OESTERREICHISCHE NATIONALBANK
EUROSYSTEM

Working Paper 281

From Tone to Trajectory: Continuous Sentiment and the Shape of Monetary Policy Communication

Martin Feldkircher, Márton Kardos,
Kristoffer Laigaard Nielbo

OESTERREICHISCHE NATIONALBANK
EUROSYSTEM

OeNB Working Paper Series
40 pages

From Tone to Trajectory: Continuous Sentiment and the Shape of Monetary Policy Communication

Central bank press conferences are scrutinised line by line. Most research reduces each statement to a single number. We measure the tone word by word across ECB and Fed press conferences and track how it rises, turns, and settles across the statement — its „sentiment arc." The shape of that arc helps predict the next rate decision and how professional forecasters revise inflation expectations, on top of the overall tone. How a statement is arranged, not just what it says, is part of the signal.

Author

Martin Feldkircher
Oesterreichische Nationalbank,
martin.feldkircher@oenb.at
Vienna School of International Studies,
martin.feldkircher@oenb.at

Márton Kardos
School of Culture and Society -
Center for Humanities Computing,
martonkardos@cas.au.dk

Kristoffer Laigaard Nielbo
School of Culture and Society -
Center for Humanities Computing
kln@cas.au.dk

JEL classification
C55, C88, E52, E58, D83

Keywords
monetary policy; introductory statement; sentiment arc

Measuring the shape, not just the tone

Central bank statements are usually boiled down to one average tone score, which ignores structure. We instead measure sentiment word by word across ECB and Fed press conferences, tracing the shape of each statement — its „sentiment arc": overall level, trend, and curvature.

The arc predicts the next rate decision

The shape of the arc helps predict the ECB's next interest rate decision — on top of the average tone, and not just because certain topics appear in certain parts of the statement. Where firm or cautious language sits, and how consistently a stance is held, carries a policy signal.

Forecasters respond to structure

Arc shape is also linked to how professional forecasters revise their inflation expectations after a press conference — evidence that the audience responds to structure, not only to headline content. For the Fed the pattern is similar but narrower, matching its more standardised style.

From Tone to Trajectory: Continuous Sentiment and the Shape of Monetary Policy Communication

MARTIN FELDKIRCHER[1,2,3], MÁRTON KARDOS[4], and KRISTOFFER LAIGAARD NIELBO[*4]

[1]Vienna School of International Studies (DA)

[2]Oesterreichische Nationalbank (OeNB)

[3]Centre for Applied Macroeconomic Analysis (CAMA)

[4]School of Culture and Society - Center for Humanities Computing (AU)

September 8, 2026

Abstract

Central bank press conferences are not merely information releases — they are structured narratives. We study whether the *shape* of sentiment within a statement, not just its average tone, carries policy-relevant signals. Constructing sentiment arcs for ECB press conferences along three dimensions — monetary policy stance, economic outlook, and uncertainty — we assess their predictive content for policy rate changes and inflation expectations. Our findings show that arc shape robustly predicts rate decisions beyond lexicon-based benchmarks: it is not merely whether a statement sounds hawkish or economically optimistic on average, but how these sentiments are sequenced, emphasized, and sustained across the statement, that carries the policy signal. Arc features also shape how professional forecasters update inflation expectations, pointing to a receiver-side effect distinct from the direct policy signal. Our results are robust to a range of additional exercises, including different sample splits as well as controlling for the content of the statement. They also carry over to a short sample of Fed press conferences. These findings suggest that the sequencing and emphasis of policy language across a statement is a first-order feature of the policy signal, not a second-order refinement.

Keywords: monetary policy; introductory statement; sentiment arc

JEL Codes: C55, C88, E52, E58, D83.

[*]Martin Feldkircher martin.feldkircher@da-vienna.ac.at (corresponding author), Márton Kardos, martonkardos@cas.au.dk, and Kristoffer Laigaard Nielbo kln@cas.au.dk. Any views expressed in this paper represent those of the authors only and do not necessarily coincide with those of the Oesterreichische Nationalbank or the Eurosystem. We would like to thank Kilian Rieder, an anonymous referee, and participants of the ZEW workshop in Mannheim, and an internal seminar at the Oesterreichische Nationalbank for helpful comments.

Non-technical summary

When a central bank changes interest rates, the press conference that accompanies the decision is scrutinised as closely as the decision itself. Every sentence of the prepared statement is drafted in the knowledge that markets, journalists, and professional forecasters will parse it word by word. Most research that measures the tone of these statements reduces each one to a single number — how hawkish or dovish, how optimistic or pessimistic it sounds on average. This discards the order in which things are said.

This paper asks whether the *shape* of a statement — how its tone rises, turns, and settles from the first word to the last — carries information that the average tone misses. We call this shape the statement's *sentiment arc*. For the full set of European Central Bank press conferences, and for a shorter set of Federal Reserve press conferences, we measure sentiment along three dimensions: the monetary policy stance (hawkish versus dovish), the assessment of economic conditions (favourable versus adverse), and the degree of uncertainty conveyed. We then summarise each statement's arc by its overall level, its trend from start to finish, and its curvature, and ask whether these features help explain the central bank's next interest rate decision and how professional forecasters revise their inflation expectations after the press conference.

We find that arc shape matters on both counts. Where a central bank places its hawkish or uncertain language within a statement, and how consistently it sustains a stance from opening to close, is systematically related to its next policy move — over and above conventional average-tone measures, and not merely because certain topics tend to appear in certain parts of the statement. Arc shape is also related to how forecasters adjust their inflation expectations. The pattern is broadly similar for the Federal Reserve, though narrower and concentrated in the policy-stance dimension, consistent with its more standardised, template-driven communication style.

For central banks, the results suggest that communication design — the sequencing and emphasis of language within a statement — is itself a policy lever, distinct from the substance of what is announced.

1 Introduction

Over the last two decades, communication has become an increasingly important aspect of monetary policy (Blinder et al., 2008, 2024). Central banks have moved away from deliberate opacity toward active management of expectations, investing heavily in forward guidance, structured press conferences, and clearly worded policy statements (Blinder et al., 2008). The European Central Bank (ECB) and the US Federal Reserve (Fed) are two prominent examples. Following each Governing Council meeting, the ECB President delivers a prepared introductory statement before opening the floor to journalists' questions; the statement and the press conference together constitute distinct information events, each with its own effect on financial markets (Altavilla et al., 2019; Ehrmann and Fratzscher, 2009). Since 2011, the Fed has held a press conference after each monetary policy meeting and published its transcript on its website. These documents are crafted with awareness that every word will be scrutinised by market participants (Hubert and Labondance, 2021).

A sizeable empirical literature documents that this communication content matters: ECB press conferences provide substantial additional information to financial markets beyond what is contained in the rate decision itself (Rosa and Verga, 2007; Ehrmann and Fratzscher, 2009; Conrad and Lamla, 2010; Parle, 2022; Gardner et al., 2022), the tone of CB statements shapes inflation expectations (Cho and Jung, 2026; Cho and Rho, 2026; Montes et al., 2016; Hoffmann et al., 2026), and communication content helps predict future policy decisions (Kanelis and Siklos, 2024; Baranowski et al., 2021; Hubert and Labondance, 2021). Much of this evidence[1] reduces a policy statement to a single aggregate tone or sentiment measure — a scalar capturing the overall hawkish or dovish inclination of the text (Picault and Renault, 2017; Baranowski et al., 2021). While this approach has proven productive, it discards the internal structure of the communication: how sentiment evolves *within* a statement, where it rises or falls, whether it accelerates or reverses, and how the statement ends.

We argue that this internal structure — the *sentiment arc* — carries information not captured by the aggregate tone. A sentiment arc represents the sequential evolution of sentiment scores across a document. Deriving sentiment arcs instead of a single aggregate sentiment score for a document can provide more granular information about local variation in affective tone. Our motivation draws on two strands of work. First, research in computational narratology shows that emotional arcs of stories are not arbitrary but dominated by a small number of basic shapes, and that particular arc shapes are associated with greater audience success (Reagan et al., 2016); this regularity extends across different genres and communication contexts, suggesting that arc shape is a robust feature of structured communication (Neugarten et al., 2025).[2] Second,

[1] A related literature studies the *delivery* rather than the content of central bank communication: Gorodnichenko et al. (2023) and Curti and Kazinnik (2023) show that vocal tone and facial expressions during press conferences move markets even after controlling for the text itself. We instead study the sequential structure of what is said, not how it is delivered.

[2] Sentiment arc features have also been shown to carry information about the perceived quality of literary narratives beyond what average sentiment conveys (Bizzoni et al., 2022).

evidence from psychology documents a peak-end bias: retrospective evaluations of sequential experiences are systematically governed by the peak and the endpoint, not the time-averaged experience (Müller et al., 2019). These regularities motivate the conjecture that position within a statement is not informationally neutral, without requiring that readers experience these documents exactly as a novel reader or a live audience would. We therefore treat the shape of the arc as an empirical question: does it convey policy-relevant signal beyond what aggregate tone captures?

To investigate this, we construct word-level sentiment scores for the full corpus of ECB introductory statements along three conceptually distinct dimensions: (i) *monetary policy sentiment* (hawkish vs. dovish), capturing the direction of the signalled policy intention, (ii) *economic sentiment* (positive vs. negative), capturing the central bank's assessment of underlying economic conditions — both motivated and assessed in Picault and Renault (2017) — and (iii) *uncertainty sentiment*, measuring the degree of epistemic uncertainty conveyed by the statement. Scores are then derived using Concept Vector Projection (CVP), a method proposed by Lyngbaek et al. (2025) that projects sentence embeddings from a large language model onto a direction in embedding space defined by manually curated seed phrases. Relative to bag-of-words dictionaries (Picault and Renault, 2017), fine-grained rule-based lexicon approaches (Consoli et al., 2022), fine-tuned classifiers (Shah et al., 2023), and large language model-based scoring (Hansen and Kazinnik, 2024; Silva et al., 2025), CVP provides continuous scores in a common, context-sensitive embedding space without requiring model fine-tuning or task-specific prompting.

Our main results can be summarized as follows. First, arc shape predicts ECB policy rate changes beyond what lexicon-based benchmark measures capture: it is not just whether a statement sounds hawkish or economically optimistic on average, but how monetary, economic, and uncertainty sentiment are sequenced and shaped across the statement that carries the policy signal. Second, this is not an artefact of which topics happen to be discussed where in the statement: arc shape remains predictive once we control directly for the section-by-section topic content of each statement. Third, arc features also predict how professional forecasters revise their inflation expectations, at both short- and long-run horizons, pointing to a receiver-side channel distinct from the policy signal itself. Finally, as a cross-check on a considerably smaller sample, the same sender- and receiver-side pattern replicates for the Federal Reserve, indicating that these results are not specific to the ECB.

These results contribute to the literature on central bank communication (Blinder et al., 2008; Ehrmann and Fratzscher, 2009; Ehrmann et al., 2024; Wabitsch, 2025) in two ways. Empirically, we show that the arc structure of policy statements encodes information about future rate decisions that aggregate sentiment scores miss. Conceptually, we draw on the computational narratology literature on emotional arcs (Reagan et al., 2016) to motivate treating arc shape as an object of study, showing empirically that the structure of textual sentiment — level, slope, and curvature — carries predictive information beyond aggregate tone.

The remainder of the paper is organised as follows. Section 2 describes our text corpus and introduces the methodology. Section 3 presents the full-sample sentiment arcs and validates the scoring approach. Section 4 reports results from predicting rate changes and inflation expectations. Finally, Section 5 concludes.

2 Data and Methodology

For this study, we utilize data from the press conferences of the European Central Bank (ECB). The ECB's Governing Council convenes regularly to decide on monetary policy, with outcomes communicated through two steps (Altavilla et al., 2019). First, at 13:45 Central European Time (CET), the ECB releases a concise statement regarding its monetary policy decision (MPD). This statement is brief and contains limited textual content. Second, at 14:30 CET, the ECB President delivers an *introductory statement* during a press conference. This carefully crafted document informs the public about the rationale behind the interest rate choices, presents the ECB's perspective on the economic situation, and provides insights into its future conduct. From 1999 to 2003, each statement followed a two-pillar template: after it opens with the policy decision, it proceeds through the *monetary analysis* (covering monetary dynamics and credit conditions) to the *economic analysis* (covering the inflation outlook and near-to-medium-term risks to price stability). The ECB's 2003 strategy review reversed this ordering: from 2003 to July 2013 the rate decision was followed by economic analysis, which was followed by the monetary analysis and finally a cross-check of the two (Holm-Hadulla et al., 2021). Since 2013, an explicit forward guidance chapter has been added to this structure. Typically, the statements end with an address to the member states of the euro area, calling for sound fiscal policies and structural reforms to reap the full benefits of the monetary union. Following another strategy review, from July 2021 the statements added a new section labelled risk assessment. The introductory statement lasts approximately fifteen minutes and is followed by a forty-five-minute session of questions and answers.

Our raw text data for the ECB contains 269 statements from the period 1998-06-09 to 2025-06-06 and we obtain the texts directly from the ECB's website (`https://www.ecb.europa.eu/press/pressconf/html/index.en.html`) using automated web scraping. Since the pre-2002 statements do not follow the same structure as the latter statements, we restrict the sample period from 2002-01-03 to 2025-06-05 which yields a final set of 235 statements. For the predictive regressions, carried out in Section 4, we also obtain policy rates. These are collected from the BIS database. For the euro area and from 12-09-2024 on, the policy rate no longer constitutes the main refinancing rate (MRO) but the deposit facility rate (DPR); policy-rate coverage is complete across the 235-statement sample.

In an extension, we also provide sentiment arcs for the US Fed. Here, we use the transcripts of the press conference that explains the rate decision and is most directly comparable to the ECB statements but available only from 2011 onward. The total sample features 88 transcripts.[3] The

[3]First, the Fed publishes the monetary policy decision with a short explanation, which is called the FOMC

transcripts are downloaded directly from the Federal Reserve's website (`https://www.federalreserve.gov/monetarypolicy/fomccalendars.htm`).

For the inflation expectations regressions in Section 4.2, we additionally use professional forecasters' one-year-ahead inflation expectations from the Consensus Economics survey. The data are converted from fixed event to fixed horizon forecasts as in Siklos (2013). We also use long-run forecasts, which are an average of 6 to 10 year ahead forecasts, also from Consensus Economics. The survey is conducted around the 10th of each month – long-run forecasts are surveyed only twice per year (April and October). To control for the persistent level of inflation expectations, we include the actual year-over-year inflation rate Δp_t^{yoy}: headline HICP for the euro area, obtained from Eurostat (dataset `prc_hicp_midx`, annual rate of change), and headline CPI for the United States, obtained from the Federal Reserve Bank of St. Louis FRED database (series `CPIAUCSL`).

2.1 Concept Vector Projection

In order to calculate sentiment arcs over the duration of a press conference, we use Concept Vector Projection (CVP) (Lyngbaek et al., 2025). In contrast to prevailing sentiment analysis pipelines, which typically treat polarity as a multi-class classification problem, using either bag-of-words or transformer representations, CVP relies on the linear representation hypothesis (Park et al., 2024) to assign continuous sentiment scores to text. Concretely, this assumes that sentences can be ordered along a single axis from negative to positive sentiment. To apply the methodology, we have to first define concepts. Following Picault and Renault (2017), we focus on two concepts, namely monetary policy sentiment (hawkish positive, dovish negative) and economic sentiment (optimistic positive, pessimistic negative). As a third dimension, we add uncertainty sentiment, where uncertainty is the positive pole. Each statement can evolve in these three dimensions, word by word. Each concept is defined by a set of positive (P^+) and negative (N^-) seed phrases that have to be specified by the researcher.

Given the embedding matrices of positive seed phrases (P^+) and negative seed phrases (N^-), this difference is computed as follows:

$$\overrightarrow{v} = \bar{p} - \bar{n}, \text{ where } \bar{p}_j = \frac{\sum_i P^+_{ij}}{|P^+|}, \text{ and } \bar{n}_j = \frac{\sum_i N^-_{ij}}{|N^-|} \tag{1}$$

$$\hat{v} = \frac{\overrightarrow{v}}{||\overrightarrow{v}||_2} \tag{2}$$

where $|P^+|$ and $|N^-|$ denote the cardinalities of the positive and negative seed sets, respectively. We refer to $\hat{v}$ as the *concept vector.*

statement. In addition to the statement, the Fed publishes the transcript of the press conference, which we analyse in this study. These documents are released immediately. In addition, it also releases meeting minutes, published three weeks later, which summarize the main discussions and reflect an evolution in Federal Reserve communication. Finally, verbatim transcripts of the meeting are released with a five-year lag. Here, we focus on the transcripts and disregard information contained in the minutes due to the three-week publication lag.

To assign a sentiment score to a given document, we first encode the document using the same transformer model to obtain an embedding vector e, and then project it onto the concept vector using the dot product:

$$s = e \cdot \hat{v} \tag{3}$$

Lyngbaek et al. (2025) apply their method in a multilingual literary context and observe substantially improved alignment with human evaluation compared to baseline approaches. CVP is particularly advantageous for our use case for several reasons:

1. It does not necessitate finetuning transformer models, only inference. Model training would be considerably more computationally intensive and requires access to large-scale data.
2. CVP, thanks to neural transformer embeddings, is context-sensitive and can deal with synonymy, polysemy, and spelling errors.
3. It allows us to compute continuous sentiment scores throughout the course of the document, thereby allowing us to compute so-called sentiment arcs.
4. CVP's flexibility allows us to measure sentiment from multiple angles.

To come up with potential seed phrases, we use a multi-layered process. First, we draw on the dictionary from Picault and Renault (2017) to gather candidate sentences from both the ECB and FOMC corpora. Following Lyngbaek et al. (2025), we target 15 seed phrases per pole. In a second step, we then remove any ECB or US Fed specific tone or wording (e.g., "federal funds rate"). To give some examples, the most hawkish phrases are *"Inflation remains too high and requires further policy tightening"* and *"Rising inflationary pressures call for higher interest rates."*. Highly dovish sentences are *"Economic conditions warrant maintaining a low policy rate."* and *"We will continue asset purchases to support favorable financing conditions."*. The hawkish and dovish phrases span topics including inflation, conventional monetary policy, and unconventional measures. To give some examples for the economic environment seed phrases, a positive environment is captured by phrases such as *"The economy is expanding at a solid pace with strong job growth."*, *"Consumer demand remains strong and supports continued growth."*. Negative sentences comprise *"Weak demand and tight credit conditions are restraining growth."*, or *"Household spending remains constrained by low income growth."*. The full list of phrases is provided in Tables 7 to 9 in the appendix. These seed phrases should reflect the two polar opposites for each dimension and, ideally, these dimensions are orthogonal to each other. Hence, as a further step, we assess the orthogonality of the embedded seed phrases using PCA and pairwise scatter plots shown in Figures 4 and 5 in the appendix. This process yielded a final seed set of 11 hawkish and 13 dovish phrases, 14 positive and 13 negative economic outlook phrases, and 15 high-uncertainty and 13 low-uncertainty phrases.

For this final set of seed phrases, we assess their validity as polar anchors by applying the CVP methodology to the seed phrases themselves. The distribution of the CVP scores for the seed phrases is shown in Figure 1, and we would expect a good separation between positive and negative sentiment for all three dimensions if the seed phrases contain meaningful information.

Figure 1: Distribution of scores for seed phrases

(a) Monetary policy sentiment scores **(b)** Economic sentiment scores **(c)** Uncertainty scores

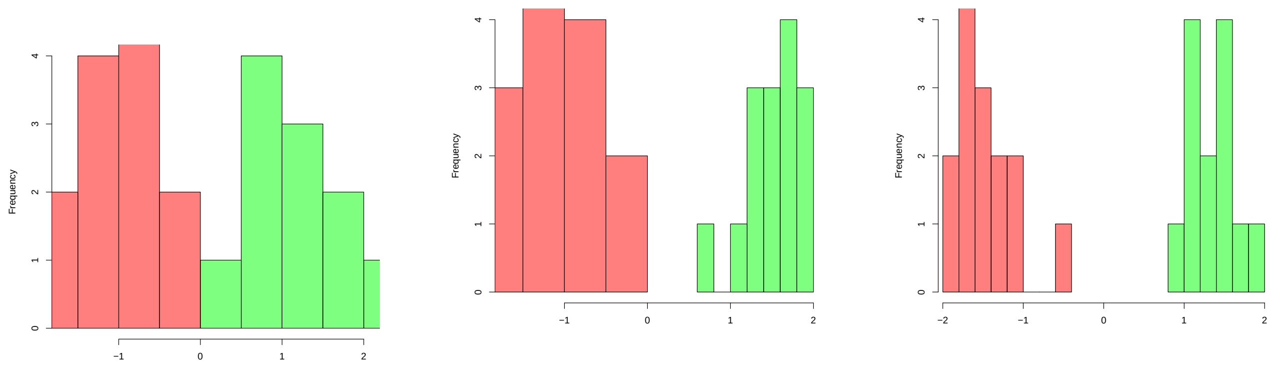

Notes: Distribution of CVP scores for the seed phrases of each dimension. Green bars show scores for the positive pole (hawkish, economic positive, uncertain); red bars show scores for the negative pole (dovish, economic negative, certain). Green and red bars overlapping indicates less separation between poles.

The plot shows that for all three dimensions we find a clear separation between positive and negative seed phrases. Two points are worth noting. First, our seed phrases are intentionally generic and institution-neutral, though the approach naturally extends to institution-specific language. Second, our approach can be easily expanded to capture further sentiment dimensions.

Finally, to derive the sentiment arc, each document must be processed at a chosen level of granularity. This determines how the arc is constructed and governs the degree to which individual segments contribute to the measured sentiment trajectory.

When applying the CVP method, we opt for two strategies.

1. A *sentence-level* sentiment arc, where we divide the document into sentences, produce separate embeddings for each sentence, then project them onto the concept vectors, and

2. A *contextual* sentiment arc, where we encode all tokens in the document at once, thereby letting the tokens fully attend to each other, and project token representations directly onto the concept vectors, without pooling. Token scores are then aggregated into equal-width position bins along the document to form the arc. For this purpose, we use the `nvidia/llama-embed-nemotron-8b` embedding model (Babakhin et al., 2025), which has a large enough context length to fit all introductory statements (32,768 tokens) and, at the time of writing, is the best-performing model on the Massive Multilingual Text Embedding Benchmark (Enevoldsen et al., 2025) when filtering for *Financial*, *Government*, and *Legal* domains. The `nvidia/llama-embed-nemotron-8b` is a mean-pooling-trained embedding model, which ensures that all tokens belong to the same embedding space.[4]

[4]Mean-pooling-trained models optimize the sentence representation as the average over all token embeddings,

We calculate contextual arcs as our benchmark and provide results for sentence-level arcs as a robustness exercise, shown in the appendix.

3 Sentiment arcs

A sentiment arc is the sequence of CVP scores read from the opening to the close of a statement, tracing how sentiment evolves as the communication unfolds. Because statements can differ in length, the position axis is normalized to the interval $[0\%, 100\%]$, mapping the opening token to 0% and the closing token to 100% regardless of document length. This normalization allows arcs from different statements to be averaged on a common scale. To reduce token-level noise while preserving within-statement structure, each document-level arc is first smoothed with a LOESS curve (span = 0.6, degree 2) evaluated on a 100-point grid, and point-wise means and 95% confidence intervals are then computed across statements.

Figure 2 shows the full-sample average arcs for the ECB, pooled across all statements in the corpus. The panels show the arcs for the three dimensions analyzed in this study: monetary policy sentiment (hawkish positive), economic sentiment (optimistic positive), and uncertainty (uncertain positive). The dotted horizontal line marks zero.

---

which induces a consistent geometric structure across tokens. CLS-pooled models, by contrast, optimize only the [CLS] token for sentence-level representation, leaving remaining token embeddings without a guarantee of mutual comparability — a requirement for geometrically meaningful token-level CVP projections.

Figure 2: Full-sample average sentiment arcs: ECB

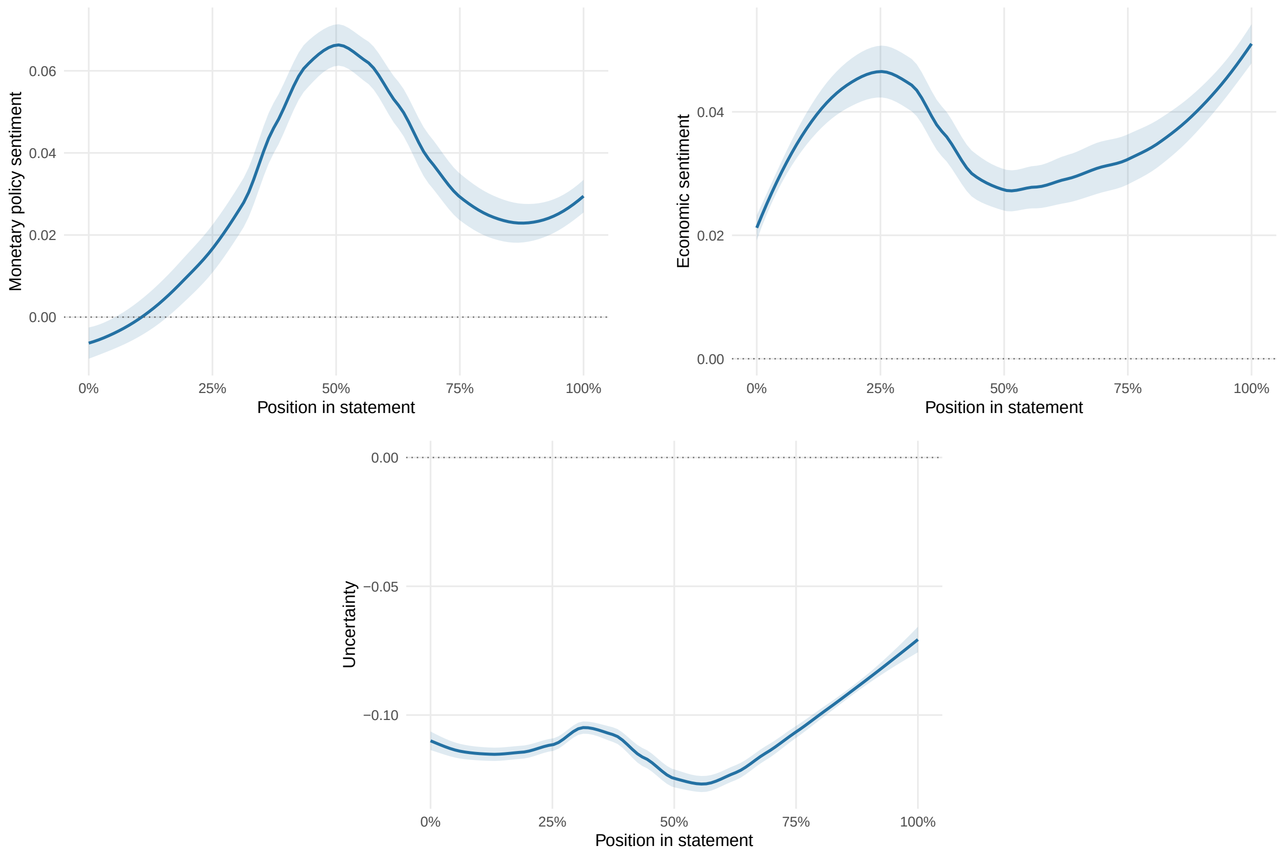

Notes: LOESS-smoothed (span = 0.6) full-sample average sentiment arcs. Panels show monetary policy sentiment (hawkish positive), economic sentiment (optimistic positive), and uncertainty sentiment (uncertain positive). The x-axis is the normalized position within the statement (0% = opening, 100% = closing token), mapping all statements to a common scale regardless of length. Shaded bands are 95% pointwise confidence intervals across statements. $N = 235$ statements, 2002–2025.

Several features of Figure 2 are worth noting. As the statement opens with a general introduction, followed by the policy decision — a factual announcement — monetary policy sentiment starts slightly below zero, on average, before rising into positive territory around the opening tenth of the statement and building to a peak around the midpoint; it then recedes gradually toward the close. Economic sentiment follows a related but distinct within-statement pattern: positive throughout, it rises quickly to an early peak around the first quarter of the statement, eases through the middle, and then rises again toward the close, tracking the statement's progression from economic assessment to the concluding remarks. That said, the average arc described here is a general pattern across all statements; deviations from it — for instance during crisis periods or ahead of rate changes — are of particular interest and will be examined more formally in the next section. The uncertainty sentiment arc is negative through most of the statement — that is, ECB language is on average relatively confident and unhedged rather than uncertain — and becomes *more* confident through the analytical middle of the statement before turning markedly less certain again toward the close.

To verify that the CVP scores capture meaningful content, Table 1 reports the highest- and lowest-scoring ECB sentences in each CVP dimension, extracted from the full corpus. These are the sentences that sit at the extreme poles of each dimension — the most hawkish, the most dovish, the most optimistic, and so on — as measured by mean token-level CVP projection.

Table 1: Highest and lowest scoring sentences by CVP dimension

Dimension	Pole	Date	Score	Sentence
MP stance	Hawkish	2022-06	+0.257	Accordingly, measures of underlying inflation have been rising further.
MP stance	Dovish	2015-09	−0.127	Our asset purchase programme continues to proceed smoothly.
Economic	Optimistic	2007-07	+0.217	Domestic demand in the euro area is also expected to maintain its relatively strong momentum.
Economic	Pessimistic	2022-09	−0.155	Fourth, uncertainty remains high and confidence is falling sharply.
Uncertainty	Uncertain	2025-04	+0.082	The economic outlook is clouded by exceptional uncertainty.
Uncertainty	Certain	2013-07	−0.258	Medium-term inflation expectations remain firmly anchored in line with price stability.

Notes: For each CVP dimension, the table reports the ECB sentence with the highest score (positive pole) and the lowest score (negative pole) across all statements in the corpus. Scores are the mean of token-level CVP projections for tokens belonging to that sentence, computed from contextual arcs. Sentences with fewer than three tokens or fewer than 40 characters are excluded.

The sentences in Table 1 are readily interpretable. On the monetary policy sentiment dimension, the highest-scoring sentence — *"Accordingly, measures of underlying inflation have been rising further"* (June 2022) — captures the broadening of inflationary pressure at the onset of the tightening cycle, while the lowest-scoring one, from September 2015, reflects the language of the ECB's asset purchase programme in its early, pre-pandemic phase. On the economic sentiment dimension, the top-scoring sentence describes pre-crisis euro area domestic demand, while the bottom-scoring sentence captures the 2022 energy-driven confidence shock. The uncertainty sentiment dimension warrants a brief note on interpretation: it scores the degree of epistemic uncertainty in language, not the valence of economic conditions. The highest-scoring sentence (positive score) uses hedged, qualified language acknowledging uncertainty, whereas the lowest-scoring sentence (negative score) uses assertive, settled language. Across all dimensions, the extreme sentences align with the historical episodes one would expect, supporting the validity of the CVP scoring approach.

4 Predictive regressions

We study whether the *shape* of sentiment arcs — not just their average level — carries predictive information, organized around two conceptually distinct perspectives on how communication design operates. From the *sender's perspective*, the central bank actively structures its narrative: different policy environments call for different arc shapes, and this deliberate structure should be recoverable in the predictive relationship between arc features and subsequent policy decisions. From the *receiver's perspective*, the sequential structure of a statement shapes how audiences process and respond to it: if market participants and households are subject to peak-end effects (Müller et al., 2019), the arc — not just its average — should affect how beliefs are updated. These two perspectives motivate two families of regressions:

$$\text{Sender (rate changes):} \quad \Delta i_{t+k} = \alpha + \boldsymbol{\beta}' \, \mathbf{x}_t^{\text{bench}} + \boldsymbol{\beta}' \, \mathbf{x}_t^{\text{arc}} + \varepsilon_t \quad k \in \{1, 2\} \tag{4}$$

$$\text{Receiver (inflation expectations):} \quad \pi_{t+1}^{e,\text{ST}} = \alpha + \boldsymbol{\beta}' \, \mathbf{x}_t^{\text{bench}} + \boldsymbol{\beta}' \, \mathbf{x}_t^{\text{arc}} + \gamma \, \Delta p_t^{\text{yoy}} + \varepsilon_t \tag{5}$$

In equation (4), Δi_{t+k} is the policy rate change decided k meetings ahead of statement t. Because ECB meetings are typically held at six-week intervals, this reflects the signaling horizon of each statement rather than a literal one-period-ahead forecast. We report $k = 1$ as the headline horizon throughout Section 4.1, while we provide results for $k = 2$ as a robustness exercise in the appendix.

In equation (5), $\pi_{t+1}^{e,\text{ST}}$ is the level of professional forecasters' one-year-ahead (short-term) inflation expectations surveyed in month $t + 1$. We also examine the analogous long-run outcome $\pi_{t+1}^{e,\text{LT}}$, the Consensus Economics 6–10-year-ahead average inflation forecast, matched to the nearest preceding statement (Section 4.2). As with $t + k$ in equation (4), the subscript $t(+1)$ here indexes timing – specifically, the Consensus Economics survey date relative to the underlying statement – rather than the forecast horizon itself; the forecast horizon (one-year-ahead vs. 6–10-year-ahead) is instead carried by the ST/LT superscript. Δp_t^{yoy} is the actual year-over-year inflation rate in month t, which controls for the persistent level of expectations. We deliberately match each statement to the *following* month's survey rather than the same month's: because the Consensus Economics survey is conducted around the 10th of the month, an ECB statement from later in month t could otherwise postdate that month's own survey, reversing the intended timing and invalidating the predictive relationship. Matching to month $t + 1$ instead guarantees that any ECB press conference in month t – regardless of its exact day – necessarily precedes the survey it is matched to, so timing is unambiguous by construction and no additional restriction on the day of the press conference is needed.

The feature vectors are constructed as follows. The benchmark $\mathbf{x}^{\text{bench}}$ contains the Picault–Renault dictionary scores (`mp_pr`, `ec_pr`) described in Section 2. The arc features consist of the Nelson–Siegel factors — level, slope, and curvature — estimated separately for each of the three arcs.[5] The *level* captures the overall average sentiment across the statement; the *slope* captures the directional trend, that is, whether sentiment rises or falls from the opening to the close of the statement; and the *curvature* captures non-linear within-statement structure,

[5]Extracting the Nelson–Siegel level, slope, and curvature factors requires fixing a decay parameter λ, which determines the exponential basis functions underlying the slope and curvature loadings. Following standard practice in the dynamic Nelson–Siegel literature (Diebold and Li, 2006), we fix $\lambda = 9.5$ across all three sentiment dimensions rather than estimating it separately for each; this keeps the factors comparable across dimensions and lets them be extracted by simple ordinary least squares. This choice places the peak of the curvature loading approximately one-fifth of the way into the statement, in the early analytical section following the policy decision.

distinguishing statements that build toward a central argument from those that open and close at higher intensity than the analytical middle. The Nelson–Siegel factors are complemented by the Hurst exponent for each dimension, which captures arc persistence (Bizzoni et al., 2022). We then build a model ladder: M1 uses $\mathbf{x}^{\text{bench}}$ alone; M2 replaces the dictionary scores with CVP document means, providing a natural midpoint between the lexicon baseline and the full arc approach; M3 adds the arc characteristics for the monetary policy sentiment dimension; M4 extends M3 with arc features of the economic sentiment dimension; the Full arc specification further adds arc features of the uncertainty sentiment dimension. This ladder approach allows us to assess which dimension contributes predictive content. All arc measures are derived from contextual arcs; Appendix A.4 confirms that contextual arcs substantially outperform sentence-level arcs.

4.1 Arc structure as policy signal

Figure 3 provides a first visual indication of whether arc shape tracks the subsequent rate decision. Sentiment arcs conditioned on the next meeting's outcome — hike, unchanged, or cut — reveal systematic differences: statements preceding rate hikes display distinctly more hawkish monetary policy sentiment and more positive economic sentiment throughout. These patterns motivate the regression framework below.

Figure 3: Sentiment arcs before rate changes: ECB

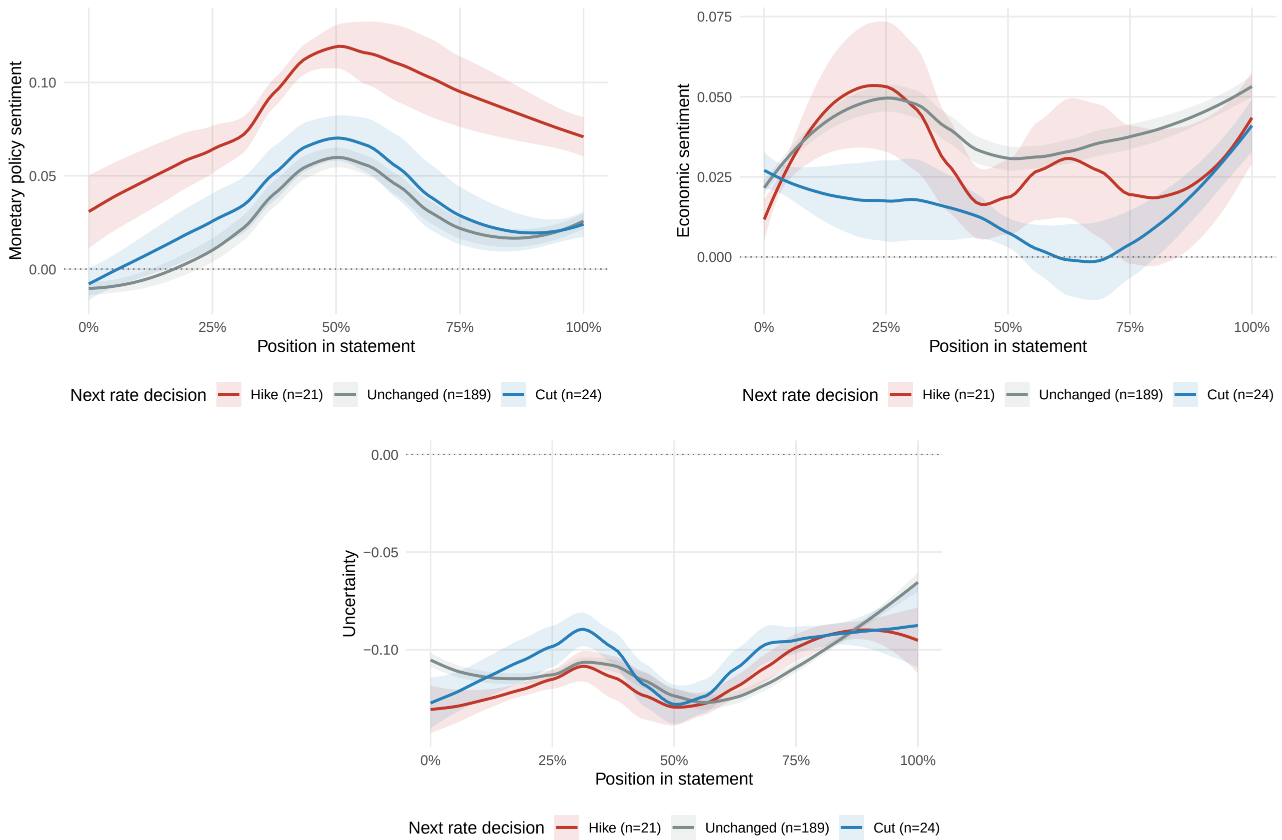

Notes: LOESS-smoothed (span = 0.6) sentiment arcs grouped by the subsequent policy rate decision: hike (red), unchanged (grey), cut (blue). Panels show monetary policy sentiment (hawkish positive), economic sentiment (optimistic positive), and uncertainty sentiment (uncertain positive). The x-axis is normalized position within statement (0% = opening, 100% = close). Shaded bands are 95% pointwise confidence intervals across statements.

Table 2 presents the model ladder for the rate change regression (4) for the ECB.

For the ECB, the benchmark M1 shows that the Picault–Renault economic sentiment score (`ec_pr`) enters with a positive and highly significant coefficient, while the hawkish–dovish indicator (`mp_pr`) is statistically insignificant, consistent with earlier findings in the literature: lexicon-based hawkishness is a noisy predictor of rate changes, while economic assessments carry cleaner signal. M2 turns to the mean of the CVP scores. Fit increases markedly, and both monetary and economic sentiment averages are highly significant, with economic sentiment entering with the larger coefficient. Introducing the arc features of the monetary policy sentiment dimension delivers the largest single gain in fit (adjusted $R^2 = 0.202$). The level of the arc (`mp_level`) enters significantly and positively, consistent with more hawkish statements preceding rate increases; slope and curvature are individually insignificant. The Hurst exponent for monetary policy sentiment (`hurst_mp`) enters with a negative and significant coefficient: statements with a more persistent and self-similar monetary policy sentiment arc — smooth rather than volatile sentiment trajectories — tend to be associated with lower subsequent rate changes. Model M4 adds arc features of the economic sentiment dimension. Here, the level, slope, and curvature of economic sentiment are all highly significant: a more positive economic

Table 2: Rate change regressions: ECB (Δi_{t+1})

	M1 (Bench.)	M2 (Mean)	M3 (MP)	M4 (MP+EC)	Full arc (MP+EC+UNC)
(Intercept)	-0.018	-0.095^{***}	0.605^{***}	0.422^{*}	0.295
	(-1.406)	(-3.761)	(3.807)	(2.378)	(1.430)
`mp_pr`	-0.017				
	(-0.607)				
`ec_pr`	0.215^{***}				
	(4.246)				
`mp_mean`		1.297^{***}			
		(3.968)			
`econ_mean`		1.398^{**}			
		(2.812)			
`mp_level`			2.251^{***}	3.800^{***}	4.839^{***}
			(5.579)	(6.992)	(8.219)
`mp_slope`			0.414	2.061^{***}	2.425^{***}
			(0.878)	(3.619)	(4.224)
`mp_curve`			-0.089	0.093	0.058
			(-0.961)	(0.883)	(0.541)
`hurst_mp`			-1.051^{***}	-0.814^{**}	-0.552^{*}
			(-4.107)	(-3.291)	(-2.229)
`econ_level`				5.270^{***}	4.239^{***}
				(4.988)	(3.706)
`econ_slope`				3.659^{***}	2.965^{***}
				(4.544)	(3.454)
`econ_curve`				0.780^{***}	0.417^{*}
				(5.220)	(2.326)
`hurst_econ`				-0.164	-0.180
				(-0.876)	(-0.971)
`unc_level`					-0.170
					(-0.183)
`unc_slope`					0.223
					(0.409)
`unc_curve`					-0.741^{***}
					(-3.581)
`hurst_unc`					-0.185
					(-0.846)
R^2	0.075	0.077	0.216	0.312	0.377
Adj. R^2	0.067	0.069	0.202	0.287	0.344
Residual SE	0.173	0.173	0.160	0.152	0.146
N			234		

Notes: OLS estimates. t-statistics in parentheses. Outcome: one-meeting-ahead policy rate change Δi_{t+1}. M1 uses Picault–Renault dictionary scores. M2 uses CVP document means. M3 adds Nelson–Siegel factors for the MP arc and `hurst_mp`. M4 extends M3 with NS factors for the economic arc and `hurst_econ`. Full arc extends M4 with NS factors for the uncertainty sentiment arc and `hurst_unc`. Significance: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $^{\cdot}p < 0.10$.

assessment (`econ_level`), an assessment that becomes progressively more optimistic over the course of the statement (`econ_slope`), and one that places greater emphasis on the analytical core (`econ_curve`) are each associated with subsequent policy tightening. Adding the economic arc features also renders the monetary arc slope significant and positive; monetary curvature remains insignificant throughout the ladder. Finally, in the Full arc specification, the curvature of the uncertainty sentiment arc (`unc_curve`) enters negatively and significantly: statements in which expressed uncertainty is concentrated at the statement's opening and close, leaving the analytical core comparatively calm, tend to be associated with smaller subsequent rate increases. The remaining uncertainty arc features — level, slope, and the Hurst exponent — are not precisely estimated. Arc shape along the monetary and economic sentiment dimensions, and along the curvature of the uncertainty dimension specifically, thus each contribute materially to rate predictions for the ECB. These results are also corroborated using the two-meeting-ahead rate change Δi_{t+2} as the dependent variable (see Appendix A.5).

Summing up, our results demonstrate that the *structure* of economic assessment within a statement — not just its average tone — encodes the forward policy signal. The next section tests this claim directly, more thoroughly controlling for the effect of topic content, which could potentially contaminate our arc-based findings.

4.1.1 Arc shape versus topic composition

A valid concern is the fact that the ECB's statement follows a stable template and that the arc features simply proxy for which topic is being discussed at which point, rather than capturing a genuinely rhetorical structure. Since a topic model would only yield a decomposition of the text for the whole document – ignoring the position in the document where a certain topic is addressed – we decided to resort to manual annotation of the text. For the ECB, this is rather straightforward, since the ECB follows standard templates that, while occasionally revised, make consistent annotation feasible.

Consequently, we categorize each statement into sections (i) Rate decision (which also contains a general introduction), (ii) Economic pillar (including inflation / price stability), (iii) Monetary pillar, (iv) Conclusion (cross-check of monetary and economic pillars) and (v) Structural. From July 2021 onward, the statement also contains a risk assessment category. To maximize our observations, we fold the risk assessment into the structural part. This allows us to map consistently onto both the pre-2003 and post-2003 template orderings. This coding is fully and consistently identifiable from 2002 onward.

In a second step, we then construct for each section the mean economic, monetary policy and uncertainty sentiment. This yields a model with 15 parameters, comparable in size to the 12-parameter Full arc specification. We then estimate three models on the common complete-case sample: the Full arc specification; a Section-average specification using these 15 topic-average regressors in place of the arc features; and a Combined specification containing both sets. Documents that could not be annotated are dropped from these three models (19 of 235, concentrated before 2021). We report the results in Table 3.

The table shows that in the combined specification, mean uncertainty sentiment plays a significant role in the rate-decision section, as does mean monetary policy sentiment in the economic section, whereas mean economic sentiment is highly significant in the structural section. Content-wise, the only sections with no significant sentiment dimensions are the conclusion and monetary sections.[6] Most importantly though, and comparing the adjusted R^2 values for the section average model and the model including only sentiment arc specifications, we see that arc features can explain more than twice as much variance as the section-average model can explain. This finding demonstrates that it is indeed the shape of the arcs, across the different dimensions, that predicts rate changes. The combined model further increases the fit, showing that content also plays a role. We also more formally test this using nested F-tests, which ask whether each set of regressors adds explanatory power once the other is already included. The continuous arc features remain significant beyond the section averages ($F(12, 187) = 6.99$, $p < 0.0001$). Section averages, in turn, add no significant explanatory power beyond the arc ($F(15, 187) = 1.71$, $p = 0.051$).

Summing up, our results thus indicate that arc features are not merely a more transparent decomposition based on the tone of different statement sections, but rather reflect policy signals embedded in the structure of the document.

4.1.2 Subsample and regime robustness

In this section we test whether the arc–rate relationship changes during specific sub-sample periods such as the zero lower bound on policy rates, or periods of acute macroeconomic stress. Consequently, we perform three checks, each augmenting the same Full arc baseline with a regime dummy, shown in Table 4.

[6] A reason for the latter could be that we did not explicitly introduce an arc dimension that reflects the monetary indicators discussed in that pillar (monetary aggregates, credit aggregates, etc.). The monetary pillar is highly ECB-specific, and our aim was to provide a more generalistic arc dimension that can be easily applied to other central banks' statements.

Table 3: Arc shape vs. topic composition, ECB: model comparison and coefficients (Δi_{t+1})

	Full arc	Section avg	Combined
(Intercept)	0.314	0.010	0.554*
	(1.410)	(0.065)	(2.194)
`mp_level`	4.803***		8.988*
	(7.736)		(2.524)
`mp_slope`	2.231***		3.249*
	(3.616)		(2.158)
`mp_curve`	0.018		0.831
	(0.160)		(1.296)
`hurst_mp`	−0.579*		−0.354
	(−2.166)		(−1.143)
`econ_level`	4.316***		13.023*
	(3.557)		(2.175)
`econ_slope`	3.044**		4.660*
	(3.312)		(2.056)
`econ_curve`	0.409*		1.912·
	(2.160)		(1.700)
`hurst_econ`	−0.251		−0.209
	(−1.221)		(−0.944)
`unc_level`	−0.250		−6.224
	(−0.245)		(−1.131)
`unc_slope`	0.006		−3.470
	(0.011)		(−1.550)
`unc_curve`	−0.778***		−2.249·
	(−3.496)		(−1.966)
`hurst_unc`	−0.163		−0.388
	(−0.699)		(−1.355)
`mp_Rate_decision`		−0.143	0.575
		(−0.116)	(0.320)
`econ_Rate_decision`		1.650	−0.140
		(0.924)	(−0.053)
`unc_Rate_decision`		1.443	6.242*
		(1.047)	(2.134)
`mp_Economic`		−2.313·	−4.697**
		(−1.656)	(−2.689)
`econ_Economic`		−1.510	−4.442·
		(−1.485)	(−1.877)
`unc_Economic`		−1.464	0.604
		(−1.153)	(0.260)
`mp_Monetary`		−0.521	−0.592
		(−0.612)	(−0.532)
`econ_Monetary`		−1.522	−1.643
		(−1.432)	(−1.318)
`unc_Monetary`		1.009	1.127
		(0.748)	(0.890)
`mp_Structural`		2.190·	−1.341
		(1.778)	(−0.905)
`econ_Structural`		0.441	−3.768*
		(0.343)	(−2.160)
`unc_Structural`		0.097	−1.240
		(0.086)	(−1.021)
`mp_Conclusion`		2.379**	0.509
		(2.887)	(0.603)
`econ_Conclusion`		4.046**	1.534
		(3.296)	(1.242)
`unc_Conclusion`		0.581	−0.042
		(0.670)	(−0.051)
R^2	0.386	0.218	0.460
Adj. R^2	0.350	0.159	0.382
N	215	215	215

Notes: OLS estimates. t-statistics in parentheses. Full arc = Nelson–Siegel level/slope/curvature + Hurst exponent for all three arc dimensions. Section avg = mean sentiment in each hand-coded category (Rate decision, Economic, Monetary, Structural, Conclusion). Combined includes both sets. All three models are re-fit on one shared complete-case sample so the nested F-tests (reported in text) compare like with like; this Full arc column is therefore not the paper's headline Full-arc table, which uses more documents (see main text). Significance: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $^{\cdot}p < 0.10$.

Table 4: Rate change regressions: subsample and regime robustness, ECB (Δi_{t+1})

	Full arc	$+D_{\mathrm{fg}}$	$+D_{\mathrm{new}}$	ZLB interact.	Excl. ZLB	Crisis interact.
(Intercept)	0.295	0.312	0.309	0.418·	0.383	0.402·
	(1.430)	(1.488)	(1.435)	(1.843)	(1.304)	(1.934)
mp_level	4.839***	4.812***	4.777***	5.698***	6.276***	4.507***
	(8.219)	(8.118)	(7.359)	(7.907)	(7.249)	(7.649)
mp_slope	2.425***	2.353***	2.364***	3.288***	3.788***	1.664**
	(4.224)	(3.954)	(3.726)	(3.936)	(3.758)	(2.867)
mp_curve	0.058	0.072	0.051	0.022	0.044	0.093
	(0.541)	(0.647)	(0.466)	(0.178)	(0.303)	(0.829)
hurst_mp	−0.552*	−0.595*	−0.562*	−0.495·	−0.360	−0.852***
	(−2.229)	(−2.252)	(−2.232)	(−1.689)	(−1.051)	(−3.522)
econ_level	4.239***	4.244***	4.302***	4.148***	5.303**	4.511***
	(3.706)	(3.704)	(3.651)	(3.415)	(3.007)	(3.908)
econ_slope	2.965***	2.972***	2.991***	2.793**	3.755**	3.158***
	(3.454)	(3.455)	(3.448)	(2.973)	(2.695)	(3.787)
econ_curve	0.417*	0.408*	0.418*	0.408*	0.491·	0.203
	(2.326)	(2.260)	(2.326)	(2.125)	(1.953)	(1.155)
hurst_econ	−0.180	−0.172	−0.176	−0.173	−0.422	−0.036
	(−0.971)	(−0.920)	(−0.942)	(−0.928)	(−1.596)	(−0.191)
unc_level	−0.170	−0.047	−0.071	0.783	−0.056	0.731
	(−0.183)	(−0.049)	(−0.070)	(0.676)	(−0.036)	(0.762)
unc_slope	0.223	0.281	0.303	0.958	1.020	0.252
	(0.409)	(0.502)	(0.468)	(1.278)	(0.945)	(0.432)
unc_curve	−0.741***	−0.728***	−0.740***	−0.678**	−0.863**	−0.418*
	(−3.581)	(−3.477)	(−3.568)	(−2.852)	(−2.789)	(−1.996)
hurst_unc	−0.185	−0.173	−0.197	−0.233	−0.272	−0.025
	(−0.846)	(−0.785)	(−0.874)	(−1.044)	(−0.963)	(−0.103)
D_fg		0.013				
		(0.471)				
D_new			0.014			
			(0.231)			
D_zlb				−0.231		
				(−0.491)		
D_zlb×mp_level				−0.823		
				(−0.788)		
D_zlb×mp_slope				−0.707		
				(−0.521)		
D_zlb×mp_curve				0.566*		
				(2.218)		
D_zlb×hurst_mp				0.366		
				(0.483)		
crisis						−0.823
						(−1.161)
mp_level×crisis						2.944
						(1.088)
mp_slope×crisis						5.484*
						(2.217)
mp_curve×crisis						0.341
						(0.894)
hurst_mp×crisis						2.751**
						(3.156)
econ_level×crisis						4.012
						(0.675)
econ_slope×crisis						6.000
						(1.222)
econ_curve×crisis						1.173
						(1.407)
hurst_econ×crisis						−0.880
						(−1.238)
unc_level×crisis						−2.040
						(−0.554)
unc_slope×crisis						3.200
						(1.281)
unc_curve×crisis						−1.212
						(−1.511)

Continued on next page

Table 4 continued

	Full arc	$+D_{\text{fg}}$	$+D_{\text{new}}$	ZLB interact.	Excl. ZLB	Crisis interact.
`hurst_unc×crisis`						-1.201^{*} (-2.073)
R^2	0.377	0.378	0.377	0.412	0.436	0.501
Adj. R^2	0.344	0.341	0.341	0.366	0.394	0.441
Residual SE	0.146	0.146	0.146	0.143	0.161	0.134
N	234	234	234	234	171	234

Notes: OLS estimates. t-statistics in parentheses. Outcome: Δi_{t+1}. Full arc: baseline model. $+D_{\text{fg}}/+D_{\text{new}}$: regime dummy ($D_{\text{fg}} = 1$ from 4 July 2013; $D_{\text{new}} = 1$ from 22 July 2021). ZLB interact.: D_{zlb} (= 1, 10 Sep 2014–9 Jun 2022) and its interaction with the monetary-arc terms. Excl. ZLB: Full arc re-estimated excluding ZLB periods. Crisis interact.: crisis dummy (EABC: GFC 2008/04–2009/03, sovereign debt crisis 2011/07–2013/03, COVID 2020/01–2020/06) and its interaction with the level/slope/curvature and Hurst-exponent arc terms. Significance: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $^{\cdot}p < 0.10$.

ECB communication-regime dummies. The ECB's communication framework underwent two significant structural shifts during our sample period: the introduction of explicit forward guidance on 4 July 2013, and the adoption of a new monetary policy strategy announced on 22 July 2021, which resulted in a slightly modified structure of the introductory statements as explained in Section 2. We augment the Full arc specification with two dummy variables: `D_fg` = 1 for all statements from 4 July 2013 onward, and `D_new` = 1 for all statements from 22 July 2021 onward (columns $+D_{\text{fg}}/+D_{\text{new}}$). The arc coefficients are stable relative to the baseline: the signs and significance of the key predictors – `mp_level`, `mp_slope`, `hurst_mp`, `econ_level`, `econ_slope`, `econ_curve`, and the curvature of the uncertainty arc – are preserved when either regime dummy is included, and neither regime dummy itself is statistically significant. Overall, these results complement the topic-content regression reported above and confirm that the arc–rate relationship reflects a stable communication pattern rather than an artefact of the two structural breaks in ECB communication design.

Zero lower bound. During ZLB periods the policy rate is mechanically constrained, so a statement may convey a clear easing or tightening signal while Δi_t is pinned near zero, potentially altering the composition of the estimated arc–rate relationship. We augment the Full arc specification with a ZLB regime dummy D_{zlb} and its interaction with the four *monetary policy-arc* terms (`mp_level`, `mp_slope`, `mp_curve`, `hurst_mp`); $D_{\text{zlb}} = 1$ for ECB meetings 10 September 2014–9 June 2022, when the policy rate going into the decision was at its effective floor (verified directly against the policy-rate series). We also re-estimate the Full arc model excluding ZLB periods entirely (column Excl. ZLB). The joint test of the five ZLB terms against the Full arc model is statistically significant ($F(5, 216) = 2.58$, $p = 0.027$). This implies that the arc–rate relationship is not fully invariant to the ZLB regime. More specifically, the interaction of monetary policy sentiment curvature with the ZLB dummy (`mp_curve`$\times D_{\text{zlb}}$) is a stronger predictor of subsequent rate changes specifically during ZLB periods. That said, all terms of the Full arc remain significant with the same sign as in the baseline estimation, with two exceptions: the Hurst exponent of monetary policy sentiment loses significance, and the curvature of economic sentiment weakens to only marginal significance once ZLB periods are excluded entirely. This holds also true when considering the model that excludes the ZLB period explicitly. Summing up, the arc–rate relationship survives ZLB conditioning for the ECB, which further lends credence to our main results.

Communication discipline under stress. Last, we examine whether the ECB communicates differently during periods of macroeconomic stress. For that purpose, we draw on data from the Euro Area Business Cycle Network (EABC) dating committee to define three crisis episodes: the Global Financial Crisis (April 2008–March 2009), the euro area sovereign debt crisis (July 2011–March 2013), and the COVID crisis (January–June 2020). Looking at the crisis arcs (Figure 6, Appendix A.6) reveals that economic sentiment arcs are consistently more negative during crisis periods than during normal periods. A more formal test of any systematic differences is carried out in Table 4, column six. Here, we introduce a crisis dummy variable and interact it with the arc features.

While the crisis dummy itself is not significant, there are some significant interaction terms. Among those, it is the persistence dimension that shows the clearest crisis-period effect: the Hurst exponent of monetary policy sentiment interacts positively and significantly with the crisis dummy, and the Hurst exponent of the uncertainty arc interacts negatively and significantly – crisis periods sharpen how strongly arc persistence along both dimensions predicts subsequent rate changes. The slope of monetary policy sentiment also strengthens significantly during crises. More importantly though, the main arc features from the Full arc specification remain significant in the crisis-interaction specification and also show the same sign. This implies that the main relationship is also present under economic stress, pointing to consistent communication by the ECB.

4.2 Arc structure and audience responses: Inflation expectations

We now turn to the receiver side: does the arc structure of a statement affect how audiences process and respond to it? We focus on professional forecasters updating inflation expectations. If the arc carries purely internal policy signals, its predictive content should be exhausted by the rate regressions above; finding effects on audience responses would instead suggest that the sequential narrative structure shapes perception beyond the policy content itself.

More specifically, we examine whether arc features predict professional inflation expectations from Consensus Economics. These forecasters feature commercial banks, research institutions or think tanks, who are more likely than households to track central bank statements closely.[7] We consider two forecast horizons to separate two possible channels. The short-run outcome is the next-month Consensus Economics one-year-ahead inflation expectation $\pi_{t+1}^{e,\mathrm{ST}}$ and the long-run outcome $\pi_{t+1}^{e,\mathrm{LT}}$ – the Consensus Economics 6–10-year-ahead average inflation forecast, which is only surveyed twice a year.

As shown in equation (5), the timing of the survey means we always match to the introductory

[7] On household attention to central bank communication, see for example Hajdini et al. (2026) using a novel survey or more indirectly Feldkircher et al. (2025) linking Google search data to central bank speech events.

statement from the preceding month. For the short-run regression, when multiple statements fall in the same month we take the latest one, yielding $N = 224$ matched observations; for the long-run regression, we have $N = 46$ matched surveys after dropping two: the October 2025 survey, whose nearest preceding statement is from June 2025, a stale match caused by our ECB corpus's June 2025 collection cutoff (Section 2) and the October 2006 survey, for which the ECB genuinely held no meeting in September 2006 (its nearest preceding statement is dated 31 August 2006). Table 5 reports both regressions side by side.

At the short-run horizon, actual inflation Δp_t^{yoy} dominates the regression, as expected given the high persistence of professional forecasters' beliefs (adj. $R^2 = 0.897$ in M1). Conditional on that level effect, the Picault–Renault scores in M1 enter positively for both monetary and economic sentiment, consistent with hawkish and optimistic ECB communication being associated with higher near-term inflation expectations. Adding the arc features raises fit further (adj. $R^2 = 0.915$), and the within-statement structure of the *monetary policy sentiment* arc is informative: level, slope, and curvature all enter positively and significantly, indicating that statements which are not just hawkish on average but also directionally and non-linearly hawkish are associated with higher expected inflation. The direction of our findings is consistent with that based on the dictionary-based benchmark indicators. For the economic sentiment arc, only the curvature factor enters significantly and positively — non-linear optimism about the economic outlook, concentrated in the analytical core of the statement, provides incremental signal about future expectations. From the uncertainty dimension, the curvature enters negatively and significantly, echoing the rate regressions: non-linear structure in how uncertainty is communicated carries incremental predictive content for the ECB, here as well as for subsequent policy. That said and at the short-term horizon, inflation expectations are generally more volatile since macroeconomic news can change abruptly. In that sense, significant arc features could reflect that the central bank has new information, which then triggers forecasters to revise their forecasts – independent of any communication design.

It is hence important to also test the effects of communication on the long-run outcome $\pi_{t+1}^{e,\text{LT}}$. Long-run inflation expectations should be anchored around the central bank's inflation target and, as the focus of credibility-oriented communication, have much less reason to move on near-term information alone. At the long-run horizon, the Full arc specification dramatically outperforms the Picault–Renault benchmark: adjusted R^2 rises from 0.132 to 0.374. Given the small sample, individual coefficients should be read with caution (32 residual degrees of freedom against 13 regressors in the Full arc column), but the overall gain in fit is itself informative: arc structure predicts long-run, credibility-anchored expectations, not only near-term ones – exactly the pattern the Delphic-effect concern would not predict. At the long-run forecast horizon, the uncertainty arc is particularly informative: its curvature enters negatively and significantly, while its level and Hurst exponent are only marginal ($p \approx 0.06$ for each). The consistency of the uncertainty-curvature channel across both horizons – one where information revelation is a live alternative explanation, one where it is much less plausible – is significant

Table 5: Inflation expectations regressions: short-run vs. long-run, ECB

	Short-run ($\pi^{e,\mathrm{ST}}_{t+1}$)		**Long-run** ($\pi^{e,\mathrm{LT}}_{t+1}$)	
	M1	Full arc	M1	Full arc
(Intercept)	0.689***	1.049*	1.923***	1.763***
	(17.864)	(2.228)	(88.989)	(6.470)
$\Delta p_t^{\mathrm{yoy}}$	0.562***	0.500***	0.016*	0.010
	(43.090)	(21.660)	(2.442)	(0.872)
`mp_pr`	0.114*		0.028	
	(2.078)		(0.912)	
`ec_pr`	0.488***		−0.103	
	(4.833)		(−1.559)	
`mp_level`		7.084***		0.703
		(4.047)		(0.751)
`mp_slope`		5.877***		0.410
		(4.350)		(0.563)
`mp_curve`		1.311***		0.122
		(5.093)		(0.815)
`hurst_mp`		−0.500		0.189
		(−0.890)		(0.536)
`econ_level`		1.490		0.346
		(0.591)		(0.265)
`econ_slope`		−0.317		1.428
		(−0.167)		(1.469)
`econ_curve`		1.228**		−0.021
		(3.095)		(−0.118)
`hurst_econ`		−0.122		0.060
		(−0.296)		(0.216)
`unc_level`		−2.007		−2.535·
		(−0.990)		(−1.987)
`unc_slope`		0.376		−0.199
		(0.319)		(−0.319)
`unc_curve`		−1.015*		−0.710*
		(−2.225)		(−2.474)
`hurst_unc`		−0.296		−0.541·
		(−0.617)		(−1.998)
Adj. R^2	0.897	0.915	0.132	0.374
Residual SE	0.340	0.310		
N	224	224	46	46

Notes: OLS estimates. t-statistics in parentheses. Short-run outcome $\pi^{e,\mathrm{ST}}_{t+1}$: next-month Consensus Economics one-year-ahead inflation expectation. Long-run outcome $\pi^{e,\mathrm{LT}}_{t+1}$: Consensus Economics 6–10-year-ahead average inflation forecast, surveyed twice yearly (April/October), matched to the most recent press statement dated strictly before the 1st of the survey month, and requiring that statement to fall in the calendar month immediately preceding the survey; two surveys (October 2006, October 2025) are dropped because no ECB statement falls in the required preceding month (see main text). $\Delta p_t^{\mathrm{yoy}}$ is actual year-over-year euro area HICP inflation in month t in both columns, controlling for the persistent level of expectations. M1 uses Picault–Renault dictionary scores; Full arc uses Nelson–Siegel factors and Hurst exponent for all three arc dimensions. Significance: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $^{\cdot}p < 0.10$.

evidence that arc structure carries something beyond near-term information content.

Finally and in analogy to the sender regressions, we test via an F-test whether these results are driven by topic content as opposed to communication design. We do this for the short-run inflation expectations regression only, due to lack of observations for the long-run specification. Our results show that arc terms remain jointly significant beyond topic composition ($F(12, 176) = 2.64$, $p = 0.0028$) on a matched sample of $N = 205$ (versus $N = 224$ for the headline specification, the same complete-case attrition as in Section 4.1.1), confirming that the arc's incremental content for short-term inflation expectations is not an artefact of topic composition.

4.3 Institutional extension: the case for the Fed

In this section, we report as an extension sentiment arcs for the Fed. Since the Fed only started in 2011 to do press conferences after FOMC meetings, the sample is considerably shorter. Despite the considerably smaller number of available observations, we include results to cross-check whether we can corroborate the main findings for the ECB, namely that arc shape predicts subsequent policy action and shapes how audiences respond – with another institution.

Table 6 reports both sides together in one compact specification – benchmark and Full arc only, no intermediate ladder steps.

On the sender side, the benchmark M1 has essentially no explanatory power (Adj. R^2 slightly negative), while the Full arc specification reaches Adj. $R^2 = 0.378$: the level and curvature of monetary policy sentiment (`mp_level`, `mp_curve`) are both significant at the 0.1% level, `mp_slope` only marginally so, and none of the economic-sentiment or uncertainty-sentiment arc features reach conventional significance – the Fed's arc signal is concentrated narrowly in the monetary policy sentiment dimension, unlike the ECB's richer multi-dimensional pattern.

On the receiver side, actual US inflation (Δp_t^{yoy}) dominates as expected (Adj. $R^2 = 0.932$ in M1 alone), but conditional on that level effect, arc shape still adds explanatory power (Adj. $R^2 = 0.944$ in the Full arc specification): the level and curvature of monetary policy sentiment are again the informative terms (`mp_level`, `mp_curve`), both entering positively – a more hawkish, more non-linearly structured monetary arc is associated with higher near-term US inflation expectations, the same direction as the ECB's short-run result and consistent with the same information-channel reading (Section 4.2). Consistent with findings for the ECB, the uncertainty dimension plays a role for short-term inflation expectations. Here, the Hurst exponent of the uncertainty arc is marginally significant, a different informative term from the ECB's uncertainty-curvature channel, but pointing to the same broad conclusion: the sequential structure of the statement carries information for forecasters beyond its average tone, replicating – in narrower form, on a much smaller sample – the paper's central receiver-side finding in a second institutional setting.

Due to the limited number of observations, we refrain from executing the sub-sample, long-run

Table 6: Fed institutional extension: sender and receiver results, corrected sample

	Sender (Δi_{t+1})		Receiver ($\pi^{e,\mathrm{ST}}_{t+1}$)	
	M1 (Bench.)	Full arc	M1 (Bench.)	Full arc
(Intercept)	0.088*	0.051	1.182***	0.556
	(2.060)	(0.106)	(14.264)	(0.783)
`mp_pr`	0.141		0.197	
	(1.275)		(1.356)	
`ec_pr`	−0.003		−0.336·	
	(−0.022)		(−1.811)	
`mp_level`		3.830***		4.152**
		(5.072)		(3.339)
`mp_slope`		1.134·		0.848
		(1.722)		(0.937)
`mp_curve`		0.635***		0.841**
		(3.939)		(3.122)
`hurst_mp`		0.124		−0.206
		(0.202)		(−0.210)
`econ_level`		1.897		3.326
		(1.271)		(1.616)
`econ_slope`		−0.471		0.966
		(−0.511)		(0.740)
`econ_curve`		0.017		−0.061
		(0.057)		(−0.146)
`hurst_econ`		−0.567		−0.081
		(−1.423)		(−0.148)
`unc_level`		−0.045		1.558
		(−0.040)		(0.951)
`unc_slope`		−0.609		0.480
		(−0.886)		(0.463)
`unc_curve`		−0.337		−0.391
		(−1.307)		(−1.065)
`hurst_unc`		0.106		1.057·
		(0.268)		(1.936)
$\Delta p_t^{\mathrm{yoy}}$			0.501***	0.468***
			(31.114)	(20.474)
R^2	0.019	0.465	0.934	0.952
Adj. R^2	-0.004	0.378	0.932	0.944
N	87	87	86	86

Notes: OLS estimates. t-statistics in parentheses. Fed, corrected press-conference-transcript sample. Sender outcome: one-meeting-ahead policy rate change. Receiver outcome: next-month Consensus Economics one-year-ahead US inflation expectation, matched the same way as the ECB regression (latest statement per month, actual YoY CPI inflation as a control); long-run Fed inflation expectations are not attempted (sample too small to match against the biannual long-run survey). Full arc: standard 12-parameter specification (Nelson–Siegel level/slope/curvature + Hurst exponent, all three dimensions). Significance: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $^{\cdot}p < 0.10$.

inflation expectations and content robustness regressions – the latter is also complicated as the Fed press conferences do not follow such consistent templates as does the ECB. We do, however, provide additional results in the appendix: descriptive statistics and the arc figures, reported in full in Appendix A.7.

5 Conclusions

This paper asks whether the *shape* of sentiment in a central bank statement – how tone builds, turns, and settles across a statement, not just its average level – carries information about monetary policy beyond what aggregate sentiment measures capture. It does, and on both sides of communication. On the sender side – how the central bank's own choices in constructing a statement relate to its next policy move – arc shape robustly predicts the ECB's subsequent rate decision, through the level and slope of monetary policy sentiment, the persistence of the hawkish–dovish signal, and the level, slope, and curvature of economic sentiment. On the receiver side – how outside audiences interpret and respond to that same communication – arc features predict how professional forecasters revise their short-term inflation expectations. At the long-run, credibility-anchored horizon, it is specifically the uncertainty dimension – and its curvature in particular – that carries predictive power, rather than arc features more broadly. Neither result is an artefact of easy alternative explanations: arc shape remains significant once we control directly for the topic content of each section of the statement, and its predictive power for inflation expectations extends to that long-run horizon, where near-term information revelation is an implausible alternative story. The ECB relationship also proves stable under stress: the key arc predictors are preserved across the ECB's two communication-regime shifts, the arc–rate relationship survives zero-lower-bound conditioning and, during episodes of acute macroeconomic stress, it is the persistence dimension that sharpens rather than weakens.

As a cross-check on a considerably smaller sample, the same sender- and receiver-side pattern replicates for the Federal Reserve. The signal there is narrower, concentrated almost entirely in the monetary policy sentiment dimension, with no comparable role for economic sentiment, uncertainty, or persistence – consistent with a more disciplined, template-driven communication style than the ECB's richer, multi-dimensional one.

Taken together, our findings point to a feature of central bank communication that the literature has largely overlooked: within-document structure. Sentiment arcs – and in particular their persistence and curvature – encode forward-looking policy signals that aggregate, whole-statement sentiment measures suppress. For practitioners, this means that communication design choices, such as where hawkish or uncertain language is placed within a statement and how consistently a stance is sustained from opening to close, carry measurable consequences for how markets and forecasters read the stance – a lever available to central banks independently of the substance of what is actually being said.

References

Altavilla, C., Brugnolini, L., Gürkaynak, R. S., Motto, R., and Ragusa, G. (2019). Measuring euro area monetary policy. *Journal of Monetary Economics*, 108:162–179.

Babakhin, Y., Osmulski, R., Ak, R., Moreira, G., Xu, M., Schifferer, B., Liu, B., and Oldridge, E. (2025). Llama-embed-nemotron-8b: A universal text embedding model for multilingual and cross-lingual tasks.

Baranowski, P., Bennani, H., and Doryń, W. (2021). Do the ECB's introductory statements help predict monetary policy? Evidence from a tone analysis. *European Journal of Political Economy*, 66:101964.

Bizzoni, Y., Nielbo, K., Peura, T., and Thomsen, M. R. (2022). The fractality of sentiment arcs for literary quality assessment: The case of Nobel laureates. In *Proceedings of the 2nd International Workshop on Natural Language Processing for Digital Humanities (NLP4DH)*, pages 31–41, Taipei, Taiwan. Association for Computational Linguistics.

Blinder, A. S., Ehrmann, M., de Haan, J., and Jansen, D.-J. (2024). Central bank communication with the general public: Promise or false hope? *Journal of Economic Literature*, 62(2):425–457.

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):910–45.

Cho, D. and Jung, J. (2026). Mind the tone: Responses of inflation expectations to central bankers' speeches. *Journal of International Money and Finance*, 160:103452.

Cho, D. and Rho, S. (2026). Can the tone of central bankers' speeches help shape inflation expectations? Evidence from Japan. *Journal of International Financial Markets, Institutions & Money*, 107:102283.

Conrad, C. and Lamla, M. J. (2010). The high-frequency response of the EUR-USD exchange rate to ECB communication. *Journal of Money, Credit and Banking*, 42(7):1391–1417.

Consoli, S., Barbaglia, L., and Manzan, S. (2022). Fine-grained, aspect-based sentiment analysis on economic and financial lexicon. *Knowledge-Based Systems*, 247:108781.

Curti, F. and Kazinnik, S. (2023). Let's face it: Quantifying the impact of nonverbal communication in fomc press conferences. *Journal of Monetary Economics*, 139:110–126.

Diebold, F. X. and Li, C. (2006). Forecasting the term structure of government bond yields. *Journal of Econometrics*, 130(2):337–364.

Ehrmann, M. and Fratzscher, M. (2009). Explaining monetary policy in press conferences. *International Journal of Central Banking*, 5(2):41–84.

Ehrmann, M., Holton, C., Kedan, D., and Phelan, G. (2024). Monetary policy communication: Perspectives from former policymakers at the ECB. *Journal of Money, Credit and Banking*, 56(4):837–864.

Enevoldsen, K., Chung, I., Kerboua, I., Kardos, M., Mathur, A., Stap, D., Gala, J., Siblini, W., Krzemiński, D., Winata, G. I., Sturua, S., Utpala, S., Ciancone, M., Schaeffer, M., Misra, D., Dhakal, S., Rystrøm, J., Solomatin, R., Çağatan, Ö. V., Kundu, A., Bernstorff, M., Xiao, S., Sukhlecha, A., Pahwa, B., Poświata, R., GV, K. K., Ashraf, S., Auras, D., Plüster, B., Harries, J. P., Magne, L., Mohr, I., Zhu, D., Gisserot-Boukhlef, H., Aarsen, T., Kostkan, J., Wojtasik, K., Lee, T., Suppa, M., Zhang, C., Rocca, R., Hamdy, M., Michail, A., Yang, J., Faysse, M., Vatolin, A., Thakur, N., Dey, M., Vasani, D., Chitale, P. A., Tedeschi, S., Tai, N., Snegirev, A., Hendriksen, M., Günther, M., Xia, M., Shi, W., Lù, X. H., Clive, J., K, G., Anna, M., Wehrli, S., Tikhonova, M., Panchal, H. S., Abramov, A., Ostendorff, M., Liu, Z., Clematide, S., Miranda, L. J. V., Fenogenova, A., Song, G., Safi, R. B., Li, W.-D., Borghini, A., Cassano, F., Hansen, L., Hooker, S., Xiao, C., Adlakha, V., Weller, O., Reddy, S., and Muennighoff, N. (2025). MMTEB: Massive multilingual text embedding benchmark. In *The Thirteenth International Conference on Learning Representations*.

Feldkircher, M., Korab, P., and Teliha, V. (2025). What do central bankers talk about? Evidence from the BIS archive. CAMA Working Paper 35/2025, Centre for Applied Macroeconomic Analysis (CAMA), Crawford School of Public Policy.

Gardner, B., Scotti, C., and Vega, C. (2022). Words speak as loudly as actions: Central bank communication and the response of equity prices and expectations. *Journal of Econometrics*, 232(2):471–490.

Gorodnichenko, Y., Pham, T., and Talavera, O. (2023). The voice of monetary policy. *American Economic Review*, 113(2):548–84.

Hajdini, I., Knotek, E. S., Mitchell, J., Pedemonte, M., and Shiroff, T. (2026). Consumers' attention to monetary policy: The importance of having "Skin in the Game". Mimeo.

Hansen, A. L. and Kazinnik, S. (2024). Can chatgpt decipher fedspeak? Available at SSRN.

Hoffmann, M., Moench, E., Pavlova, L., and Schultefrankenfeld, G. (2026). A KISS for central bank communication in times of high inflation. *Journal of International Money and Finance*, 165:103585.

Holm-Hadulla, F., Musso, A., Rodriguez-Palenzuela, D., and Vlassopoulos, T. (2021). Evolution of the ECB's analytical framework. Occasional Paper Series 277, European Central Bank.

Hubert, P. and Labondance, F. (2021). The signaling effects of central bank tone. *European Economic Review*, 133(C).

Kanelis, D. and Siklos, P. L. (2024). The ECB press conference statement: deriving a new sentiment indicator for the euro area. *International Journal of Finance & Economics*, 30:652–664.

Lyngbaek, L., Feldkamp, P., Bizzoni, Y., Nielbo, K. L., and Enevoldsen, K. (2025). Continuous sentiment scores for literary and multilingual contexts. *Anthology of Computers and the Humanities*, 3:480–497.

Montes, G. C., Oliveira, L. V., Curi, A., and Nicolay, R. T. d. F. (2016). Effects of transparency, monetary policy signalling and clarity of central bank communication on disagreement about inflation expectations. *Applied Economics*, 48(7):590–607.

Müller, U. W. D., Witteman, C. L. M., Spijker, J., and Alpers, G. W. (2019). All's bad that ends bad: There is a peak-end memory bias in anxiety. *Frontiers in Psychology*, 10:1272.

Neugarten, J., Karsdorp, F., and Bloem, J. (2025). Happily ever after: Comparing sentiment arcs in emotionally-inflected fanfiction genres across fandoms. *Anthology of Computers and the Humanities*, 3:771–793.

Park, K., Choe, Y. J., and Veitch, V. (2024). The linear representation hypothesis and the geometry of large language models. In *Proceedings of the 41st International Conference on Machine Learning*, ICML'24. JMLR.org.

Parle, C. (2022). The financial market impact of ECB monetary policy press conferences — A text based approach. *European Journal of Political Economy*, 74:102230.

Picault, M. and Renault, T. (2017). Words are not all created equal: A new measure of ECB communication. *Journal of International Money and Finance*, 79:136–156.

Reagan, A. J., Mitchell, L., Kiley, D., Danforth, C. M., and Dodds, P. S. (2016). The emotional arcs of stories are dominated by six basic shapes. *EPJ Data Science*, 5:31.

Rosa, C. and Verga, G. (2007). On the consistency and effectiveness of central bank communication: Evidence from the ECB. *European Journal of Political Economy*, 23(1):146–175.

Shah, A., Paturi, S., and Chava, S. (2023). Trillion dollar words: A new financial dataset, task & market analysis. In Rogers, A., Boyd-Graber, J., and Okazaki, N., editors, *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6664–6679, Toronto, Canada. Association for Computational Linguistics.

Siklos, P. L. (2013). Sources of disagreement in inflation forecasts: An international empirical investigation. *Journal of International Economics*, 90(1):218–231.

Silva, T. C., Moriya, K., and Veyrune, R. M. (2025). From text to quantified insights: A large-scale llm analysis of central bank communication. IMF Working Paper 2025/109, International Monetary Fund. Accessed June 9, 2025.

Wabitsch, A. (2025). The messenger matters. Discussion paper, SSRN. Available at SSRN: https://ssrn.com/abstract=5238987.

Declaration of AI Use

During the preparation of this manuscript, the authors used Claude (Anthropic) to assist with data-processing scripts, exploratory robustness checks, and editing of the manuscript text. All AI-assisted content and code were reviewed, verified against the underlying data and estimation output, and edited by the authors, who take full responsibility for the content of this publication.

A Appendix A

A.1 Seed phrases

Table 7: Seed phrases: Monetary policy sentiment

Hawkish monetary policy (11 seeds)
Inflation remains too high and requires further policy tightening.
Rising inflationary pressures call for higher interest rates.
Strong wage growth poses upside risks to inflation.
Financial conditions remain loose given current inflation dynamics.
Interest rates will need to increase further to ensure price stability.
Persistent inflation requires a firm and timely policy response.
Elevated inflation expectations warrant decisive action.
Price pressures are broadening and justify higher rates.
Monetary tightening is necessary to prevent overheating.
We are prepared to raise interest rates if inflation accelerates.
Further rate increases are likely to be appropriate in the coming meetings.
Dovish monetary policy (13 seeds)
Economic conditions warrant maintaining a low policy rate.
We will continue asset purchases to support favorable financing conditions.
Policy rates will remain at current or lower levels for an extended period.
Accommodative monetary policy remains necessary to support the recovery.
Weak demand and subdued inflation call for supportive policy measures.
Liquidity provisions will continue to safeguard smooth market functioning.
We will reinvest maturing securities for as long as needed.
Forward guidance indicates continued policy accommodation.
Low interest rates remain appropriate while inflation is below target.
Asset purchases help maintain favorable financial conditions.
The current stance supports lending and economic activity.
Liquidity operations will continue to ease financing conditions.
Maintaining accommodative policy will help restore inflation to target.

Notes: Seed phrases anchoring the monetary policy sentiment dimension of the CVP scoring. Seeds are manually curated and validated by checking that each phrase scores in the expected direction and does not cross-load heavily onto the economic sentiment dimension.

Table 8: Seed phrases: Economic sentiment

Economic sentiment: positive (14 seeds)
The economy is expanding at a solid pace with strong job growth.
Consumer demand remains strong and supports continued growth.
Labor-market conditions are strong and improving.
Business investment continues to strengthen.
Economic momentum has increased across many sectors.
Job gains remain robust and unemployment is low.
Financial conditions remain supportive of growth.
Economic indicators point to a sustained expansion.
Productivity gains continue to support economic growth.
Confidence indicators suggest improving economic outlook.
Household spending is strong and continues to improve.
Exports and external demand remain supportive.
Economic activity continues to progress at a solid pace.
Broad-based growth supports continued recovery.
Economic sentiment: negative (13 seeds)
Weak demand and tight credit conditions are restraining growth.
Household spending remains constrained by low income growth.
Financial-market stress is weighing on economic activity.
Job losses and declining wealth are weakening sentiment.
Rising energy prices are reducing purchasing power.
Industrial production and exports are weakening.
Growth remains fragile and uneven across sectors.
Economic activity is slowing and confidence is deteriorating.
Credit conditions remain tight and continue to constrain spending.
Inflation remains below target and reflects economic weakness.
Weak investment and productivity are dampening outlook.
Economic recovery remains slow and vulnerable.
Business surveys signal worsening economic conditions.

Notes: Seed phrases anchoring the economic sentiment dimension. Seeds are validated to score near zero on the monetary policy sentiment and uncertainty sentiment axes.

Table 9: Seed phrases: Uncertainty sentiment

High uncertainty (15 seeds)
The outlook remains highly uncertain.
The recovery process is likely to be uneven and subject to high uncertainty.
Economic growth is expected to remain uneven in an environment of uncertainty.
The outlook is subject to particularly high uncertainty and intensified downside risks.
The near-term economic outlook remains clouded by uncertainty.
The speed and scale of the recovery remain highly uncertain.
Pandemic-related uncertainty is likely to dampen the recovery in consumption, investment, and labour markets.
Geopolitical and financial uncertainty continue to weigh on the growth outlook.
Heightened uncertainty makes policy flexibility especially important.
Substantial uncertainty surrounds the timing and pace of the improvement.
The course of the virus and the outlook for the economy remain highly uncertain.
Significant uncertainty surrounds the strength of final demand.
Business investment remains a major source of uncertainty for the overall outlook.
The future course of the economy is subject to a marked degree of uncertainty.
Elevated uncertainty around the economic outlook justifies maintaining the policy stance.
Low uncertainty / high confidence (13 seeds)
Available information is broadly in line with the baseline scenario.
Inflation is expected to remain moderate and consistent with price stability.
Risks to the inflation outlook are broadly balanced.
Inflation expectations remain firmly anchored.
Current inflation developments are in line with previous expectations.
Wage and price developments remain subdued and consistent with price stability.
The Governing Council stands ready to act to preserve price stability and keep expectations anchored.
Inflation is expected to stabilize around the Committee's 2 percent objective over the medium term.
Longer-term inflation expectations remain well anchored.
Incoming data are broadly in line with staff expectations.
The current policy stance remains appropriate given the baseline outlook.
Resource slack is expected to keep inflation contained.
Financial markets largely anticipate no change in the policy rate.

Notes: Seed phrases anchoring the uncertainty sentiment dimension (uncertain positive, certain negative). Seeds are validated to score near zero on the monetary policy sentiment and economic sentiment axes.

A.2 Seed phrase validation: dimension orthogonality

A key requirement of the CVP approach is that the three sentiment dimensions — monetary policy sentiment, economic sentiment, and uncertainty sentiment — are sufficiently orthogonal to be treated as distinct signals. Figure 4 shows pairwise scatters of the CVP scores assigned to each seed phrase on all three axes. Seeds are coloured and shaped by category. If the dimensions are orthogonal, seeds belonging to one dimension (e.g. MP hawk/dove) should score near zero on the other two axes (econ, uncertainty), and vice versa. The scatter confirms this pattern: seeds cluster in their expected quadrant on their own dimension and remain close to zero on the other axes, with only minor cross-loading for a small number of seeds that were subsequently removed from the final set.

Figure 4: Seed phrase validation: pairwise CVP score scatter

(a) MP vs. economic sentiment **(b)** MP vs. uncertainty sentiment **(c)** Economic sentiment vs. uncertainty sentiment

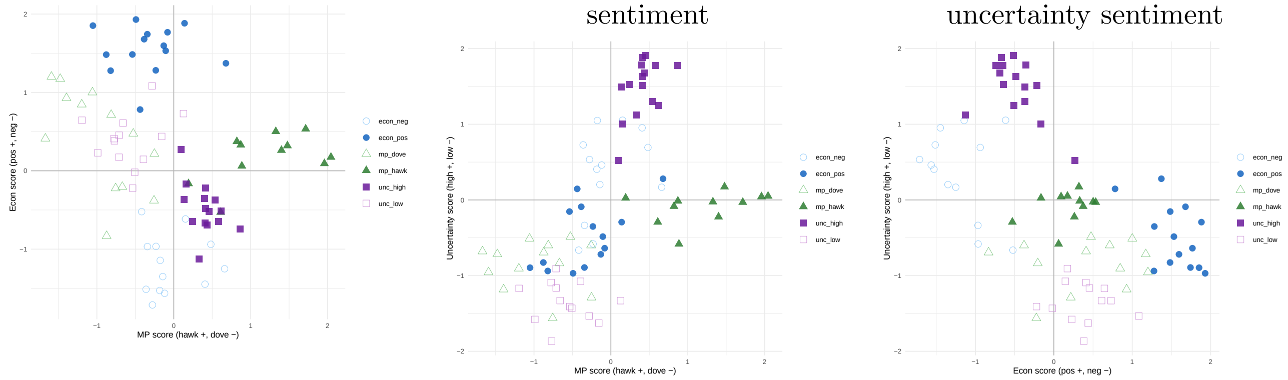

Notes: Each point is one seed phrase, coloured and shaped by category. Panel A: MP score vs. economic sentiment score. Panel B: MP score vs. uncertainty sentiment score. Panel C: Economic sentiment score vs. uncertainty sentiment score. Seeds scoring near zero on the non-own axes confirm that the three concept vectors capture largely independent dimensions of variation. Seeds with substantial cross-loading were removed from the final seed set (see main text).

As a complementary check, Figure 5 applies principal component analysis (PCA) to the sentence embeddings of all seed phrases using the paraphrase-multilingual-mpnet-base-v2 sentence transformer. The seed phrases used in the main analysis are encoded with nvidia/llama-embed-nemotron-8b – as are the texts – due to its large context length.

If the three dimensions capture genuinely distinct linguistic signals, seeds from different categories should form well-separated clusters in the low-dimensional embedding space. The plots confirm this: seeds group into six distinct clusters in the first three principal components — one cluster per pole of each dimension — with little overlap across categories. This embedding-level separation provides assurance that the three concept vectors are not merely orthogonal by construction but reflect substantively different aspects of language used in central bank communication.

Figure 5: Seed phrase validation: PCA of seed embeddings

(a) PC1 vs. PC2 (b) PC1 vs. PC3

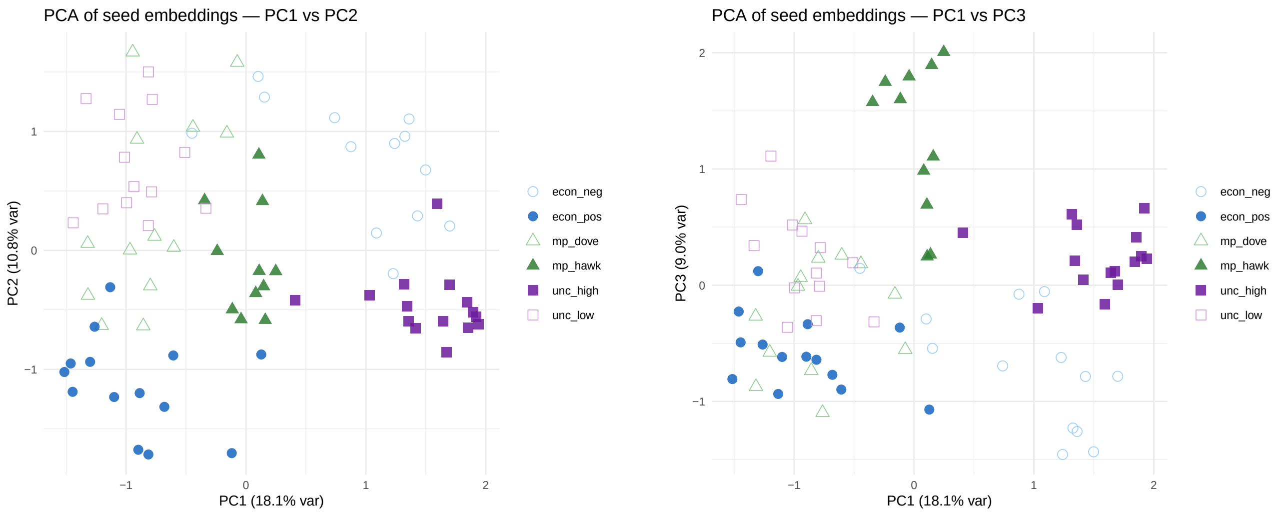

Notes: PCA applied to the 768-dimensional sentence embeddings of all seed phrases. Well-separated CVP dimensions produce six distinct clusters in PC space, one per pole.

A.3 Descriptive statistics: arc features

Table 10 reports the mean, standard deviation, minimum, median, and maximum of the Nelson–Siegel arc factors and Hurst exponents for the ECB and the Fed side by side. All features are estimated from contextual arcs using the cleaned seed set and $\lambda = 9.5$.

A.4 Contextual versus sentence-level arcs

Throughout the paper, arc features are estimated from *contextual* arcs – each token scored in the context of the full document – rather than *sentence-level* arcs, where the transformer scores each sentence independently. Table 11 validates this choice for both institutions: it compares the Full arc model estimated with each scoring approach.

For both institutions, contextual arcs outperform sentence arcs in terms of explanatory power (adjusted R^2 of 0.344 versus 0.241 for the ECB; 0.378 versus 0.277 for the Fed), confirming that context-sensitive scoring, which allows the embedding to reflect the surrounding discourse, produces more informative arc features than sentence-level scoring in isolation. For the Fed, the direction of every coefficient that is significant under contextual scoring is preserved under sentence-level scoring, generally at a larger magnitude. It is also worth noting that, given the inferior performance of the sentence arcs, they would still improve upon the benchmark measures, and considerably so. The replication across both institutions confirms that the choice of contextual over sentence-level scoring is not an ECB-specific artefact.

A.5 Predictive regressions at the two-meeting-ahead horizon

Throughout Section 4.1 and Appendix A.7, we report the one-meeting-ahead rate change Δi_{t+1} as the headline horizon and note that results at the two-meeting-ahead horizon Δi_{t+2} are

Table 10: Descriptive statistics: arc features, ECB and Fed

	ECB ($N = 235$)					**Fed** ($N = 88$)				
	Mean	SD	Min	Median	Max	Mean	SD	Min	Median	Max
NS arc factors: monetary policy sentiment										
`mp_level`	0.040	0.035	-0.022	0.036	0.152	-0.008	0.063	-0.118	-0.006	0.120
`mp_slope`	-0.074	0.035	-0.165	-0.070	0.033	-0.047	0.055	-0.193	-0.053	0.101
`mp_curve`	0.065	0.153	-0.349	0.081	0.518	0.260	0.182	-0.365	0.263	0.610
NS arc factors: economic sentiment										
`econ_level`	0.030	0.028	-0.054	0.031	0.095	0.064	0.030	-0.021	0.064	0.162
`econ_slope`	-0.004	0.034	-0.101	-0.001	0.074	0.063	0.039	-0.043	0.067	0.158
`econ_curve`	0.037	0.113	-0.259	0.050	0.318	0.032	0.154	-0.424	0.041	0.388
NS arc factors: uncertainty										
`unc_level`	-0.083	0.017	-0.142	-0.083	-0.036	-0.115	0.033	-0.178	-0.119	-0.016
`unc_slope`	-0.008	0.028	-0.093	-0.003	0.062	-0.019	0.044	-0.179	-0.019	0.139
`unc_curve`	-0.121	0.095	-0.359	-0.124	0.198	-0.005	0.144	-0.346	-0.028	0.470
Hurst exponent										
`hurst_mp`	0.632	0.042	0.490	0.638	0.723	0.650	0.039	0.533	0.655	0.718
`hurst_econ`	0.552	0.062	0.363	0.556	0.677	0.598	0.068	0.386	0.610	0.692
`hurst_unc`	0.557	0.060	0.443	0.549	0.704	0.560	0.062	0.413	0.561	0.720

Notes: Arc features estimated from contextual arcs, $\lambda = 9.5$. ECB: cleaned seed set, 235-document sample. Fed: same seeds, corrected sample of press-conference transcripts. NS factors (level, slope, curvature) are the Nelson–Siegel decomposition of the sentence-level arc. Uncertainty is scored as uncertain positive (high uncertainty = high score). Hurst exponent is the R/S statistic (`pracma::hurstexp`), with values above 0.5 indicating persistent, trending arcs.

Table 11: Full arc model: contextual vs. sentence-level arcs, ECB and Fed

	ECB ($N = 234$)		**Fed** ($N = 87$)	
	Contextual	Sentence	Contextual	Sentence
(Intercept)	0.295	0.047	0.051	0.158
	(1.430)	(0.209)	(0.106)	(0.361)
`mp_level`	4.839***	8.179***	3.830***	10.852***
	(8.219)	(6.046)	(5.072)	(5.497)
`mp_slope`	2.425***	3.003***	1.134·	3.803**
	(4.224)	(3.920)	(1.722)	(3.260)
`mp_curve`	0.058	0.889***	0.635***	2.070***
	(0.541)	(4.107)	(3.939)	(5.305)
`hurst_mp`	−0.552*	−0.067	0.124	0.543
	(−2.229)	(−0.244)	(0.202)	(0.990)
`econ_level`	4.239***	7.432***	1.897	5.461
	(3.706)	(3.633)	(1.271)	(1.402)
`econ_slope`	2.965***	5.722***	−0.471	−0.204
	(3.454)	(4.222)	(−0.511)	(−0.176)
`econ_curve`	0.417*	1.222***	0.017	0.674
	(2.326)	(3.815)	(0.057)	(0.885)
`hurst_econ`	−0.180	−0.460*	−0.567	−0.707·
	(−0.971)	(−2.385)	(−1.423)	(−1.770)
`unc_level`	−0.170	4.306***	−0.045	2.726
	(−0.183)	(3.377)	(−0.040)	(1.168)
`unc_slope`	0.223	0.626	−0.609	0.289
	(0.409)	(1.030)	(−0.886)	(0.255)
`unc_curve`	−0.741***	0.078	−0.337	0.257
	(−3.581)	(0.300)	(−1.307)	(0.603)
`hurst_unc`	−0.185	0.181	0.106	−0.234
	(−0.846)	(0.974)	(0.268)	(−0.613)
R^2	0.377	0.280	0.465	0.378
Adj. R^2	0.344	0.241	0.378	0.277
Residual SE	0.146	0.157	0.196	0.211
N	234	234	87	87

Notes: OLS estimates. t-statistics in parentheses. Outcome: one-meeting-ahead policy rate change Δi_{t+1}. Full arc model includes Nelson–Siegel factors and Hurst exponent for all three arc dimensions (monetary policy, economic sentiment, uncertainty). Contextual arcs score each token in context of the full document; sentence arcs score each sentence independently. Significance: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $^{\cdot}p < 0.10$.

qualitatively the same. Table 12 reports the results at $t + 2$ for both institutions.

A.6 Sentiment arcs during crisis periods

Figure 6 shows sentiment arcs for the ECB, grouped by whether the statement falls within a crisis period as dated in Section 4.1.2 (Global Financial Crisis, euro area sovereign debt crisis, COVID crisis) or outside one. Economic sentiment is consistently more negative during crisis periods throughout the statement; the monetary and uncertainty arcs show less visually clear-cut separation. Table 4 carries out the corresponding formal test.

Figure 6: Sentiment arcs during crisis periods: ECB

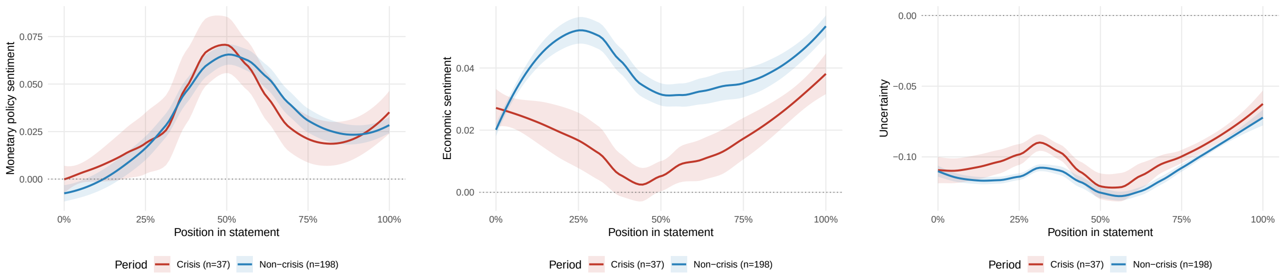

Notes: LOESS-smoothed (span = 0.6) sentiment arcs grouped by whether the statement falls within a crisis period (red) or not (grey). Panels show monetary policy sentiment (hawkish positive), economic sentiment (optimistic positive), and uncertainty sentiment (uncertain positive). The x-axis is normalized position within statement (0% = opening, 100% = close). Shaded bands are 95% pointwise confidence intervals across statements.

A.7 The Fed as an institutional extension

This appendix reports the full supporting detail behind the compact summary of Section 4.3: example sentences and the full-sample and rate-conditional arc figures (descriptive statistics are reported alongside the ECB's in Appendix A.3), all on the same corrected Fed sample of press-conference transcripts used there ($N = 87/86$ at the headline $t + 1$ horizon).

Example sentences. Table 13 reports the Fed counterpart to Table 1 (Section 3): the highest- and lowest-scoring sentences in each CVP dimension. *"Since then, inflation has again surprised to the upside... and projections for inflation this year have been revised up notably"* (June 2022) is the single most hawkish sentence in the Fed sample, from the onset of the most aggressive tightening cycle in four decades; the most dovish, from April 2012, reflects the language of the maturity-extension program ("Operation Twist"). The top-scoring economic sentiment sentence describes a strengthening US labor market, while the bottom-scoring one captures the 2011 European sovereign-debt spillover to US financial conditions – the same pattern of historically interpretable extremes found for the ECB.

Table 12: Rate change regressions: ECB and Fed (Δi_{t+2})

	ECB ($N = 233$)					**Fed** ($N = 86$)				
	M1 (Bench.)	M2 (Mean)	M3 (MP)	M4 (MP+EC)	Full arc (MP+EC+UNC)	M1 (Bench.)	M2 (Mean)	M3 (MP)	M4 (MP+EC)	Full arc (MP+EC+UNC)
(Intercept)	−0.035	−0.175***	1.339***	1.097***	0.852*	0.149·	0.001	0.415	0.889	0.598
	(−1.624)	(−3.996)	(4.987)	(3.672)	(2.417)	(1.833)	(0.008)	(0.604)	(1.077)	(0.675)
mp_pr	−0.047					0.173				
	(−0.987)					(0.828)				
ec_pr	0.401***					−0.066				
	(4.584)					(−0.229)				
mp_mean		2.460***					6.195***			
		(4.373)					(6.973)			
econ_mean		2.499**					−0.899			
		(2.908)					(−0.523)			
mp_level			4.192***	6.881***	8.482***			6.709***	7.203***	6.499***
			(6.158)	(7.530)	(8.382)			(6.361)	(6.307)	(4.707)
mp_slope			0.911	3.839***	4.493***			2.680*	2.603*	1.642
			(1.143)	(4.010)	(4.586)			(2.541)	(2.400)	(1.363)
mp_curve			−0.099	0.250	0.187			1.116***	1.323***	1.272***
			(−0.632)	(1.410)	(1.029)			(4.262)	(4.766)	(4.321)
hurst_mp			−2.285***	−1.893***	−1.549***			−0.664	−0.484	−0.414
			(−5.276)	(−4.535)	(−3.662)			(−0.621)	(−0.437)	(−0.369)
econ_level				8.722***	7.445***				1.820	2.177
				(4.881)	(3.810)				(0.843)	(0.797)
econ_slope				5.764***	5.073***				−0.747	−1.488
				(4.239)	(3.451)				(−0.597)	(−0.881)
econ_curve				1.298***	0.782*				0.019	−0.355
				(5.155)	(2.556)				(0.045)	(−0.651)
hurst_econ				−0.373	−0.374				−1.201·	−1.063
				(−1.184)	(−1.181)				(−1.696)	(−1.435)
unc_level					−0.186					−0.438
					(−0.117)					(−0.210)
unc_slope					0.672					−0.748
					(0.716)					(−0.596)
unc_curve					−1.039**					−0.794·
					(−2.925)					(−1.667)
hurst_unc					−0.119					0.172
					(−0.318)					(0.238)
R^2	0.085	0.089	0.259	0.352	0.397	0.009	0.392	0.402	0.461	0.489
Adj. R^2	0.077	0.081	0.246	0.329	0.364	-0.015	0.377	0.372	0.405	0.405
N	233	233	233	233	233	86	86	86	86	86

Notes: OLS estimates. t-statistics in parentheses. Outcome: two-meeting-ahead policy rate change Δi_{t+2}. Residual SE omitted for both panels for column consistency (not captured by the Fed pipeline's output). Significance: $^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$, $^{\cdot}p < 0.10$.

Table 13: Highest and lowest scoring sentences by CVP dimension: Fed

Dimension	Pole	Date	Score	Sentence
MP stance	Hawkish	2022-06	+0.265	Since then, inflation has again surprised to the upside, some indicators of inflation expectations have risen, and projections for inflation this year have been revised up notably.
MP stance	Dovish	2012-04	−0.143	Our program to extend the average maturity of the Federal Reserve's security holdings—announced in September—will continue as scheduled.
Economic	Optimistic	2016-03	+0.278	The labor market continues to strengthen.
Economic	Pessimistic	2011-11	−0.121	Most notably, concerns about European fiscal and banking issues have contributed to strains in global financial markets, which have likely had adverse effects on confidence and growth.
Uncertainty	Uncertain	2020-06	+0.135	There is great uncertainty about the future.
Uncertainty	Certain	2019-10	−0.263	Inflation continues to run below our symmetric 2 percent objective.

Notes: For each CVP dimension, the table reports the Fed sentence with the highest score (positive pole) and the lowest score (negative pole) across all statements in the corpus. Scores are the mean of token-level CVP projections for tokens belonging to that sentence, computed from contextual arcs. Positive scores reflect the positive pole of each dimension: hawkish for monetary policy sentiment, optimistic for economic sentiment, and uncertain for the uncertainty sentiment dimension. Sentences with fewer than three tokens or fewer than 40 characters are excluded. The uncertainty dimension scores reflect the degree of epistemic uncertainty in language: the high-scoring sentence scores as highly uncertain while the low-scoring sentence reflects low-uncertainty, assertive language. The corresponding ECB table is reported in Table 1 (Section 3).

Full-sample average sentiment arcs. Figure 7(a) shows the Fed counterpart to Figure 2.

Figure 7: Sentiment arcs: Fed

(a) Full-sample arcs

(b) Arcs before rate changes

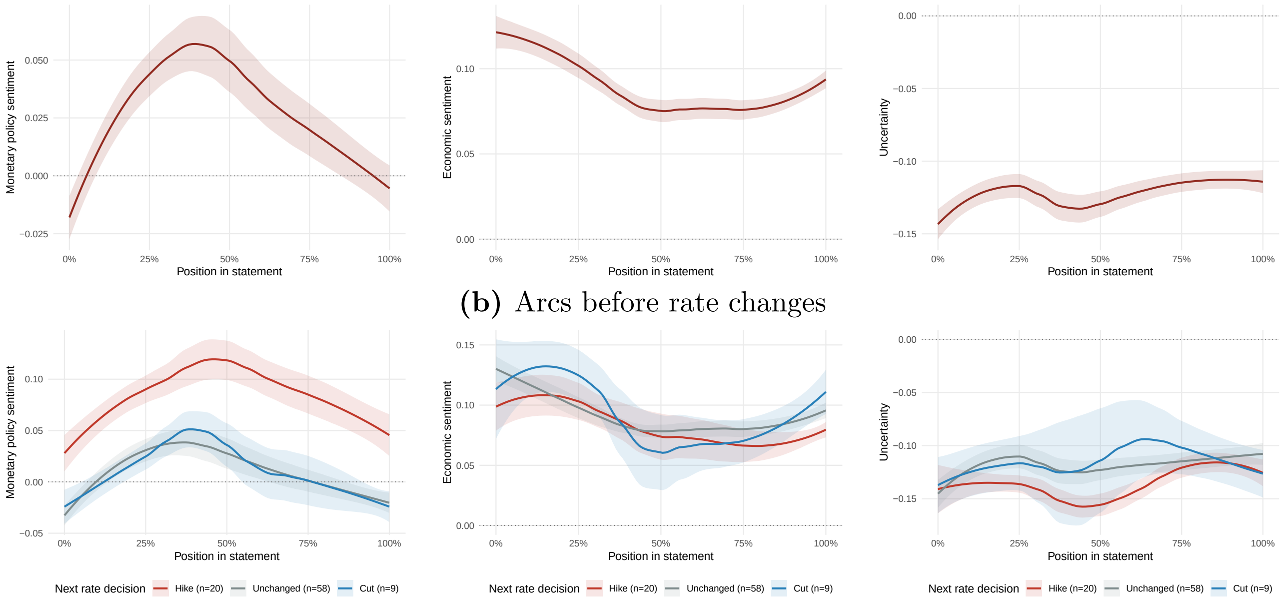

Notes: LOESS-smoothed (span = 0.6) sentiment arcs. Panel (a): full-sample averages (N = 88 transcripts, 2011–2025, corrected Fed sample); shaded bands are 95% pointwise confidence intervals across statements. Panel (b): arcs grouped by the subsequent policy rate decision (hike, red; unchanged, grey; cut, blue). Panels show monetary policy sentiment (hawkish positive), economic sentiment (optimistic positive), and uncertainty sentiment (uncertain positive). The x-axis is normalized position within statement (0% = opening, 100% = close).

On this sample, the monetary policy sentiment arc rises from a mildly dovish opening to a peak around two-fifths of the way through the statement, then declines steadily back toward zero by the close. Economic sentiment opens at its highest, declines through the first half to a shallow trough around the mid-point, then rises again toward the close. The uncertainty arc is negative throughout (confident language on average), with a mild wave: most confident around the opening and the middle of the statement, and less confident thereafter.

Sentiment arcs before rate changes. Figure 7(b) shows the Fed counterpart to Figure 3. On this sample (n = 20 Hike, 58 Unchanged, 9 Cut), statements preceding a hike show distinctly more hawkish monetary policy sentiment throughout the statement than those preceding no change or a cut, which track each other closely – this part of the earlier, larger-sample pattern holds up. The economic sentiment arc does not separate as cleanly by the next rate decision here: the confidence bands overlap throughout, and statements preceding a cut are, if anything, no less optimistic than those preceding a hike. The uncertainty arc shows a similar pattern: statements preceding a cut have a markedly elevated (less confident) uncertainty arc across the middle of the statement, though with only 9 such statements the band around that line is wide.

The Working Paper series of the Oesterreichische Nationalbank is designed to disseminate and to provide a platform for discussion of either work of the staff of the OeNB economists or outside contributors on topics which are of special interest to the OeNB. To ensure the high quality of their content, the contributions are subjected to an international refereeing process. The opinions are strictly those of the authors and do in no way commit the OeNB.

The Working Papers are also available on our website (http://www.oenb.at) and they are indexed in RePEc (http://repec.org/).

Publisher and editor Oesterreichische Nationalbank
Otto-Wagner-Platz 3, 1090 Vienna, Austria
PO Box 61, 1011 Vienna, Austria
www.oenb.at
oenb.info@oenb.at
Phone (+43-1) 40420-6666

Editor Martin Summer

Cover Design Information Management and Services Division

DVR 0031577

ISSN 2310-533X (Online)