Australian National University
Crawford School of Public Policy

# CAMA

**Centre for Applied Macroeconomic Analysis**

# Narratives and the Term Structure of Inflation Expectations

**CAMA Working Paper 29/2026**
**May 2026**

**Jonathan Benchimol**
Bank of Israel
Centre for Applied Macroeconomic Analysis, ANU

**Sathya Mellina**
Centre for Applied Macroeconomic Analysis, ANU

## Abstract

We examine how inflation language in FOMC statements and Chair press conferences maps into the breakeven inflation (BEI) term structure from two to ten years. Five indices capture stance, broad and current inflation language, and the Delphic/Odyssean decomposition. Conditional on the high-frequency rate surprise, statement language is associated with lower BEI at short-to-intermediate maturities, consistent with markets reading committee-vetted text through the policy reaction function. Press conferences differ: Delphic language loads positively on long-horizon forwards; Odyssean language compresses the short-to-intermediate segment. Intraday BEI moves positively with every press-conference index; no analogous statement-window effect survives partialling out the rate surprise.

### Keywords

central bank communication, breakeven inflation, high-frequency identification, event-study methods, Delphic and Odyssean forward guidance, natural language processing

### JEL Classification

C45, E43, E52, E58, G12, G14

**Address for correspondence:**

(E) cama.admin@anu.edu.au

**ISSN 2206-0332**

**The Centre for Applied Macroeconomic Analysis** in the Crawford School of Public Policy has been established to build strong links between professional macroeconomists. It provides a forum for quality macroeconomic research and discussion of policy issues between academia, government and the private sector.

**The Crawford School of Public Policy** is the Australian National University's public policy school, serving and influencing Australia, Asia and the Pacific through advanced policy research, graduate and executive education, and policy impact.

# Narratives and the Term Structure of Inflation Expectations*

Jonathan Benchimol† and Sathya Mellina‡

May 2026

**Abstract**

We examine how inflation language in FOMC statements and Chair press conferences maps into the breakeven inflation (BEI) term structure from two to ten years. Five indices capture stance, broad and current inflation language, and the Delphic/Odyssean decomposition. Conditional on the high-frequency rate surprise, statement language is associated with lower BEI at short-to-intermediate maturities, consistent with markets reading committee-vetted text through the policy reaction function. Press conferences differ: Delphic language loads positively on long-horizon forwards; Odyssean language compresses the short-to-intermediate segment. Intraday BEI moves positively with every press-conference index; no analogous statement-window effect survives partialling out the rate surprise.

*Keywords:* Central bank communication, breakeven inflation, high-frequency identification, event-study methods, Delphic and Odyssean forward guidance, natural language processing.

*JEL Classification:* C45, E43, E52, E58, G12, G14.

---

*The views expressed in this paper are those of the authors and do not necessarily represent the views of the Bank of Israel. The authors thank Inbar Bahat and Itamar Caspi for constructive comments.

†Bank of Israel, Research Department, Jerusalem, Israel. Corresponding author. Email: jonathan.benchimol@boi.org.il

‡Centre for Applied Macroeconomic Analysis (CAMA), Crawford School of Public Policy, Australian National University.

# 1 Introduction

Each scheduled FOMC announcement releases a bundled information set: a policy-rate decision, the committee's assessment of incoming inflation data, explicit or implicit guidance about the policy path, and, at meetings with a press conference, the Chair's elaboration of risks, contingencies, and outlook (Woodford, 2005; Gürkaynak et al., 2005). The narrative components of this bundle move inflation compensation along the term structure, and they do so in ways that depend on *which* communication channel delivers them, *how* the institutional format constrains their interpretation, and *where* along the maturity spectrum markets absorb them.[1]

We ask how inflation-related language in the two main FOMC communication objects, the post-meeting statement and the Chair's press conference, is associated with breakeven inflation (BEI) compensation from two to ten years. The two objects differ both in format and in coverage: the statement is short and committee-vetted, with a sample beginning in 2004; the press conference is longer, more conditional, permits unscripted elaboration of the inflation outlook, and is observed only from 2011 onward following the introduction of regular post-meeting press conferences. We measure both BEI yields, which average compensation across horizons, and BEI instantaneous forwards, which isolate horizon-specific repricing, and we complement the daily event study with intraday BEI changes within each narrow announcement window (Gürkaynak et al., 010b,a; Acosta et al., 2025). A central concept in what follows is *co-entailment*: the property that, in a constrained communication format, an inflation assessment and the policy response implied by that assessment are bundled and priced through a single reaction-function signal. Section 3 develops this notion formally; we return to it in interpreting the statement results.

Two design issues organize the empirical exercise. First, policy announcements bundle the target-rate surprise, information about fundamentals, and shifts in beliefs about the central bank's reaction function, and these components are not separately observed even within narrow windows (Jarociński and Karadi, 2020; Miranda-Agrippino and Ricco, 2021; Bauer and Swanson, 2023). Second, on press-conference days the announcement information arrives sequentially: the rate decision is priced at the statement release, so that the same-window target-rate surprise is near-degenerate by the time the Chair speaks. Following the event-study design in Acosta et al. (2025), we therefore distinguish a narrow statement-release window from a narrow press-conference window and align the surprise control with the window in which the dependent variable is measured.

On the measurement side, we use FOMC-RoBERTa, a domain-adapted transformer

[1] Throughout, we use "inflation compensation" to refer to breakeven inflation (BEI), which combines expected inflation, an inflation risk premium, and a relative liquidity wedge between nominal Treasuries and TIPS. We reserve "inflation expectations" for the pure expectations component, which we do not separately identify. This terminological discipline follows Acosta et al. (2025); it matters because the sign and magnitude of communication effects on the underlying components may differ (Haubrich et al., 2012; Abrahams et al., 2016; D'Amico et al., 2018). See Section 3.2 for a detailed discussion.

fine-tuned for hawkish/dovish classification (Shah et al., 2023), combined with pre-specified lexical dictionaries that identify inflation-relevant content. For each document we construct a monetary-policy stance index ($mps_t$), a broad inflation narrative index ($\pi_t^g$), and a current-inflation index ($\pi_t^{cur}$). We further decompose inflation narratives into *Delphic* ($\delta_t^{\pi}$) and *Odyssean* ($\omega_t^{\pi}$) components, the former conveying outlook information, the latter signaling intended future policy contingent on inflation (Campbell et al., 2012; Andrade et al., 2019; Andrade and Ferroni, 2021).

On the identification side, we embed these measures in a high-frequency event-study framework that purges each communication index of the component linearly associated with the high-frequency surprise in the current policy-rate decision, measured in the event window appropriate to each communication object. The resulting residual variation is best read as a text-based communication innovation conditional on the rate decision, not as a structural communication shock.[2] For statements, the relevant window is the statement-release window, in which the policy surprise displays material variation. For press conferences, we use the full announcement-day (monetary-event) window, because the press-conference-only surprise is near-degenerate by the time the Chair speaks and provides no meaningful purge of the rate-decision component. As a complementary identification layer, we exploit intraday BEI changes within the announcement windows of each communication object; these changes have a higher signal-to-noise ratio than close-to-close daily outcomes and are largely free from non-FOMC news contamination.

The empirical findings differ systematically across formats. For post-meeting statements, inflation-related language is associated with lower BEI compensation at short-to-intermediate maturities, conditional on the current rate-decision surprise; both the Delphic and Odyssean statement indices load negatively on BEI, in a pattern consistent with institutional co-entailment in the committee-vetted text, where inflation assessments and policy-response language are bundled in a format that markets price through the reaction function. For press conferences, the maturity profile differs in character. Delphic inflation language, which captures outlook and risk discussion, tracks positive long-horizon forward repricing, while Odyssean language, which captures policy-commitment language conditional on inflation developments, is associated with compression in the short-to-intermediate forward segment. The press-conference results provide a term-structure application of the Delphic/Odyssean distinction: outlook information and policy commitment are not priced at the same horizons.

[2]Following Acosta et al. (2025), we denote by $MP1_t$ the high-frequency surprise around the current FOMC meeting decision, constructed from federal funds futures in the relevant event window. Conditioning on $MP1_t$ removes only the linear component of text variation associated with the contemporaneous rate decision; it does not isolate exogenous communication variation from the remaining bundled FOMC information set, which includes policy-path news, reaction-function signals, the Summary of Economic Projections (SEP) and dot-plot information at SEP meetings, and the central bank's economic outlook (Romer and Romer, 2000; Melosi, 2017; Bu et al., 2021). The coefficients reported below should be read as associations between text-based communication indices and BEI repricing conditional on the current rate decision, not as structurally identified causal effects of individual communication dimensions.

The intraday evidence localizes the timing of these associations. Within the press-conference window, all five communication indices co-move positively with BEI changes at the available maturities. Within the narrow statement window, by contrast, text-based statement indices are not separately detectable after conditioning on the contemporaneous rate-decision surprise. The negative statement association documented in daily data therefore appears to be a broader announcement-day response rather than a narrow-window text effect realized at release.

The paper has three contributions. The first concerns the cross-format asymmetry: statement inflation language is associated with lower BEI compensation, whereas press-conference inflation language is associated with positive long-horizon forward repricing, conditional on the rate-decision surprise.

A second contribution is term-structure evidence on the Delphic/Odyssean distinction. Within press conferences, outlook-oriented language loads on long-horizon forwards while policy-commitment language loads on the short-to-intermediate forward segment. The distinction therefore has a horizon-specific counterpart in market-based inflation compensation, not only in tone aggregates.

A third contribution is the within-day localization. Intraday BEI changes show that the positive press-conference component is realized within the Chair's speaking window, whereas the statement-related negative association is a broader announcement-day phenomenon. The pricing of inflation narratives therefore depends jointly on content, institutional format, and the maturity segment that absorbs the news.

The primary estimand of the paper is the maturity-specific reduced-form loading of BEI compensation on a residualized inflation-narrative index, conditional on the high-frequency rate-decision surprise. We do not estimate a structural communication shock, an effect on inflation expectations separated from inflation risk premia, or a partial derivative holding other dimensions of FOMC information fixed. The empirical exercise is therefore best read as a horizon-specific mapping from text-based proxies of FOMC inflation narratives to inflation-compensation repricing, rather than as identification of a narrative-specific causal channel.

The remainder of the paper is organized as follows. Section 2 situates the contribution within three traditions: high-frequency event studies, the information-effect debate, and text-as-data. Section 3 develops the reduced-form framework that maps Delphic, Odyssean, and stance signals into maturity-specific BEI responses, and derives the four propositions on sign, shape, co-entailment, and within-day dynamics that guide the empirical analysis. The communication indices and the BEI yield and forward data are described in Section 4. Baseline event-study results appear in Section 5; Section 6 decomposes press-conference effects into Delphic and Odyssean components and documents the maturity segmentation. Section 7 uses intraday windows to localize the press-conference repricing and the statement-window silence. Section 8 discusses implications, and Section 9 concludes. Maturity-by-maturity coefficient estimates underlying the figures are reported in Appendix F (Tables 15 and 16), and robustness exercises are collected in Appendix E.

## 2 Related literature

The paper connects three research traditions: evidence that central bank communication (CBC) moves financial prices materially beyond the contemporaneous policy-rate decision; high-frequency event studies as the standard design for identifying monetary-policy news, together with the information-effect debate that this design has generated; and text-as-data approaches that translate qualitative narratives into quantitative, replicable measures of policy language. We position our contribution against each in turn.

A large empirical literature documents that central bank announcements move asset prices at horizons far beyond the overnight rate. Classic event-study contributions show that a single policy-rate surprise does not summarize announcement-day repricing: longer-maturity yields react materially to the informational bundle released with the decision, and a second factor, often interpreted as forward guidance or path news, explains a substantial share of medium- and long-maturity movements (Gürkaynak et al., 2005; Kuttner, 2001). Woodford (2005) formalizes the argument that managing expectations through communication is itself a primary tool of monetary policy, and one that can dominate the rate instrument when the short rate is constrained. Subsequent work shows that communication shifts not only interest rates but also risk premia, exchange rates, equities, and volatility, with effects that depend on the credibility of the sender and the channel through which information is transmitted (Ehrmann and Fratzscher, 2007; Blinder et al., 2008).

Within this literature, recent work emphasizes heterogeneity across instruments, maturities, and communication formats. Distant forward real rates move by more than a narrow expectations-hypothesis view would imply, pointing to a role for term premia and risk-taking channels (Hanson and Stein, 2015). Altavilla et al. (2025) and Ahrens et al. (2025) document heterogeneous market effects across communication objects, with press conferences and Q&A settings transmitting richer information than committee-vetted statements. Boguth et al. (2019) show that press conferences reshape attention and coordination in ways that need not map one-for-one into verbal content, a consideration that bears directly on the interpretation of our press-conference estimates. Gorodnichenko et al. (2023) reinforce this point by showing that vocal characteristics of central bankers' delivery during press conferences affect financial-market responses independently of textual content; our text-based press-conference indices may therefore capture content effects that interact with non-textual delivery components we cannot isolate. Acosta et al. (2025) document that press conferences have become the dominant source of announcement-day market repricing across a wide range of assets, and they provide the high-frequency event-study infrastructure we use to construct surprise controls. We complement that work by asking what within the press conference drives the repricing, specifically, whether it is Delphic inflation content, Odyssean commitment language, or the Chair's current-inflation assessment, and by focusing on the term structure of inflation compensation rather than on aggregate price responses.

The high-frequency event-study literature identifies monetary-policy news from asset-price changes in narrow windows around central bank announcements, using futures, OIS rates, and Treasury yields as market-based proxies for unexpected policy news (Kuttner, 2001; Gürkaynak et al., 2005; Bernanke and Kuttner, 2005; Swanson, 2021). Recent contributions broaden this approach by extracting policy-news factors from a richer cross-section of rates and event windows. Bu et al. (2021) construct a unified monetary-policy shock measure that loads on target, path, and longer-maturity components, while Acosta et al. (2025) provide event-specific surprises for FOMC statements, press conferences, monetary-event windows, and minutes. Our design is parsimonious by construction: we condition on the high-frequency surprise in the current policy-rate decision rather than on a broad policy-news factor, because the object of interest is whether textual inflation narratives contain market-relevant information *beyond* the rate decision itself. This choice preserves policy-path, outlook, and reaction-function content that may be transmitted through language, but it also implies a narrower interpretation: the residualized text indices are conditional on the current rate decision, not orthogonal to all other components of the FOMC information set.

A central interpretational caveat in this literature is that announcement-window surprises need not be pure monetary-policy shocks. They may combine changes in the intended policy stance with central bank information about fundamentals and with market learning about the systematic reaction function (Nakamura and Steinsson, 2018; Jarociński and Karadi, 2020; Miranda-Agrippino and Ricco, 2021; Bauer and Swanson, 2023; Cieslak and Schrimpf, 2019). The caveat bears directly on our setting because the same inflation language can have different meanings depending on whether markets interpret it as news about the inflation outlook or as news about the policy rule. Delphic communication conveys information about economic conditions and risks, whereas Odyssean communication conveys intended policy behavior or commitment conditional on those conditions (Campbell et al., 2012; Andrade and Ferroni, 2021). A positive BEI response to inflation language therefore need not signal a loss of credibility, and a negative response need not signal that inflation news is benign: the sign depends on whether the dominant channel is outlook revision, policy commitment, reaction-function learning, or inflation-risk compensation. This is the conceptual link between the information-effect debate and our Delphic–Odyssean decomposition.

Several contributions show that communication contains market-relevant information beyond the contemporaneous target-rate decision. Hansen et al. (2019) document that central bank communication moves asset prices through information not captured by conventional policy-rate surprises, and Lakdawala and Schaffer (2019) provide evidence that the Fed's private information is transmitted to financial markets through communication. Bauer and Swanson (2023) argue, on the other hand, that many apparent information effects are attenuated once the systematic policy reaction function is accounted for. Our design is therefore conditional rather than structural: it asks whether inflation narratives price the BEI term structure over and above the current rate decision, while recognizing that text may still proxy for path news, SEP-related information, or reaction-function learning.

The measurement window is part of the identification problem rather than a technical detail. FOMC statements and Chair press conferences release different types of information at different points within the announcement day. The event-study database of Acosta et al. (2025) explicitly separates statement, press-conference, minutes, and monetary-event windows, and shows that press conferences have become a substantial source of policy news in their own right. We build on this logic by matching the surprise-control window to the BEI window associated with each communication object, and by using intraday BEI changes to distinguish statement-release repricing from press-conference-window repricing.

BEI is widely used as a real-time measure of market-based inflation pricing, but it is not a pure measure of expected inflation. It embeds expected inflation, inflation risk premia, and, especially in early TIPS-market years or during stressed episodes, liquidity premia (Gürkaynak et al., 010b; Haubrich et al., 2012; Abrahams et al., 2016; Fleckenstein et al., 2014; D'Amico et al., 2018). For this reason, we refer to our outcome as inflation compensation rather than inflation expectations. The distinction has substantive content: a movement in BEI can reflect revisions in expected inflation, changes in the inflation risk premium, changes in TIPS liquidity premia, or a combination of these components. We return to this decomposition in Section 3.2. Acosta et al. (2025) adopt the same inflation-compensation framing and document that BEI responses to monetary-policy surprises peak around the 4–5-year forward segment. We ask which textual inflation narratives account for that term-structure repricing, and why the sign and maturity profile of the response differ between statements and press conferences. The two-object focus motivates our parallel use of BEI yields, which average compensation across horizons, and BEI instantaneous forwards, which isolate horizon-specific repricing.[3]

Text-as-data methods provide the measurement bridge between what central banks say and what markets price. Tetlock (2007) establishes a canonical framework for mapping textual tone to asset-price responses, while Loughran and McDonald (2011) and Gentzkow et al. (2019) provide methodological foundations for domain-specific dictionaries and for the distinction between prediction and interpretation in textual measurement. Lucca and Trebbi (2009) develop an automated scoring method for FOMC

[3] A broader question is how central bank communication transmits beyond financial markets to households and firms. Coibion et al. (2022) provide direct evidence on the expectations channel, while Binder (2017) and Lamla and Vinogradov (2019) show that communication format and media amplification shape transmission to non-specialist audiences. Our market-based analysis captures the first link in this broader transmission chain.

statements focused on communication about future policy-rate decisions. Hansen and McMahon (2016) combine topic modeling and tone measures to show that different dimensions of central bank language carry distinct information and have distinct market effects. Recent work moves beyond bag-of-words representations: Ahrens et al. (2025) use supervised NLP to recover economically interpretable objects such as outlook revisions from central bank communication, and Chadha et al. (2026) combine topic modeling with targeted textual indices to map communication themes into yield-curve components.

Our measurement approach builds on this literature but is tailored to the inflation-compensation object studied here. Rather than reducing FOMC communication to a single scalar measure of tone or stance, we construct inflation-specific text indices for statements and press conferences and separate broad inflation discussion, current-inflation assessment, Delphic outlook language, and Odyssean commitment language. The construction is described in Section 4.1; the conceptual point at this stage is that combining interpretable inflation dictionaries with a domain-adapted transformer preserves the economic structure of the Delphic–Odyssean distinction while embedding the resulting measures in a window-consistent event-study design.[4]

# 3 Theoretical framework

This section develops a reduced-form framework that links central bank communication to the term structure of inflation compensation. The aim is not to identify a fully structural model of communication but to discipline the empirical interpretation of the event-study coefficients. The framework addresses three issues: why the same inflation language may be priced differently when interpreted as outlook information versus policy commitment; why BEI yields and BEI forwards can display different maturity profiles; and why statement releases and Chair press conferences may generate different within-day repricing patterns.

The framework rests on three ingredients. BEI is treated as inflation compensation rather than as pure expected inflation: it combines expected inflation, inflation risk premia, and TIPS liquidity components. Communication is represented as a vector of signals containing a monetary-policy stance component, a Delphic inflation component, and an Odyssean inflation component. The institutional format of communication determines how separately these signals can be priced; press conferences allow markets to distinguish outlook information from policy commitment more easily, whereas committee-vetted statements embed inflation assessments within reaction-function language. The framework delivers the sign, maturity, co-entailment, and within-day

[4]Press-conference estimates should be interpreted as the joint effect of spoken content and the attention environment created by the event. Boguth et al. (2019) show that press conferences can reshape attention and coordination in ways that need not map one-for-one into the verbal content alone. Broader policy-news factors such as the one in Bu et al. (2021) provide a useful benchmark for robustness, but may also absorb communication content relevant to our object of interest.

predictions evaluated in Sections 5, 6, and 7.

### 3.1 A reduced-form macro-finance framework with communication

Let $t$ index scheduled FOMC announcement dates and let $\tau \in \{2, \ldots, 10\}$ denote maturity in years. Let $y_t^N(\tau)$ and $y_t^R(\tau)$ be nominal and real zero-coupon yields, and let $f_t^N(\tau)$ and $f_t^R(\tau)$ be the corresponding instantaneous forward rates. Define BEI yields and BEI forwards as

$$\mathrm{BEI}_t(\tau) \equiv y_t^N(\tau) - y_t^R(\tau), \qquad \mathrm{BEI}_t^f(\tau) \equiv f_t^N(\tau) - f_t^R(\tau). \tag{1}$$

In the term-structure literature, BEI rates are referred to as *inflation compensation* because they reflect risk-adjusted market pricing rather than a survey-based measure of expected inflation (Gürkaynak et al., 010a); following Acosta et al. (2025), we adopt this terminology throughout. The objects $\mathrm{BEI}_t(\tau)$ and $\mathrm{BEI}_t^f(\tau)$ are not pure expectations: they combine expected inflation with compensation for inflation risk and, in practice, TIPS liquidity components, and the relative importance of these components varies across maturities and regimes (Gürkaynak et al., 010b,a; D'Amico et al., 2018). Our empirical objects are announcement-day repricing $\Delta\mathrm{BEI}_t(\tau)$ and $\Delta\mathrm{BEI}_t^f(\tau)$, and the theoretical aim is to interpret their response to distinct, multidimensional communication signals across both daily and intraday identification windows.

We embed communication into a standard monetary–macro environment with forward-looking private-sector expectations and an asset-pricing representation. Let $(\pi_t, x_t)$ summarize the inflation and real-activity state relevant for market pricing at announcement date $t$. A reduced-form New Keynesian block is

$$\pi_t = \beta \mathbb{E}_t \pi_{t+1} + \kappa x_t + u_t^\pi, \tag{2}$$

$$x_t = \mathbb{E}_t x_{t+1} - \sigma^{-1}\left(i_t - \mathbb{E}_t \pi_{t+1} - r_t^n\right) + u_t^x, \tag{3}$$

where $i_t$ is the short nominal policy rate, $r_t^n$ is the natural real rate, $\sigma > 0$ is the inverse intertemporal elasticity of substitution (so that $\sigma^{-1}$ governs the sensitivity of aggregate demand to the ex-ante real-rate gap), and $u_t^\pi, u_t^x$ are shocks. The central bank sets $i_t$ using a systematic reaction function with possible deviations,

$$i_t = \bar{i} + \phi_\pi \pi_t + \phi_x x_t + \varepsilon_t^i, \tag{4}$$

where $\bar{i}$ denotes the long-run nominal policy rate and $\varepsilon_t^i$ is an unexpected policy-action component. The central bank also communicates around policy decisions through statements and press conferences. Let $\mathcal{I}_t^{CB}$ denote the information set of the central bank at the announcement time and $\mathcal{I}_{t^-}^M$ the market information set just before communication is released. The information-effect literature emphasizes that $\mathcal{I}_t^{CB}$ may contain signals about fundamentals not fully spanned by $\mathcal{I}_{t^-}^M$: the central bank may possess or synthesize information not yet incorporated in market prices.[5]

Let $\mathcal{S}_t^j$ denote the latent vector of communication signals released on announcement date $t$ through communication object $j$, where

$$j \in \mathcal{J} \equiv \{\text{post-meeting statement, Chair press conference}\}. \tag{5}$$

We suppress the superscript $j$ below; the empirical implementation constructs all signals separately for the two communication objects. We summarize this vector with three components: an overall stance signal and two inflation-specific signals corresponding to the Delphic/Odyssean distinction. First, $mps_t$ denotes an overall monetary-policy stance signal, capturing the hawkish–dovish tilt of the communication and proxying for the restrictiveness of the policy path implied by the text. Second, $\delta_t^\pi$ denotes a *Delphic* inflation signal: information about the inflation outlook, inflation persistence, and inflation-related risks. Third, $\omega_t^\pi$ denotes an *Odyssean* inflation signal: information about intended future policy behavior, namely forward guidance or reaction-function commitment, conditional on the inflation state. The separation follows the forward-guidance literature, in which Delphic guidance conveys news about the economic outlook while Odyssean guidance conveys commitment to a future policy path or reaction function (Campbell et al., 2012; Andrade and Ferroni, 2021).

For tractability, we write belief revisions as linear projections on the policy-action component and the communication signals. Let $X_t$ collect the macro-financial state relevant for inflation compensation, with $\pi_{t+h} \subset X_{t+h}$. The market's pre-announcement belief is $(\mathbb{E}_{t^-} X_{t+h})_{h\geq 0}$. On the announcement date, the market observes $(i_t, \mathcal{S}_t)$ and updates to $(\mathbb{E}_t X_{t+h})_{h\geq 0}$. Isolating the inflation component of these revisions,

$$\Delta_t \mathbb{E}\pi_{t+h} = a_\pi(h)\,\varepsilon_t^i + b_\pi(h)\,mps_t + c_\pi(h)\,\delta_t^\pi + d_\pi(h)\,\omega_t^\pi + \eta_{\pi,t}(h), \qquad h \geq 1, \tag{6}$$

where $h$ denotes the inner forecasting horizon (in the same time units as the maturity index $\tau$ used below for BEI), $\varepsilon_t^i$ denotes the unexpected policy-action component, and $\eta_{\pi,t}(h)$ collects other announcement-day news and projection error not summarized by the communication signals. Equation (6) formalizes the point of the information-effects literature: announcement-day repricing can reflect multiple belief revisions, not only the unexpected policy action. Delphic signals $\delta_t^\pi$ operate primarily through beliefs

[5]This formulation follows the Fed information-effect literature, in which central bank announcements may reveal information about fundamentals not fully incorporated in market prices (Nakamura and Steinsson, 2018; Jarociński and Karadi, 2020). It does not imply that announcement-day repricing is dominated by private-information revelation. As Bauer and Swanson (2023) emphasize, market responses may also reflect learning about the systematic policy reaction function. Our framework accommodates both interpretations: Delphic inflation signals capture communication that revises beliefs about the inflation outlook or the distribution of inflation risks, while Odyssean signals capture communication about intended policy behavior conditional on those states. In committee-vetted statements, these two dimensions are typically institutionally bundled, with inflation assessments communicated together with the implied policy response. We refer to this bundling as a co-entailment mechanism. The empirical analysis identifies which interpretation is more consistent with the sign, maturity profile, and timing of BEI repricing across communication formats, without claiming to structurally identify a unique channel.

about inflation-relevant fundamentals (e.g., shocks $u_t^{\pi}, u_t^{x}$, the natural rate, persistence), whereas Odyssean signals $\omega_t^{\pi}$ operate through beliefs about the future policy rule and the distribution of future $i_{t+h}$ conditional on states (Campbell et al., 2012; Nakamura and Steinsson, 2018; Bauer and Swanson, 2023). The stance component $mps_t$ captures the overall hawkish–dovish tilt of the communication. It is not inflation-specific and may affect inflation compensation through both expected-policy-path and credibility channels.

Asset prices translate these belief revisions into BEI yields and forwards. A standard no-arbitrage decomposition implies that breakeven inflation can be written, up to a first-order approximation, as expected average inflation over the maturity plus an inflation-compensation wedge:

$$\mathrm{BEI}_t(\tau) \approx \frac{1}{\tau}\sum_{h=1}^{\tau} \mathbb{E}_t \pi_{t+h} + \chi_t^{\pi}(\tau), \tag{7}$$

where $\chi_t^{\pi}(\tau)$ collects the non-expectational components of BEI, including inflation risk premia and relative liquidity premia in nominal Treasuries and TIPS (Gürkaynak et al., 010b,a; Haubrich et al., 2012; Abrahams et al., 2016; D'Amico et al., 2018). The corresponding forward object is

$$\mathrm{BEI}_t^{f}(\tau) \approx \mathbb{E}_t \pi_{t+\tau} + \chi_t^{\pi,f}(\tau), \tag{8}$$

where $\chi_t^{\pi,f}(\tau)$ is the forward inflation-compensation wedge. Equation (7) averages expected inflation and premia across horizons, whereas Equation (8) isolates marginal compensation at horizon $\tau$. BEI forwards therefore identify horizon-specific repricing more sharply, while BEI yields summarize the average compensation embedded in the term structure.

Combining the belief-revision equation (6) with the no-arbitrage representation, event-window repricing of BEI yields is the linear projection

$$\Delta\mathrm{BEI}_t(\tau) = B_i(\tau)\varepsilon_t^{i} + B_{mps}(\tau) mps_t + B_{\delta}(\tau)\delta_t^{\pi} + B_{\omega}(\tau)\omega_t^{\pi} + \epsilon_t(\tau). \tag{9}$$

The coefficients in Equation (9) are reduced-form loadings of inflation compensation on the policy-action component and the communication signals. For a generic signal $q \in \{i, mps, \delta, \omega\}$, the BEI loading decomposes conceptually into

$$B_q(\tau) = B_q^{\mathbb{E}\pi}(\tau) + B_q^{\chi}(\tau), \tag{10}$$

where $B_q^{\mathbb{E}\pi}(\tau)$ captures the contribution of revisions in expected inflation and $B_q^{\chi}(\tau)$ captures the contribution of non-expectational BEI wedges, including inflation risk premia and liquidity components. Our event-study regressions estimate $B_q(\tau)$, the total inflation-compensation loading. They do not separately identify $B_q^{\mathbb{E}\pi}(\tau)$ and $B_q^{\chi}(\tau)$, which would require additional structure such as a term-structure model, survey expectations, or inflation-swap-based evidence.

An analogous projection holds for the forward response $\Delta \mathrm{BEI}^{f}_{t}(\tau)$, with coefficients denoted $B^{f}_{q}(\tau)$. The distinction between yields and forwards is central to the paper. A communication signal that mainly revises beliefs about distant inflation risks may have a limited effect on BEI yields, because yields average across near and distant horizons, while generating a larger response in long-horizon forwards. Conversely, a signal that changes beliefs about the policy reaction function over the stabilization horizon may be concentrated in short- or intermediate-horizon forwards and only partially visible in yields.

The decomposition clarifies the economic interpretation of Delphic and Odyssean inflation signals. A Delphic signal that raises perceived inflation persistence or upside inflation risks can increase expected future inflation, raise compensation for inflation risk, or both; its effect should therefore be most visible at longer forward horizons, where persistence and cumulative inflation uncertainty are priced more directly. An Odyssean signal that strengthens the perceived policy reaction function can lower expected inflation and compress inflation risk compensation through a credibility channel; its effect should be concentrated at intermediate horizons, where markets price the expected stabilization of inflation back toward target (Campbell et al., 2012; Andrade and Ferroni, 2021). The stance component $mps_t$ is more general: it captures the overall hawkish–dovish tilt and may affect inflation compensation through expected-policy-path, credibility, and risk-premium channels. Its sign and maturity profile are therefore left unrestricted.

The linear projection in Equation (9) leaves the maturity profiles of the Delphic and Odyssean loadings unrestricted. We now impose minimal additional structure to obtain comparative predictions for their signs and shapes. The objective is not to derive a structural asset-pricing model but to discipline the maturity patterns that the event-study estimates should display if markets price Delphic and Odyssean inflation signals differently.

Consider first the expected-inflation loading of a Delphic signal. Delphic inflation communication is interpreted, in the standard reading, as news about the inflation outlook, inflation persistence, or the distribution of inflation risks. To illustrate the persistence channel, let inflation be expressed as a deviation from the long-run target, $\tilde{\pi}_t \equiv \pi_t - \pi^*$, and suppose that, locally,

$$\tilde{\pi}_{t+1} = \rho\, \tilde{\pi}_t + u^{\pi}_{t+1}. \tag{11}$$

A Delphic signal that raises the perceived persistence parameter by $\Delta\rho > 0$ changes the $h$-step-ahead inflation forecast by

$$\Delta_t \mathbb{E} \tilde{\pi}_{t+h} \approx h\, \rho^{h-1}\, \tilde{\pi}_t\, \Delta\rho. \tag{12}$$

When current inflation is above target, this revision is positive and is larger at more distant horizons over the relevant range. Treating $h$ as a continuous variable, $h\rho^h$ reaches its maximum at $h^* = -1/\ln\rho$; for annual persistence values $\rho \in [0.90, 0.98]$, this turning point lies between roughly 10 and 50 years. Over the 2–10-year maturities studied here, a Delphic persistence signal therefore loads relatively more strongly on long-horizon forwards than on near-horizon forwards. The prediction is not strict monotonicity at every maturity but concentration of the positive component toward the long end of the forward curve. Because BEI yields average compensation across horizons, the corresponding yield response is more muted:

$$B_\delta(\tau) \approx \frac{1}{\tau}\sum_{h=1}^{\tau} B_\delta^f(h). \tag{13}$$

This yield–forward distinction underpins the horizon-specific predictions tested in the empirical analysis.

Consider next the expected-inflation loading of an Odyssean signal. Odyssean inflation communication conveys intended policy behavior conditional on the inflation state. If $\omega_t^\pi$ tightens the perceived reaction function, by raising the perceived responsiveness to inflation or by strengthening a higher-for-longer commitment, markets revise down expected inflation over the horizon at which policy is expected to stabilize inflation. Let $H$ denote this market-perceived stabilization window. For horizons inside this window, the Odyssean loading is negative: $d_\pi(h) < 0$ for $h \leq H$. Beyond that window, the incremental loading attenuates because inflation is already expected to have returned toward target. The corresponding BEI-forward response $B_\omega^f(\tau)$ is therefore negative in the short-to-medium segment of the curve and weaker at longer horizons. This is the "belly-compression" prediction: Odyssean commitment affects most strongly the maturities at which markets price the expected return of inflation toward target.

Risk premia complicate, rather than mechanically overturn, these maturity predictions. BEI compensation combines expected inflation with non-expectational wedges, including inflation risk premia and TIPS liquidity components. A Delphic signal that raises perceived inflation persistence or upside inflation risks may increase both expected future inflation and the compensation investors require for bearing inflation risk, making the long-horizon BEI-forward response more positive. An Odyssean signal can work in the opposite direction if credible commitment lowers the perceived variance of future inflation over the stabilization window, thereby compressing both expected inflation and inflation risk compensation. The sign of the risk-premium response is ultimately an empirical matter: Acosta et al. (2025) emphasize that hawkish monetary-policy surprises may raise inflation risk premia, in which case the observed BEI decline understates the fall in true expected inflation. We therefore state the predictions at the level of total inflation compensation. Risk-premium movements may amplify or attenuate the expected-inflation channel, and our BEI regressions do not separately identify the two components.[6]

[6]If hawkish communication lowers expected inflation but raises inflation risk premia, the observed

The preceding predictions assume that markets can price Delphic and Odyssean signals separately. This assumption is more plausible for Chair press conferences, which are longer and more discursive and allow the Chair to explain inflation risks, outlook assessments, and policy contingencies in separate parts of the event. It is less plausible for post-meeting statements. Statements are short, committee-vetted, and drafted in a format that embeds inflation assessments within reaction-function language: an assessment that inflation is elevated is typically communicated together with, or immediately adjacent to, the policy response implied by that assessment.

A simple signal-extraction representation clarifies the implication. Suppose markets price the statement through a one-dimensional composite signal

$$z_t^{stmt} = \ell_\delta \delta_t^{\pi,stmt} + \ell_\omega \omega_t^{\pi,stmt}, \qquad \ell_\delta, \ell_\omega > 0, \tag{14}$$

where the composite is interpreted through the policy reaction function. If BEI repricing loads on $z_t^{stmt}$ rather than on the two components separately, then univariate regressions on either $\delta_t^{\pi,stmt}$ or $\omega_t^{\pi,stmt}$ inherit the sign and maturity profile of the same underlying policy-commitment signal. We call this property *co-entailment*: Delphic inflation assessments and Odyssean policy commitments are not necessarily identical textual objects, but in the statement format they are priced as mutually entailed parts of the same reaction-function message. Co-entailment does not require the two text indices to be perfectly correlated; it requires only that the institutional framing of the statement makes markets price both through a common policy-response signal. The empirical counterpart is developed in Section 6.

On announcement days with a post-meeting press conference, information arrives sequentially. At $T_0 \approx$ 14:00 ET, the written statement is released: markets price the current rate-decision surprise and the committee-vetted policy message. In this format, inflation assessments tend to be embedded in reaction-function language, so Delphic and Odyssean components are more likely to be priced through a common policy-response signal (Proposition 3). At $T_1 \approx$ 14:30 ET, when the Chair begins the press conference, the current rate decision has largely been incorporated into prices. The incremental information in the press conference is therefore different in character: it consists of narrative elaboration about the inflation outlook, risks, uncertainty, and the policy contingencies attached to those risks.

Let $\Delta \text{BEI}_t^{PC}(\tau)$ denote the BEI change measured in the narrow press-conference window. Because the current rate-decision surprise is largely priced at the statement release, this window contains little variation in the contemporaneous target-rate surprise. It can nevertheless contain substantial policy-path and state-contingent information. In particular, language that appears Odyssean in a daily regression may also reveal

BEI response is attenuated relative to the expected-inflation response. If credible Odyssean commitment lowers both expected inflation and inflation uncertainty, the two components reinforce each other. Distinguishing these cases would require a term-structure decomposition of BEI, survey expectations, or inflation-swap-based evidence (Haubrich et al., 2012; Abrahams et al., 2016; D'Amico et al., 2018).

the inflation state that motivates the commitment. A higher-for-longer policy posture, for example, is not only a commitment signal but also conveys information about the persistence, uncertainty, or upside risks of inflation. The press-conference window may therefore generate positive BEI repricing even for language that compresses BEI over the full daily window.

The daily BEI response is the net effect of distinct within-day components,

$$\Delta \mathrm{BEI}_t^{daily}(\tau) \approx \Delta \mathrm{BEI}_t^{stmt}(\tau) + \Delta \mathrm{BEI}_t^{PC}(\tau) + \xi_t(\tau), \tag{15}$$

where $\xi_t(\tau)$ collects residual same-day news, measurement differences between narrow and close-to-close windows, and market adjustment outside the two intraday windows. This sequential pricing structure generates a testable within-day implication: press-conference-window repricing should reveal the positive inflation-state component of the Chair's narrative, while the daily response reflects the net effect of this component and the statement-related reaction-function channel.

The four propositions that follow summarize the comparative predictions that guide the empirical analysis. They are not structural identifying assumptions; their role is to discipline the expected sign, maturity tilt, and timing of BEI repricing across communication formats.

**Proposition 1** (Maturity segmentation). *In communication formats where Delphic and Odyssean signals can be priced separately, Delphic inflation signals should load relatively more on long-horizon BEI forwards than on near-horizon forwards. The positive component of $B^f_\delta(\tau)$ is expected to be concentrated at long maturities. Odyssean inflation signals should instead generate negative responses over the market-perceived stabilization window, producing short-to-intermediate forward compression: $B^f_\omega(\tau) < 0$ for $\tau \leq H$ and attenuating effects for $\tau > H$. BEI yield responses are expected to be more muted because yields average compensation across horizons.*

**Proposition 2** (Risk-premium channel). *Risk-premium movements can amplify or attenuate the expected-inflation channel. Delphic signals that raise perceived inflation persistence or upside inflation risks may increase long-horizon inflation risk compensation, reinforcing positive long-end repricing. Odyssean signals that strengthen the credibility of the reaction function may reduce inflation uncertainty over the stabilization window, reinforcing belly compression. Because BEI combines expected inflation, inflation risk premia, and liquidity components, these predictions are stated for total inflation compensation rather than for its individual components.*

**Proposition 3** (Co-entailment in constrained formats). *In communication formats where inflation assessments and policy commitments are institutionally co-entailed, as in committee-vetted statements, Delphic and Odyssean signals are more likely to be priced through a common reaction-function message. If markets price the statement through a composite signal*

$$z_t^{stmt} = \ell_\delta \delta_t^{\pi,stmt} + \ell_\omega \omega_t^{\pi,stmt}, \qquad \ell_\delta, \ell_\omega > 0, \tag{16}$$

*then univariate regressions on either component inherit the sign and maturity profile of the same underlying policy-response signal. The prediction is that statement-based Delphic and Odyssean indices move BEI compensation in the same negative direction, with the strongest effects at the maturities where reaction-function news is priced.*

**Proposition 4** (Within-day dynamics)**.** *On announcement days that include a press conference, the press-conference window should reveal the positive inflation-state component of the Chair's narrative, because the current rate decision has already been largely priced at statement release. Statement-related communication operates primarily through the reaction-function channel and contributes to the daily BEI response without necessarily generating a detectable text effect within the narrow statement window after conditioning on the rate-decision surprise. The net daily response at each maturity reflects the balance between these components, plus any residual same-day adjustment outside the narrow intraday windows.*

## 3.2 Expected inflation versus the inflation risk premium

Propositions 1–4 are stated in terms of total inflation compensation, $\Delta \text{BEI}_t(\tau)$. This object reflects revisions in expected inflation, changes in inflation risk compensation, and a relative liquidity wedge between nominal Treasuries and TIPS (Gürkaynak et al., 010b; Fleckenstein et al., 2014). Using the notation of Equation (10), our event-study regressions estimate the total loading $B_q(\tau)$, not its separate components $B_q^{\mathbb{E}\pi}(\tau)$ and $B_q^{\chi}(\tau)$. Separating these channels would require additional structure such as an affine term-structure model (D'Amico et al., 2018; Abrahams et al., 2016), survey expectations, or inflation-swap-based evidence (Haubrich et al., 2012).

Because BEI movements can in principle arise from changes in nominal yields, in real (TIPS) yields, or in their relative liquidity, mapping our coefficients to a macroeconomic mechanism requires distinguishing among these underlying margins. We do not attempt that decomposition here; we read the empirical results as informative about the pricing of total inflation compensation and acknowledge that an extension reporting separate loadings on nominal and real yields would refine the interpretation.

The expected-inflation/risk-premium distinction matters for the interpretation of each channel separately. For the statement channel, if hawkish communication lowers expected inflation but raises inflation risk premia, the net effect on BEI is attenuated relative to the fall in expected inflation; if the risk-premium increase is large enough, the BEI response may be close to zero or even positive, so the negative statement coefficients are a lower bound on the magnitude of the expected-inflation decline rather than an unbiased estimate of it. The negative statement coefficients are therefore conservative with respect to the expected-inflation component. For the press-conference Delphic channel, positive long-horizon BEI-forward responses may reflect higher expected future inflation, higher compensation for inflation risk, or both. For the Odyssean belly-compression channel, a credibility interpretation implies that commitment language lowers expected inflation and may also reduce inflation risk compensation over the stabilization window. These component interpretations are economically distinct, but they are not separately identified by BEI data alone.

The contribution of the paper is therefore to document how total inflation compensation is repriced across communication objects, maturities, and measurement windows. Under an expected-inflation interpretation, the positive Delphic press-conference effect reflects revisions to the medium- and long-horizon inflation outlook. Under a risk-premium interpretation, it reflects higher perceived inflation uncertainty or a larger price of inflation risk at long horizons. Both interpretations are consistent with the total BEI response; distinguishing between them is beyond the identifying variation used in the baseline event-study design.

In the empirical analysis, the latent signals in $\mathcal{S}_t$ are proxied by document-level indices extracted separately from post-meeting statements and Chair press conferences. We collect these indices in a $K$-dimensional vector $CBC_t$ ($K = 5$ for the baseline indices defined in Section 4.1) and, stacking over meetings, in a $T \times K$ matrix $CBC$ with row $t$ equal to $CBC_t$. Section 4 describes the construction of the text measures, and Section 5 describes the event-study design used to estimate the maturity-specific loadings.

# 4 Data

The empirical design relies on three data inputs. The first is a textual corpus of FOMC post-meeting statements and Chair press-conference transcripts, transformed into document-level communication indices. The second is a panel of daily BEI yields and instantaneous forwards from fitted nominal Treasury and real TIPS curves over maturities from two to ten years, complemented by intraday BEI changes at selected maturities constructed from high-frequency Treasury and TIPS yield movements. The third is the set of high-frequency monetary-policy surprises around FOMC announcements, which serve as orthogonalization controls in the event-study design.

## 4.1 Communication indices from FOMC texts

The text corpus contains 188 post-meeting FOMC statements (2004–2025) and 89 Chair press-conference transcripts (2011–2025), aligned to scheduled FOMC meetings. The event-study samples are slightly smaller: $N = 185$ for statements and $N = 88$ for press conferences. The discrepancy reflects three unscheduled or special-circumstance meetings whose announcement-day BEI observations are unavailable or non-standard, and one press conference for which the daily close BEI cannot be matched to the meeting date. The intraday samples differ further: $N = 180$ for statement-window regressions and $N = 89$ for press-conference-window regressions, reflecting the availability of high-frequency Treasury and TIPS quotes in the relevant event windows. We treat each statement and each press conference as a single communication object released on the announcement date, and map each document into a small set of indices designed to capture (i) overall monetary-policy stance and (ii) inflation-relevant narratives, including Delphic and Odyssean components.

The measurement starts from FOMC-RoBERTa, a domain-adapted transformer fine-tuned to classify FOMC text as hawkish, dovish, or neutral (Shah et al., 2023). We apply the classifier to textual units obtained from each document after sentence segmentation and, when needed, model-length chunking. For each unit $i$ in document $t$, the classifier returns model-implied class probabilities $(p_{it}^H, p_{it}^D, p_{it}^N)$, with $p_{it}^H + p_{it}^D + p_{it}^N = 1$. We map these probabilities into a signed unit-level stance score and a model-implied informativeness measure:

$$m_{it} \equiv p_{it}^H - p_{it}^D, \qquad \nu_{it} \equiv 1 - p_{it}^N. \tag{17}$$

Positive values of $m_{it}$ correspond to hawkish language, negative values to dovish language, and $|m_{it}|$ reflects the model-implied tilt between hawkish and dovish classes. The scalar $\nu_{it} \in [0, 1]$ measures the extent to which the model assigns the unit to a directional (non-neutral) category and is retained as a diagnostic indicator rather than as a baseline aggregation weight.

We aggregate unit-level stance into a document-level monetary-policy stance index. For each document $t$,

$$mps_t \equiv \frac{1}{N_t} \sum_{i=1}^{N_t} m_{it}, \tag{18}$$

computed separately for statements and press conferences.

To capture inflation-relevant communication, we complement the classifier with transparent, pre-specified dictionary patterns that identify inflation-related content and distinguish different inflation narratives.

Let $g_{it}^{(k)} \in [0, 1]$ denote a theme-focus score for unit $i$ and theme $k$, constructed from interpretable pattern matches (words and short phrases). We consider four inflation-related themes aligned with the paper's indices: broad inflation-topic language ($k = \pi^g$), current and recent inflation developments ($k = \pi^{cur}$), Delphic inflation language ($k = \delta^\pi$), and Odyssean inflation language ($k = \omega^\pi$). We combine theme focus and unit-level stance using nonnegative deterministic salience weights $w_{it} \geq 0$. The baseline specification does not use classifier-confidence weighting; instead, the weight $w_{it}$ captures only transparent salience adjustments tied to the communication object itself. For FOMC statements, $w_{it}$ applies mild deterministic salience adjustments to units containing explicit policy-instrument or balance-sheet language. For press conferences, the same policy-salience component is combined with a turn-type adjustment that gives lower weight to journalist questions than to the Chair's answers. Appendix A reports the exact weighting rules. This design preserves interpretability and avoids treating classifier probabilities as calibrated certainty measures. Model-implied quantities such as $\nu_{it} \equiv 1 - p_{it}^N$ are retained as diagnostic indicators of directional informativeness but do not enter the baseline weighting rule.

For inflation narratives common to both corpora, we use an exposure-weighted aggregation (an "extensive × intensive" construction):

$$\pi_t^g \equiv \frac{\sum_{i=1}^{N_t} w_{it}\, g_{it}^{(\pi^g)}\, m_{it}}{\sum_{i=1}^{N_t} w_{it}}, \tag{19}$$

$$\pi_t^{cur} \equiv \frac{\sum_{i=1}^{N_t} w_{it}\, g_{it}^{(\pi^{cur})}\, m_{it}}{\sum_{i=1}^{N_t} w_{it}}. \tag{20}$$

The indices are dictionary-guided[7]: the inflation dictionaries are initialized from the most frequent inflation-related terms in the corpus and refined to increase precision and reduce false positives, while preserving full transparency and replicability.

Figures 1–3 display the standardized time series of each baseline index. The monetary-policy stance (Figure 1) tilts dovish during the zero-lower-bound era and turns hawkish during the 2022–2023 tightening cycle; statements and press conferences co-move but are not identical, which confirms that the two corpora encode partly distinct information. The broad inflation narrative (Figure 2) rises steeply after 2021, reflecting the post-pandemic inflation surge, while the current-inflation index (Figure 3) captures higher-frequency variation in realized-inflation language.

For Delphic and Odyssean inflation signals, we follow the forward-guidance distinction between outlook information (Delphic) and intended policy behavior/reaction function (Odyssean) (Campbell et al., 2012; Andrade and Ferroni, 2021). For press conferences, we again use the exposure-weighted aggregation:

$$\delta_t^{\pi} \equiv \frac{\sum_{i=1}^{N_t} w_{it}\, g_{it}^{(\delta^\pi)}\, m_{it}}{\sum_{i=1}^{N_t} w_{it}}, \tag{21}$$

$$\omega_t^{\pi} \equiv \frac{\sum_{i=1}^{N_t} w_{it}\, g_{it}^{(\omega^\pi)}\, m_{it}}{\sum_{i=1}^{N_t} w_{it}}. \tag{22}$$

For statements, we adopt a more conservative aggregation. Statements are short and tightly structured, so an exposure-weighted measure may load mechanically on how much of the limited statement text is devoted to Delphic/Odyssean markers rather than on the stance tilt *within* those themes. We therefore use a within-theme (conditional)

[7] Appendix C reports the lexical items underlying each inflation dictionary for both corpora (statements and press conferences) and provides a concise description of the co-occurrence rules used to construct $\pi_t^{cur}$, $\delta_t^{\pi}$, and $\omega_t^{\pi}$. In implementation, dictionary entries are matched case-insensitively using regular expressions (regex), i.e., text patterns that accommodate common variants (pluralization and minor hyphenation/spacing differences).

**Figure 1.** Time series of $mps_t$ at FOMC-meeting events

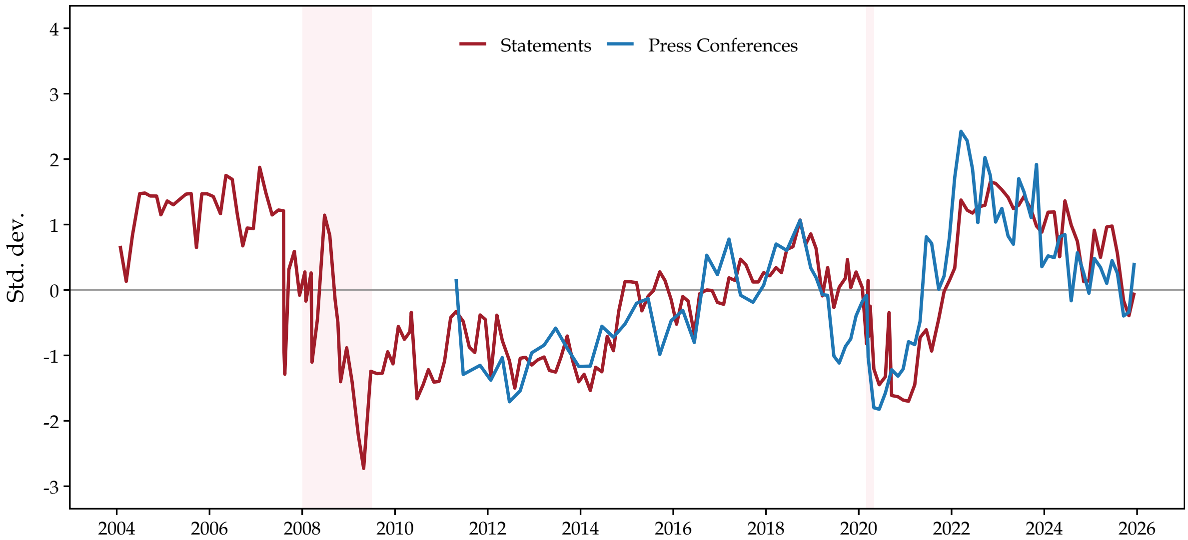


*Notes:* The index is standardized within each corpus and expressed in standard-deviation units. Shaded pink areas indicate NBER recession periods. Standardized overlays are intended to show within-corpus time variation rather than cross-corpus level differences: a one-standard-deviation movement in the statement index need not correspond to the same underlying textual shift as a one-standard-deviation movement in the press-conference index.

aggregation for statements:

$$\delta_t^{\pi} \equiv \frac{\sum_{i=1}^{N_t} w_{it}\, g_{it}^{(\delta\pi)}\, m_{it}}{\sum_{i=1}^{N_t} w_{it}\, g_{it}^{(\delta\pi)}}, \tag{23}$$

$$\omega_t^{\pi} \equiv \frac{\sum_{i=1}^{N_t} w_{it}\, g_{it}^{(\omega\pi)}\, m_{it}}{\sum_{i=1}^{N_t} w_{it}\, g_{it}^{(\omega\pi)}}. \tag{24}$$

If the denominator in the within-theme aggregation is zero, the corresponding index is set to zero before standardization and a no-hit diagnostic is retained. This convention keeps the event-study sample fixed and treats documents with no identified theme exposure as neutral with respect to that theme. Reporting no-hit rates by theme and corpus, and re-estimating the event-study regressions after dropping no-hit observations rather than imputing zero, would separate the extensive (occurrence) and intensive (intensity) margins of inflation language; we treat this as a measurement extension.

A consequence of using different aggregation formulas is that cross-channel comparisons (statements versus press conferences) involve indices constructed with different denominators. Signs and maturity profiles remain directly comparable across corpora; raw levels and cross-corpus magnitude comparisons should be read with caution. All regressions standardize each index within its own sample, so coefficient magnitudes are interpretable as basis-point responses to a one-standard-deviation move in the corresponding index. The main qualitative findings, namely negative BEI effects for

**Figure 2.** Time series of $\pi_t^g$ at FOMC-meeting events

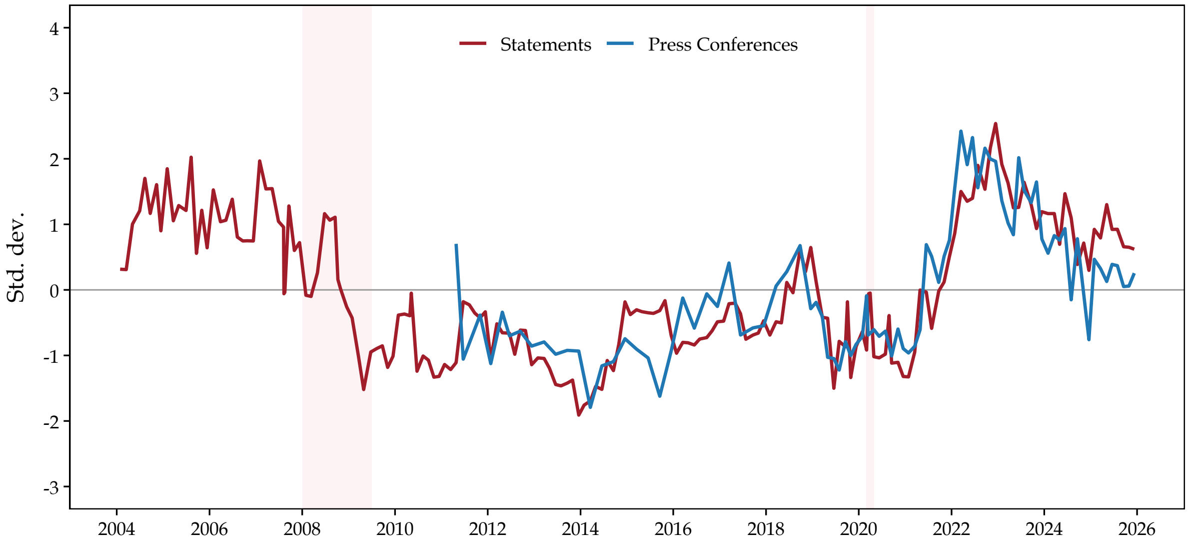


*Notes:* See notes to Figure 1.

statements and positive long-horizon forward effects for press conferences, hold across the robustness exercises that vary the surprise control (Appendix E), which indicates that the cross-channel asymmetry is not driven by the aggregation choice.

For each corpus (statements and press conferences), we collect the document-level indices in a communication index vector

$$CBC_t \equiv \left(mps_t,\ \pi_t^g,\ \pi_t^{cur},\ \delta_t^\pi,\ \omega_t^\pi\right)'. \tag{25}$$

Stacking over meetings yields the corresponding $T \times K$ matrix $CBC$ used in the empirical analysis.

For visual comparability across indices and corpora, we plot standardized versions of each document-level series (mean zero, unit variance) and report one figure per index with both statements and press conferences. Figures 4 and 5 display the Delphic and Odyssean indices. The Delphic index (Figure 4) captures outlook-driven variation, rising during the 2021–2023 episode as the FOMC discussed inflation risks and persistence. The Odyssean index (Figure 5) loads on policy-commitment language and rises during tightening cycles, particularly as the Committee emphasized "restrictive" and "higher-for-longer" postures during 2022–2024.

Tables 1 and 2 present the distributional properties and co-movement of the raw communication indices. Two features of the data bear directly on interpretation.

The first is the contrast in dispersion between corpora for the Delphic and Odyssean components. Statement-based $\delta_t^\pi$ and $\omega_t^\pi$ are roughly six to fourteen times more dispersed than their press-conference counterparts. The post-meeting statement is a collectively edited document whose inflation language shifts in deliberate, often large increments across regimes (from the explicit disinflation concerns of 2009 to the

**Figure 3.** Time series of $\pi_t^{cur}$ at FOMC-meeting events

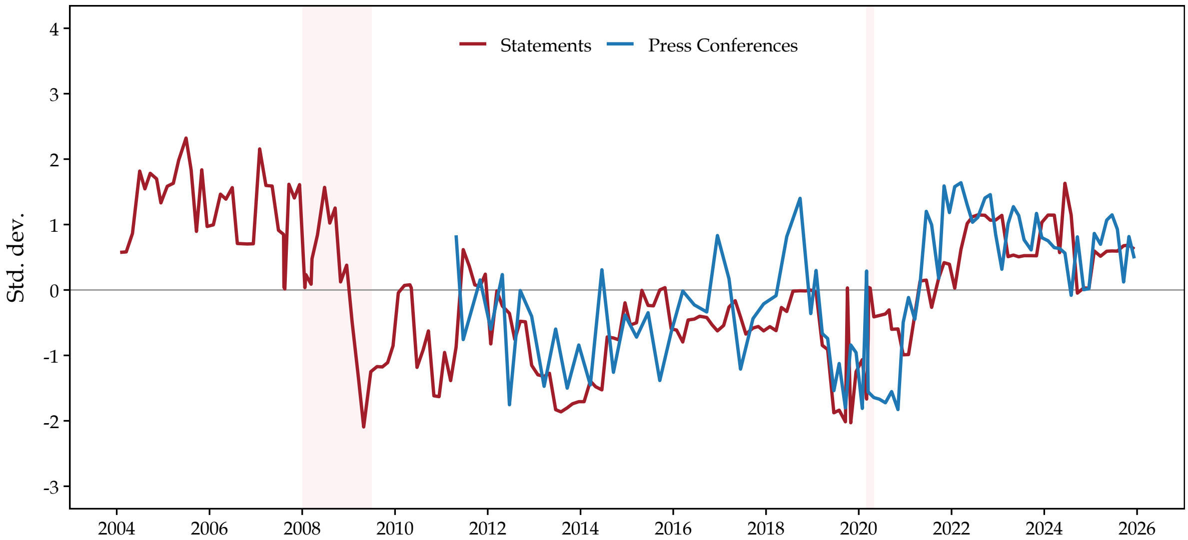


*Notes:* See notes to Figure 1.

hawkish resolve of 2022), whereas the Chair's spoken language at a press conference is modulated within a narrower register that prioritizes continuity and avoids unscripted policy commitments. The wider within-theme (conditional) construction for statements mechanically amplifies cross-meeting variation, whereas the exposure-weighted construction for press conferences compresses it.

The second feature is the heterogeneous correlation structure within each corpus. Stance, broad inflation language, and current-inflation assessment are tightly correlated within both samples, consistent with a common hawkishness component running through several dimensions of communication. The Delphic–Odyssean pair, by contrast, has a lower correlation ($\rho = 0.37$ for statements, $\rho = 0.61$ for press conferences), which provides enough within-corpus variation to compare the univariate maturity profiles of the two indices.

A textual feature deserves comment. The Delphic–Odyssean correlation is lower for statements ($\rho = 0.37$) than for press conferences ($\rho = 0.61$), even though Proposition 3 predicts that markets price statement Delphic and Odyssean indices through a common reaction-function signal. This is not a contradiction. Proposition 3 concerns market pricing, not the unconditional textual correlation of the two indices: independent classification by the dictionary gates is fully consistent with both indices being priced through a single composite signal in the constrained statement format.

The high inter-index correlations bear directly on the empirical strategy. With pairwise correlations reaching 0.90 in the statement panel and 0.89 in the press-conference panel, a multivariate specification would face collinearity that inflates standard errors and complicates the interpretation of individual coefficients as channel-specific estimates. We therefore follow Acosta et al. (2025) and estimate one index at a time; each

**Figure 4.** Time series of $\delta_t^\pi$ at FOMC-meeting events

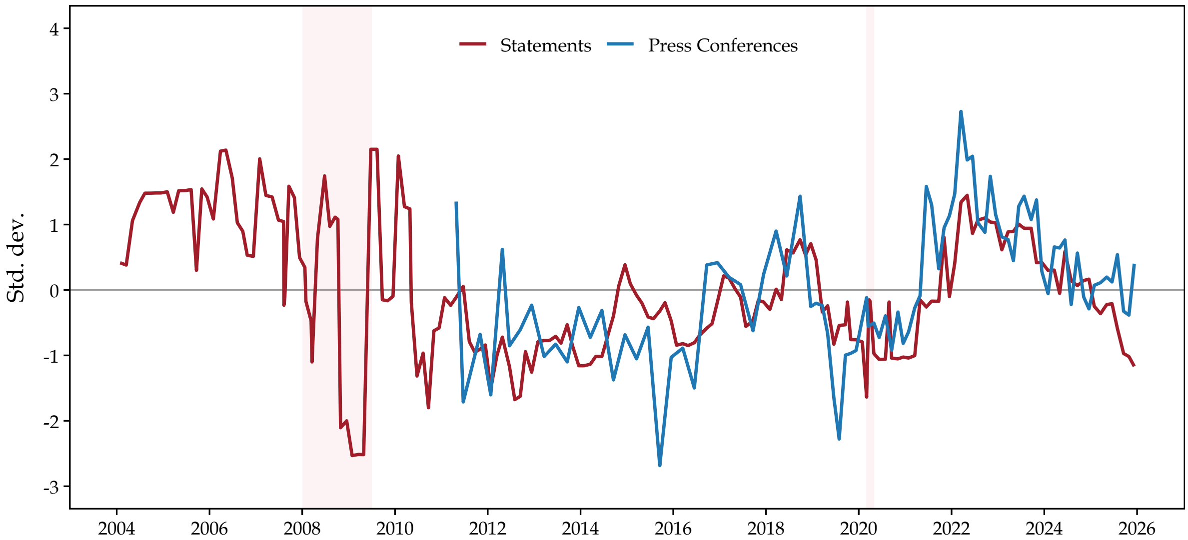


*Notes:* The index is standardized within each corpus and expressed in standard-deviation units. Shaded pink areas indicate NBER recession periods. Delphic and Odyssean indices are constructed using within-theme aggregation for statements and exposure-weighted aggregation for press conferences; cross-corpus magnitudes should therefore be interpreted with caution.

coefficient is interpreted as the total BEI response to a broad shift in communication of which the named index is a summary measure, not as a partial derivative holding other dimensions fixed.[8] The univariate design also has a clean economic interpretation: markets are unlikely to decompose a single statement or press conference into mutually orthogonal textual components before pricing it. The univariate coefficients summarize how BEI repricing covaries with each narrative dimension when that dimension is used as a stand-alone summary of the document.

## 4.2 Market data

The daily inflation-compensation outcomes are constructed from fitted nominal Treasury and real TIPS zero-coupon yield curves, using the publicly available Federal Reserve datasets that implement the Gürkaynak et al. (2007) Treasury and Gürkaynak et al. (010b) TIPS zero-coupon yield-curve methodologies, supplemented for missing observations as documented in the replication archive. For each FOMC announcement date $t$ and maturity $\tau \in \{2, \dots, 10\}$, we define the BEI yield and the BEI instantaneous

[8]This univariate design is a direct consequence of the high inter-index correlations documented in Table 2. With pairwise correlations approaching 0.90, conclusions about which specific dimension "drives" repricing should be drawn cautiously: the maturity profiles documented below characterize the association between each index and BEI conditional on the rate surprise but do not establish that a given dimension is incrementally informative after conditioning on other dimensions. Cross-index comparisons of coefficient magnitudes, significance levels, and $R^2$ values should be read as descriptive summaries of the univariate associations rather than as evidence of differential causal channels.

**Figure 5.** Time series of $\omega_t^{\pi}$ at FOMC-meeting events

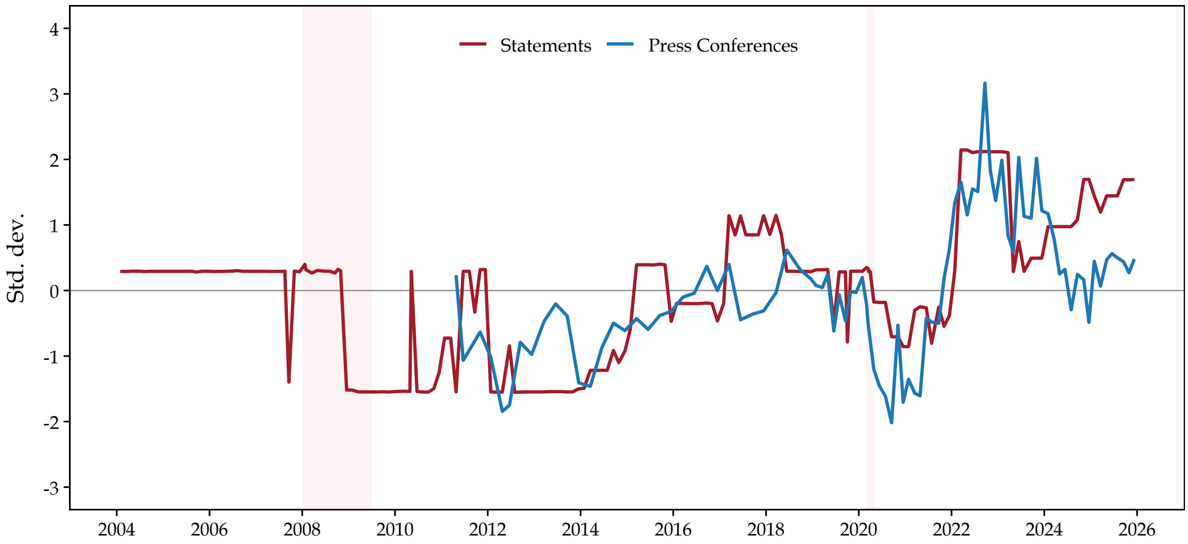


*Notes:* See notes to Figure 4.

forward as

$$\text{BEI}_t(\tau) = y_t^N(\tau) - y_t^R(\tau), \qquad \text{BEI}_t^f(\tau) = f_t^N(\tau) - f_t^R(\tau), \tag{26}$$

where nominal and real yields are evaluated at the daily close. The dependent variables in the daily event-study regressions are close-to-close changes on scheduled FOMC announcement dates, $\Delta\text{BEI}_t(\tau)$ and $\Delta\text{BEI}_t^f(\tau)$. The daily sample begins in 2004, when the TIPS market is generally considered sufficiently liquid for term-structure inference, and we assess the sensitivity of the results to stressed episodes in the robustness analysis.

The high-frequency surprise controls are taken from the event-study database of Acosta et al. (2025), which separates FOMC communication into distinct intraday event windows. The *statement window* is a 30-minute interval around the release of the post-meeting statement, beginning 10 minutes before the release and ending 20 minutes after it. The *press-conference window* is a 70-minute interval around the Chair's press conference, beginning 10 minutes before the start of the press conference and ending 60 minutes after it begins. The *monetary-event window* combines the two: on days with a press conference, it begins 10 minutes before the statement release and ends 60 minutes after the Chair begins speaking; on days without a press conference, it coincides with the statement window. Acosta et al. (2025) compute high-frequency changes using median prices or yields in pre-event and post-event windows surrounding these intervals, which reduces sensitivity to outliers and misquotes.

The baseline specification conditions on the high-frequency surprise in the current policy-rate decision, denoted $MP1_t$. This variable is constructed from federal funds futures and captures the unexpected component of the current FOMC meeting decision

**Table 1.** Descriptive statistics of communication indices

| *Panel A: Post-meeting statements* ($N = 188$) | | | | |
|---|---|---|---|---|
| Index | Mean | SD | Min | Max |
| $mps_t$ | −0.04 | 0.27 | −0.77 | 0.47 |
| $\pi_t^g$ | 0.01 | 0.16 | −0.30 | 0.41 |
| $\pi_t^{cur}$ | −0.00 | 0.11 | −0.23 | 0.24 |
| $\delta_t^\pi$ | 0.07 | 0.43 | −1.00 | 1.00 |
| $\omega_t^\pi$ | −0.16 | 0.54 | −1.00 | 1.00 |
| *Panel B: Press conferences* ($N = 89$) | | | | |
| Index | Mean | SD | Min | Max |
| $mps_t$ | −0.06 | 0.25 | −0.51 | 0.53 |
| $\pi_t^g$ | 0.03 | 0.11 | −0.18 | 0.30 |
| $\pi_t^{cur}$ | −0.05 | 0.52 | −1.00 | 0.79 |
| $\delta_t^\pi$ | −0.00 | 0.07 | −0.19 | 0.18 |
| $\omega_t^\pi$ | −0.00 | 0.04 | −0.09 | 0.13 |

*Notes:* Raw, non-standardized series.

**Table 2.** Pairwise correlations of communication indices

| | $mps_t$ | $\pi_t^g$ | $\pi_t^{cur}$ | $\delta_t^\pi$ | $\omega_t^\pi$ |
|---|---|---|---|---|---|
| $mps_t$ | 1.00 | 0.86 | 0.74 | 0.75 | 0.70 |
| $\pi_t^g$ | 0.89 | 1.00 | 0.90 | 0.71 | 0.69 |
| $\pi_t^{cur}$ | 0.81 | 0.81 | 1.00 | 0.67 | 0.56 |
| $\delta_t^\pi$ | 0.84 | 0.88 | 0.78 | 1.00 | 0.37 |
| $\omega_t^\pi$ | 0.86 | 0.83 | 0.67 | 0.61 | 1.00 |

*Notes:* Upper-triangular entries refer to post-meeting statements; lower-triangular entries refer to Chair press conferences. Diagonal entries equal one. Correlations are computed on raw, non-standardized series within each corpus separately.

in the relevant event window. Orthogonalizing each communication index with respect to $MP1_t$ removes the component of text variation linearly associated with the current rate-decision surprise, while leaving policy-path, outlook, and reaction-function content that may be transmitted through communication. The choice is parsimonious by construction: the object is not to purge all policy news from the text but to study whether inflation narratives contain information for BEI repricing beyond the current rate decision.

A critical identification detail is that the relevant surprise window differs by communication object. For post-meeting statements, we use the statement-window $MP1_t$, because the written statement and the current rate-decision surprise are released and priced in that window. For daily press-conference specifications, using the press-conference-window $MP1_t$ would provide little meaningful orthogonalization: by the time the Chair speaks, the current policy-rate decision has already been announced with the statement, so the corresponding same-window $MP1_t$ is close to degenerate. We therefore orthogonalize press-conference text indices using $MP1_t$ measured over the monetary-event window, which captures the rate-decision surprise relevant for the daily close-to-close BEI change and avoids a nearly zero control variable in the press-conference-only window.

For the intraday specifications, the logic is different. The dependent variable is measured within the same narrow window as the communication object, so the surprise control is also measured in that window. Statement-window intraday regressions therefore use statement-window $MP1_t$, whereas press-conference-window intraday regressions use press-conference-window $MP1_t$. The latter is close to zero by construction, but using it preserves window consistency: the intraday press-conference regressions ask whether communication indices explain BEI movements within the Chair's speaking window, after the current rate decision has already been largely priced.

As a robustness exercise, we replace $MP1_t$ with the broader monetary-policy news factor constructed by Acosta et al. (2025). For each event window, that factor is the first principal component of high-frequency changes in money-market futures rates covering the current meeting and the expected policy path over roughly the following year. It summarizes a broader policy-news component than $MP1_t$, including path news and forward-guidance information; it is therefore a stricter and more conservative control and may absorb genuine communication content, especially when policy-path information is transmitted through the statement or the press conference itself.

For the intraday BEI analysis, we use high-frequency changes in market-based inflation compensation from the same database. These outcomes are available at the 5-year, 10-year, and 5-year/5-year forward maturities and are constructed as the window changes in nominal Treasury yields minus the window change in TIPS yields:

$$\Delta \text{BEI}_t^w(k) = \Delta y_t^{N,w}(k) - \Delta y_t^{R,w}(k), \qquad k \in \{5\text{y}, 10\text{y}, 5\text{y}5\text{y}\}, \tag{27}$$

where $w$ denotes the relevant intraday event window. Intraday BEI changes are smaller in magnitude and variance than their daily counterparts. Narrow event windows reduce, but do not eliminate, non-FOMC news contamination: quote staleness, intraday liquidity variation, and contemporaneous order-flow effects, particularly in the TIPS market, remain. The intraday windows nevertheless provide a higher signal-to-noise setting for assessing the timing of communication effects than close-to-close daily changes.

# 5 Main results

This section presents the baseline event-study evidence. We ask whether inflation-related FOMC language is systematically associated with announcement-day changes in market-based inflation compensation, conditional on the high-frequency surprise in the current policy-rate decision. The object of interest is not a structurally identified communication shock but the component of each text index that is linearly orthogonal to the current rate-decision surprise. As discussed in Sections 4.2 and 3, the resulting coefficients are reduced-form associations conditional on the rate decision, not causal effects of isolated narrative components. The analysis combines the document-level communication indices defined in Equation (25), daily BEI yields and forwards constructed from nominal Treasury and real TIPS curves, and high-frequency monetary-policy surprises from the FOMC event-study database.

For each maturity $\tau \in \{2, \ldots, 10\}$ and each outcome, we estimate a separate event-study regression. Let $\Delta\text{BEI}_t(\tau)$ denote the daily change in the BEI yield of maturity $\tau$ on FOMC announcement date $t$, and let $\Delta\text{BEI}_t^f(\tau)$ denote the corresponding change in the instantaneous BEI forward. For each scalar component $CBC_{k,t}$ of $CBC_t$, we first orthogonalize the raw text index with respect to the rate-decision surprise:

$$CBC_{k,t} = a_k + \gamma_k MP1_t + u_{k,t}, \qquad CBC_{k,t}^{\perp} \equiv u_{k,t}. \tag{28}$$

We then standardize the residualized index within the relevant sample:

$$\widetilde{CBC}_{k,t}^{\perp} \equiv \frac{CBC_{k,t}^{\perp} - \overline{CBC_k^{\perp}}}{\text{sd}(CBC_k^{\perp})}. \tag{29}$$

The baseline maturity-by-maturity specification is

$$\Delta Y_t(\tau) = \alpha_k(\tau) + \beta_k(\tau)\widetilde{CBC}_{k,t}^{\perp} + \varepsilon_{k,t}(\tau), \tag{30}$$

where $\Delta Y_t(\tau)$ is either $\Delta\text{BEI}_t(\tau)$ or $\Delta\text{BEI}_t^f(\tau)$. Because the regressor is standardized, $\widehat{\beta}_k(\tau)$ is measured in basis points and represents the BEI co-movement associated with a one-standard-deviation shift in the target-orthogonalized text index. Statistical inference uses a wild bootstrap with $B = 10{,}000$ replications and Rademacher weights, which is robust to heteroskedasticity and applicable in small samples (Cameron et al., 2008).[9]

[9] The analysis estimates many maturity-by-maturity specifications across two communication formats and five text indices, generating on the order of one hundred eighty pointwise tests in the daily-frequency tables alone. The wild bootstrap addresses heteroskedasticity at each pointwise test but does not address the family-wise error rate, and adjacent maturities are mechanically correlated because the underlying fitted yield-curve inputs are highly correlated across maturities. We therefore interpret the evidence primarily through sign consistency across the curve, the shape of the maturity profile, and robustness across alternative controls and samples, rather than through isolated pointwise significance at any single maturity. A formal joint inference exercise, using max-$t$ confidence bands or a Romano–Wolf stepdown

## 5.1 Statements

We begin with the post-meeting statement. Figure 6 reports the term structure of BEI responses to three text indices available for both communication objects: the overall monetary-policy stance index ($mps_t$), the broad inflation narrative index ($\pi_t^g$), and the current-inflation index ($\pi_t^{cur}$). Each index is orthogonalized with respect to the statement-window $MP1_t$, aligning the surprise control with the window in which the written statement and the current rate decision are released.

The empirical strategy follows a hierarchy of specifications, moving from least to most conservative with respect to the rate-and-path information set. The baseline reported in this section conditions only on the current rate-decision surprise $MP1_t$ and is the specification that retains the largest volume of identifying variation in the text indices. Appendix E additionally reports (i) the unconditional specification that omits any intraday surprise control, capturing the total announcement-day BEI co-movement with each index; (ii) a stricter specification that purges each index with respect to the first principal component of the Acosta et al. (2025) high-frequency surprises, which absorbs path-news content and is therefore an upper bound on the purging of rate-related variation; (iii) sample-restriction checks that exclude the Global Financial Crisis window and calendar year 2020; and (iv) an SEP/non-SEP interaction specification for press conferences. The qualitative sign and maturity patterns reported in this section are preserved across these alternative specifications, although precision varies with sample size and the breadth of the surprise control. None of these specifications fully purges path news, SEP-related outlook information, reaction-function learning, or other information-effect components of the announcement-day FOMC bundle (Nakamura and Steinsson, 2018; Jarociński and Karadi, 2020; Bauer and Swanson, 2023; Bu et al., 2021; Romer and Romer, 2000; Melosi, 2017).

The point estimates are negative at every maturity for the three baseline statement indices, with statistical significance concentrated at the short-to-intermediate end and attenuating beyond seven years. The pattern holds for the overall stance index and for the broad inflation index, and it is weaker but still negative for the current-inflation index. Statement language with a more hawkish or inflation-focused tilt is therefore priced as reducing inflation compensation rather than raising it.

The forward-rate estimates sharpen the maturity interpretation. Statement indices co-move with negative BEI-forward responses at the short-to-intermediate end of the curve, while responses at longer forward horizons are small and statistically indistinguishable from zero. The yield–forward contrast is informative: the statement does not generate the positive long-horizon repricing that would be expected if markets interpreted the text primarily as Delphic news about a deteriorating inflation outlook. The response is concentrated at horizons where markets price the expected stabilization of inflation through the policy reaction function.

procedure across maturities, would sharpen the shape-based claims and is a natural extension. We discuss this further in Section 9.

**Figure 6.** Term structure of inflation compensation sensitivity to FOMC statements

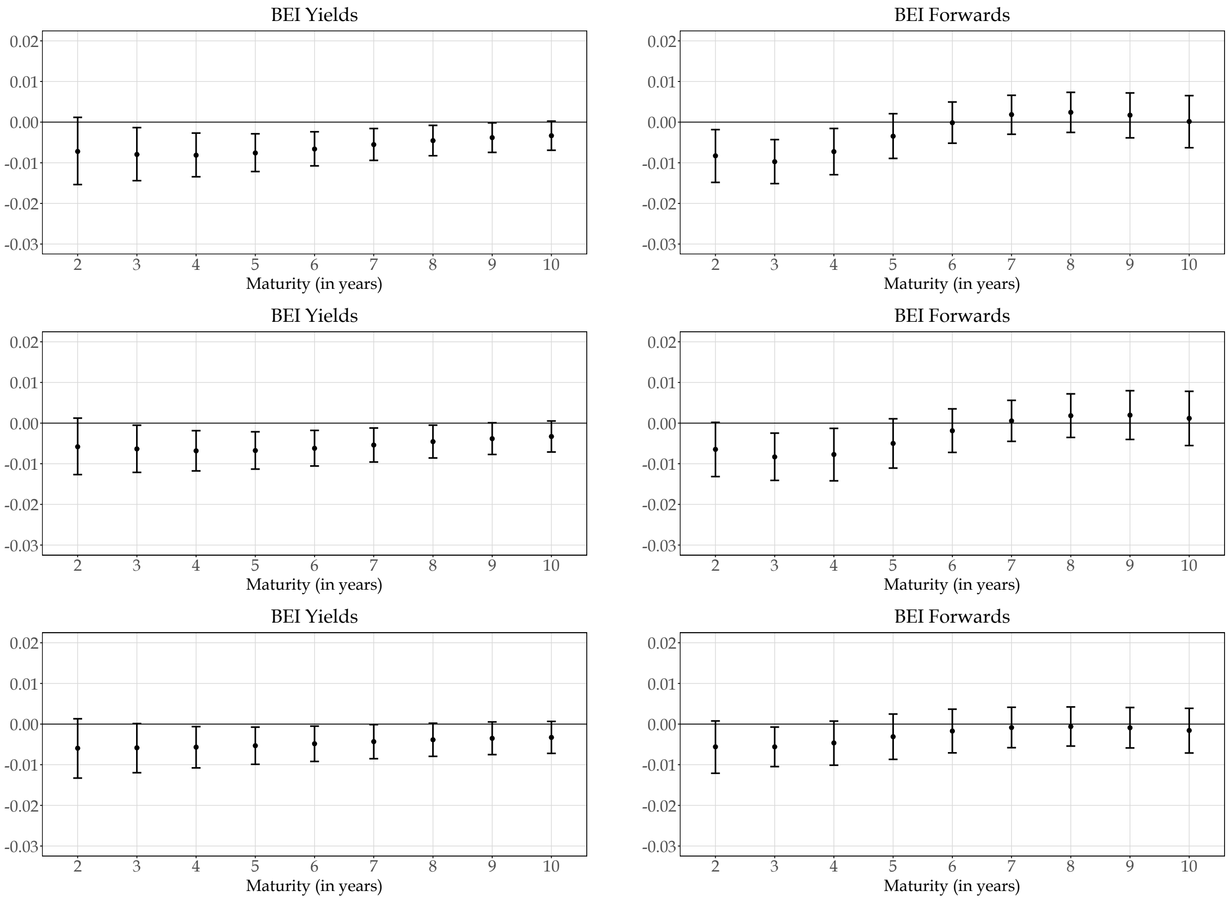


*Notes:* Event-study regressions of daily changes in BEI yields (left column) and BEI forwards (right column) on target-orthogonalized FOMC statement indices. The top, middle, and bottom rows report responses to $mps_t$, $\pi_t^g$, and $\pi_t^{cur}$, respectively. The sample comprises $N = 185$ scheduled FOMC announcements from 2004 to 2025. Dots denote OLS estimates for maturities $\tau \in \{2, \ldots, 10\}$; error bars are 90% confidence intervals based on wild bootstrap ($B = 10{,}000$, Rademacher weights).

This interpretation is consistent with the institutional role of the post-meeting statement. The statement is short, committee-vetted, and formulaic; inflation assessments in this format are typically communicated together with the implied policy response. Markets appear to read statement inflation language through the reaction function: language that acknowledges inflation pressure is priced as a signal of policy resolve, compressing BEI compensation at the maturities where stabilization news is most relevant. This is the empirical counterpart of the co-entailment mechanism in Proposition 3.

The reaction-function reading is an interpretation of the daily co-movement, not a high-frequency identification of statement text being priced through the reaction function within the statement window itself. As Section 7 documents, text-based statement indices do not explain a detectable component of BEI repricing within the narrow statement-release window once the contemporaneous target-rate surprise is partialled out. The negative daily statement coefficient therefore reflects broader-session pricing of the FOMC announcement bundle: possibly delayed digestion of the committee-vetted text, possibly the interaction between the statement and subsequent press-conference information on press-conference days, and possibly residual same-day information unrelated to the text. We return to this timing limitation in Sections 7 and 9.

The pattern is consistent with the high-frequency monetary-policy literature. Acosta et al. (2025) document that monetary-policy surprises have negative effects on BEI, with the strongest responses at intermediate maturities, and interpret these patterns as consistent with policy news affecting beliefs about the conduct of monetary policy. Our evidence is complementary: rather than using market-based policy surprises directly, we show that the text of the statement contains inflation-related variation that covaries with BEI in the same direction, conditional on the current rate-decision surprise.

## 5.2 Press conferences

We turn to Chair press conferences. The press-conference sample begins in 2011 and contains $N = 88$ scheduled FOMC announcements. For daily press-conference specifications, we orthogonalize each text index with respect to $MP1_t$ measured over the monetary-event window rather than the press-conference-only window. As discussed in Section 4.2, the current rate decision has already been announced by the time the Chair speaks, so the same-window target-rate surprise contains little variation and provides no meaningful purge of the rate-decision component for daily BEI outcomes. Figure 7 reports the results.

The press-conference results differ from the statement results. BEI yield responses are muted and imprecisely estimated, especially at longer maturities. Yields average compensation across horizons, and the press conference appears to generate offsetting movements across different parts of the curve: intermediate-horizon responses are slightly negative for some indices, while longer-horizon forward responses move in the opposite direction.

The clearest evidence appears in BEI forwards. Inflation-related press-conference indices co-move with positive responses at the long end of the forward curve. The broad inflation index and the current-inflation index display their largest positive responses around the eight-to-ten-year segment, while the corresponding BEI yield responses remain small. The yield–forward wedge is informative for interpretation: the Chair's inflation discussion does not move average inflation compensation uniformly across all horizons but loads on the part of the curve where persistence, anchoring, and long-run inflation-risk compensation are priced.

The pattern is consistent with the press conference conveying information distinct from the committee-vetted statement. The Chair's answers are longer, more conditional, and more explicit about risks, uncertainty, and the inflation outlook, so inflation language in the press conference can be priced less as an immediate reaction-function signal and more as information about the future distribution of inflation outcomes. The

**Figure 7.** Term structure of inflation compensation sensitivity to FOMC press conferences

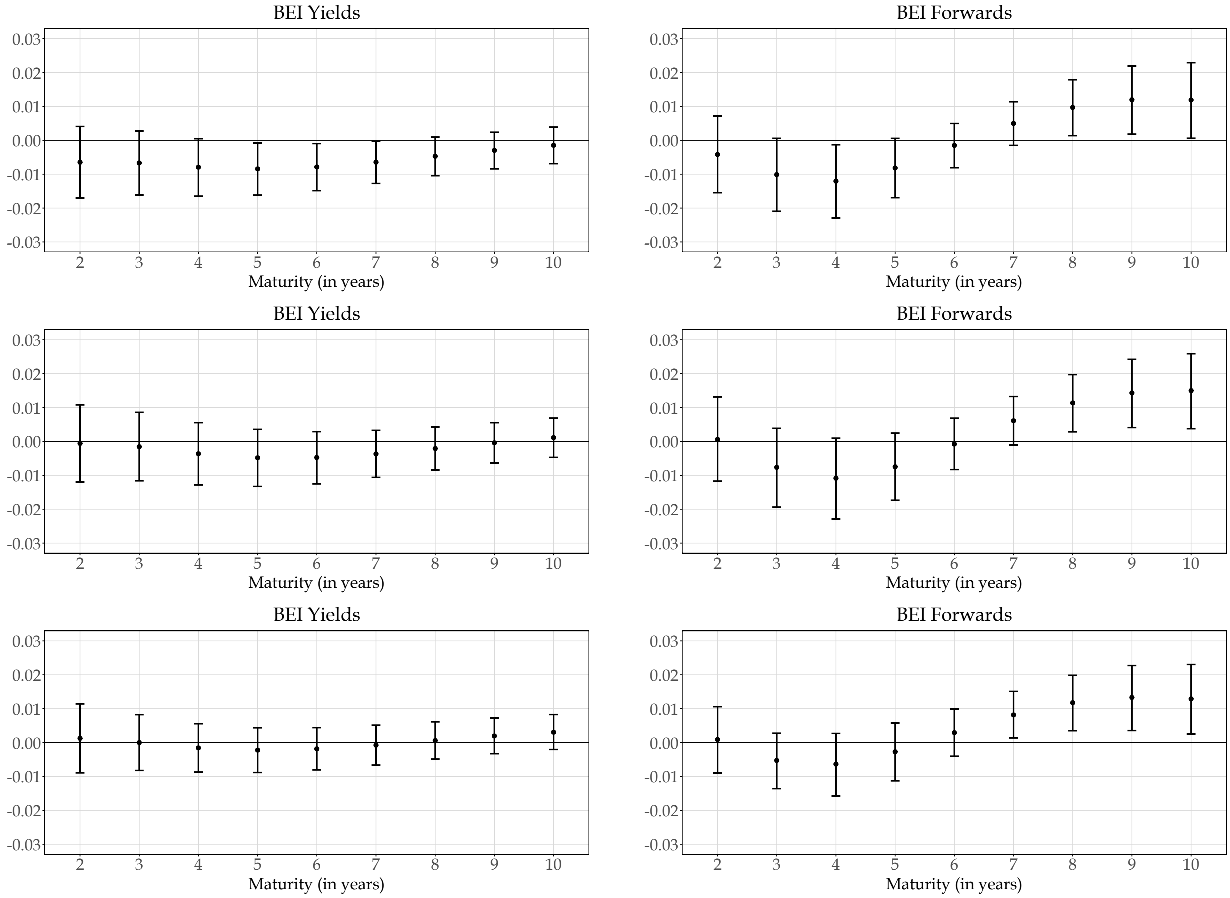


*Notes:* Event-study regressions of daily changes in BEI yields (left column) and BEI forwards (right column) on target-orthogonalized Chair press-conference indices. The top, middle, and bottom rows report responses to $mps_t$, $\pi_t^g$, and $\pi_t^{cur}$, respectively. The sample comprises $N = 88$ scheduled FOMC announcements from 2011 to 2025. Dots denote OLS estimates for maturities $\tau \in \{2, \ldots, 10\}$; error bars are 90% confidence intervals based on wild bootstrap ($B = 10{,}000$, Rademacher weights).

positive long-horizon forward response is therefore consistent with a Delphic interpretation: press-conference inflation narratives raise compensation for inflation-linked exposure at distant horizons, either through revisions in expected inflation, through changes in inflation risk premia, or through both.

The sign reversal relative to statements is the main result of the baseline analysis. Statement language is associated with lower BEI compensation, especially at short-to-intermediate horizons, while press-conference inflation language tracks higher long-horizon forward compensation. The same broad class of inflation-related language is therefore priced differently depending on the communication object. This is consistent with the view that institutional format matters: the statement bundles inflation assessments with the policy response, while the press conference allows the Chair to elaborate on risks and contingencies in a way that markets can price as inflation-outlook information.

A potential confound is the Summary of Economic Projections. Until January 2019, post-meeting press conferences were held only after meetings with an SEP release, and the Chair often discusses SEP projections during the Q&A. After January 2019, press conferences became routine following every scheduled meeting. The institutional change generates two distinct identification concerns that we cannot fully resolve with the present sample. Before 2019, the press-conference treatment bundles Chair language with SEP projections and dot-plot interpretation, and any inflation-outlook signal in the press-conference text may proxy for SEP information rather than independent narrative content. Second, the post-2019 sub-sample includes the COVID episode, the 2021–2023 inflation surge, and the associated tightening cycle, while the pre-2019 sub-sample is dominated by the zero-lower-bound era; a pooled press-conference coefficient therefore mixes substantively different communication regimes and macroeconomic states. The baseline Delphic dictionaries exclude explicit SEP labels, but generic projection and outlook language is retained. Orthogonalization with respect to $MP1_t$ removes the component of press-conference language linearly associated with the rate-decision surprise, and the broader Acosta et al. (2025) factor provides a stricter control for policy-path news; neither approach can fully purge inflation-outlook information that is economically related to the SEP, nor can it isolate the post-2019 routine press-conference regime from the pre-2019 SEP-only regime. Appendix E reports SEP-interaction specifications and a sample-restriction check excluding calendar year 2020. The qualitative press-conference pattern is preserved, which indicates that the long-horizon forward response is not driven solely by SEP meetings; the small SEP/non-SEP cell sizes preclude a sharper conclusion. A pre-2019 versus post-2019 split is a natural follow-up exercise, but with $N = 88$ press-conference observations it is at the limit of statistical power; we flag it as a priority for the next revision rather than report it here as an underpowered diagnostic.

The next section decomposes inflation communication into Delphic and Odyssean components. The decomposition clarifies whether the long-horizon press-conference response is associated with outlook language and whether policy-commitment language instead loads on the intermediate segment of the forward curve, as predicted by Proposition 1.

# 6 Delphic and Odyssean inflation narratives

This section examines whether the baseline inflation-language effects conceal different maturity profiles for outlook information and policy-commitment language. We use the two inflation-specific indices introduced in Section 4.1: the Delphic index $\delta_t^{\pi}$, which captures inflation-outlook language, and the Odyssean index $\omega_t^{\pi}$, which captures policy-commitment language conditional on inflation developments (Campbell et al., 2012; Andrade and Ferroni, 2021). The theoretical framework predicts that, when these dimensions can be priced separately, Delphic language should load more strongly on long-horizon forward compensation, whereas Odyssean language should be associated with compression in the short-to-intermediate segment of the forward curve (Proposition 1). In constrained formats such as committee-vetted statements, the same dimensions may instead be priced through a common reaction-function signal (Proposition 3).

We estimate the same maturity-by-maturity event-study specification as in Section 5. Each index is standardized and orthogonalized with respect to $MP1_t$ in the relevant event window. For press conferences, $\delta_t^\pi$ and $\omega_t^\pi$ are constructed using exposure-weighted aggregation, reflecting the longer and more discursive nature of the communication object. For statements, they are constructed using within-theme aggregation, reflecting the shorter and more formulaic structure of the post-meeting text.

**Figure 8.** Term structure of inflation compensation sensitivity to FOMC statements: Delphic and Odyssean indices

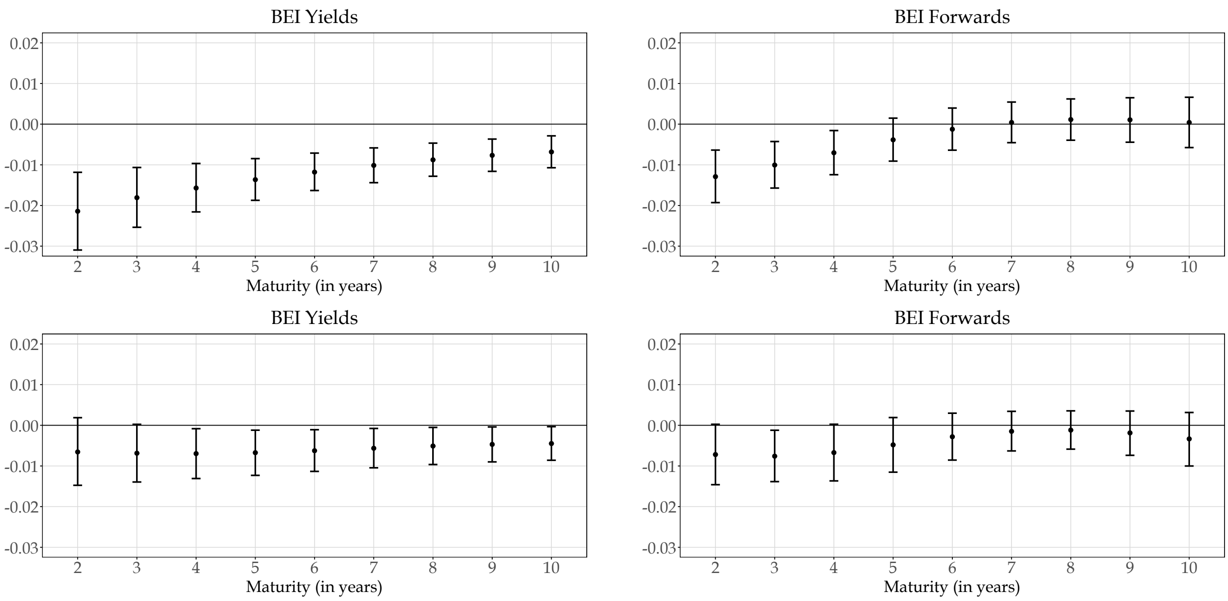


*Notes:* Event-study regressions of daily changes in BEI yields (left column) and BEI forwards (right column) on target-orthogonalized FOMC statement indices. The top row reports responses to $\delta_t^\pi$ and the bottom row to $\omega_t^\pi$. The sample comprises $N = 185$ scheduled FOMC announcements from 2004 to 2025. Dots denote OLS estimates for maturities $\tau \in \{2, \dots, 10\}$; error bars are 90% wild bootstrap confidence intervals.

Figure 8 reports the statement results. Both Delphic and Odyssean statement indices co-move with lower BEI compensation. The pattern is strongest for the Delphic index, whose yield response is negative across the maturity spectrum and whose forward response is concentrated at short-to-intermediate horizons before attenuating at the long end. The Odyssean index displays a more moderate but similarly negative profile.

The sign and maturity profile of the statement Delphic index are not consistent with an interpretation in which this measure captures a pure outlook-news shock. If statement Delphic language were priced primarily as adverse information about future inflation, one would expect positive long-horizon forward repricing. The estimates are instead negative and concentrated at maturities where markets price the expected stabilization of inflation. The pattern is consistent with the co-entailment mechanism in Proposition 3: in the post-meeting statement, inflation assessments are typically embedded in, or placed adjacent to, language about the policy response. Statement inflation language is therefore priced through the reaction function rather than as a separately identifiable Delphic signal.

This interpretation is consistent with the institutional role of the post-meeting statement. The statement is short, committee-vetted, and drafted to summarize both the Committee's inflation assessment and its policy implications. References to elevated inflation, inflation risks, or price stability therefore do not arrive as stand-alone outlook information; they arrive in a format that simultaneously communicates how the Committee intends to respond. The negative statement responses should accordingly be interpreted as evidence that statement inflation narratives are priced as part of a policy-response message, rather than as evidence that Delphic inflation language is intrinsically disinflationary.

The press-conference format is institutionally distinct. It is longer, more conditional, and allows the Chair to discuss inflation risks, uncertainty, and policy contingencies in a less compressed format. Figure 9 reports the corresponding Delphic/Odyssean profiles.

For press conferences, the two inflation narratives display more distinct maturity profiles. The Delphic index co-moves with positive BEI-forward responses at the long end of the curve, while the corresponding yield responses remain small. The yield–forward wedge is economically informative: because yields average compensation across horizons, repricing concentrated at distant maturities is more visible in forwards than in yields. The pattern is therefore consistent with Chair inflation-outlook language raising compensation at maturities where persistence, anchoring, and long-horizon inflation-risk compensation are priced, the channel anticipated in Proposition 2.

The Delphic forward profile is not uniformly increasing across all maturities. The positive response is concentrated at the long end, while short- and intermediate-horizon estimates are close to zero or modestly negative. At individual maturities, only the nine-year forward is statistically distinguishable from zero at conventional levels in Table 16; we therefore emphasize the term-structure shape rather than pointwise inference, consistent with the multiple-testing footnote in Section 5. The shape is consistent with the comparative implication of Proposition 1: the relevant prediction is not strict monotonicity at every maturity but a stronger loading of Delphic inflation language on long-horizon forwards than on near-horizon forwards when markets can price outlook information separately. The press-conference estimates satisfy this comparative implication.

The Odyssean press-conference index displays a complementary profile. Its forward

**Figure 9.** Term structure of inflation compensation sensitivity to Delphic and Odyssean inflation communication in FOMC press conferences

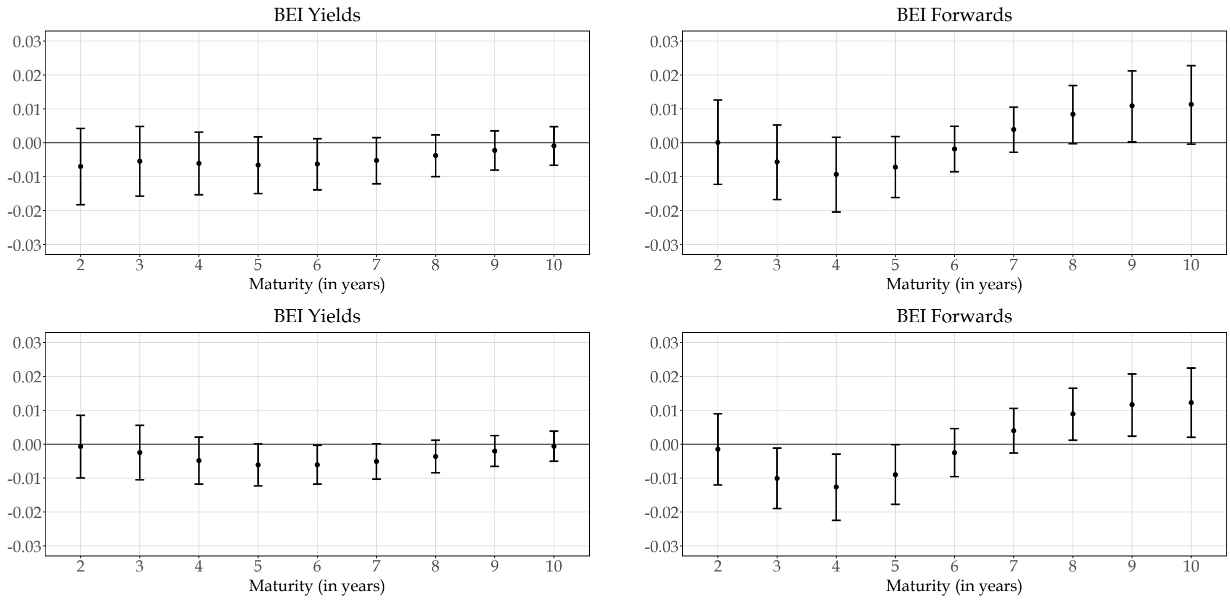


*Notes:* Event-study regressions of daily changes in BEI yields (left column) and BEI forwards (right column) on target-orthogonalized Chair press-conference indices. The top row reports responses to $\delta_t^\pi$ and the bottom row to $\omega_t^\pi$. The sample comprises $N = 88$ scheduled FOMC announcements from 2011 to 2025. Dots denote OLS estimates for maturities $\tau \in \{2, \dots, 10\}$; error bars are 90% wild bootstrap confidence intervals.

response is negative in the short-to-intermediate segment of the curve and attenuates at longer horizons, consistent with the stabilization-window mechanism: policy-commitment language is priced most strongly at maturities where markets expect policy to bring inflation back toward target. Longer-horizon estimates should be interpreted cautiously, because the smaller press-conference sample and the state-contingent nature of the Chair's communication may allow inflation-outlook information to offset part of the pure commitment component.

Long-horizon point estimates are positive for both Delphic and Odyssean press-conference indices, and statistically significant for $\omega_t^\pi$ at the nine- and ten-year maturities (1.2 bp, $p < 0.05$; Table 16, Panel B). The pattern anticipates the intraday evidence in Section 7: state-contingent commitment language appears to convey information about the inflation environment that motivates the commitment, a component visible at long horizons in the daily forwards and across all available maturities within the press-conference window.

The press-conference estimates therefore provide evidence of maturity segmentation. In the univariate profiles, Delphic inflation language tracks long-horizon forward repricing, whereas Odyssean language is associated with short-to-intermediate forward compression. This does not imply that markets observe two orthogonal communication shocks; it shows that the two text indices summarize different dimensions of the same communication event and that these dimensions covary with different segments of the BEI forward curve.

The contrast with statements is central. In statements, Delphic and Odyssean indices both load negatively on BEI, consistent with institutional co-entailment and reaction-function pricing. In press conferences, the two indices display different maturity profiles, consistent with a format in which outlook information and policy commitment can be elaborated more separately. The asymmetry provides an empirical counterpart to the Delphic/Odyssean distinction in the forward-guidance literature while preserving the reduced-form interpretation of the estimates.

The results also inform the interpretation of announcement-day information effects. The statement evidence is best read through the policy-reaction-function channel: inflation language is associated with lower inflation compensation, rather than with the higher compensation that would arise if adverse outlook news dominated the response. The press-conference Delphic profile differs in character. Its positive long-horizon forward response is consistent with the Chair conveying information about the inflation outlook, inflation persistence, or inflation risk compensation beyond what is summarized by the current rate decision. Because BEI combines expected inflation and risk premia, the evidence does not distinguish between these components; it does show that the format in which inflation language is delivered matters for how the BEI term structure is repriced.

# 7 High-frequency BEI responses

The daily evidence in Sections 5–6 measures close-to-close BEI repricing on FOMC announcement days. On days with both a post-meeting statement and a Chair press conference, the daily window aggregates information released at different points within the trading day. Following the event-window definitions in Acosta et al. (2025), the statement window is a 30-minute interval beginning 10 minutes before and ending 20 minutes after the release of the post-meeting statement, while the press-conference window is a 70-minute interval beginning 10 minutes before and ending 60 minutes after the start of the Chair's press conference. In typical scheduled FOMC meetings, the statement is released first, together with the policy decision, and the press conference begins shortly thereafter. The daily response therefore combines the immediate market reaction to the statement and rate decision, the incremental reaction during the press conference, and any additional adjustment outside the narrow high-frequency windows.

This section uses intraday BEI changes to examine the timing of these responses. The objective is not to decompose the daily effect mechanically into fully identified statement and press-conference shocks; it is to ask whether the sign and strength of the communication–BEI association differ when the dependent variable is measured within the narrow event window associated with each communication object. The exercise is especially useful for press conferences, because the current rate decision has already been announced when the Chair begins speaking. The press-conference window therefore provides a high-frequency setting for studying the incremental pricing of the Chair's narrative after the rate decision has been largely incorporated into market prices.

We use intraday BEI changes from Acosta et al. (2025), constructed as the difference between nominal Treasury and TIPS yield changes within event window $w$:

$$\Delta \mathrm{BEI}_t^w(k) \equiv \Delta y_t^{N,w}(k) - \Delta y_t^{R,w}(k). \tag{31}$$

The available maturities are $k \in \{5\mathrm{y}, 10\mathrm{y}, 5\mathrm{y}5\mathrm{y}\}$, where the 5-year/5-year forward is approximated as

$$\Delta \mathrm{BEI}_t^w(5\mathrm{y}5\mathrm{y}) = 2\Delta \mathrm{BEI}_t^w(10\mathrm{y}) - \Delta \mathrm{BEI}_t^w(5\mathrm{y}). \tag{32}$$

The approximation is exact under continuous compounding and equal-weight averaging across the 5–10y segment, but it amplifies measurement error in the 5- and 10-year inputs by construction: under independent zero-mean quote noise with equal variances, the variance of the approximated 5y5y change is roughly five times the variance of the underlying 5- and 10-year changes. We therefore report 5y5y intraday estimates alongside the 5y and 10y estimates rather than as a stand-alone object, and we read the qualitative pattern of signs across the three maturities rather than the magnitude of the 5y5y coefficient in isolation. The event-study specification mirrors Equation (30), with the daily outcome replaced by the corresponding intraday change. The communication index is orthogonalized with respect to $MP1_t$ measured in the same event window as the dependent variable, as described in Section 4.2. The matched samples contain $N = 180$ statement events and $N = 89$ press conferences.[10]

Table 3 reports the intraday statement-window results. The estimates are small and statistically indistinguishable from zero for almost all indices and maturities. The only marginal exception is the 5-year response to $mps_t$, which is positive at the 10% level. The text-based statement indices therefore do not explain a detectable component of BEI repricing within the narrow statement-release window once the contemporaneous target-rate surprise is accounted for. The marginal positive 5-year $mps_t$ coefficient does not contradict Proposition 4, which states that statement-related communication contributes to the daily response without necessarily generating a detectable text effect within the narrow statement window; the broader-session interpretation of the negative daily statement coefficient is therefore consistent with both a small early micro-response

[10] The matched samples differ from the daily samples because some FOMC dates lack matched high-frequency Treasury and TIPS quotes in the relevant 30- or 70-minute intraday window. The intraday statement sample is five observations smaller than the daily statement sample, reflecting these missing intraday quotes. The intraday press-conference sample is one observation larger than the daily press-conference sample because one FOMC date has valid intraday Treasury and TIPS observations but cannot be matched to the daily BEI closing data used in the baseline daily regressions.

and the more substantial reaction-function pricing observed over the remainder of the announcement day.

This result does not contradict the negative daily statement responses documented in Section 5; it clarifies their timing. The daily effects appear to be broader-session responses to the statement and the associated policy narrative rather than immediate high-frequency reactions fully realized within the short statement window. The interpretation is consistent with the institutional nature of the statement: the policy decision and the statement arrive simultaneously, and the immediate price reaction is dominated by the rate-decision and policy-news bundle. The additional textual component associated with inflation language may require more time to be incorporated into BEI compensation, especially when the object of interest is a fitted Treasury–TIPS term-structure measure rather than a highly liquid short-rate contract.

**Table 3.** FOMC statements: high-frequency BEI responses

| | 5y | 10y | 5y5y |
|---|---|---|---|
| $mps_t$ | 0.28* | 0.13 | −0.02 |
| | (0.16) | (0.17) | (0.25) |
| $R^2$ | 0.02 | 0.01 | 0.00 |
| $\pi_t^g$ | 0.06 | −0.02 | −0.09 |
| | (0.16) | (0.13) | (0.16) |
| $R^2$ | 0.00 | 0.00 | 0.00 |
| $\pi_t^{cur}$ | 0.07 | −0.01 | −0.08 |
| | (0.15) | (0.15) | (0.20) |
| $R^2$ | 0.00 | 0.00 | 0.00 |
| $\delta_t^\pi$ | 0.21 | 0.11 | 0.02 |
| | (0.16) | (0.19) | (0.26) |
| $R^2$ | 0.01 | 0.00 | 0.00 |
| $\omega_t^\pi$ | 0.05 | 0.02 | −0.01 |
| | (0.17) | (0.15) | (0.19) |
| $R^2$ | 0.00 | 0.00 | 0.00 |

*Notes:* Event-study regressions of intraday BEI changes on target-orthogonalized, standardized FOMC statement indices. Wild bootstrap standard errors ($B = 10{,}000$, Rademacher weights) are in parentheses. Coefficients and standard errors are in basis points. $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Sample: $N = 180$.

Table 4 reports the corresponding press-conference-window estimates. The contrast with the statement window is pronounced. All five communication indices are positively associated with intraday BEI changes at the 5-year, 10-year, and 5-year/5-year maturities, and the estimates are more precisely measured than in the statement-window regressions. At the 10-year maturity, the coefficients range from 0.59 basis points for $\pi_t^{cur}$ and $\delta_t^\pi$ to 0.74 basis points for $mps_t$. At the 5-year/5-year forward, the estimates range from 0.43 to 0.55 basis points.

The sign of the press-conference-window response is informative. In the daily regressions, press-conference inflation language exhibits a more heterogeneous maturity profile, with positive long-horizon forward responses and some negative intermediate-horizon responses for Odyssean language. Within the narrow press-conference window, by contrast, the dominant association is positive across all available maturities and indices. The pattern indicates that the immediate market reaction to the Chair's narrative is primarily an inflation-state or inflation-risk response. Once the rate decision has already been priced, the Chair's discussion of inflation, risks, uncertainty, and policy contingencies co-moves with higher compensation for inflation exposure. Monetary-policy uncertainty itself, as measured by Husted et al. (2020), is a potentially omitted state variable that could simultaneously raise the intensity of FOMC inflation language and move inflation-compensation pricing; our coefficients are therefore interpreted as conditional on the rate-decision surprise, not on the full uncertainty state.

The positive coefficient on $\omega_t^{\pi}$ is especially useful for interpretation. Odyssean language is designed to capture policy-commitment content, and in the daily press-conference regressions it co-moves with short-to-intermediate forward compression. In the narrow press-conference window, however, the same textual dimension is positively associated with BEI at all available maturities. This does not imply that commitment language is inflationary; it indicates that state-contingent commitment language also conveys information about the inflation environment that motivates the commitment. At high frequency, the state-information component appears to dominate; over the daily window, the commitment and reaction-function components become more visible in the intermediate segment of the forward curve.

**Table 4.** FOMC press conferences: high-frequency BEI responses

| | 5y | 10y | 5y5y |
|---|---|---|---|
| $mps_t$ | 0.95** | 0.74*** | 0.53*** |
| | (0.40) | (0.24) | (0.20) |
| $R^2$ | 0.10 | 0.14 | 0.08 |
| $\pi_t^g$ | 0.88*** | 0.71*** | 0.55*** |
| | (0.34) | (0.22) | (0.21) |
| $R^2$ | 0.08 | 0.13 | 0.08 |
| $\pi_t^{cur}$ | 0.70* | 0.59** | 0.48** |
| | (0.41) | (0.24) | (0.22) |
| $R^2$ | 0.05 | 0.09 | 0.06 |
| $\delta_t^{\pi}$ | 0.75** | 0.59** | 0.43* |
| | (0.38) | (0.26) | (0.22) |
| $R^2$ | 0.06 | 0.09 | 0.05 |
| $\omega_t^{\pi}$ | 0.76*** | 0.62*** | 0.47*** |
| | (0.28) | (0.19) | (0.19) |
| $R^2$ | 0.06 | 0.10 | 0.06 |

*Notes:* Event-study regressions of intraday BEI changes on target-orthogonalized, standardized Chair press-conference indices. Wild bootstrap standard errors ($B = 10{,}000$, Rademacher weights) are in parentheses. Coefficients and standard errors are in basis points. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Sample: $N = 89$.

The intraday evidence refines the interpretation of the daily results in two ways. First, the negative daily statement effects are not visible as immediate text-driven

BEI movements in the narrow statement window; they appear instead to be incorporated over the broader announcement day, after the simultaneous policy decision and committee-vetted statement have been processed. Second, the press-conference window contains a positive communication component that is visible across all available BEI maturities and across all text indices, including the Odyssean index.

These findings support the broader timing interpretation in Proposition 4, with one qualification. The evidence does not imply an exact additive decomposition of the daily response into a negative statement-window component and a positive press-conference component; the daily window includes trading outside the two narrow windows, possible delayed adjustment, and other same-day information. What the intraday estimates show is more specific: the Chair's press conference contains a positive high-frequency inflation-compensation component, whereas the text-based statement component is not separately detectable in the narrow statement-release window after conditioning on the target-rate surprise.

The distinction matters for the interpretation of the main results. Daily BEI responses summarize the net repricing of inflation compensation over the announcement day; intraday responses identify when part of that repricing occurs. The press-conference evidence indicates that the Chair's narrative is priced immediately as information about the inflation state or inflation-risk environment. The statement evidence indicates that the negative reaction-function interpretation of statement language is a broader announcement-day phenomenon rather than a narrow-window text effect. The combination reinforces the central conclusion: the pricing of inflation language depends not only on what is said but also on the institutional format and timing through which the information is released.

# 8 Policy implications

The findings have implications for the interpretation of central bank communication that we state conditionally, given the identification caveats discussed in Sections 3 and 5.[11]

Statements and press conferences exhibit distinct patterns of BEI repricing within the FOMC communication toolkit. Statement inflation language is associated with lower BEI compensation, particularly at short-to-intermediate maturities, in a manner consistent with markets pricing the committee-vetted text through the policy reaction function. Press-conference inflation language tracks positive high-frequency BEI repricing and positive long-horizon forward responses in the daily regressions. The same broad class of inflation language is therefore priced differently across the two formats,

[11]The communication indices may still capture path news, SEP-related content, reaction-function information, monetary-policy uncertainty, or other components of the FOMC information set not removed by $MP1_t$ orthogonalization. The interpretations in this section are therefore suggestive associations rather than prescriptive guidance.

conditional on the rate-decision surprise.

The Delphic/Odyssean distinction maps into different maturity profiles in market-based inflation compensation. In press conferences, Delphic inflation language loads positively on long-horizon forward repricing, while Odyssean language is associated with negative responses in the short-to-intermediate segment of the forward curve. The maturity at which markets reprice inflation compensation therefore depends on whether communication is read as outlook information, as policy commitment, or as a combination of both.

We do not infer prescriptive guidance from these associations. Establishing that a given calibration of inflation language would change inflation expectations or anchoring at policy-relevant horizons would require a structural decomposition of BEI into expected inflation and risk premia, external validation of the text indices, and an identification strategy that isolates exogenous communication variation. We discuss these as priorities for future work in Section 9.

The methodology nonetheless provides a replicable framework for monitoring how markets price central bank communication over time. Changes in the sign, magnitude, or maturity profile of the estimated responses may signal shifts in how markets interpret the central bank's reaction function, inflation outlook, or credibility. While the analysis focuses on the Federal Reserve, the framework extends to other institutional settings that combine committee-vetted written statements with subsequent oral explanation.

# 9 Conclusion

This paper studies how inflation-related FOMC communication is priced into the maturity structure of market-based inflation compensation. The main finding is that the institutional format through which a given inflation message is delivered conditions how breakeven inflation reprices around that message. In post-meeting statements, inflation-related language is associated with lower BEI compensation, especially at short-to-intermediate horizons, in a pattern consistent with markets reading the committee-vetted text through the policy reaction function: inflation assessments and policy-response language are bundled in a format that links outlook information to stabilization intent. In Chair press conferences, by contrast, inflation language tracks positive long-horizon forward repricing, and the intraday evidence shows a positive BEI response within the press-conference window itself.

The Delphic/Odyssean decomposition sharpens the interpretation. In statements, both Delphic and Odyssean inflation indices load negatively on BEI compensation, consistent with institutional co-entailment. In press conferences, the two indices display different maturity profiles: Delphic language loads positively on long-horizon forwards, whereas Odyssean language is associated with short-to-intermediate forward compression. The maturity segmentation provides a term-structure application of the Delphic/Odyssean distinction in the forward-guidance literature and indicates that communication format matters for the pricing of inflation compensation.

These results should be read as reduced-form event-study associations, not as structurally identified causal effects of isolated communication channels. Orthogonalization with respect to the current rate-decision surprise removes only the linear component of the text indices associated with $MP1_t$; it does not purge path news, SEP-related information, reaction-function learning, or other components of the FOMC information-effect bundle. Several additional caveats bear emphasis.

Inference is limited along two dimensions. High inter-index correlations (up to $\rho \approx 0.89$–$0.90$ within corpus) preclude a credible multivariate decomposition, so coefficient comparisons across indices should be read as descriptive summaries of univariate associations rather than as evidence of differential causal channels. The press-conference sample contains $N = 88$ scheduled meetings, so pointwise inference at any single maturity must be interpreted alongside the term-structure shape of the estimates, and the wild bootstrap addresses heteroskedasticity but not the family-wise error rate across the maturity spectrum. A formal joint-inference exercise (max-$t$ confidence bands, Romano–Wolf stepdown, or functional bootstrap tests of shape restrictions) is a natural extension. A related point applies to the statement coefficient. Its negative daily value is not visible in the narrow statement-window intraday regressions once $MP1_t$ is partialled out, so we interpret it as a broader-session phenomenon rather than as a high-frequency reaction-function effect, with corresponding limits on what can be inferred about the channel through which statement text moves daily BEI.

Sample composition introduces two further concerns. The press-conference sample spans an institutional regime change in early 2019, mixes SEP-only meetings (pre-2019) with routine press conferences (post-2019), and includes the 2021–2023 inflation surge. A pre/post-2019 split, leave-one-year-out diagnostics, labeled influence checks, and Chair-specific subsamples (Bernanke, Yellen, Powell) would discipline the press-conference results further but require more statistical power than the present sample affords. The text indices themselves have not been validated against an independent hand-coded benchmark; while time-series patterns and event-study co-movements support face validity, a stratified human-coded validation exercise reporting inter-coder agreement, confusion matrices, and precision-recall metrics for Delphic and Odyssean labels, following the template of Loughran and McDonald (2011) and Baker et al. (2016), is a priority for revision, as is external validation of the Delphic index against SEP inflation revisions and of the Odyssean index against path-surprise factors.

Measurement granularity bounds what the present design can identify. The intraday analysis decomposes the announcement day into two windows (the full announcement day and the press-conference window of Acosta et al. (2025)) rather than into a finer four-window partition isolating statement-release, interim, press-conference, and post-conference segments; the latter would sharpen the within-day attribution but requires high-frequency BEI quotes for windows outside the present dataset. The dependent variable is $\Delta \text{BEI}_t(\tau)$ throughout. Estimating the same specifications on the underlying nominal-yield change $\Delta y_t^N(\tau)$ and real-yield change $\Delta y_t^R(\tau)$ would distinguish the nominal- and real-yield channels and is a natural extension. The text indices are aggregated at the document level; a sentence-level adjacency indicator separating inflation language co-entailed with policy-response language from inflation language that is not would directly test the measurement-side counterpart of Proposition 3. The consistency of sign, maturity, and timing patterns across statements, press conferences, daily responses, and intraday windows indicates that inflation narratives contain market-relevant information beyond the contemporaneous rate decision; the caveats above describe the boundary of what the present design identifies.

Two extensions would further sharpen the analysis. First, decomposing BEI responses into model-implied expected inflation, inflation risk premia, and liquidity components using a term-structure framework (Abrahams et al., 2016; Haubrich et al., 2012; D'Amico et al., 2018) would clarify whether the Delphic press-conference effect reflects revisions to expected inflation, higher compensation for inflation risk, or both. Second, embedding the communication indices in a richer identification framework that separates target, path, and information components of FOMC announcements (Bu et al., 2021) would move the interpretation closer to structural channels. The implication of the present analysis is that the effect of inflation communication on market-based inflation compensation depends not only on what the central bank says, but also on the institutional format and the maturity at which the information is priced.

# References


Abrahams, M., T. Adrian, R. K. Crump, E. Moench, and R. Yu (2016). Decomposing real and nominal yield curves. *Journal of Monetary Economics 84*, 182–200.

Acosta, M., A. Ajello, M. D. Bauer, F. Loria, and S. Miranda-Agrippino (2025). Financial market effects of FOMC communication: Evidence from a new event-study database. Working Paper 2025-30, Federal Reserve Bank of San Francisco.

Ahrens, M., D. Erdemlioglu, M. McMahon, C. J. Neely, and X. Yang (2025). Mind your language: Market responses to central bank speeches. *Journal of Econometrics 249*, 105921.

Altavilla, C., R. S. Gürkaynak, L. Laeven, and T. Kind (2025). Monetary transmission with frequent policy events. Working Paper Series 3157, European Central Bank.

Andrade, P. and F. Ferroni (2021). Delphic and odyssean monetary policy shocks: Evidence from the Euro Area. *Journal of Monetary Economics 117*, 816–832.

Andrade, P., G. Gaballo, E. Mengus, and B. Mojon (2019). Forward guidance and heterogeneous beliefs. *American Economic Journal: Macroeconomics 11*(3), 1–29.

Baker, S. R., N. Bloom, and S. J. Davis (2016). Measuring economic policy uncertainty. *Quarterly Journal of Economics 131*(4), 1593–1636.

Bauer, M. D. and E. T. Swanson (2023). An alternative explanation for the "Fed information effect". *American Economic Review 113*(3), 664–700.

Bernanke, B. S. and K. N. Kuttner (2005). What explains the stock market's reaction to Federal Reserve policy? *Journal of Finance 60*(3), 1221–1257.

Binder, C. (2017). Fed speak on main street: Central bank communication and household expectations. *Journal of Macroeconomics 52*, 238–251.

Blinder, A. S., M. Ehrmann, M. Fratzscher, J. De Haan, and D.-J. Jansen (2008). Central bank communication and monetary policy: A survey of theory and evidence. *Journal of Economic Literature 46*(4), 910–945.

Boguth, O., V. Grégoire, and C. Martineau (2019). Shaping expectations and coordinating attention: The unintended consequences of FOMC press conferences. *Journal of Financial and Quantitative Analysis 54*(6), 2327–2353.

Bordo, M. D. and K. Istrefi (2023). Perceived FOMC: The making of hawks, doves and swingers. *Journal of Monetary Economics 136*, 125–143.

Bu, C., J. Rogers, and W. Wu (2021). A unified measure of Fed monetary policy shocks. *Journal of Monetary Economics 118*, 331–349.

Cameron, A. C., J. B. Gelbach, and D. L. Miller (2008). Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics 90*(3), 414–427.

Campbell, J. R., C. L. Evans, J. D. Fisher, and A. Justiniano (2012). Macroeconomic effects of Federal Reserve forward guidance. *Brookings Papers on Economic Activity 43*(1 (Spring)), 1–80.

Chadha, J. S., C. Macchiarelli, S. Goel, A. Hantzsche, and S. Mellina (2026). Deciphering Delphic guidance: The Bank of England and geopolitical uncertainty. *Journal of International Money and Finance 162*, 103530.

Cieslak, A., S. Hansen, M. McMahon, and S. Xiao (2023). Policymakers' uncertainty. NBER Working Papers 31849, National Bureau of Economic Research.

Cieslak, A. and M. McMahon (2024). Tough talk: The Fed and the risk premium. CEPR Discussion Papers 19313, Centre for Economic Policy Research.

Cieslak, A. and A. Schrimpf (2019). Non-monetary news in central bank communication. *Journal of International Economics 118*, 293–315.

Coibion, O., Y. Gorodnichenko, and M. Weber (2022). Monetary policy communications and their effects on household inflation expectations. *Journal of Political Economy 130*(6), 1537–1584.

D'Amico, S., D. H. Kim, and M. Wei (2018). TIPS from TIPS: The informational content of treasury inflation-protected security prices. *Journal of Financial and Quantitative Analysis 53*(1), 395–436.

Ehrmann, M. and M. Fratzscher (2007). Communication by central bank committee members: Different strategies, same effectiveness? *Journal of Money, Credit and Banking 39*(2-3), 509–541.

Fleckenstein, M., F. A. Longstaff, and H. Lustig (2014). The tips-treasury bond puzzle. *Journal of Finance 69*(5), 2151–2197.

Gentzkow, M., B. Kelly, and M. Taddy (2019). Text as data. *Journal of Economic Literature 57*(3), 535–574.

Gorodnichenko, Y., T. Pham, and O. Talavera (2023). The voice of monetary policy. *American Economic Review 113*(2), 548–584.

Gürkaynak, R. S., A. Levin, and E. Swanson (2010a). Does inflation targeting anchor long-run inflation expectations? Evidence from the U.S., UK, and Sweden. *Journal of the European Economic Association 8*(6), 1208–1242.

Gürkaynak, R. S., B. Sack, and E. Swanson (2005). Do actions speak louder than words? The response of asset prices to monetary policy actions and statements. *International Journal of Central Banking 1*(1), 55–93.

Gürkaynak, R. S., B. Sack, and J. H. Wright (2007). The U.S. Treasury yield curve: 1961 to the present. *Journal of Monetary Economics 54*(8), 2291–2304.

Gürkaynak, R. S., B. Sack, and J. H. Wright (2010b). The TIPS yield curve and inflation compensation. *American Economic Journal: Macroeconomics 2*(1), 70–92.

Hansen, S. and M. McMahon (2016). Shocking language: Understanding the macroeconomic effects of central bank communication. *Journal of International Economics 99*(S1), 114–133.

Hansen, S., M. McMahon, and M. Tong (2019). The long-run information effect of central bank communication. *Journal of Monetary Economics 108*, 185–202.

Hanson, S. G. and J. C. Stein (2015). Monetary policy and long-term real rates. *Journal of Financial Economics 115*(3), 429–448.

Haubrich, J. G., G. Pennacchi, and P. Ritchken (2012). Inflation expectations, real rates, and risk premia: Evidence from inflation swaps. *Review of Financial Studies 25*(5), 1588–1629.

Husted, L., J. Rogers, and B. Sun (2020). Monetary policy uncertainty. *Journal of Monetary Economics 115*, 20–36.

Jarociński, M. and P. Karadi (2020). Deconstructing monetary policy surprises: The role of information shocks. *American Economic Journal: Macroeconomics 12*(2), 1–43.

Kuttner, K. N. (2001). Monetary policy surprises and interest rates: Evidence from the Fed funds futures market. *Journal of Monetary Economics 47*(3), 523–544.

Lakdawala, A. and M. Schaffer (2019). Federal Reserve private information and the stock market. *Journal of Banking & Finance 106*, 34–49.

Lamla, M. J. and D. V. Vinogradov (2019). Central bank announcements: Big news for little people? *Journal of Monetary Economics 108*, 21–38.

Loughran, T. and B. McDonald (2011). When is a liability not a liability? Textual analysis, dictionaries, and 10-ks. *Journal of Finance 66*(1), 35–65.

Lucca, D. O. and F. Trebbi (2009). Measuring central bank communication: An automated approach with application to FOMC statements. NBER Working Papers 15367, National Bureau of Economic Research.

Melosi, L. (2017). Signalling effects of monetary policy. *Review of Economic Studies 84*(2), 853–884.

Miranda-Agrippino, S. and G. Ricco (2021). The transmission of monetary policy shocks. *American Economic Journal: Macroeconomics 13*(3), 74–107.

Nakamura, E. and J. Steinsson (2018). High-frequency identification of monetary non-neutrality: The information effect. *Quarterly Journal of Economics 133*(3), 1283–1330.

Romer, C. D. and D. H. Romer (2000). Federal Reserve information and the behavior of interest rates. *American Economic Review 90*(3), 429–457.

Shah, A., S. Paturi, and S. Chava (2023, 7). Trillion dollar words: A new financial dataset, task & market analysis. In A. Rogers, J. Boyd-Graber, and N. Okazaki (Eds.), *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Toronto, Canada, pp. 6664–6679. Association for Computational Linguistics.

Swanson, E. T. (2021). Measuring the effects of Federal Reserve forward guidance and asset purchases on financial markets. *Journal of Monetary Economics 118*, 32–53.

Tetlock, P. C. (2007). Giving content to investor sentiment: The role of media in the stock market. *Journal of Finance 62*(3), 1139–1168.

Woodford, M. (2005). Central bank communication and policy effectiveness. In *Proceedings - Economic Policy Symposium - Jackson Hole*, pp. 399–474.

# A NLP implementation details

This appendix documents the text-measurement layer used to construct the CBC inflation indices. The strategy combines a transformer-based stance classifier with transparent regex-based dictionaries: the transformer provides contextual directional classification at the textual-unit level, while the dictionaries discipline *which* units belong to each inflation narrative in an interpretable and auditable way.

**Transformer-based stance model.** For each text unit $i$ in document $t$, FOMC-RoBERTa (Shah et al., 2023) returns model-implied class probabilities $(p_{it}^H, p_{it}^D, p_{it}^N)$ over hawkish, dovish, and neutral classes, from which we construct the signed stance score

$$m_{it} \equiv p_{it}^H - p_{it}^D. \tag{33}$$

Positive values correspond to relatively hawkish language; negative values to relatively dovish language. We retain $\nu_{it} \equiv 1 - p_{it}^N$ as a diagnostic measure of directional informativeness, but the baseline aggregation does not impose classifier-confidence weighting. Document-level indices are constructed using deterministic salience weights tied to observable features of the communication object: explicit policy-language upweighting in both corpora, and additional question downweighting in press conferences. The choice keeps the baseline transparent and avoids treating $1 - p_{it}^N$ as a calibrated certainty measure.

**Text units and corpus-specific segmentation.** The two corpora differ materially in format.

For *FOMC statements*, the basic unit is the sentence. Statements are first cleaned of implementation boilerplate (voting records, operational sections), whitespace is normalized, and the text is split at sentence boundaries. Because statements are short and tightly structured, the inflation anchor and the policy-guidance language may appear in adjacent sentences; the pipeline therefore allows a small within-document window when applying inflation-specific gates.

For *Chair press conferences*, the basic unit is a turn-aware chunk. The transcript is parsed into question and answer turns using explicit `Q.`/`A.` markers or speaker labels; long turns are then subdivided into overlapping chunks, preserving Q/A boundaries. Turn-aware chunking is preferable to a naive fixed-window split because it prevents a journalist's question and the Chair's answer from being merged into the same unit. In the baseline implementation, question chunks are downweighted relative to answer chunks, reflecting the identification object of interest.

The replication archive reports document-level descriptive statistics for each corpus, including the average number of text units per document, the distribution of inflation-dictionary hits across the four narrative themes, and the number of nonzero events per index per corpus. These descriptive statistics support the validity-of-construction discussion in Appendix C.

We use the official Federal Reserve transcripts of the Chair's press conferences. These are released shortly after each press conference and incorporate post-event editorial corrections that may not have been available in real-time market processing. In principle, this introduces a small wedge between the textual content as scored by our pipeline and the textual content as parsed by markets in the live event window. We treat this wedge as a measurement-error component absorbed by the residual term in our event-study regressions; constructing real-time text indices from contemporaneous wire-service feeds is a refinement left for future work.

The deterministic salience weights are fixed before the event-study analysis. Units containing explicit policy-instrument language receive a multiplier of 1.15, and units containing balance-sheet or QE/QT language receive an additional multiplier of 1.10. In press conferences, Chair answer chunks receive unit weight, journalist question chunks receive weight 0.35, and mixed chunks receive weight 0.70. These weights reflect observable features of the communication object; they are not estimated from market outcomes.

**Dictionary scoring.** For each theme $k$, let $h_{it}^{(k)}$ denote the number of dictionary patterns matched by unit $i$ in document $t$. The bounded salience score is

$$g_{it}^{(k)} \equiv 1 - \exp\Big(-\frac{h_{it}^{(k)}}{\max(n_{it}/60, 0.5)}\Big), \tag{34}$$

where $n_{it}$ is the unit word count. The denominator scales the match count by unit length, while the lower bound prevents extremely short units from generating unbounded salience. The exponential transformation maps the score into $[0, 1]$ and gives additional matches diminishing marginal influence.

**Narrative gates.** The dictionaries define *candidate* topical mass; corpus-specific gating rules enforce the conceptual separation between Delphic and Odyssean language. A unit is classified as *Odyssean* only when it combines inflation-relevant language with policy linkage and reaction-function content (instrument references, action verbs, conditionality, or strong-guidance templates such as *need to see*, *greater confidence*, or *not confident yet*). A unit is classified as *Delphic* only when it contains inflation-relevant language together with outlook, risk, expectation, or component language, and is not simultaneously classified as a true Odyssean unit. The ordering prevents reaction-function language from mechanically contaminating the Delphic measure. It does not impose co-entailment in the textual classification: even when Delphic and Odyssean units are classified separately, the statement format may lead markets to price them through a common reaction-function signal. Co-entailment is therefore an interpretation of market pricing in constrained communication formats, not a coding rule.

**Interpretation.** The dictionaries do not assign direction by themselves; they identify *which* units belong to each narrative. Direction comes from the transformer stance score $m_{it}$, which is combined with the thematic salience scores to form document-level indices. The division of labor is central to the design: the dictionaries preserve interpretability and auditability, while the transformer captures local context, polarity, and the semantic direction of policy language.

# B Worked examples

To complement the aggregate validation in Appendix C (§C.2), Tables 5 and 6 report representative textual units drawn from the pipeline output, together with the FOMC-RoBERTa stance score and the narrative classification assigned by the dictionary gates. The selection covers four cells of the Delphic/Odyssean $\times$ hawkish/dovish matrix for each corpus, plus one example of a text unit that is *excluded* by the gating rules despite containing inflation-relevant vocabulary. The format follows the validation approach of Cieslak et al. (2023), who illustrate their hawk-dove textual measures with annotated examples from FOMC transcripts, and of Bordo and Istrefi (2023), who classify FOMC members as hawks, doves, and swingers using narrative records from media archives, validating the measure against voting dissents, preferred rates, and macroeconomic forecasts.

**Table 5.** Worked examples: post-meeting statements

| Date | Stance | Sentence | Classification |
|---|---|---|---|
| *Panel A: Delphic inflation narrative ($\delta_t^\pi$)* | | | |
| Mar 2005 | Hawkish (+1.00) | "Though longer-term inflation expectations remain well contained, pressures on inflation have picked up in recent months and pricing power is more evident." | Inflation anchor + outlook/risk. Hawkish: rising pressures flagged, no policy action committed. |
| Sep 2020 | Dovish (−1.00) | "The ongoing public health crisis will continue to weigh on economic activity, employment, and inflation in the near term, and poses considerable risks to the economic outlook over the medium term." | Inflation anchor + risk/outlook. Dovish: downside risks to inflation outlook dominate. |
| *Panel B: Odyssean inflation narrative ($\omega_t^\pi$)* | | | |
| Jan 2024 | Hawkish (+0.88) | "The Committee does not expect it will be appropriate to reduce the target range until it has gained greater confidence that inflation is moving sustainably toward 2 percent." | Inflation anchor + strong-guidance template (*greater confidence* + *until*). Hawkish: Committee signals resistance to easing. |
| Dec 2021 | Dovish (−0.67) | "With inflation having exceeded 2 percent for some time, the Committee expects it will be appropriate to maintain this target range until labor market conditions have reached levels consistent with the Committee's assessments of maximum employment." | Inflation anchor + conditionality (*appropriate to maintain ... until*). Dovish: conditional on labor-market progress, no tightening. |
| *Panel C: Excluded by gating (false positive avoided)* | | | |
| Jun 2009 | n/a | "However, substantial resource slack is likely to dampen cost pressures, and the Committee expects that inflation will remain subdued for some time." | Raw $\delta^\pi$ score positive; gated score = 0.00. Inflation anchor and forecast-like language are present, but the phrase is treated by the gate as insufficiently distinct from boilerplate forward-guidance language to qualify as a standalone Delphic inflation-outlook signal. |

*Notes:* "Stance" is the FOMC-RoBERTa signed score $m_{it} = p_{it}^H - p_{it}^D$, where +1.00 (−1.00) indicates the model assigns probability 1.00 to hawkish (dovish). The classification column states which dictionary gate fires (+ denotes co-occurrence requirement) and the economic interpretation of the directional score.

**Table 6.** Worked examples: Chair press conferences

| Date | Stance | Chunk (Chair answer turn) | Classification |
|---|---|---|---|
| *Panel A: Delphic inflation narrative ($\delta_t^\pi$)* | | | |
| May 2022 | Hawkish (+1.00) | "[ . . . ] we can't allow inflation expectations to become unanchored. It's just something that we can't allow to happen, and so we'll look at it that way." | Inflation-expectations anchor + risk language. Hawkish: Chair frames unanchored expectations as an outcome to prevent, signalling vigilance. |
| Apr 2011 | Dovish (−0.94) | "[ . . . ] the Committee continues to project inflation to return to mandate-consistent levels in the medium term [ . . . ] the short-term increase in inflation has not prompted the Committee to tighten policy at this juncture." | Inflation anchor + projection language. Dovish: Chair projects normalization without urgency. |
| *Panel B: Odyssean inflation narrative ($\omega_t^\pi$)* | | | |
| Jun 2021 | Hawkish (+1.00) | "And if we see inflation moving [ . . . ] to a level, or persistently enough, you know, we would be prepared to use our tools to address that." | Inflation anchor + commitment (*prepared to use our tools*). Hawkish: action conditioned on inflation persistence. |
| Jun 2020 | Dovish (−0.99) | "We are committed to using our full range of tools to support the economy and to help assure that the recovery from this difficult period will be as robust as possible." | Commitment language with nearby inflation anchor in the same local press-conference window. Dovish: the Chair signals continued accommodation rather than restraint. |
| *Panel C: Excluded by gating (question chunk)* | | | |
| Apr 2011 | n/a | "Given what you know about the pace of the economy now, what is your best guess for how soon the Fed needs to begin to withdraw its extraordinary stimulus?" | Journalist's question turn. Segment type flagged as `question`; weight = 0.35 applied. The segment is retained with reduced weight so that journalist framing receives less influence than the Chair's own language in document-level aggregation. |

*Notes:* Chunks are contiguous Chair answer-turn passages of 2–5 sentences assigned the same narrative gate. Stance scores as in Table 5. Square brackets indicate abbreviated text; full passages are available in the replication archive.

# C Lexical dictionaries and measure validation

This appendix reports the lexical dictionaries used to identify inflation narratives in FOMC post-meeting statements and Chair press conferences, together with a validation of the resulting measures.

## C.1 Dictionaries

The dictionaries underpin the construction of the CBC inflation indicators $\pi_t^g$ (broad inflation narrative; Table 7), $\pi_t^{cur}$ (current-inflation narrative; Table 8), $\delta_t^\pi$ (Delphic inflation narrative; Table 9), and $\omega_t^\pi$ (Odyssean inflation narrative; Table 10). The full dictionaries are provided in machine-readable form in the replication archive (regex-by-regex, with case-insensitivity flags, stemming and pluralization rules, hyphenation variants, and tokenization parameters explicitly recorded), so that the formatted tables in this appendix can be exactly reproduced from the underlying patterns. The replication archive additionally contains the raw document list, the FOMC-event-date matching file, the text-unit segmentation code, the FOMC-RoBERTa scoring routine, the constructed indices, the BEI outcome dataset description with sources and vintages, and all scripts producing the tables and figures in the main text and appendices.

**Step 1: Candidate terms from the corpus.** For each corpus (statements; press conferences), we compute term and phrase frequencies on the original documents to obtain a candidate vocabulary that is empirically relevant in the FOMC context. We then curate the candidates into four thematic lists aligned with the inflation narratives of interest.

**Step 2: Pattern matching via regular expressions.** Each lexical item is matched using regular expressions (regex): rule-based text patterns that accommodate common morphological variants without manual enumeration. Matching is case-insensitive and applied to lightly preprocessed text. For example, a single regex simultaneously matches *expectation* and *expectations*, or *two percent* and *2 percent*.

**Step 3: Co-occurrence rules for narrative identification.** A text unit is a sentence for statements and a turn-aware chunk for press conferences, typically consisting of two to five sentences from a single speaker turn. The broad inflation narrative $\pi_t^g$ flags units containing any broad inflation term. The context-dependent indicators use local co-occurrence rules. For statements, inflation and policy anchors may appear in the same sentence or in a local adjacent-sentence window of $\pm 2$ sentences. For press conferences, where the Chair may state the inflation object and the policy implication in nearby but not identical chunks, anchor matching is evaluated within a local window of $\pm 6$ chunks inside the same meeting. Within this windowed structure, $\pi_t^{cur}$ requires an inflation anchor and a timing, release, or realized-dynamics cue; $\delta_t^\pi$ requires inflation language together with outlook, risk, expectation, or projection language; and $\omega_t^{\pi}$ requires inflation language together with policy-intentions, conditionality, reaction-function, or forward-guidance language. The rules ensure that the dictionaries identify *which* units belong to each narrative, while the FOMC-RoBERTa classifier measures the directional (hawkish or dovish) stance of those units.

**Table 7.** Dictionary for broad inflation narrative $\pi_t^{g}$

| *Dictionary block* | *Statements* | *Press conferences* |
|---|---|---|
| Price-stability mandate and target | price stability; stable prices; 2 percent; two percent; inflation objective; inflation goal; inflation target | price stability; stable prices; 2 percent; two percent; inflation objective; inflation goal; inflation target |
| Inflation measures and price indices | inflation; core inflation; headline inflation; underlying inflation; PCE; CPI; consumer price index; PCE inflation; CPI inflation; PCE price index; CPI price index; core PCE; core PCE prices | inflation; core inflation; headline inflation; underlying inflation; PCE; CPI; consumer price index; consumer prices; PCE inflation; CPI inflation; PCE prices; CPI price index; PCE price index; core PCE; core PCE prices |
| Expectations and inflation compensation | inflation expectations; longer-term inflation expectations; inflation compensation; measures of inflation compensation; anchored; well anchored; unanchored; de-anchored; breakevens; TIPS | inflation expectations; inflation compensation; market-based measures of inflation compensation; survey-based measures; anchored; well anchored; unanchored; de-anchored; breakevens; TIPS |
| Pressures, risks, and disinflation progress | inflation pressures; inflationary pressures; price pressures; cost pressures; inflation risks; inflation developments; bringing inflation back to 2 percent; return inflation; disinflation; disinflationary | inflationary pressures; price pressures; inflation remains elevated; above target; progress toward price stability; disinflation process; disinflation; last mile; final mile |

*Notes:* The first column groups lexical items into blocks solely for readability; operationally, a unit is flagged for $\pi_t^{g}$ if it matches any item in the union of blocks. Matching is case-insensitive and implemented via regex.

**Table 8.** Dictionary for current-inflation narrative $\pi_t^{cur}$

| *Dictionary block* | *Statements* | *Press conferences* |
|---|---|---|
| Inflation anchors | inflation; prices; price stability; PCE; CPI; core inflation; headline inflation; inflation expectations; inflation compensation; price pressures; inflationary pressures; disinflation | inflation; prices; price stability; PCE; CPI; core inflation; headline inflation; inflation expectations; inflation compensation; price pressures; inflationary pressures; disinflation |
| Timing / recentness cues | recently; in recent months; over recent months; latest; most recent; monthly; month-to-month | recently; in recent months; over recent months; last few months; latest; most recent; monthly; month-to-month |
| Release / reading language | came in; came out; readings | came in; printed; readings |
| Realized-inflation verbs and descriptors | eased; moderated; declined; fell; slowed; softened; picked up; increased; rose; accelerated; rebounded; stalled; plateaued; elevated; subdued; sticky; persistent | eased; moderated; declined; fell; slowed; softened; picked up; increased; rose; accelerated; rebounded; stalled; plateaued; elevated; subdued; sticky; persistent |
| Key components / drivers | energy prices; food prices; gasoline prices; goods inflation; services inflation; shelter/housing/rent inflation | energy prices; food prices; gasoline prices; goods inflation; services inflation; shelter/housing/rent inflation |

*Notes:* A unit contributes to $\pi_t^{cur}$ when it matches (i) at least one inflation-anchor term and (ii) at least one current-dynamics cue (timing/recentness, release/reading language, or realized-inflation verbs/descriptors) within the same unit.

**Table 9.** Dictionary for Delphic inflation narrative $\delta_t^\pi$

| *Dictionary block* | *Statements* | *Press conferences* |
|---|---|---|
| Inflation anchors | inflation; prices; price stability; PCE; CPI; core inflation; headline inflation; 2 percent; two percent | inflation; prices; price stability; PCE; CPI; core inflation; headline inflation; 2 percent; two percent |
| Outlook / projections | outlook; forecast; projection(s); baseline; expected; participants' projections; staff forecast | outlook; forecast; projection(s); baseline; expected; participants' projections; staff forecast |
| Risks / uncertainty | risks; balance of risks; upside risks; downside risks; uncertainty; highly uncertain | risks; balance of risks; upside risks; downside risks; uncertainty; highly uncertain |
| Drivers and components (inflation-relevant) | energy prices; food prices; financial conditions; credit conditions; supply chain; unit labor costs; wage growth; goods inflation; services inflation; shelter/housing/rent | energy prices; food prices; financial conditions; credit conditions; supply chain; unit labor costs; wage growth; goods inflation; services inflation; shelter/housing/rent |
| Expectations (informational) | inflation expectations; longer-term inflation expectations; anchored; well anchored; unanchored; de-anchored | inflation expectations; longer-term inflation expectations; anchored; well anchored; unanchored; de-anchored |

*Notes:* A unit contributes to $\delta_t^\pi$ when it matches an inflation-anchor term and co-occurs with outlook/projection or risk/uncertainty language within the same unit. The baseline Delphic dictionaries exclude explicit SEP labels such as "SEP" and "Summary of Economic Projections" to avoid mechanical loading on SEP-release meetings. We retain generic projection and outlook language, and we separately assess SEP-related heterogeneity through interaction specifications.

**Table 10.** Dictionary for Odyssean inflation narrative $\omega_t^\pi$

| *Dictionary block* | *Statements* | *Press conferences* |
|---|---|---|
| Inflation anchors | inflation; prices; price stability; PCE; CPI; core inflation; headline inflation; 2 percent; two percent | inflation; prices; price stability; PCE; CPI; core inflation; headline inflation; 2 percent; two percent |
| Policy instruments | federal funds rate; target range; policy rate; monetary policy; stance of monetary policy | federal funds rate; target range; policy rate; monetary policy |
| Policy actions | raise/increase; lower/cut; hold/maintain; pause; tighten; ease; reduce; additional tightening; further increases | raise/increase; lower/cut; hold/maintain; pause; tighten; ease; reduce; additional tightening; further increases |
| Intentions / commitment language | it will be appropriate; will be appropriate; prepared to; be prepared to; committed to; plan to | it will be appropriate; will be appropriate; prepared to; be prepared to; committed to; plan to |
| Conditionality / state dependence | data dependent; incoming data; meeting by meeting; conditional on; contingent on; provided that; as warranted; will depend on | data dependent; incoming data; meeting by meeting; conditional on; contingent on; provided that; as warranted; will depend on |
| Forward-guidance stance descriptors | restrictive; sufficiently restrictive; appropriately restrictive; higher for longer; for some time; extended period | restrictive; sufficiently restrictive; appropriately restrictive; higher for longer; for some time; extended period |

*Notes:* A unit contributes to $\omega_t^\pi$ when it matches an inflation-anchor term and co-occurs with policy intentions/conditionality, instrument references, or action language within the same unit.

## C.2 Validation of the communication measures

The communication indices combine two components: the FOMC-RoBERTa classifier and the lexical co-occurrence rules. We assess the measures along three dimensions. We rely on the documented out-of-sample performance of the trained classifier for the directional hawkish/dovish/neutral component; we inspect the time-series behavior of the indices around well-known FOMC episodes as a face-validity check; and we use the event-study evidence as a market-price relevance check. The third exercise should not be interpreted as independent validation of the dictionaries: it shows whether the resulting text measures covary systematically with FOMC-day BEI repricing in economically interpretable ways.

**Layer 1: Classifier performance.** The directional stance of each text unit is assigned by FOMC-RoBERTa (Shah et al., 2023), a RoBERTa-base model fine-tuned on human-annotated FOMC sentences for hawkish, neutral, and dovish classification. Their training-and-evaluation pipeline constitutes a dataset-based validation of the stance component that is independent of the lexical dictionaries.

**Layer 2: Episode-based face validity.** The time-series of the communication indices displays the directional patterns one would expect from a measure of inflation communication that tracks the policy record. The Delphic index $\delta_t^\pi$ reaches its most negative values during the 2009 disinflationary trough, when the FOMC explicitly discussed the risk of unwelcome disinflation; it rises through 2021–2022 as the Fed flagged a deteriorating inflation outlook, and peaks in late 2022, coinciding with the tightening cycle of that period. The Odyssean index $\omega_t^\pi$ likewise rises in 2022 and remains elevated through 2023, consistent with the explicit "higher for longer" forward guidance of that period. These patterns are consistent with the narrative record of Fed communication and with the qualitative chronologies documented by Bauer and Swanson (2023) and Acosta et al. (2025). Episode-based face validity is a comparatively weak validation layer relative to the human-coding-and-confusion-matrix template developed for text-based policy measures by Baker et al. (2016); Loughran and McDonald (2011); we return to this limitation in Section 9 as a priority for revision.

**Layer 3: Market-price relevance.** The event-study evidence provides an economic relevance check rather than an independent validation exercise. Text-based measures of central bank communication are commonly assessed by asking whether they covary with high-frequency asset-price responses or policy-surprise measures in economically interpretable ways (Lucca and Trebbi, 2009; Hansen and McMahon, 2016; Cieslak et al., 2023; Cieslak and McMahon, 2024). In the same spirit, if our indices captured only textual variation unrelated to market pricing, the orthogonalized measures would not display systematic sign and maturity patterns in FOMC-day BEI changes. The evidence in the main text shows that the indices are market-relevant in this sense: they covary with inflation-compensation repricing in patterns consistent with the institutional distinction between statements and press conferences and with the Delphic/Odyssean interpretation developed in the theoretical framework. The evidence does not, however, establish that the dictionaries uniquely identify structural communication channels. We therefore interpret the event-study results as reduced-form associations between transparent text measures and BEI repricing, not as validation of a fully structural decomposition of communication.

# D Target-rate surprise

We report the time series of the target-rate surprise $MP1_t$ constructed from federal funds futures in narrow announcement windows around scheduled FOMC meetings, following the approach of Acosta et al. (2025). As described in Section 4.2, the relevant window differs by corpus. For statements, we use the surprise measured in the statement-release window, which captures the rate-action component as it is initially priced. For press conferences in the daily event study, we use the surprise measured in the full announcement-day window, because the narrow press-conference-window surprise is near-degenerate (the rate decision is announced with the statement and is fully priced before the Chair speaks). The series is transformed into standard-deviation units (mean zero, unit variance) within the respective sample. Standardization facilitates direct magnitude comparisons with the standardized communication indices throughout the event-study analysis, while preserving the original timing and sign of the underlying high-frequency shock.

**Figure 10.** Time series of standardized $MP1_t$ at FOMC meetings

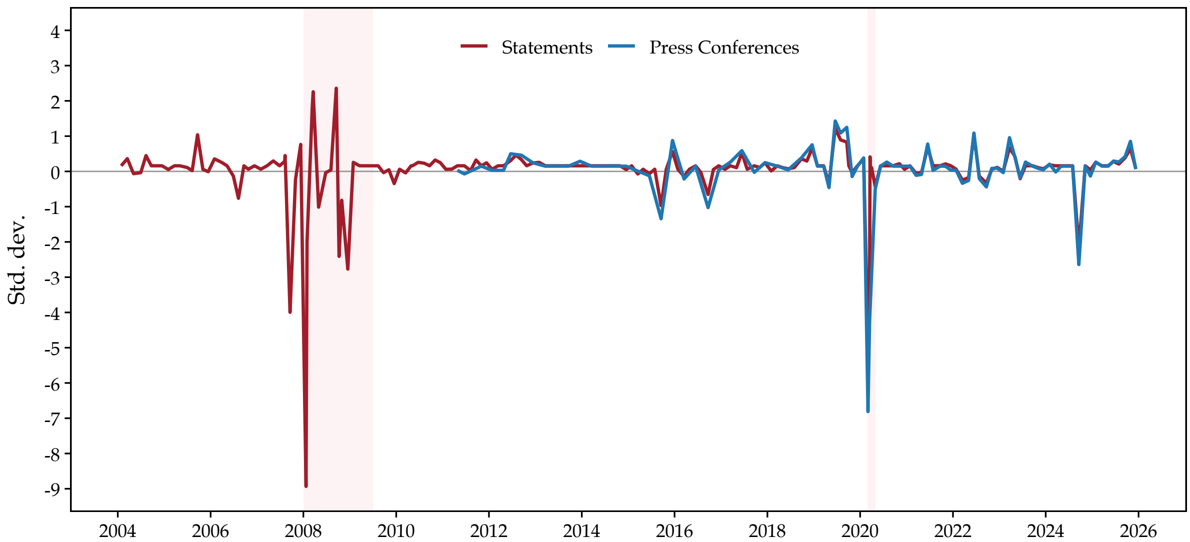


*Notes:* Original target-rate surprise $MP1_t$ series from Acosta et al. (2025). See notes to Figure 1.

Figure 10 reports the full time series of the standardized $MP1_t$ for the statements sample (2004–2025). The series is centered on zero, with the largest positive (contractionary) surprises clustered in the 2022–2023 hiking cycle and the largest negative (expansionary) surprises associated with the March 2020 emergency cuts and the 2008–2009 crisis period. The variation is broadly symmetric, with no structural breaks in the sign or scale of surprises that would raise concerns about heteroskedasticity-driven identification in the baseline regressions. For the press-conference corpus (2011–2025), the full-announcement-day $MP1_t$ has a standard deviation of approximately 0.043 percentage points, substantially larger than the press-conference-only window (0.005 pp); the contrast confirms the empirical relevance of the window-selection argument made in the main text.

# E Other robustness checks

**Regressions without target-rate control.** The baseline conditions on the target-rate surprise $MP1_t$ by orthogonalizing each communication index in the channel-specific event window. As a first robustness check, we estimate the same regressions without any intraday surprise adjustment, so that the reported coefficients capture the total announcement-day BEI co-movement with each index. Table 11 reports the results for statements and press conferences at maturities 2, 5, and 10 years. The core sign patterns and term-structure shapes are preserved: the Delphic index delivers monotone BEI compression for statements, and the long-end forward rise for press conferences survives the omission of the rate control.

**Orthogonalization with respect to the Acosta et al. (2025) surprise.** We next adopt a stricter identification that purges each communication index with respect to the first principal component of the high-frequency monetary-policy surprises constructed by Acosta et al. (2025). This factor summarizes the full cross-asset announcement-day pricing response and represents a more conservative upper bound on the purging of rate-related variation. Table 12 reports the results. The sign and maturity patterns are qualitatively similar to the baseline, which indicates that the findings are not an artifact of the choice of surprise measure used for orthogonalization.

**Alternative sample windows.** Table 13 addresses two distinct concerns about exceptional market conditions. The left panel replicates the statement baseline after excluding the Global Financial Crisis window (August 2007–June 2009, $N = 176$), removing a period in which BEI compensation was distorted by liquidity premia. The right panel augments the press-conference baseline with interaction terms between each orthogonalized index and a dummy for meetings that coincide with a Summary of Economic Projections (SEP) release. The main-effect rows correspond to non-SEP meetings; the interaction rows capture any incremental effect on SEP dates. The main effects in the right panel remain consistent with the baseline; the SEP interactions are imprecisely estimated, providing no evidence that the presence of an SEP release systematically alters the inflation-communication channel. At the ten-year forward, however, the non-SEP main-effect coefficients (Table 13, Panel B) are roughly twice the pooled-baseline values, with SEP-interaction coefficients of opposite sign that are individually imprecise. The pattern is consistent with non-SEP press conferences carrying additional long-horizon information, although the small SEP/non-SEP cell sizes preclude a sharper conclusion.

**Excluding calendar year 2020.** Table 14 drops the COVID year from both samples. The left panel covers statements ($N = 173$) and the right panel covers press conferences ($N = 80$). The sign patterns and maturity profiles are stable, with some reduction in precision expected from the smaller samples; the COVID episode, characterized by extreme BEI volatility and unconventional policy, is therefore not the main driver of the baseline results. Across these exercises, the qualitative sign and maturity patterns emphasized in the main text are preserved, although precision declines in smaller subsamples. The qualitative patterns documented in Sections 5–6, namely monotone BEI compression for hawkish statement language and long-end BEI rise for press-conference inflation communication, survive changes in $MP1$ control, orthogonalization strategy, and sample composition, with the expected reduction in precision in smaller windows.

# F Baseline results: full maturity tables

The main text presents the baseline event-study results graphically. This appendix provides the corresponding numerical results for all maturities (two to ten years), using the same specification as Sections 5–6. Table 15 covers post-meeting statements and Table 16 covers press conferences.

**Table 11.** BEI responses without target-rate control

| | Post-meeting statements | | | Press conferences | | |
|---|---|---|---|---|---|---|
| | 2y | 5y | 10y | 2y | 5y | 10y |
| *Panel A: BEI yields* | | | | | | |
| $mps_t$ | −0.8 | −0.8*** | −0.4 | −0.7 | −0.9* | −0.2 |
| | (0.5) | (0.3) | (0.2) | (0.6) | (0.5) | (0.3) |
| $R^2$ | 0.01 | 0.04 | 0.01 | 0.01 | 0.05 | 0.00 |
| $\pi_t^g$ | −0.6 | −0.7** | −0.3 | −0.1 | −0.5 | 0.1 |
| | (0.4) | (0.3) | (0.2) | (0.7) | (0.5) | (0.4) |
| $R^2$ | 0.01 | 0.03 | 0.01 | 0.00 | 0.02 | 0.00 |
| $\pi_t^{cur}$ | −0.6 | −0.5* | −0.3 | 0.1 | −0.2 | 0.3 |
| | (0.4) | (0.3) | (0.2) | (0.6) | (0.4) | (0.3) |
| $R^2$ | 0.01 | 0.02 | 0.01 | 0.00 | 0.00 | 0.01 |
| $\delta_t^\pi$ | −2.2*** | −1.4*** | −0.7*** | −0.7 | −0.7 | −0.1 |
| | (0.6) | (0.3) | (0.2) | (0.7) | (0.5) | (0.4) |
| $R^2$ | 0.11 | 0.12 | 0.04 | 0.02 | 0.03 | 0.00 |
| $\omega_t^\pi$ | −0.7 | −0.7** | −0.5* | −0.1 | −0.6 | −0.1 |
| | (0.5) | (0.3) | (0.3) | (0.6) | (0.4) | (0.3) |
| $R^2$ | 0.01 | 0.03 | 0.02 | 0.00 | 0.03 | 0.00 |
| *Panel B: BEI forwards* | | | | | | |
| $mps_t$ | −0.8** | −0.4 | −0.0 | −0.5 | −0.8 | 1.1* |
| | (0.4) | (0.3) | (0.4) | (0.7) | (0.5) | (0.7) |
| $R^2$ | 0.02 | 0.01 | 0.00 | 0.01 | 0.03 | 0.04 |
| $\pi_t^g$ | −0.7 | −0.5 | 0.1 | 0.1 | −0.7 | 1.5** |
| | (0.4) | (0.4) | (0.4) | (0.8) | (0.6) | (0.7) |
| $R^2$ | 0.01 | 0.01 | 0.00 | 0.00 | 0.03 | 0.08 |
| $\pi_t^{cur}$ | −0.6 | −0.3 | −0.1 | 0.1 | −0.3 | 1.3** |
| | (0.4) | (0.3) | (0.3) | (0.6) | (0.5) | (0.6) |
| $R^2$ | 0.01 | 0.00 | 0.00 | 0.00 | 0.00 | 0.05 |
| $\delta_t^\pi$ | −1.3*** | −0.4 | −0.0 | −0.0 | −0.7 | 1.1 |
| | (0.4) | (0.3) | (0.4) | (0.8) | (0.5) | (0.7) |
| $R^2$ | 0.06 | 0.01 | 0.00 | 0.00 | 0.02 | 0.04 |
| $\omega_t^\pi$ | −0.7 | −0.5 | −0.3 | −0.1 | −0.9* | 1.2** |
| | (0.5) | (0.4) | (0.4) | (0.6) | (0.5) | (0.6) |
| $R^2$ | 0.02 | 0.01 | 0.00 | 0.00 | 0.04 | 0.05 |

*Notes:* Entries report coefficients from horizon-by-horizon regressions of daily BEI changes on one standardized communication index at a time, without conditioning on $MP1_t$. *Left*: statements, 2004–2025 ($N = 185$). *Right*: press conferences, 2011–2025 ($N = 88$), monetary-event window. Wild bootstrap s.e. ($B = 10{,}000$) in parentheses. Coefficients in basis points. $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

**Table 12.** BEI responses with Acosta et al. (2025) orthogonalization

| | Post-meeting statements | | | Press conferences | | |
|---|---|---|---|---|---|---|
| | 2y | 5y | 10y | 2y | 5y | 10y |
| *Panel A: BEI yields* | | | | | | |
| $mps_t$ | −0.5 | −0.6** | −0.2 | −0.7 | −0.9* | −0.2 |
| | (0.5) | (0.3) | (0.2) | (0.6) | (0.5) | (0.3) |
| $R^2$ | 0.01 | 0.02 | 0.00 | 0.01 | 0.05 | 0.00 |
| $\pi_t^g$ | −0.4 | −0.6** | −0.2 | −0.1 | −0.6 | 0.1 |
| | (0.4) | (0.3) | (0.2) | (0.7) | (0.5) | (0.4) |
| $R^2$ | 0.00 | 0.02 | 0.00 | 0.00 | 0.02 | 0.00 |
| $\pi_t^{cur}$ | −0.5 | −0.5* | −0.3 | 0.2 | −0.2 | 0.3 |
| | (0.4) | (0.3) | (0.2) | (0.6) | (0.4) | (0.3) |
| $R^2$ | 0.01 | 0.01 | 0.01 | 0.00 | 0.00 | 0.01 |
| $\delta_t^\pi$ | −1.8*** | −1.1*** | −0.5** | −0.7 | −0.6 | −0.1 |
| | (0.5) | (0.3) | (0.2) | (0.7) | (0.5) | (0.4) |
| $R^2$ | 0.08 | 0.08 | 0.02 | 0.01 | 0.02 | 0.00 |
| $\omega_t^\pi$ | −0.5 | −0.6* | −0.3 | −0.1 | −0.6* | −0.1 |
| | (0.5) | (0.3) | (0.2) | (0.6) | (0.4) | (0.3) |
| $R^2$ | 0.01 | 0.02 | 0.01 | 0.00 | 0.03 | 0.00 |
| *Panel B: BEI forwards* | | | | | | |
| $mps_t$ | −0.7* | −0.2 | 0.1 | −0.5 | −0.8 | 1.1* |
| | (0.4) | (0.3) | (0.4) | (0.7) | (0.5) | (0.7) |
| $R^2$ | 0.02 | 0.00 | 0.00 | 0.01 | 0.03 | 0.04 |
| $\pi_t^g$ | −0.6 | −0.4 | 0.2 | 0.0 | −0.8 | 1.5** |
| | (0.4) | (0.4) | (0.4) | (0.8) | (0.6) | (0.7) |
| $R^2$ | 0.01 | 0.01 | 0.00 | 0.00 | 0.03 | 0.08 |
| $\pi_t^{cur}$ | −0.5 | −0.3 | −0.1 | 0.1 | −0.2 | 1.3** |
| | (0.4) | (0.3) | (0.3) | (0.6) | (0.5) | (0.6) |
| $R^2$ | 0.01 | 0.00 | 0.00 | 0.00 | 0.00 | 0.05 |
| $\delta_t^\pi$ | −1.2*** | −0.2 | 0.2 | −0.0 | −0.6 | 1.1 |
| | (0.4) | (0.3) | (0.4) | (0.8) | (0.5) | (0.7) |
| $R^2$ | 0.05 | 0.00 | 0.00 | 0.00 | 0.02 | 0.04 |
| $\omega_t^\pi$ | −0.7 | −0.4 | −0.2 | −0.2 | −0.9* | 1.2** |
| | (0.4) | (0.4) | (0.4) | (0.6) | (0.5) | (0.6) |
| $R^2$ | 0.01 | 0.01 | 0.00 | 0.00 | 0.04 | 0.05 |

*Notes:* Each index is standardized and orthogonalized with respect to the first principal component of Acosta et al. (2025) high-frequency surprises. *Left*: post-meeting statements, 2004–2025 ($N = 185$). *Right*: Press conferences, 2011–2025 ($N = 88$). Wild bootstrap s.e. ($B = 10{,}000$) in parentheses. Coefficients in basis points. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

**Table 13.** Alternative sample: statements (ex-GFC) and press conf. (SEP interaction)

| | Statements, ex-GFC | | | Press conf., SEP interaction | | |
|---|---|---|---|---|---|---|
| | 2y | 5y | 10y | 2y | 5y | 10y |
| *Panel A: BEI yields* | | | | | | |
| $mps_t$ | −0.2 | −0.6** | −0.2 | −0.1 | −0.6 | 0.3 |
| | (0.4) | (0.2) | (0.2) | (0.6) | (0.5) | (0.3) |
| $mps_t \times SEP_t$ | | | | −1.0 | −0.3 | −0.7 |
| | | | | (1.1) | (0.8) | (0.6) |
| $R^2$ | 0.00 | 0.03 | 0.00 | 0.02 | 0.05 | 0.01 |
| $\pi_t^g$ | −0.3 | −0.5* | −0.2 | 0.0 | −0.4 | 0.6 |
| | (0.4) | (0.3) | (0.2) | (0.8) | (0.6) | (0.4) |
| $\pi_t^g \times SEP_t$ | | | | −0.2 | −0.0 | −0.7 |
| | | | | (1.2) | (1.0) | (0.6) |
| $R^2$ | 0.00 | 0.02 | 0.00 | 0.00 | 0.02 | 0.01 |
| $\pi_t^{cur}$ | −0.2 | −0.3 | −0.1 | 0.0 | −0.1 | 0.5 |
| | (0.3) | (0.2) | (0.2) | (0.4) | (0.4) | (0.3) |
| $\pi_t^{cur} \times SEP_t$ | | | | 0.2 | −0.2 | −0.3 |
| | | | | (1.0) | (0.7) | (0.6) |
| $R^2$ | 0.00 | 0.01 | 0.00 | 0.00 | 0.01 | 0.01 |
| $\delta_t^\pi$ | −0.7** | −0.6** | −0.2 | 0.3 | −0.1 | 0.7* |
| | (0.3) | (0.2) | (0.2) | (0.7) | (0.6) | (0.4) |
| $\delta_t^\pi \times SEP_t$ | | | | −1.5 | −0.7 | −1.1 |
| | | | | (1.1) | (0.9) | (0.6) |
| $R^2$ | 0.02 | 0.03 | 0.01 | 0.03 | 0.04 | 0.03 |
| $\omega_t^\pi$ | −0.3 | −0.4 | −0.3 | −0.5 | −1.0** | −0.1 |
| | (0.4) | (0.3) | (0.3) | (0.9) | (0.5) | (0.3) |
| $\omega_t^\pi \times SEP_t$ | | | | 0.6 | 0.7 | 0.0 |
| | | | | (1.2) | (0.7) | (0.5) |
| $R^2$ | 0.00 | 0.02 | 0.01 | 0.00 | 0.03 | 0.00 |
| *Panel B: BEI forwards* | | | | | | |
| $mps_t$ | −0.7* | −0.4 | 0.2 | 0.1 | −1.0 | 2.7** |
| | (0.4) | (0.3) | (0.4) | (0.8) | (1.0) | (1.3) |
| $mps_t \times SEP_t$ | | | | −0.8 | 0.4 | −2.4 |
| | | | | (1.3) | (1.2) | (1.4) |
| $R^2$ | 0.02 | 0.01 | 0.00 | 0.01 | 0.04 | 0.10 |
| $\pi_t^g$ | −0.5 | −0.4 | 0.3 | 0.5 | −0.7 | 3.1** |
| | (0.4) | (0.4) | (0.4) | (1.0) | (1.2) | (1.4) |
| $\pi_t^g \times SEP_t$ | | | | −0.7 | 0.1 | −2.4 |
| | | | | (1.4) | (1.4) | (1.6) |
| $R^2$ | 0.01 | 0.01 | 0.00 | 0.00 | 0.04 | 0.12 |
| $\pi_t^{cur}$ | −0.3 | −0.2 | 0.1 | 0.7 | −0.3 | 2.0* |
| | (0.4) | (0.3) | (0.3) | (0.6) | (0.8) | (1.1) |
| $\pi_t^{cur} \times SEP_t$ | | | | −1.1 | 0.3 | −1.2 |
| | | | | (1.0) | (1.0) | (1.3) |
| $R^2$ | 0.00 | 0.00 | 0.00 | 0.01 | 0.02 | 0.07 |
| $\delta_t^\pi$ | −0.7** | −0.2 | 0.3 | 0.8 | −0.6 | 2.9* |
| | (0.3) | (0.3) | (0.3) | (0.9) | (1.2) | (1.5) |
| $\delta_t^\pi \times SEP_t$ | | | | −1.1 | −0.0 | −2.6 |
| | | | | (1.3) | (1.3) | (1.7) |
| $R^2$ | 0.02 | 0.00 | 0.00 | 0.01 | 0.04 | 0.09 |
| $\omega_t^\pi$ | −0.4 | −0.5 | −0.1 | −0.1 | −1.4 | 2.5* |
| | (0.4) | (0.4) | (0.4) | (1.0) | (1.0) | (1.4) |
| $\omega_t^\pi \times SEP_t$ | | | | −0.1 | 0.9 | −1.9 |
| | | | | (1.2) | (1.2) | (1.5) |
| $R^2$ | 0.01 | 0.01 | 0.00 | 0.00 | 0.05 | 0.08 |

*Notes: Left*: post-meeting statements, 2004–2025 excluding the Global Financial Crisis window (August 2007–June 2009), $N = 176$; baseline specification with wild bootstrap s.e. ($B =$ 10,000, Rademacher weights). *Right*: Press conferences, 2011–2025 ($N = 88$), baseline augmented with an SEP interaction. The interaction terms (rows $\cdot \times SEP_t$) are reported for descriptive purposes; the inferential focus of the paper is the univariate specifications that underpin the main-text figures, for which the wild bootstrap is implemented. Heteroskedasticity-robust White HC0 standard errors are reported for both the main-effect and interaction rows of this auxiliary specification, so that all entries within the right panel are based on a common standard-error construction. Coefficients in basis points. $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

**Table 14.** BEI responses excluding calendar year 2020

| | Post-meeting statements | | | Press conferences | | |
|---|---|---|---|---|---|---|
| | 2y | 5y | 10y | 2y | 5y | 10y |
| *Panel A: BEI yields* | | | | | | |
| $mps_t$ | −0.7 | −0.8** | −0.4 | −0.5 | −0.8 | −0.1 |
| | (0.6) | (0.3) | (0.2) | (0.7) | (0.6) | (0.4) |
| $R^2$ | 0.01 | 0.03 | 0.01 | 0.01 | 0.04 | 0.00 |
| $\pi_t^g$ | −0.5 | −0.7** | −0.3 | 0.0 | −0.4 | 0.1 |
| | (0.4) | (0.3) | (0.2) | (0.7) | (0.6) | (0.4) |
| $R^2$ | 0.01 | 0.03 | 0.01 | 0.00 | 0.01 | 0.00 |
| $\pi_t^{cur}$ | −0.6 | −0.5* | −0.3 | 0.3 | −0.1 | 0.4 |
| | (0.5) | (0.3) | (0.3) | (0.7) | (0.5) | (0.4) |
| $R^2$ | 0.01 | 0.02 | 0.01 | 0.00 | 0.00 | 0.02 |
| $\delta_t^\pi$ | −2.2*** | −1.4*** | −0.7*** | −0.6 | −0.6 | −0.1 |
| | (0.6) | (0.3) | (0.2) | (0.7) | (0.6) | (0.4) |
| $R^2$ | 0.11 | 0.12 | 0.04 | 0.01 | 0.03 | 0.00 |
| $\omega_t^\pi$ | −0.6 | −0.7* | −0.4* | 0.2 | −0.5 | 0.0 |
| | (0.5) | (0.3) | (0.3) | (0.6) | (0.4) | (0.3) |
| $R^2$ | 0.01 | 0.03 | 0.02 | 0.00 | 0.02 | 0.00 |
| *Panel B: BEI forwards* | | | | | | |
| $mps_t$ | −0.8* | −0.4 | −0.0 | −0.5 | −0.8 | 1.3* |
| | (0.4) | (0.4) | (0.4) | (0.8) | (0.6) | (0.7) |
| $R^2$ | 0.02 | 0.01 | 0.00 | 0.01 | 0.03 | 0.05 |
| $\pi_t^g$ | −0.6 | −0.5 | 0.1 | 0.1 | −0.8 | 1.6** |
| | (0.4) | (0.4) | (0.4) | (0.8) | (0.7) | (0.7) |
| $R^2$ | 0.01 | 0.01 | 0.00 | 0.00 | 0.02 | 0.08 |
| $\pi_t^{cur}$ | −0.6 | −0.3 | −0.1 | 0.1 | −0.2 | 1.4** |
| | (0.4) | (0.4) | (0.3) | (0.7) | (0.6) | (0.7) |
| $R^2$ | 0.01 | 0.00 | 0.00 | 0.00 | 0.00 | 0.06 |
| $\delta_t^\pi$ | −1.3*** | −0.4 | 0.0 | 0.0 | −0.7 | 1.2 |
| | (0.4) | (0.3) | (0.4) | (0.8) | (0.6) | (0.7) |
| $R^2$ | 0.06 | 0.01 | 0.00 | 0.00 | 0.02 | 0.04 |
| $\omega_t^\pi$ | −0.7 | −0.5 | −0.3 | −0.1 | −0.9 | 1.4** |
| | (0.5) | (0.4) | (0.4) | (0.7) | (0.6) | (0.7) |
| $R^2$ | 0.02 | 0.01 | 0.00 | 0.00 | 0.03 | 0.06 |

*Notes:* Each index is standardized and orthogonalized with respect to $MP1_t$, after removing calendar year 2020 from both samples. *Left*: post-meeting statements, $N = 173$. *Right*: Press conferences, monetary-event window, $N = 80$. Wild bootstrap s.e. ($B = 10{,}000$) in parentheses. Coefficients in basis points. $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

**Table 15.** FOMC statements: BEI yields and forwards, full maturity spectrum

| | 2y | 3y | 4y | 5y | 6y | 7y | 8y | 9y | 10y |
|---|---|---|---|---|---|---|---|---|---|
| *Panel A: BEI yields* | | | | | | | | | |
| $mps_t$ | −0.7 | −0.8** | −0.8** | −0.8*** | −0.7** | −0.6** | −0.5** | −0.4* | −0.3 |
| | (0.5) | (0.4) | (0.3) | (0.3) | (0.3) | (0.2) | (0.2) | (0.2) | (0.2) |
| $R^2$ | 0.01 | 0.02 | 0.03 | 0.03 | 0.03 | 0.02 | 0.02 | 0.01 | 0.01 |
| $\pi_t^g$ | −0.6 | −0.6* | −0.7** | −0.7** | −0.6** | −0.5** | −0.5* | −0.4 | −0.3 |
| | (0.4) | (0.3) | (0.3) | (0.3) | (0.3) | (0.3) | (0.2) | (0.2) | (0.2) |
| $R^2$ | 0.01 | 0.01 | 0.02 | 0.03 | 0.03 | 0.02 | 0.02 | 0.01 | 0.01 |
| $\pi_t^{cur}$ | −0.6 | −0.6 | −0.6* | −0.5* | −0.5* | −0.4* | −0.4 | −0.4 | −0.3 |
| | (0.4) | (0.4) | (0.3) | (0.3) | (0.3) | (0.3) | (0.2) | (0.2) | (0.2) |
| $R^2$ | 0.01 | 0.01 | 0.02 | 0.02 | 0.02 | 0.01 | 0.01 | 0.01 | 0.01 |
| $\delta_t^\pi$ | −2.1*** | −1.8*** | −1.6*** | −1.4*** | −1.2*** | −1.0*** | −0.9*** | −0.8*** | −0.7*** |
| | (0.6) | (0.4) | (0.4) | (0.3) | (0.3) | (0.3) | (0.2) | (0.2) | (0.2) |
| $R^2$ | 0.10 | 0.12 | 0.12 | 0.11 | 0.09 | 0.08 | 0.06 | 0.05 | 0.04 |
| $\omega_t^\pi$ | −0.7 | −0.7 | −0.7* | −0.7** | −0.6** | −0.6* | −0.5* | −0.5* | −0.4* |
| | (0.5) | (0.4) | (0.4) | (0.3) | (0.3) | (0.3) | (0.3) | (0.3) | (0.2) |
| $R^2$ | 0.01 | 0.02 | 0.02 | 0.03 | 0.03 | 0.02 | 0.02 | 0.02 | 0.02 |
| *Panel B: BEI instantaneous forwards* | | | | | | | | | |
| $mps_t$ | −0.8** | −1.0*** | −0.7** | −0.4 | 0.0 | 0.2 | 0.2 | 0.2 | 0.0 |
| | (0.4) | (0.3) | (0.3) | (0.3) | (0.3) | (0.3) | (0.3) | (0.3) | (0.4) |
| $R^2$ | 0.02 | 0.04 | 0.02 | 0.01 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 |
| $\pi_t^g$ | −0.6 | −0.8** | −0.8** | −0.5 | −0.2 | 0.0 | 0.2 | 0.2 | 0.1 |
| | (0.4) | (0.3) | (0.4) | (0.4) | (0.3) | (0.3) | (0.3) | (0.4) | (0.4) |
| $R^2$ | 0.01 | 0.03 | 0.03 | 0.01 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 |
| $\pi_t^{cur}$ | −0.6 | −0.6* | −0.5 | −0.3 | −0.2 | −0.1 | −0.1 | −0.1 | −0.2 |
| | (0.4) | (0.3) | (0.3) | (0.3) | (0.3) | (0.3) | (0.3) | (0.3) | (0.3) |
| $R^2$ | 0.01 | 0.01 | 0.01 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 |
| $\delta_t^\pi$ | −1.3*** | −1.0*** | −0.7** | −0.4 | −0.1 | 0.0 | 0.1 | 0.1 | 0.0 |
| | (0.4) | (0.3) | (0.3) | (0.3) | (0.3) | (0.3) | (0.3) | (0.3) | (0.4) |
| $R^2$ | 0.05 | 0.05 | 0.02 | 0.01 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 |
| $\omega_t^\pi$ | −0.7 | −0.8** | −0.7 | −0.5 | −0.3 | −0.2 | −0.1 | −0.2 | −0.3 |
| | (0.5) | (0.4) | (0.4) | (0.4) | (0.3) | (0.3) | (0.3) | (0.3) | (0.4) |
| $R^2$ | 0.02 | 0.03 | 0.02 | 0.01 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 |

*Notes:* Post-meeting statements, 2004–2025 ($N = 185$). All indices standardized and orthogonalized with respect to $MP1_t$ in the statement-release window from Acosta et al. (2025). Wild bootstrap s.e. ($B = 10{,}000$) in parentheses. Coefficients in basis points. $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

**Table 16.** Chair press conferences: BEI yields and forwards, full maturity spectrum

| | 2y | 3y | 4y | 5y | 6y | 7y | 8y | 9y | 10y |
|---|---|---|---|---|---|---|---|---|---|
| *Panel A: BEI yields* | | | | | | | | | |
| $mps_t$ | −0.6 | −0.7 | −0.8 | −0.8* | −0.8* | −0.6* | −0.5 | −0.3 | −0.2 |
| | (0.6) | (0.6) | (0.5) | (0.5) | (0.4) | (0.4) | (0.3) | (0.3) | (0.3) |
| $R^2$ | 0.01 | 0.02 | 0.04 | 0.05 | 0.05 | 0.04 | 0.02 | 0.01 | 0.00 |
| $\pi_t^g$ | −0.1 | −0.2 | −0.4 | −0.5 | −0.5 | −0.4 | −0.2 | 0.0 | 0.1 |
| | (0.7) | (0.6) | (0.6) | (0.5) | (0.5) | (0.4) | (0.4) | (0.4) | (0.3) |
| $R^2$ | 0.00 | 0.00 | 0.01 | 0.02 | 0.02 | 0.01 | 0.00 | 0.00 | 0.00 |
| $\pi_t^{cur}$ | 0.1 | 0.0 | −0.2 | −0.2 | −0.2 | −0.1 | 0.1 | 0.2 | 0.3 |
| | (0.6) | (0.5) | (0.4) | (0.4) | (0.4) | (0.3) | (0.3) | (0.3) | (0.3) |
| $R^2$ | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.01 |
| $\delta_t^\pi$ | −0.7 | −0.5 | −0.6 | −0.7 | −0.6 | −0.5 | −0.4 | −0.2 | −0.1 |
| | (0.7) | (0.6) | (0.6) | (0.5) | (0.5) | (0.4) | (0.4) | (0.3) | (0.3) |
| $R^2$ | 0.02 | 0.01 | 0.02 | 0.03 | 0.03 | 0.02 | 0.01 | 0.01 | 0.00 |
| $\omega_t^\pi$ | −0.1 | −0.2 | −0.5 | −0.6 | −0.6* | −0.5 | −0.4 | −0.2 | −0.1 |
| | (0.6) | (0.5) | (0.4) | (0.4) | (0.3) | (0.3) | (0.3) | (0.3) | (0.3) |
| $R^2$ | 0.00 | 0.00 | 0.01 | 0.03 | 0.03 | 0.02 | 0.01 | 0.00 | 0.00 |
| *Panel B: BEI instantaneous forwards* | | | | | | | | | |
| $mps_t$ | −0.4 | −1.0 | −1.2* | −0.8 | −0.2 | 0.5 | 1.0* | 1.2** | 1.2* |
| | (0.7) | (0.7) | (0.7) | (0.5) | (0.4) | (0.4) | (0.5) | (0.6) | (0.7) |
| $R^2$ | 0.01 | 0.04 | 0.05 | 0.03 | 0.00 | 0.02 | 0.05 | 0.05 | 0.05 |
| $\pi_t^g$ | 0.1 | −0.8 | −1.1 | −0.8 | −0.1 | 0.6 | 1.1** | 1.4** | 1.5** |
| | (0.8) | (0.7) | (0.7) | (0.6) | (0.5) | (0.4) | (0.5) | (0.6) | (0.7) |
| $R^2$ | 0.00 | 0.03 | 0.04 | 0.03 | 0.00 | 0.02 | 0.06 | 0.08 | 0.08 |
| $\pi_t^{cur}$ | 0.1 | −0.5 | −0.6 | −0.3 | 0.3 | 0.8** | 1.2** | 1.3** | 1.3** |
| | (0.6) | (0.5) | (0.6) | (0.5) | (0.4) | (0.4) | (0.5) | (0.6) | (0.6) |
| $R^2$ | 0.00 | 0.01 | 0.01 | 0.00 | 0.01 | 0.04 | 0.07 | 0.07 | 0.06 |
| $\delta_t^\pi$ | 0.0 | −0.6 | −0.9 | −0.7 | −0.2 | 0.4 | 0.8 | 1.1* | 1.1 |
| | (0.8) | (0.7) | (0.7) | (0.6) | (0.4) | (0.4) | (0.5) | (0.6) | (0.7) |
| $R^2$ | 0.00 | 0.01 | 0.03 | 0.02 | 0.00 | 0.01 | 0.03 | 0.05 | 0.04 |
| $\omega_t^\pi$ | −0.2 | −1.0* | −1.3** | −0.9* | −0.2 | 0.4 | 0.9* | 1.2** | 1.2** |
| | (0.6) | (0.5) | (0.6) | (0.5) | (0.4) | (0.4) | (0.5) | (0.6) | (0.6) |
| $R^2$ | 0.00 | 0.04 | 0.06 | 0.04 | 0.00 | 0.01 | 0.04 | 0.05 | 0.05 |

*Notes:* Chair press conferences, 2011–2025 ($N = 88$). All indices standardized and orthogonalized with respect to $MP1_t$ in the monetary-event window (full announcement-day window) from Acosta et al. (2025). Wild bootstrap s.e. ($B = 10{,}000$) in parentheses. Coefficients in basis points. $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.