

Decoding Central Bank Communication: A Transformer-Based Analysis of Monetary Policy Stance

Virginie Gautier
AFSE 2026 – June 18th




Introduction

Why does Fed communication matter?

- **Transparency paradigm:** Since 1994, the Fed's systematic disclosure of decisions has shifted communication from a mere byproduct to a primary tool for anchoring expectations.
- **Massive volume:** A pluralistic strategy combining official Board statements with a constant flow of individual speeches from Governors and Regional Presidents.
- **Semantic complexity:** The sheer frequency and diversity of interventions paradoxically increase the "noise," making it harder to extract a single, clear policy signal.

Number of monetary policy related speeches and press releases



This paper proposes an automated tool that identifies the global signal and isolates individual policymaker positions to measure internal consensus or “cacophony”.

The FOMC: Structure and communication

- The FOMC has 12 voting members deciding policy rates at 8 regular meetings per year (simple majority):
 - 7 Board of Governors (permanent voting rights)
 - President of the Federal Reserve Bank of New York (permanent voting right)
 - 4 seats rotating annually among 11 regional bank presidents

Communication channel	Issuer	Nature of the signal	Strategic role
Press release	FOMC	Official & decisive	Direct announcement of policy decisions; highly standardized and concise.
Speech	Individual members	Exploratory & diverse	High-granularity signals; reflects personal views and forward guidance.
Testimony	Fed Chair	Institutional & political	Democratic accountability; justification of the Fed's mandate before Congress.

From word counts to Transformers: A literature evolution

Dictionary-based approaches: Loughran & McDonald (2011), Bordo & Istrefi (2018)

⇒ *Unable to handle negation or semantic shift*

Latent Dirichlet Allocation (LDA): Hansen & McMahon (2016)

⇒ *Identifies what the CB talks about, but not the direction of its intent*

The Transformer revolution: Vaswani et al. (2017)

⇒ *Bidirectional self-attention enables dynamic, context-aware representations; resolves polysemy*

- CentralBankRoBERTa (Pfeifer & Marohl, 2023): Domain fine-tuned on documents from 21 central banks; multi-class classification by economic agent within a sentence.
- Trillion Dollar Words (Shah et al., 2023): Expert-labeled sentences (hawkish/neutral/dovish); RoBERTa-large outperforms BERT and DistilBERT on F1-score.
- MILA, Bundesbank (Geiger et al., 2025): Prompt-chaining to distinguish intensity (moderate vs. strong signals) based on presence of explicit risk justifications.

Our contribution: fine-tune RoBERTa-small on augmented Shah et al. labels; aggregate to Board- and voter-level indices; integrate into an augmented Taylor rule.

Transformer architecture & BERT


Transformer architecture & BERT

- Transformer-based models are pre-trained on masked word prediction and next-sentence prediction
- BERT belongs to the encoder-only family of Transformers. It reads the full sentence bidirectionally (each word interpreted in light of all surrounding words), resolving negation, polysemy, and hedging. It is particularly suited for classification tasks like stance detection.
- RoBERTa is an improved variant trained with more data and epochs (better robustness for textual analysis).

How it works:


- Input embedding: text split into tokens; each token becomes a dense vector encoding meaning and position.
- N encoder layers: multi-head self-attention + feedforward layer + residual connections & layer norm (stacked N times)
- Output: [CLS]token embedding (sentence-level representation) $\rightarrow$ linear + softmax $\rightarrow$ probabilities (dovish/neutral/hawkish)

Transformer encoder




Methodology

From text to stance index



From text to stance index



$$S_{Doc} = \frac{1}{n} \times \sum_{i=1}^n S_{Doc,i}$$

where n is the number of relevant sentences within a single document

$$i_t = r_t^* + \phi_\pi(\pi_t - \pi_t^*) + \phi_u(u_t - u_t^*) + \phi_S S_{Board,t}$$

where i_t is the nominal policy rate, r_t^* the equilibrium real interest rate, π_t the inflation rate, π_t^* the inflation target and $u_t - u_t^*$ the NAIRU gap (Fed dual mandate).


Extension: Lexical evolution

- **Motivation:** What is considered “hawkish” or “dovish” language can evolve with macroeconomic regimes and global context.
- **Real time RoBERTA fine-tuning:** Fine-tuning RoBERTa solely on dovish/neutral/hawkish labeled sentences from Board communications over a given period, then applying it to Fed communications across the entire 1996–2026 timeframe to identify how and when the Fed's wording shifted:
 - Expanding train window: 1996–2000, 1996–2005, 1996–2010, 1996–2015, 1996–2020
 - Comparison of different Board Stance indicators obtained from RoBERTa models fine-tuned on varying subsets of labeled sentences, put into perspective with the monthly Fed funds rate.
- **Preliminary findings:** Identification of notable characteristic periods:
 - Pre-2000s: Lack of standardized communication → No impact on the policy rate
 - 2000s: Implicit signaling, moderate tone → Moderate stance but linked to the policy rate
 - 2010s: Explicit adoption of forward guidance (calendar, thresholds, dot plots) → More pronounced stance and linked to the policy rate
- Real-time pipeline ensures the indicator adapts to semantic evolution of Fed discourse.

Results

Board stance as a leading indicator of monetary cycles


Board stance and Fed monetary cycle



- Board stance tracks all major monetary cycles (1996–2026). Directional consistency across rate-hiking and easing phases.
- Stance moves ahead of the first effective rate change.
→ Leading indicator of policy direction


From consensus to cacophony: individual stances and internal dissent

Board Stance vs. Average voter stance



- Board stance and the average voter stance exhibit the same directional pattern, but the average voter stance is less pronounced during hawkish/dovish phases. **Official communiqués smooth internal tensions.**

Monetary cacophony: Dissents among voting members



- Monetary cacophony index: rolling std. dev. (3m avg) of individual voter stances. Measures the degree of FOMC internal disagreement.


- Cacophony peaks appear to precede major cycle turning points. **Consensus fractures in communication before it fractures in the vote.**

Augmented Taylor rule: stance as a determinant of policy rates

Augmented Taylor rule - Estimation

Dependent variable: Fed funds rate (freq: monthly)	Model 1 Fundamentals, historical	Model 2 Fundamentals	Model 3 Fundamentals & Board stance	Model 4 Board stance
Intercept	0.48 ** (0.16)	2.27 *** (0.19)	1.90 *** (0.16)	1.59 *** (0.08)
Inflation (π_{t-4})	0.29 ** (0.10)	-0.39 *** (0.10)	-0.63 *** (0.08)	
Inflation (π_{t-12})	1.09 *** (0.09)	0.90 *** (0.10)	0.97 *** (0.08)	
NAIRU Gap ($u_t - u_t^*$)	-1.10 *** (0.14)	-0.72 *** (0.11)	-0.38 *** (0.09)	
NAIRU Gap ($u_{t-4} - u_{t-4}^*$)	0.32 * (0.14)	-0.16 . (0.11)	-0.22 ** (0.09)	
Board Stance ($S_{Board,t-4}$)			7.57 *** (0.57)	14.23 *** (0.64)
Estimation period	1980-2026	1997-2026	1997-2026	1997-2026
Adjusted R2	0.75	0.69	0.80	0.59
Standard errors in parentheses. Significance levels: ***p<0.001, **p<0.01, .p<0.1.				


Augmented Taylor rule - Fitted



Communication is not a commentary on fundamentals. It adds genuinely new information to the reaction function.

Communication as an instrument: the forward guidance effect

Board stance and Fed Funds rate



The grey shaded area represents the period from which the labeled sentences were extracted to train the model.

- **Unstandardized and sparse language (1990s):** Statements were published only when rates changed, brief and factual, with no forward-looking language.
- **Lack of polarization (2000s):** Early models (BS 2005, 2010) do detect directional cycles, but the Board stance oscillates in a narrow, compressed band.
- **Forward guidance era (2010s):** The Dot plot, explicit thresholds and calendar commitments transformed Fed language from descriptive to prescriptive, producing a strongly polarized stance signal that tracks rate amplitude far more closely.

Conclusions

- **Communication leads policy:** The Board stance tracks all monetary cycles and anticipates rate moves ahead of formal decisions. Adding it to a standard Taylor rule raises adj. R^2 from 0.69 to 0.80, with a highly significant coefficient.
→ Language is not a commentary on policy; it is a component of it.
- **Cacophony as early warning:** The cross-sectional dispersion of individual voter stances peaks ahead of major cycle turning points. Consensus fractures in speeches before it crystallizes in formal votes, offering a high-frequency signal of impending policy change.
- **A structural break around 2012:** Pre-forward-guidance language was descriptive and tonally flat; post-2012 communication became prescriptive and strongly polarized. The vocabulary itself became a policy instrument.
→ Lexical drift is part of the signal: a continuously updating pipeline is a requirement.

Thank you for your attention

Appendix

Panel A1	Panel B1	Panel C
inflation expectation, interest rate, bank rate, fund rate, price, economic activity, inflation, employment	unemployment, growth, exchange rate, productivity, deficit, demand, job market, monetary policy	weren't, were not, wasn't, was not, did not, didn't, do not, don't, will not, won't
Panel A2	Panel B2	
anchor, cut, subdue, decline, decrease, reduce, low, drop, fall, fell, decelerate, slow, pause, pausing, stable, nonaccelerating, downward, tighten	ease, easing, rise, rising, increase, expand, improve, strong, upward, raise, high, rapid	


- 1. **Sentence filtering:** A word from Panel A1 or B1 must appear for the sentence to be kept.
- 2. **Sentence splitting:** Contrastive connectors (Panel C) are used as split points
- 3. **Post-split filtering:** Fragments are kept only if they contain a word from Panel A1, A2, B1, or B2.

Note: Gorodnichenko et al. (2021) originally designed these panels for a dictionary-based stance classifier:


- Dovish if (A1+A2) or (B1+B2)
- Hawkish if (A1+B2) or (A2+B1)
- Reversed if Panel C.

Rule-based dictionary vs RoBERTA model


Board stance – Rule-based dictionary



Board stance – RoBERTa (Transformer)



Board stance and Fed Funds rate



BS 2000: Board stance computed from a RoBERTa model fine-tuned exclusively on labeled sentences from Board communications published between 1996 and 2000 (6m avg.).

- **Methodology:** Five RoBERTa models trained on expanding windows starting in 1996 and ending respectively in 2000, 2005, 2010, 2015 and 2020, each then applied to the full 1996–2026.

→ **How is the Fed's terminology evolving?**

- **Old vocabulary, weak signal:** Models trained only on pre-2010s speeches produce a board stance score that bears no systematic relationship with the Fed Funds rate across the entire sample period.
- **A break in communication:** Models trained with speeches after 2015 track Fed Funds rate dynamics significantly more closely.

Type of document	Press releases	Speeches
Speaker	Board	Governors (Fed "Monetary policy" filter)
Number of documents	548	1365

Type of document	Speeches (no filtering during the extraction phase)											
Speaker	New-York	Chicago	St Louis	San Francisco	Richmond	Atlanta	Boston	Dallas	Cleveland	Philadelp hia	Kansas City	Minneapo lis
Number of documents	652	280	576	288	293	206	266	269	180	211	160	298