174		4	SPF

Note: This table presents the different series used to compute interest rate expectations. Futures come from Bloomberg, and the FFR Forecast comes from the Blue Chip Financial Forecasts. The 3-month T-bill comes from the Survey of Professional Forecasters. For the survey series, *First obs* and *Last obs* refer to survey collection dates, not to the dates being forecast; the estimation samples are stated in each regression table's note.

TABLE A.3—SOURCES FOR MACRO VARIABLES

Variable	Source	Mnemonics	First obs	Last obs	Nobs
Cfnai	Fred	CFNAI	1967-04-30	2024-12-31	693
Inflation	Fred	CPIAUCSL	1961-06-30	2024-12-31	763
Emp	Fred	PAYEMS	1961-06-30	2024-12-31	763
Fedfunds	Fred	FEDFUNDS	1961-06-30	2024-12-31	763

Note: This table presents the reference macroeconomic series used in the empirical analysis, with their FRED mnemonics. The forecasting regressions use the real-time (first-release) counterparts of these series from the Federal Reserve Bank of Philadelphia real-time database and the ALFRED vintage archive, as described in Section 1.3 and Appendix A.3; the revised FRED series listed here enter only where noted (revised data produce nearly identical results).

A.2 MSS Dictionary

TABLE A.4—MSS DICTIONARY

TERMS	HAWKISH MODIFIER	DOVISH MODIFIER
spending	accelerat, boom, boost, brisk, buoyant, climb, expan, heat, high, increas, overheat, pick, pickup, rise, rising, robust, rose, soar, solid, strength, strong, surg, upside, upswing, uptick, upturn	contracted, contracting, contraction, cool, curtail, decelerat, declin, decreas, deteriorat, dip, downside, downturn, drop, ease, easing, fall, fell, fragil, low, reced, retrench, shrank, shrink, slack, slow, slugg, soft, subdu, tepid, weak
economic growth	accelerat, boom, boost, brisk, buoyant, climb, expan, heat, high, increas, overheat, pick, pickup, rise, rising, robust, rose, soar, solid, strength, strong, surg, upside, upswing, uptick, upturn	contracted, contracting, contraction, cool, curtail, decelerat, declin, decreas, deteriorat, dip, downside, downturn, drop, ease, easing, fall, fell, fragil, low, reced, retrench, shrank, shrink, slack, slow, slugg, soft, subdu, tepid, weak
demand	accelerat, boom, boost, brisk, buoyant, climb, expan, heat, high, increas, overheat, pick, pickup, rise, rising, robust, rose, soar, solid, strength, strong, surg, upside, upswing, uptick, upturn	contracted, contracting, contraction, cool, curtail, decelerat, declin, decreas, deteriorat, dip, downside, downturn, drop, ease, easing, fall, fell, fragil, low, reced, retrench, shrank, shrink, slack, slow, slugg, soft, subdu, tepid, weak
employment	accelerat, climb, expand, gain, heat, high, improv, increas, overheat, pick, pickup, rais, rise, rising, robust, rose, soar, solid, strength, strong, surg, tight	contracted, contracting, contraction, cool, declin, deteriorat, dip, drop, ease, easing, fall, fell, fragil, reduc, shrank, shrink, slack, slow, slugg, soft, subdu, weak
productivity	declin, dip, drop, fall, fell, low, reduc, slow, slugg, soft, weak	elevat, rose
labor market	elevat, firmer, firming, heat, hot, overheat, pressur, robust, solid, strain, strong, surg, tight	cool, deteriorat, ease, eased, easing, fragil, loos, loose, slack, slow, soft, subdu, weak
activity	accelerat, climb, expand, gain, heat, high, improv, increas, overheat, pick, pickup, rais, rise, rising, robust, rose, soar, solid, strength, strong, surg, tight	contracted, contracting, contraction, cool, declin, deteriorat, dip, drop, ease, easing, fall, fell, fragil, reduc, shrank, shrink, slack, slow, slugg, soft, subdu, weak
unemployment	declin, dip, drop, ease, easing, fall, fell, improv, low, reced, reduc, tumbl	elevat, rose
job	accelerat, climb, expand, gain, heat, high, improv, increas, overheat, pick, pickup, rais, rise, rising, robust, rose, soar, solid, strength, strong, surg, tight	contracted, contracting, contraction, cool, declin, deteriorat, dip, drop, ease, easing, fall, fell, fragil, reduc, shrank, shrink, slack, slow, slugg, soft, subdu, weak
price	accelerat, boost, broaden, climb, elevat, escalat, firmer, firming, heat, high, increas, jump, mount, overheat, persist, pick, pickup, pressur, rise, rising, rose, soar, spik, sticky, strong, surg, unanchor, upside, uptick, upward	abate, anchor, contain, cool, decelerat, declin, decreas, dip, disinflat, down, drop, ebb, ease, easing, fall, fell, low, moderat, muted, reced, reduc, slow, soft, stable, subdu, tame, tepid, transitory, weak, weaken
inflation	accelerat, boost, broaden, climb, elevat, escalat, firmer, firming, heat, high, increas, jump, mount, overheat, persist, pick, pickup, pressur, rise, rising, rose, soar, spik, sticky, strong, surg, unanchor, upside, uptick, upward	abate, anchor, contain, cool, decelerat, declin, decreas, dip, disinflat, down, drop, ebb, ease, easing, fall, fell, low, moderat, muted, reced, reduc, slow, soft, stable, subdu, tame, tepid, transitory, weak, weaken
wage	accelerat, boost, broaden, climb, elevat, escalat, firmer, firming, heat, high, increas, jump, mount, overheat, persist, pick, pickup, pressur, rise, rising, rose, soar, spik, sticky, strong, surg, unanchor, upside, uptick, upward	abate, anchor, contain, cool, decelerat, declin, decreas, dip, disinflat, down, drop, ebb, ease, easing, fall, fell, low, moderat, muted, reced, reduc, slow, soft, stable, subdu, tame, tepid, transitory, weak, weaken
production	accelerat, climb, expand, gain, heat, high, improv, increas, overheat, pick, pickup, rais, rise, rising, robust, rose, soar, solid, strength, strong, surg, tight	contracted, contracting, contraction, cool, declin, deteriorat, dip, drop, ease, easing, fall, fell, fragil, reduc, shrank, shrink, slack, slow, slugg, soft, subdu, weak

Note: This table shows the keywords and modifiers used to construct the MSS: the thirteen effective topic terms with their hawkish and dovish modifier lists. Plural forms are matched by the scoring algorithm but not shown. For *productivity* and *unemployment*, the modifier columns reflect the inverted directional mapping described below: weak productivity growth and falling unemployment are hawkish.

Negation terms (35, common to all keywords, each reversing the classification of its subsentence; two or more leave it unoriented): *ain't*, *aren't*, *can't*, *cannot*, *couldn't*, *daren't*, *despite*, *didn't*, *doesn't*, *don't*, *hadn't*, *hasn't*, *haven't*, *isn't*, *mightn't*, *mustn't*, *needn't*, *never*, *none*, *nope*, *nor*, *not*, *nothing*, *nowhere*, *oughtn't*, *rarely*, *seldom*, *shan't*, *shouldn't*, *wasn't*, *weren't*, *without*, *won't*, *wont*, *wouldn't*.

Dictionary Composition. Of the eleven Apel et al. (2022) topics, eight are retained together with about four-fifths of their modifiers; of the omitted topics, *pressure* survives as a directional modifier, *labour* is the British-spelling duplicate of *labor*, and *resource utilization* is switched off in the baseline configuration—as are further candidate topics carried by the dictionary file (payroll, income, sentiment, and hours). To these I add the topics through which U.S. policymakers discuss the same subjects: wages, productivity, demand, jobs, and production. Productivity enters with inverted directional signs, since weak productivity growth raises unit labor costs and is therefore hawkish. Unemployment enters the same way, since falling unemployment is hawkish. I also add the FOMC's own inflation vocabulary (*anchored*, *unanchored*, *disinflation*, *sticky*, *transitory*, *overheating*) and common cyclical qualifiers (*robust*, *solid*, *slack*, *fragile*). Keyword occurrences are split roughly one-half on prices, one-third on activity, and the remainder on employment (Figure A.1, Panel B).

FIGURE A.1—THE CORPUS: COVERAGE AND TOPIC COMPOSITION BY INSTITUTIONAL GROUP

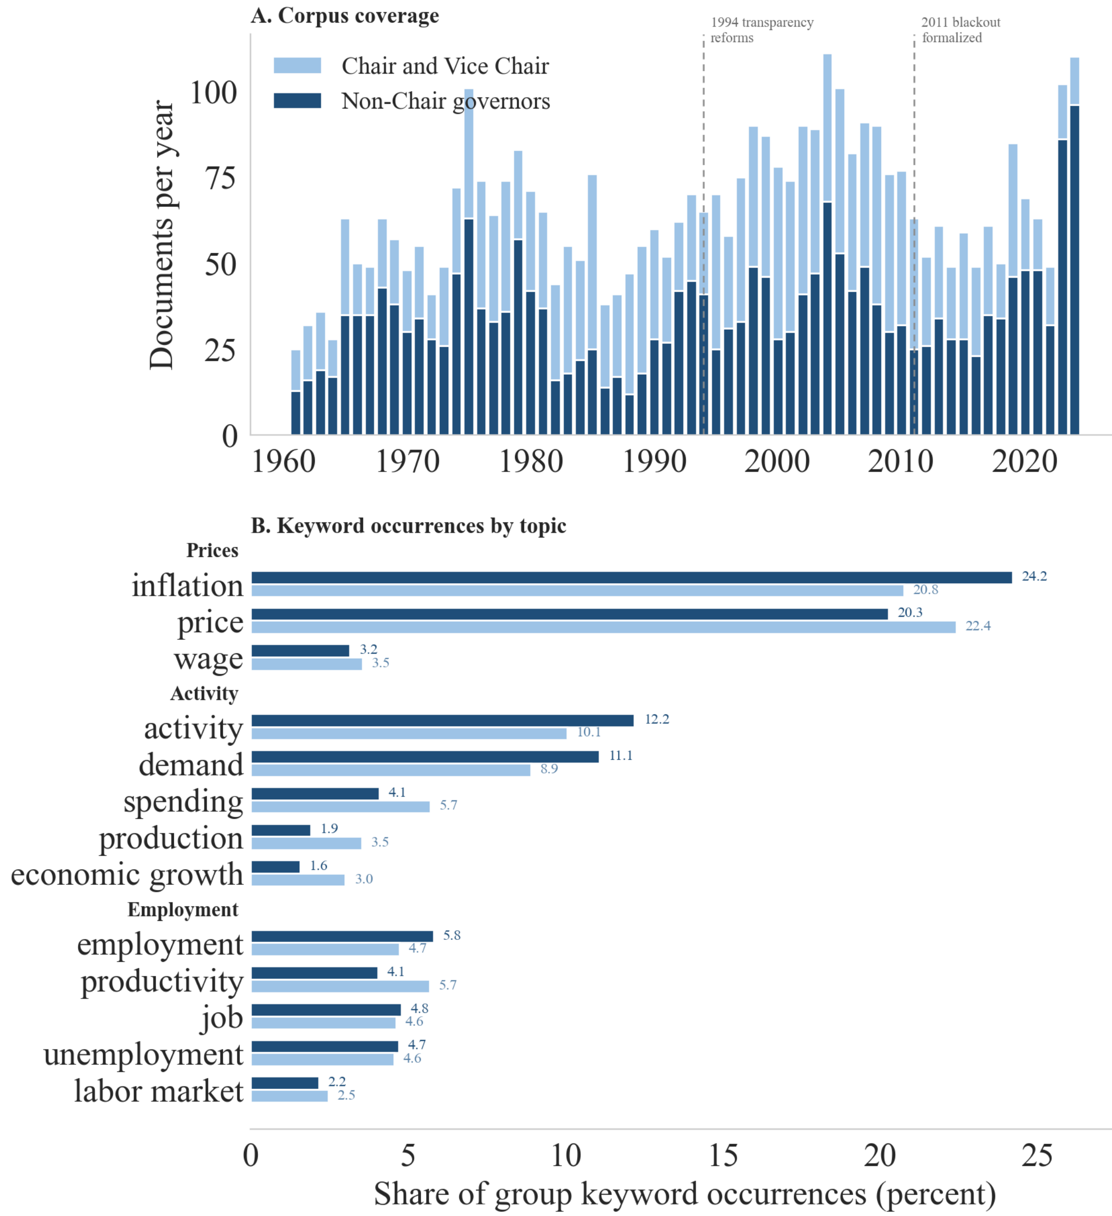

Note: Panel A plots the number of documents (speeches and congressional testimonies) per year in the corpus, by the speaker's institutional position at the time of delivery (4,177 documents, 1961–2024); FOMC meeting transcripts are not shown. The two dashed verticals mark the 1994 transparency reforms and the 2011 formalization of the pre-meeting blackout. Panel B plots the share of each group's dictionary-keyword occurrences accounted for by each topic term, over the same corpus; occurrences are keyword matches identified by the scoring algorithm, with plural forms folded into the singular, and the thirteen terms shown are the effective dictionary of Appendix Table A.4, grouped by category.

A.3 Data Construction Details

Interest-Rate Expectations. For federal funds futures, the two-month contract begins in December 1988 and one-year-ahead contracts enter the sample in January 2004; prices are sampled on the final trading day of each month, and the MSS is evaluated three days before month-end. The *Blue Chip Financial Forecasts* survey covers approximately forty-five major financial institutions each month; to obtain non-overlapping observations I use forecasts reported in the middle month of each quarter. The survey is typically conducted between the 23rd and 27th of each month and published on the first day of the following month (January surveys are generally conducted during the second half of December); each round is dated at its first-of-month publication date, and the MSS is evaluated three days before the opening of the response window, on the 20th of the response month. No speech after the evaluation date enters the signal, so every speech in the signal is public at least three days before the documented response window opens. Evaluating the MSS between the 18th and the 22nd of the response month yields similar results. SPF forecasts are quarterly forecasts of the three-month Treasury-bill rate, with the MSS evaluated three days before the collection dates published by the Federal Reserve Bank of Philadelphia.

Bond Excess Returns and Yield Factors. Annual holding-period excess returns are $rx_t^{(n)} = ni_t^{(n)} - (n-1)i_{t+12}^{(n-1)} - i_t^{(1)}$, where $i_t^{(n)}$ is the continuously compounded yield on an n-year zero-coupon bond, t indexes months, and n is measured in years, so $i_{t+12}^{(n-1)}$ is the $(n-1)$-year yield one year after purchase. Because data availability varies across maturities, the principal components of yields are constructed on both a 1–7-year panel (available from June 1961) and a 1–30-year panel (complete from November 1985); the two specifications provide robustness to differences in maturity coverage.

Real-Time Macroeconomic Variables. Employment growth and inflation are computed from the most recent first release available at each observation date—the release for month $t-1$ in the month-end-dated designs (futures, SPF, bond returns), and the release for month $t-2$ for BCFF survey rounds, which are dated at their first-of-month publication date—together with the revised value twelve months earlier. The CFNAI enters in level as its most recent first release, taken from the Chicago Fed real-time file, which begins with the December 2000 observation; earlier observations use the revised series. Apart from that fallback and the twelve-month revised base, the real-time controls use only information available at each date; the trend-inflation predictor of Cieslak and Povala (2015) and the unemployment rate are computed from the revised series. In the ALFRED/FRED release panel, each series retains the value initially published, together with its actual publication date; each release is carried forward until superseded and matched to the appropriate weekly or meeting-level observation by publication date.

Decision Regression: News Controls and the Decision Variable. The data-news panel entering the decision and surprise regressions of main-text Section 4.3 comprises the real-time ALFRED/FRED first releases of nonfarm payrolls, unemployment, headline and core CPI, headline and core PCE, real GDP, and industrial production, matched by publication date; it absorbs macroeconomic news arriving between meetings so the recovered attention reflects forecasters' reading of policymaker stances rather than their response to data releases. The decision variable is the change in the federal funds target (the upper bound of the target range after 2008) at each scheduled meeting, constructed from the official daily target series and matched to meetings through the announcement calendar, so no-change meetings enter the estimation with a decision of zero; intermeeting moves are not attributed to scheduled meetings.

B Additional Empirical Results

B.1 Demand-Side Attention: Search Interest Around Speeches

Figure B.1 reports the demand-side attention measure discussed in Section 2.1: Google Trends search interest in each policymaker's name around Board-member speeches, by the speaker's status, over 2006–2024 (the modern era in which search data are informative). Search interest jumps in the week of a Chair or Vice Chair speech and returns to baseline; around a non-Chair speech, it barely moves. Pooled across four sample eras, each anchored on the sitting Chair, the speech-week search spike averages $+2.6$ index points for principals against $+0.3$ for non-Chairs ($t = 4.4$), and is positive in every era; the principal–non-Chair gap sharpens once FOMC-announcement weeks are excluded (Chair appearances sometimes fall in FOMC weeks), leaving the non-Chair response indistinguishable from zero ($+2.4$ versus 0.0, $t = 5.1$). The pattern belongs to the office rather than to individual prominence: within person, searches for "Jerome Powell" were negligible and unmoved by his speeches while he was a governor, and jumped sharply around his speeches once he became Chair.[1] Demand-side attention, like press coverage, follows authority.

[1] Search interest is weekly Google Trends for the speaker's name (United States), normalized within each query so speakers share a Chair-anchored scale; the speech-week spike is the search index in the speech's week minus a local baseline (the mean of the weeks two and three before and after). The series is available only post-2004, so this exhibit complements rather than extends the three-decade press-coverage panel.

FIGURE B.1—SEARCH INTEREST AROUND SPEECHES, BY SPEAKER STATUS

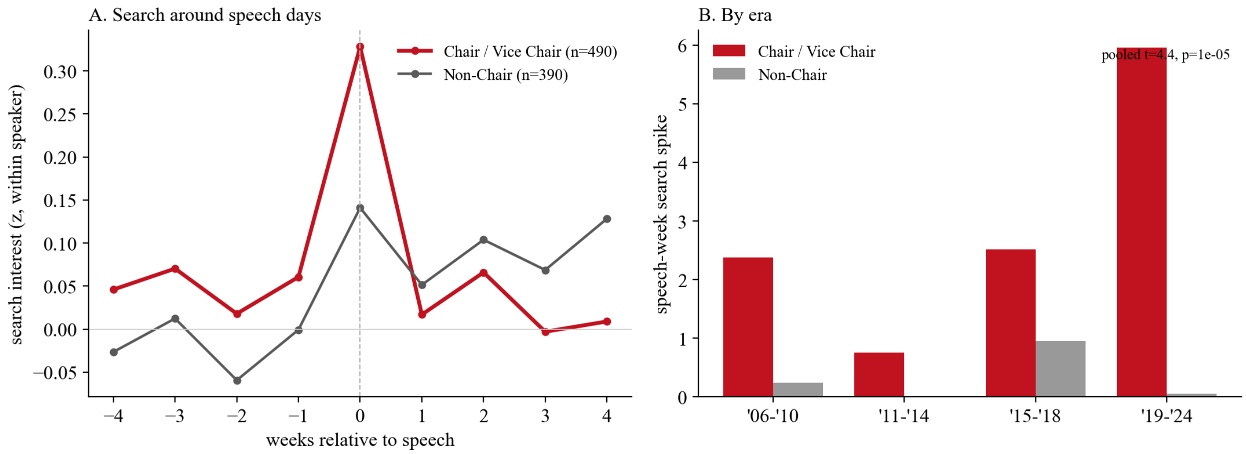

Note: The figure plots Google Trends weekly search interest in policymakers' names (United States), 2006–2024, around Board-member speech dates, classified by the speaker's status on the speech date (principal = Chair or Vice Chair). Panel A: average search interest in event time (weeks relative to the speech), standardized within speaker; the principal series peaks in the speech week ($t = 0$) while the non-Chair series does not. Panel B: the average speech-week search spike (defined in the text) for principals and non-Chairs, by era, with the sitting Chair anchoring each era's Trends scale; the pooled principal–non-Chair difference is significant by a two-sample Welch test ($t = 4.4$, $p \approx 10^{-5}$), positive in every era, and survives excluding FOMC-announcement weeks.

B.2 Same-Day Repricing on Coverage Tone

Table B.1 reports the pricing side of the attention measure discussed in Section 2.1. Each trading day's change in the front of the short-rate curve—the federal funds futures two, four, and six months ahead and the one- and two-year Treasury yields—is regressed on the day's summed *Reuters* coverage tone for each authority group, with calendar-year fixed effects and Newey–West standard errors. The estimation window is the federal funds futures window, August 1991 to December 2024, restricted to trading days on which every rate in the ladder is observed, and the two groups share the same trading days, so the columns differ only in whose coverage carries the tone. Prices move with the tone the press attributes to principals at every rung of the maturity ladder, between 0.05 and 0.24 basis points per standard deviation of coverage tone—at the 5 percent level or better at every rung but the two-year yield, which is significant at the 10 percent level—and never with the tone it attributes to non-Chairs, whose coefficients are never statistically distinguishable from zero. The regression is contemporaneous, so it identifies the attention prices pay to principal coverage rather than a causal response to the articles; the supply side of Table 1 and the pricing side here allocate attention the same way.

TABLE B.1—SAME-DAY REPRICING OF THE SHORT-RATE CURVE ON COVERAGE TONE, BY AUTHORITY

	Principal (Chair, Vice Chair)	Non-Chair (other governors)
ff_2	0.047**	0.019
	(0.019)	(0.019)
ff_4	0.099***	0.019
	(0.033)	(0.031)
ff_6	0.237**	0.061
	(0.092)	(0.051)
1Y	0.135**	0.044
	(0.054)	(0.049)
2Y	0.138*	0.043
	(0.071)	(0.070)
N (trading days)	7951	

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 trading day and regresses the same-day change in each front-of-curve rate—the fed-funds futures ff_h (h months ahead; roll-adjusted implied rate) and the 1-year and 2-year Treasury constant-maturity yields—on the day's summed coverage tone within the group, with calendar-year fixed effects and Newey and West (1987) standard errors (2 daily lags) in parentheses. Its coefficients are basis points of the same-day rate change per one standard deviation of summed tone, estimated on all trading days within the fed-funds futures window (August 1991 through December 2024) on which every rate in the ladder is observed, with a day carrying zero summed tone when no article matched that group; both columns share the same sample, so N is the number of trading days. The association is contemporaneous, so it is read as the attention prices pay to principal coverage rather than a causal response to it. *, **, and *** indicate significance at the 10, 5, and 1 percent levels.

B.3 Bond Excess Returns Across Promotion

Table B.2 is the return-side counterpart of Panel C of main-text Table 6, on the wider sample: it runs the excess-bond-return regression of Section 2.2 on all thirteen members whose careers run from non-Chair to Chair or Vice Chair, without the restriction that the news file cover them on both sides of the promotion, and enters the same individuals' pre-promotion and post-promotion tone jointly. Before promotion, their tone predicts returns at every maturity—the two-year coefficient is -0.66—with the strength declining from the three-year maturity outward, as in the full cross-section. After promotion, the coefficients are indistinguishable from zero, and a Wald test rejects equality of the pre- and post-promotion coefficients at every maturity.

TABLE B.2—THE PREDICTIVE CONTENT OF A SPEAKER'S TONE BEFORE AND AFTER PROMOTION

	$n = 2Y$	$n = 3Y$	$n = 5Y$	$n = 7Y$	$n = 10Y$
MSS_t^n (before promotion, governor)	-0.663***	-0.679***	-0.609***	-0.508***	-0.381**
	(0.148)	(0.167)	(0.178)	(0.179)	(0.180)
MSS_t^c (after promotion, Chair/Vice)	0.029	0.044	0.068	0.082	0.090
	(0.096)	(0.088)	(0.086)	(0.090)	(0.095)
Wald: equal coefficients (p)	0.000	0.000	0.000	0.001	0.015
N (months)	468	468	468	468	468

Note: The dependent variable is the n-year bond excess return, standardized by maturity, regressed jointly on the switchers' pre- and post-promotion tone with the first three yield principal components (not shown); Newey and West (1987) standard errors (18 monthly lags) in parentheses. Tone is standardized by the full-sample non-Chair standard deviation, so coefficients read as standard deviations of the return per standard deviation of tone. The sample is the 13 members whose career runs non-Chair $\rightarrow$ Chair/Vice (435 pre-promotion and 828 post-promotion speeches; the counts cover each member's full archive, so they exceed those of the reliability table, which imposes the paper's 1961 corpus start), over the 1985–2024 months in which a switcher signal is active, and both components enter jointly with a Wald test of equal coefficients. Seven of these members are covered by the post-1992 news file on both sides of promotion and form the sample of main-text Table 6. *, **, and *** indicate significance at the 10, 5, and 1 percent levels, respectively.

B.4 Attention or Aggregation?

Table B.3 reports the Monte Carlo restrictions of the speech set discussed in Section 2.4: *single-voice* draws keep, within each period, only the speeches of one randomly drawn speaker, while the *volume-matched placebo* keeps the same number of speeches drawn at random across speakers. The speaker-noise share measures the fraction of the signal attributable to cross-speaker averaging.

TABLE B.3—ATTENTION VERSUS AGGREGATION: SINGLE-VOICE SIGNALS AGAINST A VOLUME-MATCHED PLACEBO

	Single voice	Volume-matched placebo
A. Non-Chair governors → BCFF 4Q forecast errors (150 draws)		
Standardized beta (median [p10, p90])	-0.230 [-0.360, -0.091]	-0.218 [-0.352, -0.091]
t-statistic (median)	-1.96	-1.86
Share of draws negative	0.97	0.98
Share of draws $t < -1.96$	0.50	0.46
Speaker-noise share of signal	0.00 (Mann–Whitney $p = 0.33$)	
B. Non-Chair governors → 2Y excess bond returns (100 draws)		
Standardized beta (median [p10, p90])	-0.139 [-0.240, -0.051]	-0.130 [-0.233, -0.022]
t-statistic (median)	-1.51	-1.18
Share of draws negative	0.95	0.95
Share of draws $t < -1.96$	0.27	0.20
Speaker-noise share of signal	0.00 (Mann–Whitney $p = 0.51$)	

Note: Each Monte-Carlo draw restricts the speech set before it enters the paper's signal construction, which is otherwise unchanged. *Single voice*: within each period, only the speeches of one randomly drawn speaker are kept. *Volume-matched placebo*: the same number of speeches per period is kept, drawn at random across speakers. Under the aggregation alternative—the signal predicts because it averages several quasi-independent voices—the single-voice draws must be attenuated relative to the placebo; the speaker-noise share reports one minus the ratio of their median standardized betas, floored at zero (a negative value—single-voice draws no weaker than the placebo—is evidence against aggregation, not a negative noise share), with a Mann–Whitney test of distributional equality. The periods are quarters. Both panels rebuild the non-Chair MSS through Equations (1)–(2) (60-day half-life, zero-fill, sampled three days before the forecast date) and re-estimate the specifications of Tables 3 (4Q horizon, Chair tone included) and 5 ($n = 2Y$, PC_1–PC_3 controls). Speeches kept per draw: A: 34%, B: 33%. Newey and West (1987) standard errors (6 quarterly lags in Panel A, 18 monthly lags in Panel B) in parentheses.

B.5 Communication Invariance Across Status Changes

Table B.4 reports the communication-style regressions discussed in Section 2.4: for the promotion transition, four features of a speech—its stance m, the absolute stance $|m|$, its dictionary-keyword density, and its length—are regressed on the status indicator, with speaker fixed effects and standard errors clustered by speaker. Stance, its magnitude, and keyword density are statistically indistinguishable across the transition; the only feature that moves is speech length, which falls by roughly

620 words.

TABLE B.4—COMMUNICATION INVARIANCE ACROSS STATUS CHANGES

Speech feature	Promotion pre vs post (switchers)
m (stance)	0.001
	(0.002)
$\lvert m \rvert$ (absolute stance)	0.003
	(0.004)
Keyword density (/1,000 words)	0.536
	(3.686)
Speech length (words)	-623.599**
	(251.080)

Note: Each row is a within-speaker OLS regression of the speech feature on the post-promotion indicator, with speaker fixed effects and standard errors clustered by speaker; the coefficient (with its standard error below) measures how the feature shifts when institutional status changes with the speaker held fixed. m is the per-speech stance (unsmoothed, hence independent of the half-life); $|m|$ its magnitude; keyword density is the number of dictionary-keyword tokens per 1,000 words; length is the word count. The sample is the thirteen status switchers with speeches in the 1961–2024 corpus; members whose only recorded speeches predate the corpus start are excluded. Stance, its magnitude, and keyword density are statistically indistinguishable across the transition; the only feature that moves is the length of a promoted governor's speeches, which the stance measure does not load on. The regressions pool 402 pre-promotion and 779 post-promotion speeches, over 13 speaker clusters. *, **, ***: significance at 10, 5, 1%.

C Robustness

C.1 Press Coverage: Normalizations and Event Windows

Table C.1 estimates the principal coverage premium under four normalizations of the attention measure—log articles, articles per speech delivered by the speaker that year, articles per thousand words of the speech, and the speaker's within-year share of *Reuters*' Federal Reserve coverage—and under two event windows. The premium is positive and significant in every cell.

TABLE C.1—MEDIA ATTENTION AND INSTITUTIONAL STATUS: NORMALIZATIONS AND EVENT WINDOWS

	Full sample		Status switchers	
	Wide $(-5, 15)$	Tight $(-1, 5)$	Wide $(-5, 15)$	Tight $(-1, 5)$
$\log(1 + \text{Mentions})$	1.556***	1.266***	1.234***	0.972***
	(0.268)	(0.209)	(0.421)	(0.347)
Mentions/Speeches	2.684***	1.064***	4.138***	1.673***
	(0.719)	(0.305)	(1.294)	(0.537)
Mentions/1,000 words	32.83***	12.43***	38.40**	14.07***
	(7.930)	(2.919)	(16.046)	(5.271)
Coverage share	0.0250***	0.0223***	0.0333***	0.0294***
	(0.005)	(0.004)	(0.010)	(0.009)
N	2090	2090	716	716

Note: Each cell is a separate OLS regression of a media-attention measure on a *Principal* indicator—one for speeches delivered while the speaker served as Chair or Vice Chair, zero for ordinary Board governors—and controls: the speech's own tone (its MSS score), log speech length, days to the nearest FOMC meeting, and aggregate Fed-news intensity in the window, with year fixed effects. Attention is measured at the speech level over the event window in the column head, in calendar days relative to the speech date. *Mentions* is the number of distinct *Reuters News* articles in the window that name the speaker or report the speech (diary/market-wrap round-ups excluded); Mentions/Speeches scales by the speeches the speaker delivered that year; coverage share is the speaker's within-year share of Reuters' Fed coverage. The *Full sample* columns use all Board members; the *Status switchers* columns restrict to members observed first as governors and later as Chair or Vice Chair. Neither carries speaker fixed effects, so neither is a within-person estimate; for that see main-text Figure 3. Standard errors in parentheses are clustered by speaker. *, **, and *** indicate significance at the 10, 5, and 1 percent levels. The sample covers 1992–2024. Speeches by the Chair and Vice Chair account for 82.3 percent of the matched article coverage.

C.2 Do Reports Reflect the Speech's Tone?

Table C.2 asks whether the stance a speech expresses survives into the *Reuters* coverage of it. The two are positively and significantly related at every cut, but the association is weak. In the tight window the correlation is 0.219 for principals and 0.125 for governors; widening the window leaves the principal figure unchanged and lowers the governor figure to 0.102. Coverage therefore transmits the direction of a speech only partially. The comparison across offices is not clean: principals draw several times more articles per speech, so their coverage tone is averaged over more text and measured with less noise, which inflates the correlation independently of any difference in how faithfully the two offices are reported.

TABLE C.2—DO REUTERS REPORTS REFLECT THE TONE OF THE SPEECH?

	Speeches N	Articles N	Correlation ρ	Pass-through $\hat{\beta}$
A. Small $(-1, 5)$ *event window*				
Principal (Chair + Vice)	847	19,550	0.219^{***}	0.172^{***}
				(0.032)
Governor	1,076	5,393	0.125^{***}	0.161^{***}
				(0.040)
B. Extended $(-5, 15)$ *event window*				
Principal (Chair + Vice)	852	32,937	0.219^{***}	0.127^{***}
				(0.025)
Governor	1,134	6,734	0.102^{***}	0.111^{***}
				(0.031)

Note: Each row is a speaker group. The unit is a speech: for every Board-member speech that Reuters directly reported, the speech's own tone is compared with the mean tone of the Reuters articles covering it, both scored with the same MSS dictionary. *Principal* pools speeches delivered as Chair or Vice Chair; *Governor* is ordinary Board members. *Articles* counts the distinct strong-match Reuters articles behind those speeches (a direct attribution to the speaker or the speaker named in the headline, within the event window; diary/market-wrap round-ups excluded). The window is in calendar days relative to the speech date. *Correlation* is the Pearson correlation between speech tone and mean coverage tone; *Pass-through* is the slope $\hat{\beta}$ from an OLS regression of coverage tone on speech tone (HC1 standard errors in parentheses; the reported significance is that of $\hat{\beta}$, equivalently of ρ). Matched speeches 1992–2024. *, **, and *** indicate significance at the 10, 5, and 1 percent levels.

C.3 SPF Regressions

Table C.3 repeats the headline forecast-error regressions on the *Survey of Professional Forecasters*' quarterly forecasts of the three-month Treasury-bill rate; the non-Chair pattern of Section 2.2 reappears at every horizon.

Table C.3—Predictive Regressions of SPF Three-Month Treasury-Bill Expectations

	h = 1Q	h = 2Q	h = 3Q	h = 4Q
A. MSS_t^n and MSS_t^c				
MSS_t^n	-2.755***	-5.445***	-8.852***	-11.554***
	(0.898)	(1.471)	(2.032)	(2.851)
MSS_t^c	0.696	1.179	2.231	3.206*
	(0.582)	(1.123)	(1.605)	(1.927)
R2adj	0.113	0.157	0.212	0.219
B. MSS_t, FFR_t, ΔCPI_t, $CFNAI_{t-1}$ and ΔEMP_t				
MSS_t^n	-2.408***	-4.481***	-7.326***	-9.324***
	(0.915)	(1.391)	(1.720)	(2.076)
MSS_t^c	0.442	0.704	1.529	2.265
	(0.565)	(1.090)	(1.509)	(1.753)
R2adj	0.156	0.239	0.312	0.341
N (quarters)	144	144	144	144

Note: The table presents predictive regressions of SPF three-month Treasury-bill forecast errors at horizon h, with h ranging from 1 to 4 quarters ahead; forecast errors are measured in percentage points of the three-month Treasury-bill 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 $CFNAI_{t-1}$. All regressions include a constant. *, **, and *** indicate significance at the 10, 5, and 1 percent levels, respectively. The data cover the period from 1989Q1 to 2024Q4 for a total of 144 observations.

C.4 Long-Maturity Regressions

Table C.4 re-estimates the BCFF forecast-error regressions, with the non-Chair MSS entering alone, across Treasury maturities from three months to ten years: the non-Chair predictability holds through two years, strengthening with the forecast horizon, and attenuates at ten, consistent with a policy-path signal. The series behind the headline forecast-error regression are plotted in main-text Figure 2.

TABLE C.4—PREDICTIVE REGRESSIONS OF BCFF TREASURY EXPECTATIONS ACROSS MATURITIES

	h = 1Q	h = 2Q	h = 3Q	h = 4Q
A. MSS_t^n				
3M	-1.684**	-3.917***	-6.942***	-9.557***
	(0.744)	(1.385)	(2.059)	(2.904)
R2adj	0.056	0.105	0.163	0.184
6M	-1.919**	-4.309***	-7.307***	-9.894***
	(0.890)	(1.528)	(2.261)	(3.135)
R2adj	0.057	0.110	0.164	0.184
1Y	-3.549***	-5.980***	-8.823***	-11.132***
	(1.312)	(1.774)	(2.370)	(3.177)
R2adj	0.056	0.104	0.153	0.177
2Y	-2.017**	-4.531***	-6.888***	-8.907***
	(0.850)	(1.463)	(2.109)	(2.948)
R2adj	0.042	0.102	0.146	0.159
10Y	-0.770	-1.944*	-2.966**	-4.265**
	(0.626)	(1.070)	(1.481)	(1.946)
R2adj	0.002	0.024	0.044	0.071
B. MSS_t^n, FFR_t, ΔCPI_t, $CFNAI$ and ΔEMP_t				
3M	-1.710**	-3.443***	-6.060***	-8.118***
	(0.755)	(1.304)	(1.662)	(2.177)
R2adj	0.180	0.230	0.302	0.323
6M	-1.754**	-3.641***	-6.182***	-8.340***
	(0.693)	(1.199)	(1.606)	(2.239)
R2adj	0.219	0.255	0.318	0.331
1Y	-3.631***	-5.459***	-7.772***	-9.712***
	(1.300)	(1.399)	(1.626)	(2.152)
R2adj	0.133	0.195	0.266	0.298
2Y	-1.704**	-3.639***	-5.398***	-7.156***
	(0.713)	(1.123)	(1.646)	(2.434)
R2adj	0.053	0.137	0.210	0.223
10Y	-1.130	-2.045*	-2.429	-3.311*
	(0.736)	(1.158)	(1.504)	(1.908)
R2adj	-0.019	-0.000	0.037	0.080
N (quarters)	136	136	136	136

Note: The table presents predictive regressions of Blue Chip Financial Forecasts (BCFF) forecast errors at horizon h, with h ranging from 1 to 4 quarters ahead, for Treasury expectations at the 3-month, 6-month, 1-year, 2-year, and 10-year maturities; forecast errors are measured in percentage points of the corresponding Treasury yield. Each row reports the coefficient on the non-Chair MSS, MSS_t^n, with Newey and West (1987) standard errors (6 quarterly lags) in parentheses below and the adjusted R^2. These regressions estimate Equation (3), where X_t denotes the Panel B controls and is empty in Panel A. 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.

C.5 Out-of-Sample Regressions

Out-of-sample forecasts are one step ahead, generated recursively from an expanding window whose parameters are first estimated on the initial 25 percent of each sample, and forecast improvements are evaluated with Clark and West (2007) statistics (Tables C.5, C.6, and C.7); the table notes give the design. For the survey errors this design is demanding in its early years: the expanding window first estimates the MSS coefficient on the pre-2000 sample, when Federal Reserve communication was scarce and the relationship least precisely estimated, and the forecast improvements accordingly concentrate in the later part of the evaluation window, mirroring the rolling-window pattern of Figure C.1.

TABLE C.5—OUT-OF-SAMPLE PREDICTIVE REGRESSIONS FOR FFR FUTURES

Predictors	h=2M	h=4M	h=6M	h=8M
Constant	0.000	0.002^{*}	0.018^{**}	0.060^{*}
	(1.031)	(1.513)	(1.671)	(1.536)
ΔEMP_t	0.000^{*}	0.002^{**}	0.016^{**}	0.051^{**}
	(1.461)	(1.871)	(1.860)	(1.659)
ΔCPI_t	0.000^{***}	0.005^{***}	0.022^{**}	0.059^{**}
	(2.577)	(2.352)	(1.940)	(1.728)
$CFNAI_{t-1}$	0.000^{**}	0.006^{**}	0.027^{**}	0.070^{**}
	(2.146)	(2.133)	(1.894)	(1.757)

Note: The table reports differences in out-of-sample mean squared prediction error for FFR futures forecast errors estimated with and without the MSS included among the predictors. A positive value indicates that the MSS improves forecast accuracy. In each successive row, the predictor listed in the leftmost column is added to the set of predictors included in the previous rows. The estimation sample spans the period from February 1989 to December 2024. Out-of-sample forecasts are one step ahead from an expanding window: parameters are first estimated by OLS on the initial 25 percent of the sample and re-estimated each period on the data available to that point, yielding 294 out-of-sample forecasts. Forecast errors are measured in percentage points, so the differences are in squared percentage points. The Clark and West (2007) statistics are reported in parentheses. *, **, and *** indicate significance at the 10, 5, and 1 percent levels of the one-sided test, respectively. To maintain a sufficiently long sample that is comparable across prediction horizons, the out-of-sample exercise is restricted to shorter-maturity fed funds futures.

TABLE C.6—OUT-OF-SAMPLE PREDICTIVE REGRESSIONS FOR BCFF SHORT-TERM RATE EXPECTATIONS

Predictors	h=1Q	h=2Q	h=3Q	h=4Q
Constant	0.001*	0.026*	0.120*	0.294*
	(1.343)	(1.563)	(1.514)	(1.406)
ΔEMP_t	0.002**	0.023**	0.098**	0.208*
	(1.734)	(1.776)	(1.790)	(1.597)
ΔCPI_t	0.005***	0.031**	0.110**	0.222*
	(2.415)	(2.048)	(1.894)	(1.621)
CFNAI	0.017*	0.099*	0.272*	0.471**
	(1.424)	(1.490)	(1.643)	(1.703)

Note: The table reports differences in out-of-sample mean squared prediction error for Blue Chip Financial Forecasts (BCFF) forecast errors estimated with and without the MSS included among the predictors. A positive value indicates that the MSS improves forecast accuracy. In each successive row, the predictor listed in the leftmost column is added to the set of predictors included in the previous rows. The estimation sample spans the period from March 1991 to December 2024. Out-of-sample forecasts are one step ahead from an expanding window: parameters are first estimated by OLS on the initial 25 percent of the sample and re-estimated each period on the data available to that point, yielding 102 out-of-sample forecasts. Forecast errors are measured in percentage points, so the differences are in squared percentage points. The Clark and West (2007) statistics are reported in parentheses. *, **, and *** indicate significance at the 10, 5, and 1 percent levels of the one-sided test, respectively.

TABLE C.7—OUT-OF-SAMPLE PREDICTIVE REGRESSIONS FOR BOND RETURNS

Predictors	n=2Y	n=3Y	n=4Y	n=5Y	n=6Y	n=7Y
Constant	0.068^{**}	0.114^{**}	0.128^{**}	0.128^{**}	0.121^{***}	0.113^{***}
	(2.001)	(2.151)	(2.256)	(2.319)	(2.349)	(2.356)
PCs	0.071^{**}	0.105^{**}	0.106^{**}	0.095^{**}	0.081^{**}	0.069^{**}
	(1.917)	(1.872)	(1.837)	(1.796)	(1.751)	(1.705)
ΔEMP_t	0.063^{**}	0.094^{**}	0.096^{**}	0.087^{**}	0.075^{**}	0.065^{**}
	(2.257)	(2.107)	(2.011)	(1.931)	(1.859)	(1.793)
τ_t^{CPI}	0.049^{***}	0.075^{**}	0.076^{**}	0.068^{**}	0.058^{**}	0.048^{**}
	(2.327)	(2.274)	(2.217)	(2.164)	(2.114)	(2.064)
$CFNAI_{t-1}$	0.051^{**}	0.077^{**}	0.078^{**}	0.070^{**}	0.059^{**}	0.049^{**}
	(2.269)	(2.180)	(2.114)	(2.056)	(2.001)	(1.946)

Note: The table reports differences in out-of-sample mean squared prediction error for bond return forecasts estimated with and without the MSS included among the predictors. A positive value indicates that the MSS improves forecast accuracy. In each successive row, the predictor listed in the leftmost column is added to the set of predictors included in the previous rows. The estimation sample spans the period from April 1967 to December 2024. Out-of-sample forecasts are one step ahead from an expanding window: parameters are first estimated by OLS on the initial 25 percent of the sample and re-estimated each period on the data available to that point, yielding 520 out-of-sample forecasts. The dependent variable is the annual excess return divided by bond maturity, in percentage points, so the differences are in squared percentage points. The Clark and West (2007) statistics are reported in parentheses. *, **, and *** indicate significance at the 10, 5, and 1 percent levels of the one-sided test, respectively.

C.6 Subsample and Rolling-Window Stability

Table C.8 re-estimates the headline forecast-error specification over the subsamples discussed in Section 2.2: the predictive content survives outside the zero-lower-bound years and strengthens at every horizon beyond the shortest; in the pre-2008 sample on its own the survey slopes are significant at the 5 percent level or better at every horizon, and the futures slopes at every horizon (the eight-month cell at the 10 percent level).

TABLE C.8—SUBSAMPLE STABILITY OF THE NON-CHAIR PREDICTIVE CONTENT

	A. BCFF forecast errors (quarterly)			
	$h = 1Q$	$h = 2Q$	$h = 3Q$	$h = 4Q$
Full sample	-1.374***	-3.773***	-6.824***	-10.216***
	(0.519)	(1.040)	(1.575)	(2.295)
N	136	136	136	136
Pre-2008	-1.940**	-4.387***	-6.247***	-8.366**
	(0.805)	(1.615)	(2.406)	(3.678)
N	68	68	68	68
Excl. 2008–2015	-1.272***	-3.987***	-7.394***	-11.186***
	(0.365)	(1.000)	(1.591)	(2.367)
N	104	104	104	104
Post-2015	-0.744**	-2.960**	-6.415***	-10.336***
	(0.348)	(1.277)	(1.990)	(2.592)
N	36	36	36	36
	B. FFR-futures forecast errors (monthly)			
	$h = 4M$	$h = 8M$	$h = 12M$	
Full sample	-1.893***	-5.159***	-9.527***	
	(0.573)	(1.301)	(1.930)	
N	432	368	252	
Pre-2008	-2.459***	-3.383*	-9.205***	
	(0.694)	(1.761)	(2.894)	
N	228	164	48	
Excl. 2008–2015	-1.848***	-5.409***	-9.788***	
	(0.430)	(1.186)	(2.007)	
N	336	272	156	
Post-2015	-0.987*	-4.589***	-8.795***	
	(0.532)	(1.523)	(2.270)	
N	108	108	108	

Note: Coefficients on the non-Chair MSS from the joint Panel-A specification of Tables 4 and 3 ($y_t = \gamma_0 + \gamma_1 MSS_t^n + \gamma_2 MSS_t^c + \varepsilon_t$; the Chair coefficient, insignificant throughout, is not shown), re-estimated over the indicated subsamples. Forecast errors are measured in percentage points of the federal funds rate; the full-sample rows reproduce Panel A of the referenced tables, whose notes state the sample windows. Newey and West (1987) standard errors (6 quarterly lags in Panel A, 18 monthly lags in Panel B) in parentheses. *Excl. 2008–2015* drops the zero-lower-bound years. *, **, and *** indicate significance at the 10, 5, and 1 percent levels.

Figure C.1 reports the non-Chair MSS slope from the joint Panel-A specification of Table 3, estimated over a rolling ten-year window, together with 90 percent Newey–West confidence bands. The slope is negative in more than nine-tenths of the windows at both horizons, and significantly negative at the 10 percent level in about three-fifths of them at the one-quarter horizon and just over half at the one-year horizon. The relationship first attains sustained significance in windows ending in 2004, holds through 2016, lapses over 2016–2021, and returns from late 2021 to the end of the sample. In the earliest windows the slope is smaller and, at the one-year horizon, far less precisely estimated; both sharpen over the sample, consistent with the growing role of central bank communication.

FIGURE C.1—ROLLING PREDICTIVE COEFFICIENT OF THE NON-CHAIR MSS

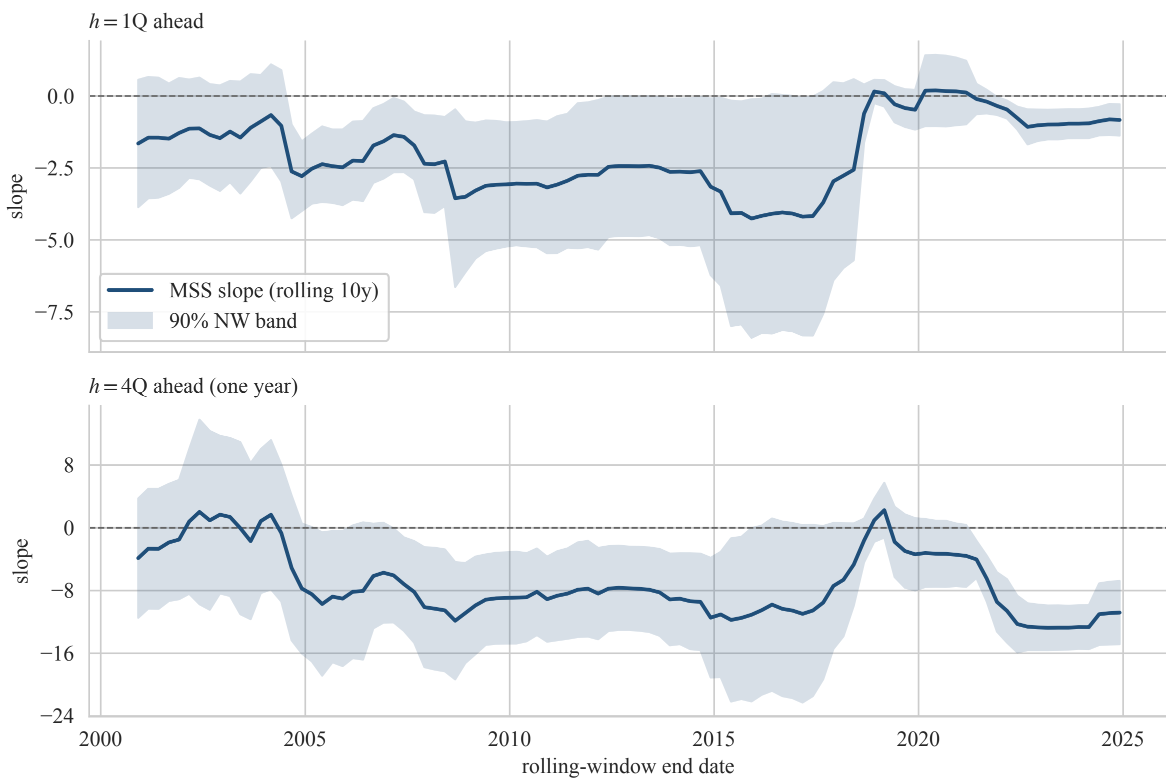

Note: Each line shows the slope coefficient on the non-Chair MSS from regressing BCFF federal-funds-rate forecast errors on the non-Chair and Chair MSS jointly (the Panel-A specification of Table 3; the Chair coefficient is not shown), estimated over a rolling ten-year (40-quarter) window—97 windows ending 2000:Q4 through 2024:Q4. The upper panel uses the one-quarter-ahead forecast error and the lower panel the one-year-ahead (four-quarter) error. Shaded areas are 90 percent confidence bands computed with Newey–West standard errors (six quarterly lags). The MSS is on the same scale as the predictive-regression tables, and the horizontal axis marks the end date of each rolling window.

C.7 Joint Robustness to Smoothing and Dictionary Choice

The headline excess-bond-return result relies on two researcher choices that the predictive regressions hold fixed: the half-life parameterizing the EWMA smoothing weight ω of Equation (2), and the set of keywords used to score speeches. To verify that the result does not hinge on a specific combination of the two, I re-estimate the two-year excess-bond-return regression

$$rx_t^{(2)} = \gamma_0 + \gamma_1 MSS_t^j + \gamma_2' PC_t + u_t, \qquad j \in \{c, n\},$$

where PC_t collects the first three principal components of yields. The grid varies both dimensions jointly: the smoothing half-life ranges over 5, 10, 20, 30, 50, 60, 90, and 120 days, and the dictionary is perturbed by leaving out each single keyword (leave-one-out) and every pair of keywords (leave-two-out), which together with the full thirteen-term baseline gives 92 dictionary variants. The baseline cell is the sixty-day half-life with the full dictionary, and the grid spans $8 \times 92 = 736$ specifications for each speaker group.

Figure C.2 displays the t-statistic on the speech signal in every cell. The non-Chair result holds in every cell: the coefficient is negative in all 736 specifications and statistically significant at the 5 percent level in every one, with $|t|$ ranging from 2.0 to 3.8. The Chair signal, used here as a null control, is never significant in any cell ($|t| \leq 1.9$ throughout). The result therefore does not hinge on either choice, and the Chair/non-Chair contrast survives their joint perturbation.

FIGURE C.2—JOINT ROBUSTNESS OF THE TWO-YEAR PREDICTIVE RESULT

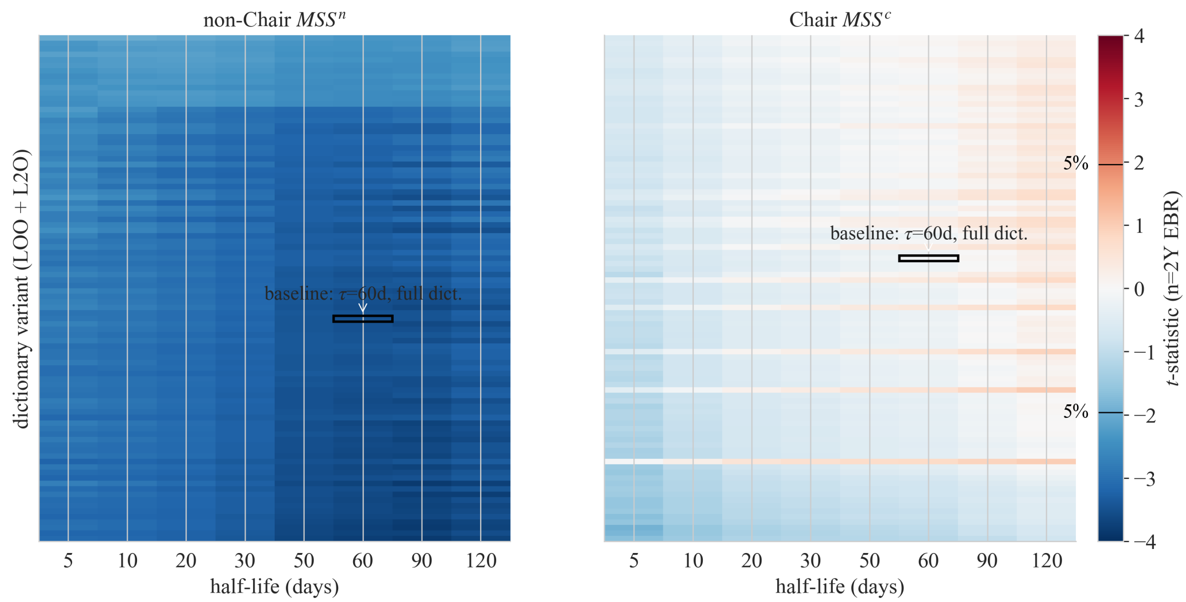

Note: Each cell reports the t-statistic on the speech signal in the two-year excess-bond-return regression $rx_t^{(2)} = \gamma_0 + \gamma_1 MSS_t^j + \gamma_2' PC_t + u_t$, with PC_t the first three yield principal components, estimated with Newey and West (1987) standard errors (18 monthly lags). The horizontal axis varies the half-life (in days) of the EWMA smoothing weight ω of Equation (2); the vertical axis varies the scoring dictionary across leave-one-out and leave-two-out perturbations of the baseline keyword set (92 variants, including the baseline). The left panel uses the non-Chair MSS^n and the right panel the Chair MSS^c; within each panel, rows are ordered by mean $|t|$, and the black rectangle marks the baseline cell (sixty-day half-life, full dictionary). Black reference marks on the colorbar denote the ± 1.96 (5 percent) significance thresholds.

The perturbation exercise extends beyond the single headline regression to the paper's other tracked results: the forecast-error regressions of Tables 3 and 4, the rule-deviation and surprise regressions of Tables 7 and 8, and the switcher excess-return regression of Table B.2. Figure C.3 removes a random 5, 10, or 15 percent of the pooled dictionary tokens—keywords and hawkish/dovish modifier stems alike—re-scores the full corpus, and re-estimates each tracked regression, with 1000 Monte Carlo draws per fraction. Each result that is significant under the current dictionary remains so in 73 to 99 percent of draws even with 15 percent of the dictionary removed, and the designated null control stays insignificant in at least 97 percent of draws at every fraction. At the 5 percent fraction the current dictionary sits between the 25th and 73rd percentile of every draw distribution, not at a favorable edge; at larger fractions the draws attenuate toward zero, the direction classical measurement error implies. Across these draws, no tracked estimate depends on the presence of any particular part of the dictionary.

FIGURE C.3—DICTIONARY ROBUSTNESS OF THE HEADLINE t-STATISTICS

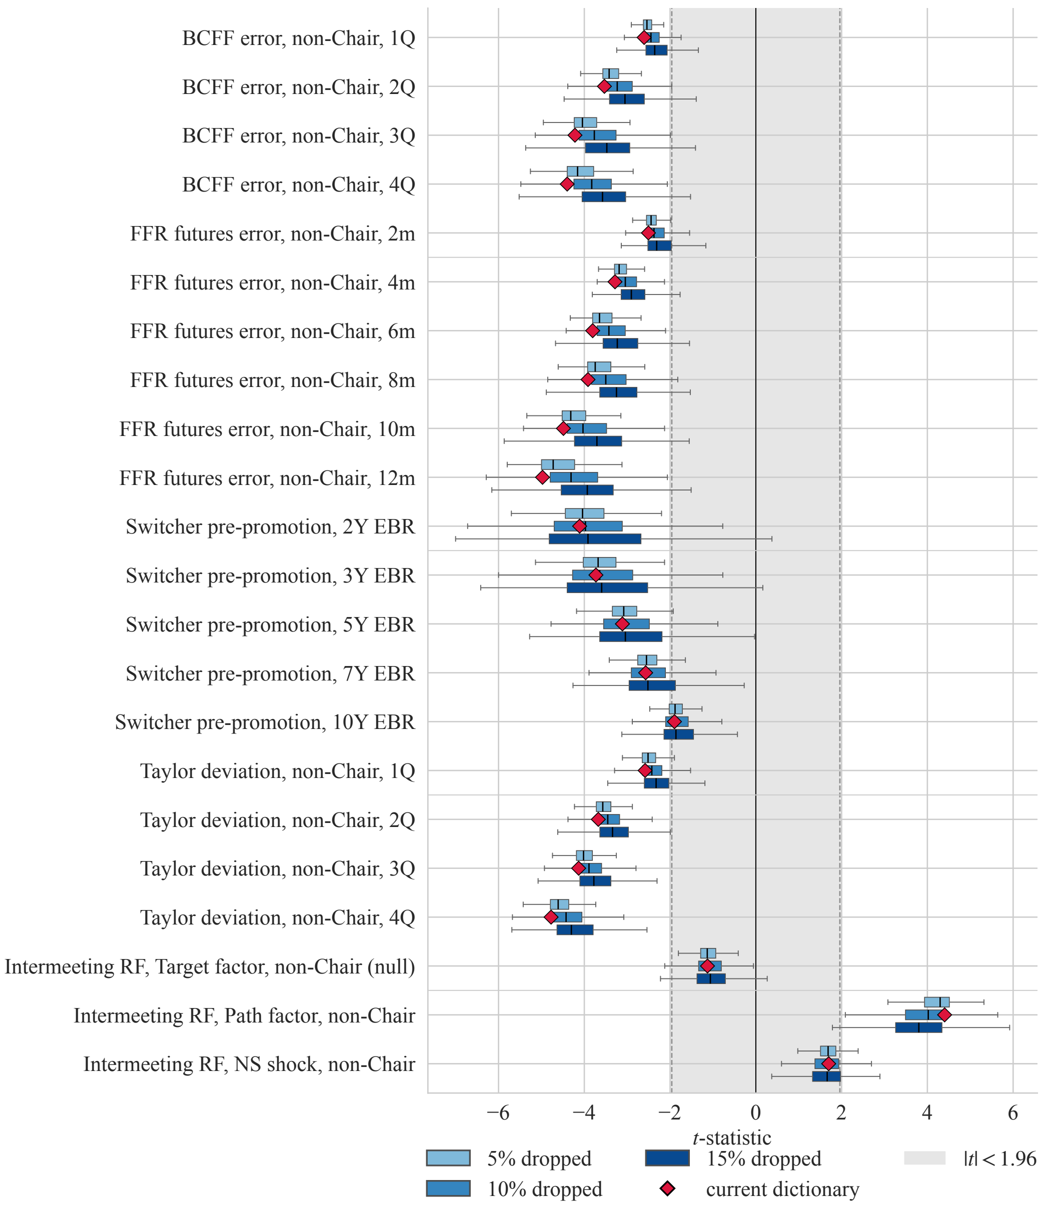

Note: Each row is a tracked t-statistic, labeled by the outcome and regressor of its regression. A random fraction (5, 10, or 15 percent) of the pooled dictionary tokens—keywords and hawkish/dovish modifier stems—is removed, the full corpus is re-scored, and the regression is re-estimated through the identical pipeline; each box summarizes the distribution of the t-statistic across 1000 Monte Carlo draws (boxes span the interquartile range, whiskers the Tukey range). The red diamond marks the current dictionary, evaluated with the same re-scorer at zero tokens dropped. The shaded band is the $|t| < 1.96$ non-significance region. The row marked (null)—the target-factor surprise—is a control cell expected to remain insignificant.

Table C.9 varies the construction of the daily signal itself rather than its inputs: a carry-forward rule that holds a member's last observed score until the next speech instead of decaying toward zero on quiet days, and a whole-sample linear detrend of each group's aggregated series. The non-Chair coefficient of main-text Table 3 stays negative and significant at every horizon under both constructions.

TABLE C.9—NON-CHAIR PREDICTIVE COEFFICIENT UNDER ALTERNATIVE SIGNAL CONSTRUCTIONS

	1Q	2Q	3Q	4Q
Baseline (decay to zero on quiet days)	-1.373***	-3.767***	-6.815***	-10.217***
	(0.520)	(1.042)	(1.577)	(2.295)
Carry-forward fill (hold last score)	-0.147**	-0.303***	-0.480***	-0.630***
	(0.064)	(0.115)	(0.161)	(0.211)
Whole-sample linear detrend	-1.365***	-3.751***	-6.797***	-10.217***
	(0.522)	(1.049)	(1.591)	(2.309)

Note: Coefficient on the non-Chair MSS in the joint Panel-A specification of Table 3 (both tones enter; the Chair coefficient is not shown), with the daily tone series rebuilt under each listed construction. *Carry-forward fill* holds a member's last observed score until the next speech instead of decaying toward zero on quiet days. *Whole-sample linear detrend* removes a linear time trend, fit on the full sample, from each group's aggregated daily series. Coefficients are not comparable across rows because each construction changes the signal's scale; the sign and significance pattern is the object. Forecast errors are measured in percentage points of the federal funds rate; $N = 136$ survey rounds at every horizon. Newey and West (1987) standard errors (6 quarterly lags) in parentheses. *, **, and *** indicate significance at the 10, 5, and 1 percent levels, respectively.

C.8 Robustness to Externally Designed Dictionaries

The leave-one-out and leave-two-out perturbations above vary my own keyword set. Here, I test whether the result survives measures built by other researchers. I re-score the entire corpus with three external measures and pass each through the identical pipeline—same smoothing half-life, same Chair/non-Chair split, same regressions—so that the scoring rule is the only thing that changes.[2] Two are policy-stance dictionaries. The dictionary of Apel et al. (2022) pairs a rich

[2]Each external list enters the pipeline encoded hawkish-positive in the baseline dictionary format; the Gorodnichenko et al. (2023) index, whose published convention takes $+1$ as dovish, is sign-flipped accordingly, and for Loughran–McDonald I treat positive financial sentiment as hawkish (strong-economy language). The expected predictive sign is therefore negative throughout.

set of macroeconomic topics—including employment and resource utilization—with topic-specific modifier lists; and the hawkish/dovish word list of Gorodnichenko et al. (2023) (their Appendix Table D1) is built specifically on FOMC communication. As a discriminant control I also include the financial-sentiment lexicon of Loughran and McDonald (2011), which scores generic positive versus negative tone rather than policy stance. The three differ substantially from mine: at the daily frequency the non-Chair series correlate with my baseline at $\rho = 0.63$ (Apel et al., 2022), $\rho = 0.53$ (Gorodnichenko et al., 2023), and $\rho = -0.22$ (Loughran and McDonald, 2011).

Table C.10 shows two things. First, the two policy-stance dictionaries reproduce the predictive content. Under the Apel et al. (2022) lists the non-Chair result is reproduced at every horizon but the shortest: the two-year excess-bond-return coefficient is negative and significant at the 5 percent level ($t = -2.2$), and the federal-funds-futures forecast-error coefficient is negative throughout and strengthens with the horizon, from $t = -1.6$ at four months to $t = -3.6$ at one year. The FOMC-specific list of Gorodnichenko et al. (2023) likewise reproduces the headline excess-return result ($t = -2.1$ at two years) and the one-year forecast-error coefficient ($t = -2.4$), though it is weaker at the shorter forecast horizons. Second, the generic Loughran and McDonald (2011) sentiment lexicon—which scores optimism rather than policy stance—has no predictive content: the non-Chair coefficient is insignificant at every horizon and even changes sign, consistent with its weak negative correlation with the MSS ($\rho = -0.22$). This sharpens the interpretation of the result: what predicts forecast errors is the *hawkish/dovish* orientation of speech, not its general positivity.

Throughout, the Chair signal serves as the null control: it is insignificant in the large majority of cells, and the few Chair coefficients that do reach significance (the Gorodnichenko et al. excess-return cell and the one-year forecast errors under the baseline and Loughran–McDonald lists) come from specifications in which the Chair signal enters alone. In the joint specifications of the main text, where MSS^c and MSS^n appear together (Table 4), the Chair coefficient is insignificant beyond the shortest futures horizons (Section 2.2). Measures of policy stance designed by others, applied without tuning, thus reproduce the direction, the horizon pattern, and the speaker asymmetry of the main results, whereas a generic sentiment lexicon does not.

TABLE C.10—DICTIONARY ROBUSTNESS: EXTERNAL STANCE AND SENTIMENT DICTIONARIES

	Non-Chair MSS^n				Chair MSS^c			
	Baseline	ABH (2022)	VMP	LM	Baseline	ABH (2022)	VMP	LM
Panel A: excess bond return (NW, 18 lags)								
rx_t^{2Y}	-44.801***	-32.172**	-28.442**	-4.167	-11.170	-17.111	-31.801*	-11.762
	(13.231)	(14.849)	(13.660)	(11.599)	(15.484)	(13.962)	(17.414)	(17.122)
Panel B: FFR-futures forecast error (NW, 18 lags)								
$h = 4$M	-0.081***	-0.045	-0.013	0.022	0.017	0.004	0.003	-0.051
	(0.031)	(0.028)	(0.029)	(0.031)	(0.038)	(0.035)	(0.033)	(0.045)
$h = 8$M	-0.247***	-0.202**	-0.108	-0.024	-0.059	-0.093	-0.105	-0.142
	(0.084)	(0.083)	(0.076)	(0.078)	(0.104)	(0.096)	(0.089)	(0.115)
$h = 12$M	-0.484***	-0.409***	-0.282**	-0.125	-0.281***	-0.215	-0.241	-0.348**
	(0.134)	(0.114)	(0.116)	(0.119)	(0.098)	(0.161)	(0.156)	(0.164)
Panel C: agreement with LLM-classified reference sample								
Sample corr. ρ	0.17	0.16	0.12	0.05	0.17	0.16	0.12	0.05
Sample bal. accuracy	0.46	0.42	0.44	0.37	0.46	0.42	0.44	0.37
VoMP-gold corr. ρ	0.27	0.29	0.22	0.16	0.27	0.29	0.22	0.16
VoMP-gold bal. accuracy	0.42	0.42	0.46	0.38	0.42	0.42	0.46	0.38

Note: Each cell reports the coefficient on the speech signal per one-standard-deviation increase in the (hawkish-positive) signal, with Newey and West (1987) standard errors (18 monthly lags) in parentheses. Each dictionary's MSS series is standardized to unit standard deviation over its estimation sample, so coefficients are comparable across columns despite the word lists' different native scales (t-statistics unchanged). The *Baseline* columns use the paper's dictionary; the other columns re-score the identical corpus and pipeline with external word lists—*ABH (2022)*: the extended Apel et al. (2022) lists; *VMP*: Gorodnichenko et al. (2023) (Appendix Table D1); *LM (Loughran–McDonald)*: the Loughran and McDonald (2011) financial lexicon (net positive minus negative tone, positive = hawkish)—all signed hawkish-positive so the predictive sign is comparable. Panel A is the $n = 2$Y excess-bond-return regression $rx_t^{2Y} = \gamma_0 + \gamma_1 MSS_t^j + \gamma_2' PC_t$ (as in Table 5), estimated on $N = 482$ monthly observations, 1985:11–2025:12; Panel B regresses FFR-futures forecast errors at horizon h on the MSS (as in Table 4), estimated on $h = 4$M: $N = 432$, 1988:12–2024:12; $h = 8$M: $N = 368$, 1989:01–2024:12; $h = 12$M: $N = 252$, 2004:01–2024:12. At the daily frequency the external non-Chair series correlate with the baseline MSS^n at $\rho =$ ABH (2022) 0.63, VMP 0.53, LM -0.22. Panel C compares each method's signed per-sentence score with the LLM-classified labels of a 5,000-sentence sample of the corpus (Hawkish $= +1$, Neutral $= 0$, Dovish $= -1$): Pearson correlation of the raw score, and balanced accuracy (mean per-class recall) of the predicted class sign(score). Balanced accuracy is used because most sentences are Neutral; the panel's rows are method properties, identical across the two groups. The *VoMP-gold* rows repeat the comparison on the gold-labelled sentences of Gorodnichenko et al. (2023). *, **, ***: significance at the 10%, 5%, 1% level.

A frontier-LLM audit of the sentence-level classifications. As a further validation, I use a frontier large language model (Anthropic's Claude Opus 4.8) as an independent stance rater to re-classify a stratified random subsample of the LLM-classified reference sample of Table C.10, Panel C: 150 sentences, 50 per class, drawn with a fixed seed and classified without access to the stored reference labels. The audit yields two results. First, the dictionary's balanced accuracy on the subsample is 0.45, matching its full-sample agreement with the reference labels (0.46, Panel C)—the subsample is representative. Second, the dictionary is a sparse, high-precision filter: it assigns a non-zero score to only 17 percent of these (stance-enriched) sentences, but conditional on scoring, its sign agrees with the frontier-LLM rater on 81 percent and with the reference labels on 77 percent of sentences—it abstains on most stance-bearing language and rarely mis-signs the language it does score. Because the reference sample is itself LLM-classified, this exercise checks the dictionary against LLM stance judgments rather than human annotation; the independent human-annotated benchmark is the external *Voice of Monetary Policy* (VoMP) gold sample of Panel C. Missed sentences act as classical measurement error, which attenuates the predictive coefficients toward zero.

C.9 Group-Specific Reliability of the Tone Series

The attenuation just described puts the cross-group (Chair versus non-Chair) contrast in doubt only if the Chair series is the noisier of the two: a noisier or lower-variance Chair series would be mechanically more attenuated, and the Chair "null" partly a signal-to-noise artifact. Table C.11 measures the relevant object by performing a split-half reliability of the aggregated series, computed by splitting each speech's scored subsentences by parity and re-running both halves through the paper's aggregation procedure. The ordering runs against the artifact: the principal series is more reliable than the non-Chair series ($r_c^{SB} = 0.84$ vs. $r_n^{SB} = 0.78$), and within the promotion switchers the post-promotion series is more reliable than the pre-promotion one (0.88 vs. 0.77)—the significant coefficients are, if anything, the more attenuated ones. Disattenuating the Chair predictive coefficient by its own reliability rescales it by $1/r_c^{SB} \approx 1.2$ and leaves its t-statistic unchanged, so the Chair null survives correction for its own measurement error.

TABLE C.11—SPLIT-HALF RELIABILITY OF THE MSS, BY GROUP AND PROMOTION STATUS

	Non-Chair	Chair & Vice	Switchers, pre	Switchers, post
Scored subsentences per speech	8.6	10.6	10.6	10.6
Split-half correlation	0.645	0.728	0.630	0.779
Spearman–Brown reliability r_j^{SB}	0.784	0.843	0.773	0.876
N (speeches)	2277	1900	402	779
N (months)	766	765	748	765

Note: Each speech's dictionary-scored subsentences are split by parity of their within-document order (odd vs. even); each half is aggregated into a monthly signal through the paper's aggregation procedure (identical zero-fill, smoothing, and decay), and the two resulting series are correlated. The decomposition is exact: the two halves sum to the speech's baseline score. Spearman–Brown adjusts the half-length correlation to the reliability of the full-length series used in the regressions. N (months) counts the calendar months of the aggregated series entering the split-half correlation: each leg is aggregated into one monthly series per half, and the two halves are joined on the month. The monthly series run from February 1961 through November 2024. The switcher columns restrict to the governors later promoted to Chair or Vice Chair, before and after promotion.

D Model

D.1 Scope of the Model and the Measurement Exercise

This appendix derives the rational-inattention model set out in Section 4 of the main text, and records the qualifications the measurement exercise inherits from it. The setup, the index x of Equation (6), and Proposition 1 are stated in the main text; Section D.2 derives the optimal attention allocation and Section D.3 proves the proposition.

The parameter space of Proposition 1 is one-sided: for any group with positive influence, $\beta_j = \nu_j(1-\lambda_j)$ lies in $(0, \nu_j]$—full attention drives the loading to zero, never below it. Overreaction—a loading of the sign opposite to the underreaction prediction, as in the significant short-horizon Chair cells of Table 4, where the forecast-error convention inverts the surprise sign—lies outside the model, which allocates finite capacity to a signal but never amplifies it. The quantitative exercise of main-text Section 4.3 accordingly conditions on the underreaction sign. A non-positive surprise loading would imply $1-\lambda_j \leq 0$, outside the model's parameter space; both estimated loadings are positive (0.13 for the Chair, 0.51 for non-Chairs; main-text Table 9), so the exercise lies inside it. The wrong-signed short-horizon cells should therefore be read as a boundary of the framework—consistent with overreaction to attended communication (Bordalo et al., 2020)—rather than as additional evidence of the Chair null.

Three features of the measurement exercise follow from the model's setup. First, the non-Chair variance σ_n^2 entering the index x of Equation (6) is that of the Chair-orthogonalized series—the same regressor whose loading defines ν_n—so influence and variance refer to a single object. Orthogonalization mechanically lowers σ_n^2, which raises x and with it optimal Chair attention; using the raw non-Chair variance would lower x and *widen* the over-attention gap, so the choice is conservative for the result. Second, the ratios recovered in main-text Equation (9) are unit-free—any common rescaling of a tone series cancels between the surprise and the decision regression—

but they inherit the scale of the surprise series itself, which shifts the bit *levels* while leaving the realized-versus-optimal gap $\kappa_c - \kappa_c^*$ and the perceived-influence ratio invariant. The main text accordingly rests its argument on those invariant objects. Third, the model treats stances as serially independent, whereas measured tone is persistent; λ_j then reflects prior knowledge as well as within-period attention, so the recovered κ_j overstate per-period capacity, leaving the cross-group comparison intact insofar as the two groups' tones are similarly persistent.

D.2 The Forecaster's Problem and Optimal Attention

This section derives the forecaster's optimization problem and its solution. The forecaster predicts the interest-rate decision (5); the forecast, written i_t^e, is the conditional expectation given the information set $\mathscr{H}_t = \{S_t^c, S_t^n, X_t\}$,

$$i_t^e \equiv \mathbb{E}(i_t \mid \mathscr{H}_t) = \nu_c \, \mathbb{E}(\alpha_t^c \mid \mathscr{H}_t) + \nu_n \, \mathbb{E}(\alpha_t^n \mid \mathscr{H}_t) + g(X_t), \tag{D.1}$$

where the speech signals are $S_t^j = \alpha_t^j + \eta_t^j$ with η_t^c, η_t^n Gaussian white noise, mutually independent and independent of the stances. The forecaster is subject to the information-flow constraint

$$\mathscr{I}(\{\alpha_t^c, \alpha_t^n\}, \{S_t^c, S_t^n\}) \leq \kappa, \tag{D.2}$$

which states that the average per-period information the signals carry about the stances cannot exceed κ. Attention is allocated to minimize the expected squared forecast error

$$\min_{\kappa_c, \kappa_n} \quad \mathbb{E}\left[(i_t^e - i_t)^2\right], \tag{D.3}$$

subject to (D.2) and the definitions in (5) and (D.1).

The latent stances α_t^c, α_t^n follow stationary Gaussian white-noise processes. Following Sims

(2003) and Maćkowiak and Wiederholt (2009), with Gaussian states and a quadratic objective it is optimal to acquire Gaussian signals of the form "true state plus white-noise error,"

$$S_t^c = \alpha_t^c + \eta_t^c, \tag{D.4}$$

$$S_t^n = \alpha_t^n + \eta_t^n, \tag{D.5}$$

with η_t^c, η_t^n of variances $\tilde{\sigma}_c^2, \tilde{\sigma}_n^2$, orthogonal to the latent states. The information-processing constraint is then

$$\underbrace{\tfrac{1}{2}\log_2\left(1+\tfrac{\sigma_c^2}{\tilde{\sigma}_c^2}\right)}_{\kappa_c} + \underbrace{\tfrac{1}{2}\log_2\left(1+\tfrac{\sigma_n^2}{\tilde{\sigma}_n^2}\right)}_{\kappa_n} \leq \kappa, \tag{D.6}$$

so the signal-to-noise ratios are $\sigma_j^2/\tilde{\sigma}_j^2 = 2^{2\kappa_j} - 1$. Given the Gaussian environment, the optimal forecast is the conditional expectation

$$i_t^e = v_c \frac{\sigma_c^2}{\sigma_c^2 + \tilde{\sigma}_c^2} S_t^c + v_n \frac{\sigma_n^2}{\sigma_n^2 + \tilde{\sigma}_n^2} S_t^n + g(X_t), \tag{D.7}$$

with Kalman gains $\sigma_j^2/(\sigma_j^2 + \tilde{\sigma}_j^2) = 1 - 2^{-2\kappa_j}$, so that

$$i_t^e = v_c \left(1 - 2^{-2\kappa_c}\right)\left(\alpha_t^c + \eta_t^c\right) + v_n \left(1 - 2^{-2\kappa_n}\right)\left(\alpha_t^n + \eta_t^n\right) + g(X_t). \tag{D.8}$$

Substituting into the objective and using the orthogonality of the latent states and noise terms gives $v_c^2\sigma_c^2 2^{-2\kappa_c} + v_n^2\sigma_n^2 2^{-2\kappa_n} + \sigma_i^2$. Because the constraint binds at the optimum, $\kappa_n = \kappa - \kappa_c$, and dropping the shock variance σ_i^2 (independent of κ_c) the problem reduces to

$$\min_{\kappa_c \in [0,\kappa]} \left[v_c^2\sigma_c^2 2^{-2\kappa_c} + v_n^2\sigma_n^2 2^{-2(\kappa-\kappa_c)} \right]. \tag{D.9}$$

Assuming $\nu_n \neq 0$, the unique solution is

$$\kappa_c^* = \begin{cases} \kappa & \text{if } x \geq 2^{2\kappa}, \\ \frac{1}{2}\kappa + \frac{1}{4}\log_2(x) & \text{if } x \in \left[2^{-2\kappa}, 2^{2\kappa}\right], \\ 0 & \text{if } x \leq 2^{-2\kappa}, \end{cases} \tag{D.10}$$

where $x = (\sigma_c^2/\sigma_n^2)(\nu_c/\nu_n)^2$ as in (6). Optimal Chair attention κ_c^* is increasing in x: the forecaster concentrates on the group whose stance is more influential or more variable.

D.3 Endogeneity of Monetary Policy Surprises

After the decision is announced, the surprise is

$$Surprise_t \equiv i_t - \mathbb{E}(i_t \mid \mathscr{H}_t) = \nu_c(1-\lambda_c)\alpha_t^c + \nu_n(1-\lambda_n)\alpha_t^n - \nu_c\lambda_c\eta_t^c - \nu_n\lambda_n\eta_t^n + \varepsilon_t^i, \tag{D.11}$$

where $\lambda_j = \sigma_j^2/(\sigma_j^2 + \tilde{\sigma}_j^2) = 1 - 2^{-2\kappa_j^*}$ is the share of group j's stance incorporated into expectations before the announcement. When attention κ_j^* is high, $\lambda_j \to 1$ and group j's stance contributes little to the surprise; when attention is low, $\lambda_j \to 0$ and the stance drives a predictable surprise.

Proof of Proposition 1. The loading on α_t^j in (D.11) is $\nu_j(1-\lambda_j)$; substituting $\lambda_j = 1 - 2^{-2\kappa_j^*}$ gives $\beta_j = \nu_j 2^{-2\kappa_j^*}$, so $|\beta_j|$ is strictly decreasing in κ_j^* with limits $|\nu_j|$ at $\kappa_j^* = 0$ and 0 as $\kappa_j^* \to \infty$. Since the processing errors η_t^c, η_t^n are private to the forecaster and orthogonal, like ε_t^i, to the communicated stances, α_t^j is the only term in (D.11) correlated with a text-based measure of group j's communication, so the part of $Surprise_t$ predictable from measured tone scales with β_j. (The surprise is orthogonal to the forecaster's *own* signal S_t^j by the orthogonality of forecast errors; the model's testable content therefore concerns the public communication, which proxies α_t^j.) $\square$

D.4 Model Validity

This appendix documents the two facts that support the meeting-tone proxies of Section 4. Table D.1 shows that the FOMC's rate decision loads significantly and positively—a hawkish meeting tone preceding a tighter decision—on both groups' meeting tones; this is the sign-and-significance basis for the influence weights ν_j, whose controlled estimates Table 9 reports. Table D.2 reports the individual-stance influence estimates in the manner of Chappell et al. (2004), measuring influence over the rate *decision*—the object the surprise is defined against and the one that sets the model's ν_j; they place the Chair's implied share of committee influence 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; non-Chair dominance holds across specifications. Table D.3 reports a small but consistently positive correlation between a member's public speech tone and the same member's stance inside subsequent meetings; the alignment strengthens mildly for non-Chairs while decaying for Chairs as the window lengthens. Appendix Tables D.5–D.9 illustrate the foreshadowing with paired speech and transcript excerpts.

Table D.4 addresses the residualization choice behind the influence weights. The baseline regresses the decision on raw Chair tone and Chair-orthogonalized non-Chair tone, assigning the Committee's common tone component entirely to the Chair—a total effect set against a partial one. Re-estimating the full system under the reverse allocation (raw non-Chair, non-Chair-orthogonalized Chair), with both series raw, and with both orthogonalized, shows that the baseline is the allocation most favorable to the Chair: the influence ordering $\nu_n > \nu_c$ holds at the point estimates under every alternative, and by a wider margin than in the baseline; the both-orthogonalized column, the only alternative with an admissible gap, puts the gap at 0.63 bits against the baseline's 0.57.

TABLE D.1—REGRESSION OF MEETING INTEREST-RATE DECISIONS ON TRANSCRIPT TONE BY STATUS

	Rate decision ΔFFR_t^{fomc}	
	A. MSS_t only	B. + controls
$MSS_t^{F,n}$	5.437***	5.651***
	(1.118)	(1.124)
$MSS_t^{F,c}$	3.931***	3.581***
	(0.989)	(0.953)
R2adj	0.195	0.208
N (meetings)	185	185

Note: The table presents regressions of the FOMC interest-rate decision ΔFFR_t^{fomc} on meeting transcript tone by institutional status. Here, $MSS_t^{F,n}$ and $MSS_t^{F,c}$ are computed as the average tone of FOMC members participating in the meetings from group non-Chair and Chair, respectively. The tone for non-Chair members is purged of its Chair component to generate an orthogonal series. Newey and West (1987) standard errors (18 lags at the meeting frequency) in parentheses. Column A uses only the MSS; column B adds controls. The full set of controls in column B includes current month FFR_t, annual growth rate of employment ΔEMP_t, year-on-year headline CPI inflation ΔCPI_t, and $CFNAI_{t-1}$. All regressions include a constant. *, **, and *** indicate significance at the 10, 5, and 1 percent levels, respectively. The estimation covers the common sample of meetings with a scored statement, from February 1994 to December 2020 for a total of 185 observations. The FOMC has issued a post-meeting statement only since February 1994, which sets the start of this common sample.

TABLE D.2—CHAIR INFLUENCE: INDIVIDUAL-STANCE (CHAPPELL-STYLE) ESTIMATES

	No controls	Macro controls
Current group-mean estimator (Chair-orthogonalized non-Chair)		
Chair-group weight ν_c	3.53	3.23
Non-Chair weight ν_n (orth.)	3.98	4.26
Chappell-style individual-stance estimator: implied Chair influence share		
Mean-voter + Chair	31%	34%
Median-voter + Chair	32%	35%
Symmetric partition, Chair group vs median(n)	29%	22%

Note: Influence of the Chair recovered from the realized target-rate change regressed on individual FOMC members' meeting-intervention tone (1988–2020, all-voting Board panel; $N = 224$ meetings, on average 6.5 members per meeting). The upper block reproduces the paper's current estimator (Chair-group tone plus *Chair-orthogonalized* non-Chair tone). The lower block follows Chappell et al. (2004): the committee mean or *median* desired stance (majority rule), with the Chair entering additionally, and reports the Chair's implied share of committee influence (including the Chair's $1/N$ contribution to the aggregate). Standardized on the estimation sample. The table reports point estimates; standard errors are omitted. Chappell et al. (2004) estimate the Chair at 40 to 50 percent of the voting weight for the Burns era; because Burns chaired an unusually centralized committee, that figure is an *upper bound*, and the shares here are to be read against the 60% recovered attention share of Table 9 and the 82% mention share of Board communication in the underlying news-attention panel.

TABLE D.3—CORRELATION BETWEEN SPEECH AND TRANSCRIPT MSS

Status	30 Days	60 Days	90 Days	120 Days	150 Days	180 Days
Chairs	0.142***	0.12***	0.123***	0.109***	0.082**	0.079**
Non-Chairs	0.089**	0.098***	0.102***	0.115***	0.116***	0.115***

Note: The table presents correlations between individual members' speech tone, averaged over the indicated window before each FOMC meeting, and the same member's tone in the meeting transcript. Rows split members by institutional status; columns are the length of the pre-meeting window. *, **, and *** indicate significance at the 10, 5, and 1 percent levels, respectively. The sample is the full FOMC transcript corpus, 1976–2020. Each cell pools the member-meeting observations with at least one speech in the window: as the window widens from 30 to 180 days, N runs from 621 to 842 for Chairs and 833 to 1,537 for Non-Chairs. Stars are from the two-sided test of zero Pearson correlation (Student-t with $N-2$ degrees of freedom).

TABLE D.4—INFLUENCE ORDERING AND OVER-ATTENTION GAP: SYMMETRIC SPECIFICATIONS

	Baseline	Reverse	Both raw	Both orth.
Influence ν_c	3.23	0.40	0.40	4.78
Influence ν_n	4.26	4.46	4.26	6.60
Over-attention gap $\kappa_c - \kappa_c^*$ (bits)	0.57	—	—	0.63

Note: The table re-estimates the influence weights and the over-attention gap of Table 9 under four allocations of the common component of committee meeting tone. *Baseline* (the paper's specification) regresses on raw Chair tone and Chair-orthogonalized non-Chair tone, assigning the common component entirely to the Chair; *Reverse* assigns it entirely to the non-Chair side (raw non-Chair, non-Chair-orthogonalized Chair); *Both raw* enters both series unorthogonalized, so each coefficient is a partial effect; *Both orth.* enters each series residualized on the other. Everything else—samples, real-time data-news controls, and HAC inference—is identical to Table 9. Under the reverse and both-raw allocations the recovered Chair inattention leaves the admissible range (the Chair decision loading is statistically indistinguishable from zero), so the gap is not defined (—).

D.5 Speeches and Subsequent Meeting Interventions

Tables D.5 through D.9 pair a member's public speech with the same member's interventions at a subsequent FOMC meeting, illustrating the foreshadowing quantified in Table D.3. The five pairs are drawn from across the transcript sample (1976–2020) and cover five different members; in each, the left column reproduces stance-bearing passages of the speech and the right column the member's remarks at that subsequent meeting.

TABLE D.5—WALLICH, SPEECH ON 23 APRIL 1979 AND MEETING ON 11 JULY 1979

Speech	Transcript
...inflation has once more accelerated into the double-digit range	Here we have 13 percent CPI inflation...
We delude ourselves if we believe that a cure [for inflation] will not be painful and costly.	Bringing inflation down under these deteriorating circumstances is not just a question of taking more time to do it but it will be more costly...
Allow for somewhat greater slack in the economy in the form of more excess capacity and unemployment...	And the cost of that real income means less growth, higher unemployment.
...reject speculative calculations of the amount of output that we would be sacrificing by such a policy [against inflation].	...it [the cost of bringing inflation down] does mean less output.
We should pursue a strong anti-inflationary monetary policy...	...my concern is more that we should achieve some lasting effect on inflation... less willingness this time around to accommodate... monetarily.

Note: Left: stance-bearing passages of the speech; right: the member's interventions at the subsequent FOMC meeting. Bracketed text is the author's clarification.

TABLE D.6—PRESTON MARTIN, SPEECH ON 12 APRIL 1984 AND MEETING ON 17 JULY 1984

Speech	Transcript
On the other hand, the apparent vigor of rising intentions in 1984 to put plant and equipment in place pose a real question... the durability and depth of business fixed investment and R&D spending... These factors have contributed to a strong expansion in some types of investment...	There has been a good deal of investment in business equipment... Business fixed investment coming on stream... it is equipping; it is retrofitting; it is productivity oriented...
...management and labor alike have turned their efforts... toward... more flexible work rules and hiring practices... Union leadership... have gone beyond concessions... to changes in work rules... Nonunion labor has become more important... in construction.	Labor and management have negotiated quite a few work rule changes... Construction union? Not only the threat, but the utilization by developers of non-union workers.
...policymakers are faced with a basic dilemma... objective... to encourage further economic expansion... On the other hand, the apparent vigor... pose a real question... whether financing... will ignite boom-expectations and a market perspective and behavior leading to significantly higher rates of inflation.	That always raises the question of whether we try to do too much in too short a period of time... Is this expansion too fast in the sense of being inflationary?

Note: Left: stance-bearing passages of the speech; right: the member's interventions at the subsequent FOMC meeting.

TABLE D.7—MEYER, SPEECH ON 8 SEPTEMBER 1996 AND MEETING ON 17 DECEMBER 1996

Speech	Transcript
First, in the absence of policy adjustment, is the economy slowing or likely to slow to trend quickly enough to stabilize the unemployment rate at its current level?	Economic growth has now slowed to trend.
All [private sector forecasts] have growth near 2% in the second half of 1996 and through 1997.	On a fourth-quarter-to-fourth-quarter basis, the range for next year is from a low of 2.0 percent to a high of 2.4 percent.
Second, is the current unemployment already so low that remaining at this level would trigger a steady increase in the rate of inflation?	The unemployment rate remains locked in a low but narrow band.
Are we already below NAIRU?	A second justification for policy change would be the conviction that we are already below NAIRU...
A further decline in the unemployment rate, from its current level, would, in turn, increase the risk of an acceleration of inflation.	A second justification for policy change would be the conviction that we are already below NAIRU and not likely to move back to it quickly enough to prevent an uptick in inflation.
...stable to declining inflation over the past two years.	I read core measures of inflation as stable to still declining.
...a prudent central bank must exercise heightened surveillance of the inflationary risks and stand ready to respond if necessary.	Nevertheless, the major risk seems to be a rise in the inflation rate...

Note: Left: stance-bearing passages of the speech; right: the member's interventions at the subsequent FOMC meeting. Bracketed text is the author's clarification.

TABLE D.8—KOHN, SPEECH ON 19 OCTOBER 2005 AND MEETING ON 1 NOVEMBER 2005

Speech	Transcript
In sum, I see risks on both sides of my expectations that the growth of economic activity will slow modestly on balance over the next year or so, leaving the economy producing at about its sustainable potential.	Still, conditions do seem to be in place for a moderation in growth over the intermediate term.
Moreover, according to the employment cost index, which is based on a survey of firms, and to the growth in average hourly earnings in the payroll employment report, compensation pressures remained quite subdued.	Business labor costs are probably not putting much upward pressure on prices.
Measures of core consumer inflation eased some over the late spring and early summer, even as headline inflation was pushed higher by rising energy prices.	Core inflation has been damped, despite substantial increases in energy prices before the hurricanes.
...persistently high energy prices also raise the possibility that inflation expectations could ratchet higher, even if core inflation remains fairly low in the short run.	...higher energy prices do threaten to feed through to core prices and inflation expectations...
So far, market-based measures of required inflation compensation beyond the next few years have changed little on balance since late August.	...long-run forward measures of inflation compensation in the market have risen only modestly...

Note: Left: stance-bearing passages of the speech; right: the member's interventions at the subsequent FOMC meeting.

TABLE D.9—BRAINARD, SPEECH ON 7 MARCH 2016 AND MEETING ON 16 MARCH 2016

Speech	Transcript
I am pleased with the continued strength in the U.S. labor market... In February, the unemployment rate was 4.9 percent...	The labor market continued to strengthen last month, ... the unemployment rate holding steady at 4.9 percent...
The prime-age employment-to-population ratio remains 1-¾ percentage points below levels prevailing prior to the financial crisis...	The employment-to-population ratio for prime-age workers remains 1¾ percentage points below pre-crisis levels.
...a relatively large share of employees who are working part time would prefer to work full time. In addition, wage growth remains relatively slow.	Involuntary part-time employment is still elevated, and wage growth has shown little sign of acceleration.
And there likely is some room to go...	...there may be reasons to believe there is some room for the labor market to improve still further.
Prices for personal consumption expenditures (PCE) have increased 1.3 percent over the 12 months through January...	The 12-month change in headline PCE inflation moved up to 1.3 percent in January...
...core PCE inflation has been stubbornly stuck in the vicinity of 1-¼ to 1-½ percent since 2012.	...core inflation has remained stuck in the vicinity of 1¼ to 1½ percent for the past three years.
In this regard, core PCE inflation increased 1.7 percent over the 12 months ending in January...	...core PCE inflation moved up to 1.7 percent. [Context implies January data]
However, households' inflation expectations appear to have moved down somewhat recently. Five-to-10-year inflation expectations in the University of Michigan Surveys of Consumers... the level in February nearly ½ percentage point below the 10-year average.	...household expectations have moved lower. In February, 5-to-10-year-ahead inflation expectations in the Michigan survey were 2.5 percent, ½ percentage point below the series longer-run average.
Notably, market-based measures of inflation compensation... At the five-year, five-year-ahead horizon, inflation compensation is 1 percentage point lower than mid-2014 levels. Declines in swap-based measures... have been similar.	More notably, both the TIPS- and swaps-based measures of inflation compensation are down about 1 percentage point from mid-2014 and remain at very low levels...
Domestic activity continues to grow at a moderate pace.	...the recent data point to continued moderate growth on balance.
The pace of consumer spending, after slowing some at the end of last year, looks to have picked up in January... I expect to see growth to continue... based on solid job and income growth...	The pace of consumer spending looks to have slowed a bit recently, but solid job and income growth... should support consumption growth.
The housing sector has also contributed steadily to growth over the past year.	And the housing sector has also contributed steadily to growth over the past year...
In contrast, sectors of the economy that are sensitive to energy prices or international demand have been a drag.	In contrast, sectors of the economy exposed to the strength of the dollar or the fall in energy prices have struggled.
...net exports subtracted a little more than ½ percentage point from GDP growth in 2014 and 2015...	Net exports subtracted about a ½ percentage point from GDP growth in each of the past two years...
More broadly, sources of robust demand around the globe are few, and sources of weakness relatively greater...	More broadly, sources of economic strength across the globe are few, and sources of weakness relatively greater.

Note: Left: stance-bearing passages of the speech; right: the member's interventions at the subsequent FOMC meeting. Bracketed text is the author's clarification.

References

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.

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

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.

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.

Gorodnichenko, Yuriy, Tho Pham, and Oleksandr Talavera. 2023. "The Voice of Monetary Policy." *American Economic Review* 113 (2): 548–584. 10.1257/aer.20220129.

Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright. 2007. "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.

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.

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.

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

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