

VOTING RIGHTS AND THE VOLATILITY CONTENT OF FEDERAL RESERVE COMMUNICATION

Fuli Yang¹ | Tingting Qiu^{1,2} | Jing Huang¹

¹ GARCH Institute, Harvard Square, Cambridge, MA 02138, USA

² Fujian Vocational College of Agriculture, Fuzhou, Fujian, China

Corresponding author: Fuli Yang (dr@garch.ac) | SSRN working paper | 7 September 2026

Abstract

Formal voting authority can matter even when it does not visibly change what a central banker says. We exploit the mechanical annual rotation of voting rights among Federal Reserve Bank presidents and ask whether otherwise comparable public language receives a different loading in a persistent component of U.S. equity variance when it carries voting authority. The sample contains 819 eligible speeches by 21 presidential officeholders over 2015-2023 and 71 scheduled meeting cycles. Using an archived Yahoo Finance S&P 500 price tape and strictly lagged realized variance, the interaction between uncertainty/modal communication and statutory voting exposure is negative ($\delta = -0.461$; likelihood-ratio $p = 0.0007$). Meeting-cycle clustering yields $p = 0.037$, while a 9,999-draw restricted wild-score bootstrap yields $p = 0.084$ two-sided and $p = 0.042$ for the prespecified negative direction. The estimate is unchanged in a fixed equal-weight 66-day benchmark and remains negative across 44, 66, 88, and 132 return lags. A 500-draw random-group placebo reaches the observed absolute interaction in 4.6 percent of draws. The interaction is null at individual speech dates. A zero-inclusive selection audit finds no extensive-margin or timing shift, although voting officeholders give about 0.19 more speeches per cycle (wild-cluster $p = 0.075$). Voting status does not robustly change language under equal-weight officeholder-year inference, and rotating-district unemployment does not explain collective statements. The contribution is therefore an aggregation-specific authorization result with borderline finite-cluster inference, not a claim that rotation mechanically rewrites the Committee's statement.

JEL Classification: E52; E58; G17; C58

Keywords: FOMC; voting-right rotation; central bank communication; GARCH-MIDAS; long-run variance; textual uncertainty

1. Introduction

Macroeconomic information does not arrive with an official interpretation attached. The U.S. employment report released on 4 September 2026 offered a timely illustration: a strong payroll headline, revisions to prior months, a 4.1 percent unemployment rate, and moderating wage growth could support more than one policy narrative. The Federal Open Market Committee's scheduled statement on 16 September will transform the public data into an authorized committee message. That

episode motivates the paper, but it does not identify the paper's effect. The communications blackout begins immediately after the data release, preventing a clean post-release comparison of voting and nonvoting speakers.

The empirical question is narrower and more general. Does the formal authority attached to a speaker's language change the portion of that language that markets carry into the long-run component of equity variance? A standard communication regression treats text as if it were written by a single, institution-level author. The FOMC is instead a committee with a mechanically changing subset of Reserve Bank presidents who hold voting rights. This institutional feature allows the same president to be observed in voting and nonvoting years while permanent members, national conditions, and the committee calendar are separated from the rotation treatment.

The paper's estimand is delta: the incremental long-run-variance loading associated with communication carrying formal voting authority. It is not the average tone coefficient, a federal funds rate response, or a same-day event-window reaction. The distinction matters because three adjacent literatures already occupy those objects. Fos and Xu use rotation to study policy-rate decisions and their pricing. Ehrmann, Tietz, and Visser study how voting status changes speech behavior and daily Treasury-market reactions. Riboni and Ruge-Murcia study turnover, disagreement, and policy inertia. Tsang and Yang recover hidden dissent from transcripts and minutes. This paper instead asks whether voting authority changes the horizon-specific variance content of public language.

Table 1. Contribution boundary

Object	Existing literature	This paper
Policy choice	Rotation and the federal funds target	Not an outcome
Daily yield reaction	Vote discount around speeches	Not the main horizon
Internal disagreement	Turnover, dissent, transcripts, and minutes	Used only as context or controls
Long-run equity variance	Institution-level language loadings	Authority-contingent loading, delta

The results imply a specific aggregation pattern. At the scheduled-meeting-cycle horizon, uncertainty/modal language associated with voting presidents receives a lower long-run variance loading than comparable language associated with nonvoting presidents. At individual speech dates, the interaction is close to zero. The finding is therefore not that every voting speech is stabilizing. Formal authority becomes informative only when intermeeting communication is organized as the information set leading into the next scheduled decision.

Two negative results discipline this interpretation. First, voting status does not robustly change regional-language density, national-language density, sentiment, or hawkishness under equal speaker-year inference. Second, the unemployment composition of the four rotating voting districts does not explain the uncertainty/modal content of collective statements. The paper therefore rejects its own stronger writing channel and does not estimate a generated-regressor mediation model.

2. Institutional setting and related research

The FOMC includes the members of the Board of Governors, the president of the Federal Reserve Bank of New York, and four of the remaining eleven Reserve Bank presidents. The four rotating seats are allocated across statutory groups on an annual schedule. Nonvoting presidents continue to attend meetings and contribute to deliberations, so rotation changes formal voting authority without excluding presidents from the information set or the discussion. This creates repeated within-president variation in a publicly observable institutional status. The broader communication literature emphasizes that public language can convey both policy guidance and information about economic conditions (Blinder et al., 2008; Hansen and McMahon, 2016), while automated FOMC text measures show that statements contain information beyond the contemporaneous policy rate (Lucca and Trebbi, 2009).

Fos and Xu (2025) use 472 meetings from 1969-2019 to show that local conditions in voting districts affect the federal funds target and are priced before meetings. Their endpoint is the policy rate. Ehrmann, Tietz, and Visser (2026) find that presidents' speeches move Treasury markets less in voting years and relate this vote discount to regional content. Their endpoint is a daily market response. Riboni and Ruge-Murcia (2025) model changes in the committee median and show that disagreement is associated with dissents and policy inertia, while turnover does not raise the volatility of rate adjustments. Tsang and Yang (2025) infer hidden dissent from transcripts and minutes and show that its disclosure affects stocks, volatility expectations, and bond markets.

This paper is complementary rather than substitutive. It takes neither the policy rate nor the announcement-window change as its dependent variable. It uses a long-run variance component and asks whether the loading of language itself depends on formal authority. It thereby extends Yang, Qiu, and Huang (2026), who document that Federal Reserve text enters long-run variance but treat the Committee as a single author.

3. Empirical design

3.1 Frozen estimand and communication units

Let X_j denote a frozen textual score for communication j , V_j an indicator that the Reserve Bank president holds a statutory vote in that year, and s_t the share of communications in the relevant information set delivered by voting presidents. For a meeting cycle, the design aggregates all eligible speeches mapped to the next scheduled meeting. It separately constructs the mean language score, $C_t = \text{mean}(X_j)$, and the mean authority-weighted language score, $A_t = \text{mean}(X_j V_j)$. This distinction avoids replacing the average of individual interactions with the product of two aggregate averages.

The statutory annual schedule is the primary authority assignment. Realized participation, audited from each statement's voter paragraph, is a robustness treatment because it can reflect vacancies, alternates, or other meeting-specific facts known only after the schedule is set.

3.2 GARCH-MIDAS variance equation

Daily returns are modeled with a short-run GARCH(1,1) component and a multiplicative long-run component, following the mixed-frequency volatility literature (Ghysels, Santa-Clara, and Valkanov, 2006; Ghysels, Sinko, and Valkanov, 2007) and the GARCH-MIDAS application of Engle, Ghysels, and Sohn (2013). The working long-run equation is

$$\log(\tau_{it}) = m + \theta_{RV} \log[RV_t(\omega)] + \theta_X C_{it} + \gamma s_{it} + \delta A_{it}.$$

RV_t uses only lagged squared returns. The MIDAS kernel is a one-parameter Beta(1, ω) weight over 66 trading-day lags. Each communication becomes effective on the next trading day, preventing the text measure from using a return that occurred before the language was public. The null model omits A_{it} ; the likelihood-ratio test therefore evaluates one frozen interaction parameter.

Important technical interpretation. Across the primary fits, ω is approximately 1.000. The estimated kernel is therefore close to equal weighting over 66 lags. The draft interprets δ as an authority difference in a roughly quarterly long-run variance component; it does not claim to identify a sharply declining MIDAS profile.

3.3 Identification and inferential limits

Mechanical rotation supplies predetermined variation in V_j . It does not mechanically determine whether a president gives a speech, the topic selected, or how many speeches enter a meeting cycle. Accordingly, the present estimate is best described as an authority-contingent loading under observed speech selection, not as an unconditional causal effect of assigning a vote to an otherwise fixed communication stream. Speaker fixed effects in the expression analysis address stable speaker heterogeneity; the variance model relies on the rotation schedule, identical text measurement, and the frozen timing rule. A separate selection audit forms officeholder-by-meeting-cycle cells, retains zero-speech cycles within officeholder-years for which statutory status is observed, and estimates vote effects with officeholder and meeting fixed effects. This audit diagnoses, but cannot eliminate, endogenous topic and appearance choices.

The paper uses a gate structure. The expression gate asks whether status changes language. The authorization gate asks whether status changes the long-run loading of language. The statement-aggregation gate tests a stronger writing mechanism. Mediation is permitted only if the expression or statement-writing first stage is sufficiently strong. It is not.

4. Data and measurement

The speech discovery corpus is CBS Dataset v1.0, covering 2015-2023. After eligibility screening, the working sample contains 819 speeches, 98 officeholder-years, and 21 distinct presidential officeholders at 11 Reserve Banks. The officeholder count exceeds the number of Banks because presidencies turn over during the sample. The archived daily S&P 500 price file is sourced from Yahoo Finance and spans

2000-2026; this paper uses observations from 2011 for filter initialization and estimates the communication model from 29 January 2015 through 29 December 2023, yielding 2,302 daily observations. On the 1,842 overlapping dates from September 2016 through December 2023, its closes match the archived FRED series exactly.

Text is scored using a frozen 2020 Loughran-McDonald dictionary reconstruction and a fixed hawk-dove index. The primary measure is the percentage of uncertainty plus modal terms. Construct checks separate Loughran-McDonald uncertainty from modal language, use net sentiment, and apply the fixed hawk-dove index. All scoring rules were frozen before the authorization run; alternative dictionaries, dates, and event definitions are not searched after observing the result.

The statement mechanism test covers 71 scheduled meetings from 2015-2023. It uses unemployment information dated at least two months before each statement. The primary mechanism diagnostic uses seasonally adjusted state LAUS series apportioned to Federal Reserve districts using the Kansas City Fed district-county crosswalk. A prespecified county-count confirmation is reported in Appendix A; it does not change the conclusion and is not used to select outcomes or weights.

4.1 Data provenance chain

The Yahoo Finance CSV is preserved byte-for-byte in the replication archive with SHA-256 hash 9c9791d9407ec1d6ff2814ee06552e3239a4d118e924c5f0d5445be377c70d3d6. It contains daily S&P 500 OHLCV observations from 3 January 2000 through 30 June 2026. The estimation input is generated directly from this file and has SHA-256 hash c16c06f9af26ed85da2816fa8568efe33ea74d308c12765c129321e7beccbc4c. On 1,842 overlapping dates with the archived FRED export, closing prices and log returns are identical to machine precision. The earlier submission packages did not include optimizer source code; this study therefore uses a self-contained implementation whose inputs, multi-start audit, estimation scripts, and result files are supplied together.

5. Results

5.1 Voting status and speech content

The equal-weight speaker-year specification provides no robust evidence that voting status changes the primary language score. For uncertainty/modal language, the voting coefficient is 0.049 percentage points with a speaker-clustered p-value of 0.395. A speech-count-weighted speaker-year model and a speech-level model are positive, but this evidence is not sufficient under the frozen equal-speaker inferential rule. Regional-term density, national-term density, net sentiment, and the fixed hawk-dove index are also null in the equal speaker-year models. The expression channel is therefore provisional rather than established.

5.1.1 Speech-selection audit

Table 2. Voting status and intermeeting speech supply

Outcome	Vote beta	Cluster SE	Cluster p	Wild p	N
Speech count per cycle	0.193	0.087	0.027	0.075	774
Any speech (linear probability)	0.052	0.051	0.314	0.343	774
Mean days to meeting speech	0.072	0.717	0.920	0.925	491

The zero-inclusive grid contains 774 officeholder-cycle cells with observable annual statutory status. Conditional on officeholder and meeting fixed effects, voting status does not significantly change the probability of giving any speech (beta = 0.052; wild-cluster p = 0.343) or the mean distance between a speech and the next meeting among speaking cells (beta = 0.072 days; p = 0.925). It is associated with 0.193 additional speeches per cycle. Conventional officeholder-clustered inference gives p = 0.027, but a 9,999-draw restricted wild-cluster t bootstrap across 21 officeholders gives p = 0.075. The evidence therefore suggests a possible intensive-margin supply response, while providing no evidence of an extensive-margin or timing response. Topic choice remains unobserved selection.

5.2 Formal authority and long-run variance

Table 3. Primary uncertainty/modal authorization estimates

Aggregation	delta	SE	LR p	theta nonvote	theta vote	Events
Speech-event dates	0.027	0.069	0.6933	-0.059	-0.032	629
Scheduled meeting cycles	-0.461	0.137	0.0007	0.260	-0.201	71
Realized-voter audit	-0.551	0.136	0.0001	0.232	-0.319	71

At the scheduled-meeting-cycle horizon, the statutory assignment yields delta = -0.461 (Hessian SE = 0.137; LR p = 0.0007). The implied text loading is positive for nonvoting communication (theta = 0.260) and negative for voting communication (theta = -0.201). Replacing the statutory schedule with realized meeting participation strengthens the contrast: delta = -0.551 (SE = 0.136; LR p < 0.0001), with theta = 0.232 for nonvoting communication and theta = -0.319 for voting communication.

All 40 finite solutions converge within 0.0001 negative-log-likelihood units of the best solution for each mapping. This rules out a simple local-optimum explanation for the sign and magnitude of delta within the archived implementation.

5.3 Aggregation, not an event-date effect

At the individual speech-date frequency, delta is 0.027 (SE = 0.069; LR p = 0.693). The implied nonvoting and voting loadings are both small and negative. Thus, the authorization result is not a generic effect attached to the timestamp of each speech. It appears only after the eligible communication set is aggregated toward the next scheduled decision.

This null is informative. It rules out the stronger story that markets place a different persistent-variance weight on every utterance merely because the speaker votes that year. The supported object is the authorized information set surrounding a committee decision. Because speech selection and

aggregation remain observational, this pattern is consistent with, but does not prove, a committee-level information-integration mechanism.

5.3.1 Random-group aggregation placebo

To address the possibility that aggregation merely compresses noise, we randomly permute the observed communication-mass profile across 71 contiguous chronological groups, preserving the number of groups, strict timing, and the underlying speech-level votes and scores. Across 500 fixed-seed draws, the mean placebo interaction is -0.157 (SD 0.183); 4.6% of draws reach an absolute interaction at least as large as the observed statutory-cycle estimate. The placebo distribution is therefore centered near zero relative to the observed -0.461 contrast. This randomization addresses the aggregation-only alternative; finite-cluster inference is evaluated separately with the wild-score bootstrap.

Placebo statistic	Value
Observed statutory-cycle delta	-0.461
Random-group mean delta	-0.157
Random-group SD	0.183
Share $ \text{delta} \geq \text{observed}$	4.6%
Share $\text{delta} < 0$	81.4%

5.4 Inference and lag-length stability

Table 4. Primary inference and lag-length audit

Audit	K	delta	SE	p	Basis
observed Hessian	66	-0.461	0.137	0.0008	Inference
meeting-cycle cluster sandwich	66	-0.461	0.222	0.0373	Inference
meeting-cycle HAC(2) sandwich	66	-0.461	0.196	0.0183	Inference
restricted wild-score bootstrap	66	-0.461	0.267	0.0837	9,999 draws
Lag length K = 44	44	-0.558	0.170	0.0009	LR test
Lag length K = 66	66	-0.461	0.137	0.0007	LR test
Lag length K = 88	88	-0.451	0.137	0.0009	LR test
Lag length K = 132	132	-0.439	0.140	0.0016	LR test

Meeting-cycle clustering raises the standard error from 0.137 to 0.222 and yields $p = 0.037$; a meeting-cycle HAC(2) sandwich yields $p = 0.018$. The more conservative restricted Rademacher wild-score bootstrap, based on 9,999 draws over 71 meeting-cycle clusters, yields a two-sided p-value of 0.084. Its lower-tail p-value is 0.042 for the prespecified negative direction. Finite-cluster evidence is therefore directional but borderline under a two-sided test. Across the frozen K set of 44, 66, 88, and 132 trading days, delta ranges from -0.558 to -0.439 and every likelihood-ratio p-value is below 0.002.

5.4.1 Equal-weight long-run variance benchmark

Because the estimated MIDAS shape is at the equal-weight boundary, we fix the 66 lag weights to $1/66$ and re-estimate the short-run GARCH and all language coefficients. The benchmark produces $\text{delta} = -0.461$ (LR $p = 0.0007$), numerically identical to the free-kernel estimate. Its log likelihood differs from the free MIDAS model by less than 0.000001, while its AIC is 5871.2, compared with 5873.2 for the

extra-parameter model. The authorization contrast is therefore not created by an estimated MIDAS decay profile. The persistent horizon matters; the flexible kernel does not add fit in this sample.

5.5 Construct dependence

Table 5. Scheduled-cycle construct checks

Measure	delta	LR p	Relative to primary
Uncertainty + modal	-0.461	0.0007	Same sign
LM uncertainty	-0.273	0.0262	Same sign
Modal terms	-0.455	0.0009	Same sign
LM net sentiment	0.004	0.9697	Opposite/null
Fixed hawk-dove index	0.204	0.0824	Opposite/null

The primary result decomposes cleanly into negative uncertainty and modal interactions. Net sentiment is null, while the hawk-dove index reverses direction. Voting authority therefore does not have a text-invariant effect. The paper's claim is restricted to the uncertainty/modal construct and treats measurement sensitivity as substantive evidence about which dimension of language markets encode.

5.6 The statement-writing mechanism does not pass

Table 6. State-apportioned voting-district unemployment mix and collective statement text

Statement measure	beta	HC3 SE	p	Holm p	Rotation p
Uncertainty + modal	0.604	0.463	0.192	0.383	0.681
LM net sentiment	-0.449	0.394	0.254	0.383	0.785
Fixed hawk-dove index	-0.359	0.220	0.103	0.308	0.095

The unemployment gap of the four rotating voting districts does not explain the primary statement score. The uncertainty/modal coefficient is 0.604 (HC3 p = 0.192; Holm p = 0.383), and its p-value under 432 legal rotation assignments is 0.681. Neither equal-seat versus labor-force weighting nor land-share versus county-count apportionment changes the conclusion. The paper therefore does not claim that voting regions write the collective statement and does not run the proposed generated-regressor mediation model.

Table 7. County-level geography confirmation

Statement measure	beta	HC3 SE	p	Holm p	Rotation p
Uncertainty + modal	0.020	0.365	0.956	1.000	0.988
LM net sentiment	0.070	0.322	0.828	1.000	0.956
Fixed hawk-dove index	-0.326	0.169	0.054	0.161	0.074

As a geographic confirmation, we replace state apportionment with BLS LAUS county-level unemployed and labor-force counts merged to the Kansas City Fed county crosswalk. Under the same t-2 information rule and equal-seat mix, the uncertainty/modal coefficient is 0.020 (HC3 p = 0.956; 432-assignment rotation p = 0.988). The hawk-dove coefficient is -0.326 (HC3 p = 0.054; rotation p = 0.074; Holm p = 0.161). County-level measurement therefore confirms the null for the primary uncertainty/modal

construct; the marginal hawk-dove pattern is not robust to the prespecified rotation randomization or multiple-outcome adjustment.

6. Interpretation

The supported result is an authorization discount in long-run variance. Formal voting status changes the market meaning of a measured language dimension even though the same status does not robustly alter average language under the most conservative speaker-year inference. This separation is economically plausible: market participants may condition on who is institutionally empowered to convert views into policy, even when lexical differences are small.

The negative cycle-level interaction is also consistent with authority reducing unresolved policy uncertainty. Language from nonvoting presidents may be read as an additional, potentially competing signal about the committee's future path. Once similar language is attached to members who will vote at the next scheduled decision, part of that uncertainty may be absorbed into a more bounded policy interpretation. This is a market-interpretation account, not a demonstrated communication-strategy mechanism.

The event-date null prevents a simple speaker-label story. Voting status does not alter the persistent-variance loading of each speech as it arrives; the contrast emerges only in the communication set leading into a scheduled decision. The object is therefore not a constant tone coefficient but an aggregation- and institution-dependent loading.

7. Limitations and remaining checks

- Market-data provenance. The archived Yahoo Finance price file and transformed estimation input are hashed, and their overlap with FRED is exact. The present self-contained implementation, rather than an earlier unsupplied optimizer, is the reproducible object evaluated in this paper.
- Speech selection. The zero-inclusive fixed-effects audit finds no extensive-margin or timing effect, but it suggests 0.193 additional speeches per cycle in voting years (cluster $p = 0.027$; wild-cluster $p = 0.075$). Rotation may therefore affect communication intensity, and topic choice remains unobserved. The variance estimate is conditional on the resulting communication set.
- Kernel boundary. At $K = 66$ the MIDAS shape parameter is at its lower bound, producing almost equal weights. The interaction remains negative across the frozen K set, but the data do not identify a sharp decay profile.
- Inference. Meeting-cycle sandwich and HAC inference reject at 5 percent, but the 9,999-draw restricted wild-score bootstrap gives $p = 0.084$ two-sided and $p = 0.042$ for the prespecified negative direction. The paper treats this as borderline finite-cluster evidence.
- Aggregation placebo. The 500-draw random-group placebo yields an empirical two-sided tail share of 4.6%; retain the fixed seed and disclose that the placebo randomizes contiguous group lengths rather than speaker selection.

- **Benchmark.** The fixed equal-weight 66-day model reproduces delta and slightly improves AIC by dropping the unidentified shape parameter. The paper therefore does not claim that a flexible MIDAS decay is empirically necessary.
- **Measurement.** The result is concentrated in uncertainty/modal language and reverses for the hawk-dove index. No universal “Fed text” parameter should be claimed.
- **Geography.** The county-level BLS LAUS confirmation preserves the primary uncertainty/modal null (HC3 $p = 0.956$; rotation $p = 0.988$). A marginal hawk-dove coefficient is not robust to rotation randomization or Holm adjustment. Geography therefore does not reopen the statement-writing mechanism.

8. Conclusion

Mechanical rotation of FOMC voting rights provides more than a source of variation in policy preferences. It also indicates that the estimated variance content of public language depends on formal authority. In the present sample, uncertainty/modal language associated with voting presidents receives a lower loading in the persistent component of equity variance when communications are aggregated by scheduled meeting cycle. The interaction is null at individual speech dates, making committee-level aggregation central to the empirical pattern. The two-sided wild-score p-value of 0.084 nevertheless requires the magnitude to be interpreted as directional rather than definitive finite-cluster evidence.

The data do not support the stronger claim that rotating voters mechanically rewrite collective statements. Nor do they establish a broad change in what presidents say. The contribution is narrower: the estimated loading of the same textual dimension varies with the authority and aggregation structure behind it. In that limited sense, the volatility content of Federal Reserve communication is a committee object rather than an institutional constant.

Declarations

Declaration of interest

The authors declare no financial or personal relationships that could have influenced the research reported in this manuscript.

Funding

This research received no external funding.

Ethics approval

Ethics approval was not required. The study uses publicly available aggregate market, macroeconomic, and central-bank communication data and does not involve human participants or identifiable personal data.

Data and code availability

The replication archive supplies the archived market file, transformed inputs, source code, estimation outputs, county-level LAUS confirmation, and SHA-256 hashes. Public source files are identified by their originating agency and retrieval URL in the archive README. No confidential or restricted data were used.

Generative AI disclosure

Generative AI tools were used for language editing, code orchestration, and document quality assurance. The authors designed the study, verified all sources and computations, and take responsibility for the final content.

References

- Blinder, A. S., Ehrmann, M., Fratzscher, M., De Haan, J., and Jansen, D.-J. (2008). Central bank communication and monetary policy: A survey of theory and evidence. *Journal of Economic Literature*, 46(4), 910-945. <https://doi.org/10.1257/jel.46.4.910>
- Engle, R. F., Ghysels, E., and Sohn, B. (2013). Stock market volatility and macroeconomic fundamentals. *Review of Economics and Statistics*, 95(3), 776-797. https://doi.org/10.1162/REST_a_00300
- Ehrmann, M., Tietz, R., and Visser, B. (2026). Voting right rotation, speeches, and financial market reactions: Evidence from the U.S. Federal Open Market Committee. *Journal of International Money and Finance*, 165, 103574. <https://doi.org/10.1016/j.jimonfin.2026.103574>
- Fos, V., and Xu, N. R. (2025). When Do FOMC Voting Rights Affect Monetary Policy? NBER Working Paper No. 33762.
- Ghysels, E., Santa-Clara, P., and Valkanov, R. (2006). Predicting volatility: Getting the most out of return data sampled at different frequencies. *Journal of Econometrics*, 131(1-2), 59-95. <https://doi.org/10.1016/j.jeconom.2005.01.003>
- Ghysels, E., Sinko, A., and Valkanov, R. (2007). MIDAS regressions: Further results and new directions. *Econometric Reviews*, 26(1), 53-90. <https://doi.org/10.1080/07474930600972467>
- Hansen, S., and McMahon, M. (2016). Shocking language: Understanding the macroeconomic effects of central bank communication. *Journal of International Economics*, 99, S114-S133. <https://doi.org/10.1016/j.jinteco.2015.12.008>
- Lucca, D. O., and Trebbi, F. (2009). Measuring central bank communication: An automated approach with application to FOMC statements. NBER Working Paper No. 15367.
- Riboni, A., and Ruge-Murcia, F. (2025). Membership Turnover and Policy Disagreement at the FOMC. CEPR Discussion Paper No. 20613. Revised June 2026.
- Tsang, K. P., and Yang, Z. (2025). Agree to disagree: Measuring hidden dissent in FOMC meetings. *Journal of Economic Dynamics and Control*, 180, 105197. <https://doi.org/10.1016/j.jedc.2025.105197>
- Yang, F., Qiu, T., and Huang, J. (2026). Central Bank Uncertainty and Long-Run Financial Volatility: Evidence from FOMC Communication and a GARCH-MIDAS Framework. Working paper.

Author information and contributions

Fuli Yang (ORCID 0009-0005-4055-7561) is the corresponding author. Tingting Qiu: ORCID 0009-0009-2442-5621. Jing Huang: ORCID 0009-0004-8120-1652.

CRedit statement. Fuli Yang: Conceptualization, Methodology, Software, Formal analysis, Data curation, Writing - original draft, Visualization, Project administration. Tingting Qiu: Validation, Investigation, Writing - review and editing. Jing Huang: Validation, Investigation, Writing - review and editing.

Appendix A. Robustness and claim boundaries

Table A1. Empirical scope of the reported evidence

Component	Evidence	Interpretive boundary
Speech expression	Equal-speaker result is null	No general claim that voting rights change average wording
Speech selection	No extensive-margin/timing shift; count wild $p = 0.075$	Possible intensive-margin response; topic selection remains
Meeting-cycle authorization	$\delta = -0.461$; cluster $p = 0.037$; wild $p = 0.084$	Directional evidence; borderline two-sided finite-cluster inference
Random grouping	500-draw tail share = 4.6%	Rejects a simple aggregation-only account within the stated placebo
Equal-weight benchmark	δ unchanged; lower AIC	Persistent horizon matters; flexible MIDAS decay is unnecessary
Event-date authorization	Interaction is null	No per-speech authorization effect
Statement-writing channel	Rotation test is null	No claim that voting-district unemployment changes statement text
County geography confirmation	Uncertainty/modal $p = 0.956$; rotation $p = 0.988$	Confirms the primary state-apportionment null; hawk-dove marginal pattern is non-robust
Mediation	Not estimated	No vote-predicted-text mechanism claim
Market provenance	Archived tape and transformed input are hashed	Current pipeline is reproducible; earlier optimizer is not reconstructed

Appendix B. Reproducibility details

- The replication archive contains the archived Yahoo price file, transformed daily input, estimation code, result files, and SHA-256 hashes.
- The statutory rotation schedule is the primary treatment; realized meeting participation is reported as an audit under the identical timing rule.
- The archive includes the 9,999-draw restricted wild-score bootstrap, the 500-draw random-group placebo, the fixed equal-weight benchmark, and the frozen $K = 44, 66, 88, \text{ and } 132$ sensitivity set.
- The speech-selection audit supplies the zero-inclusive officeholder-cycle grid, fixed-effects estimates for speech count, any speech, and conditional timing, and 9,999-draw restricted wild-cluster t tests across 21 officeholders.
- The county geography confirmation uses the BLS LAUS county file `la.data.64.County` (SHA-256 50502ab8d3a53c0b74e3cf7684dce5c2d4a51903ff89924b54c884a7427babc1) and the fixed Kansas City Fed crosswalk under the same $t-2$ and equal-seat rules.

Appendix C. Prospective September 2026 illustration

The 16 September 2026 statement is reserved as a prospective descriptive observation. Before release, the frozen question is whether the Committee translates the 4 September employment report into less hedging or more hawkish language. The statement will be scored with the existing dictionaries and compared with the immediately preceding scheduled statement. It will not be added to the confirmatory 2015-2023 discovery sample, used to choose a construct, or presented as a causal test.