UNIVERSITY OF WARWICK

Economics

Embracing the Future: Tense Patterns and Forward-Looking Central Bank Communication

Andrea Guerrieri D'Amati and Gavin Hassall

July 2026
No: 1619

ISSN 2059-4283 (online)
ISSN 0083-7350 (print)

Warwick Economics Research Papers

Embracing the Future: Tense Patterns and Forward-Looking Central Bank Communication

Andrea Guerrieri D'Amati*[1] and Gavin Hassall†‡[2]

[1]University of Warwick
[2]University of Bath

This version: July 1, 2026
Most updated version available **here**

Abstract

This paper studies how forward-looking language in the Federal Open Market Committee (FOMC) minutes affects market expectations of future interest rates. We analyse the text of the FOMC minutes from 1997 to 2023 with structural topic modelling combined with LLM-based tone and tense analysis. We estimate market reactions in an event study that exploits the fact that the release of the minutes involves no policy change, ensuring any market response reflects pure expectation revisions. We show that forward-looking information about certain topics has systematically moved private sector expectations of future interest rates. In particular, hawkish forward-looking inflation language raises 2-, 5- and 10-year Treasury yields. We interpret these findings through a model where the private sector does not observe the central bank's responsiveness to its inflation outlook, and learns about it via a signal extraction problem. We argue that communication effectiveness depends not only on what topics are discussed but on how they are temporally framed.

Keywords: Central Bank Communication, Monetary Policy, Market Expectations.
JEL Classification: E52, E58, E59.

*Email: andrea.guerrieri-d-amati.1@warwick.ac.uk
†Email: gh913@bath.ac.uk
‡We would like to thank Roger Farmer, Christine Braun, and Leonardo Melosi for their helpful guidance, as well as Roberto Pancrazi, Thijs Van Rens, Juri Marcucci, Amy Handlan, Michael McMahon, Johannes Zahner, and participants of the 2025 Monash-Paris-Warwick-Zurich-CEPR Text-as-Data Workshop, of the 2025 Warwick PhD Conference, of the 2025 PhD Economics Virtual Seminar Series, of the 2025 RGS Doctoral Conference in Economics, of the 2024 ECONDAT Conference on Nontraditional Data, Machine Learning and Natural Language Processing for Macroeconomics, of the 2024 Royal Economics Society PhD Conference, of the 2024 Tor Vergata PhD Conference, of the 2024 Liverpool Text-as-data in Economics Workshop, and of the 2024 Royal Economics Society Annual Conference for their advice and feedback.

1 Introduction

How does the temporal orientation of central bank communication influence financial markets, and does this matter for monetary policy effectiveness? This paper examines how the Federal Open Market Committee (FOMC) uses forward-looking language across different topics to communicate monetary policy. Central bank communication is a powerful tool in modern monetary policy (Blinder, 2018), yet how the temporal dimension of communication shapes market expectations has been relatively less-well studied. Financial markets form expectations about future interest rates from central bank signals, but they cannot directly observe the central bank's "true" policy function.

This information problem is central to our analysis. Market participants observe policy outcomes, not whether decisions rely on current, past, or expected values of inflation or output. Communication therefore gives the central bank a way to reveal how it interprets the economy and how it may react in future. Think of Mario Draghi's "Whatever it takes" speech in 2012: by committing to protect the economy at all costs, he signalled that future monetary policy would remain highly accommodative. Market reactions to communication therefore reflect changes in the private sector's beliefs about the policy function itself.

We combine structural topic modelling (Roberts et al., 2013), sentence-level tone and tense classification with large language models, and an event study design to measure how the temporal framing of FOMC minutes affects market expectations. We focus on the FOMC minutes as they are released several weeks after each meeting and do not coincide with a policy change. So, any market reaction reflects revisions in expectations, and it is whether the temporal aspect of communication influences market expectations is what we focus on. This setting lets us identify which dimensions of communication systematically surprise markets and which information the FOMC does not fully convey through announcements, press conferences, or intermeeting speeches.

Our empirical analysis, covering 1997 to 2023, shows that certain forward-looking topics generate systematic surprises. Hawkish forward-looking inflation language in the minutes increased 2-, 5-, and 10-year Treasury yields by 0.8, 0.6, and 0.5 basis points per one-standard-deviation increase, respectively. By contrast, forward-looking discussions of interest rates were systematically more dovish than markets expected, especially over the zero lower bound period.

We interpret our findings with a model of imperfect information where private sector agents cannot directly observe the central bank's inflation expectations nor how it responds to such expectations and must infer these variables from the observed interest rate. This inference problem results in persistent forecast errors that affect the entire yield curve, as a wedge develops between forecasts done under full-information and partial-information. In a simulation, we show that this type of imperfect information has persistent effects along the yield curve. We interpret the release of FOMC minutes as providing news about these unobserved variables, explaining why financial markets respond to central bank communications even after interest rate decisions have been announced. Our empirical findings are consistent with an interpretation where the central bank's responsiveness to expected inflation has increased over time. Private sector agents are rational Bayesians who gradually learn the reaction function, but can be persistently surprised in one direction over finite periods of time.[1] While our theoretical model focuses specifically on inflation expectations, our empirical analysis captures the multidimensional nature of central bank communication across various topics.

Related Literature Communications are complicated, and it is not always clear what information they release, nor how markets react to this information. Early studies argued that forward guidance primarily works through signalling future interest rates (Gurkaynak et al., 2005), and later papers found so-called information effects about economic conditions (Nakamura and Steinsson, 2018). There are methodological challenges in isolating communication effects from policy actions in an empirical setting. Research often focuses on policy announcements or on first-order tonal measures of communication content, leaving a gap in our understanding of how further dimensions of central bank language, particularly a temporal orientation, can influence market expectations. Our paper contributes to this discussion by developing a novel methodology that combines structural topic modelling with tone and tense analysis through large language models. This approach allows us to identify how forward-looking language across different topics shapes market reactions. We analyse FOMC minutes and estimate market responses using high-frequency regressions, providing evidence on how the temporal framing of central bank communication affects different asset classes.

[1] Our treatment of imperfect information is related to papers in the literature that study a Kalman filtering problem of model agents such as Pearlman et al. (1986), Del Negro and Eusepi (2011), Levine et al. (2012) and Blanchard et al. (2013).

More generally, our paper contributes to a growing body of research that examines how central bank communications influence financial markets and economic outcomes. This literature often focuses on major policy announcements, speeches, and official statements only (Hansen, McMahon and Prat, 2018; Blinder, Ehrmann, Fratzscher, De Haan and Jansen, 2008). However, our focus on the meeting minutes provides a richer framework for textual analysis and event study approaches. We also contribute to the literature using text to analyse how a central bank sets and communicates its policy. For example, Bennani et al. (2020) demonstrate that the communication of ECB officials with respect to their monetary policy preferences helps to better understand the future direction of monetary policy, and highlight that the timing of such communication is very important for understanding monetary policy decisions. However, we implement an entirely data-driven framework rather than hand-coded dictionaries. Prior research has shown that forward-looking language in central bank communications shapes interest rate expectations (Larsen, Thorsrud and Zhulanova, 2021; Leombroni, Vedolin, Venter and Whelan, 2021); we relate such communication in the minutes to stock market reactions, expanding on the literature leveraging investor reactions to Fed communications using intra-day tick data. We build on the intuition that information regarding the future is important for explaining stock market movements (Byrne et al., 2023*a*), and we expand on this by showing the direct effects of changing the temporal and topical structure of communications on markets. To do so, we construct a new measure of forward-looking communication using a combination of topic models and large language models. Finally, our methodology corrects for estimation uncertainty coming from the topic model with simulations. Corrections for generated topic covariates have been discussed in different forms in the computer science literature (Mcauliffe and Blei, 2007; Wilcox, Jacobucci, Zhang and Ammerman, 2023), however there are relatively few papers in economics that address this (Battaglia et al., 2024).

Market reactions to communication events imply the existence of valuable information for the private sector, whether that may be on the state of the economy or on the preferences of the central bank. The empirical challenge then is to extract this information. The literature has focused on two main approaches: a market-based approach, where high-frequency data on federal funds futures and other interest rate contracts to identify the multidimensional nature of monetary policy surprises; and a text-based approach, where information content directly from the text of central bank communications is analysed to show how narratives influence market expectations and media coverage.

The first approach, pioneered by Kuttner (2001) and Gurkaynak et al. (2005), begins with collecting high-frequency data on federal funds futures and eurodollar futures contracts surrounding FOMC announcements. Researchers measure the surprise component of policy decisions by examining changes in these contracts within narrow windows, typically 30 minutes, around the announcement time. This tight window ensures that the measured changes primarily reflect the impact of monetary policy news rather than other economic developments. Following data collection, they apply a factor analysis to determine the dimensionality of monetary policy surprises. These factors are then rotated to provide economically meaningful interpretations. After appropriate rotation, the first factor can be interpreted as a "target factor" representing the surprise component of the current federal funds rate decision, while the second emerges as a "path factor" capturing information about the future trajectory of monetary policy that is orthogonal to the current rate decision[2]. The factor-based methodology offers quantitative measures of policy surprises that can be directly related to asset price movements, however they are only looking at the receiver of the signal (i.e. the private sector) and not at the sender of the signal (i.e. the central bank).

The second approach emphasises instead the textual information that the central bank sends. Text-based approaches extract information content directly from central bank communications through sophisticated natural language processing (NLP) techniques. Researchers apply NLP and machine learning algorithms to central bank statements, minutes, or speeches to quantify their informational content. These methods extract measures of policy stance, sentiment, or topics from the textual data, identifying how changes in communication narratives influence financial markets or media coverage. One of the first examples of this approach is Hansen and McMahon (2016), who find that the information released by the FOMC on the state of economic conditions, as well as the guidance the FOMC provides about future monetary policy decisions. Employing these measures within a FAVAR framework, they find that shocks to forward guidance are more important than the FOMC communication of current economic conditions in terms of their effects on market and real variables. Handlan (2022) created a "text shock" series using neural network methods for text analysis on FOMC statements and internal meeting materials, finding that the wording of FOMC statements

[2] If further used as an instrument in a VAR, these surprises have to be orthogonalised as they may be correlated with other economic variables. Bauer and Swanson (2022) and Miranda-Agrippino and Ricco (2021) provide two approaches for which variables to use.

accounts for four times more variation in federal funds futures prices than direct announcements of changes in the target federal funds rate. This paper belongs to the second approach, and we add to this by looking at the temporal dimension of communication. The closest work to this paper would be Byrne et al. (2023*a*), however, while they focus on measuring temporal orientation and showing that central bank assessments of the state of the economy can be the source of the public's information deficit, we also look into its consequences on financial markets.

The remainder of our paper is structured as follows. Section 2 outlines a theoretical foundation for the empirical results. Section 3.1 provides an overview of the data sources. Section 3.2 details the computational linguistic approach used to analyse the text data. Section 3.2.4 outlines the econometric framework, the results of which are given in Section 4. Finally, Section 5 concludes.

2 Model

In this section, we present a stylised model to illustrate how incomplete information about how forward-looking a central bank is can affect market interest rate expectations and provides a role for communications in reducing this imperfection. We first outline a Fisherian model of price determination, then introduce information imperfections.

Variables with a hat denote log deviations around that variable's non-stochastic steady-state. We assume the central bank sets interest rates in response to current inflation, $\hat{\Pi}_t$ and their expectations of future inflation, $\mathbb{E}_t^{CB}[\hat{\Pi}_{t+1}]$. Our use of a forward-looking reaction function is chosen to reflect how monetary policy is likely set in practice.[3] For example, central banks may choose to respond pre-emptively to expected inflation if there are lags in the monetary policy transmission mechanism. The gross interest rate, $\hat{R}_t$ is set according to

$$\hat{R}_t = \phi\hat{\Pi}_t + \Lambda_t\mathbb{E}_t^{CB}[\hat{\Pi}_{t+1}] + \sigma_R\varepsilon_t^R, \tag{1}$$

where Λ_t is the central bank's time varying responsiveness to their economic outlook and the monetary policy shock is $\varepsilon_t^R \sim N(0,1)$. The time varying responsiveness to expected

[3] This is similar to Clarida et al. (2000), and micro-foundations for forward-looking monetary policy reaction functions from a central bank loss function can be found in Clarida et al. (1999) or Svensson (1997).

inflation is modelled with the logistic function

$$\Lambda_t = \frac{1}{1+e^{-\lambda_t}}, \tag{2}$$

which maps the aggregated preferences of the monetary policy decision makers, λ_t, to the interval $(0, 1)$, and follows the unit root process

$$\lambda_t = \lambda_{t-1} + \sigma_\lambda \varepsilon_t^\lambda, \tag{3}$$

where $\varepsilon_t^\lambda \sim N(0, 1)$. An increase (decrease) in λ_t represents a more hawkish (dovish) forward-looking central bank. Asset market clearing is given by the Fisher equation,

$$\hat{R}_t = \hat{r}_t + \mathbb{E}_t^{PS}\left[\hat{\Pi}_{t+1}\right], \tag{4}$$

where $\mathbb{E}_t^{PS}\left[\hat{\Pi}_{t+1}\right]$ is the private sector's inflation expectation and $\hat{r}_t^*$ is the real (natural) interest rate that evolves according to

$$\hat{r}_t^* = \rho \hat{r}_{t-1}^* + \sigma_r \varepsilon_t^r, \tag{5}$$

where $\rho \in (0, 1)$ and $\varepsilon_t^r \sim N(0, 1)$. All shocks are uncorrelated. The full information rational expectation for inflation, where both the central bank and private sector have complete information is, up to a first-order, given by

$$E_t^{FIRE}\left[\hat{\Pi}_{t+1}\right] = \frac{\rho \hat{r}_t^*}{\phi - \rho(1-\Lambda_t)}. \tag{6}$$

We deviate from full information with the following assumptions about the information sets of the central bank and the private sector. The private sector does not observe shocks to the central bank's variables, ε_t^λ and ε_t^R, but observes all other variables, the structural parameters and the structure of the model equations. The model therefore focuses on a 'Fed response to news' effect (Bauer and Swanson, 2023) where the reaction to the common knowledge state of the economy is the source of imperfect information.[4] The central bank, in addition, observes ε_t^λ and ε_t^R. We assume for simplicity that the central bank sets policy

[4] It is straight-forward to introduce the natural rate as an additional unobserved variable inferred from the setting of policy, which would also capture 'signalling effects.' We abstract from this here, but note that our empirical work likely contains both effects.

as if the economy has full information such that $E_t^{CB}\left[\hat{\Pi}_{t+1}\right] = E_t^{FIRE}\left[\hat{\Pi}_{t+1}\right]$.[5]

The private sector solves a signal extraction problem to estimate λ_t from the observed interest rate setting behaviour of the central bank. The reaction function acts as the observation equation,

$$y_t^{obs} = \Lambda_t \frac{\rho \hat{r}_t^*}{\phi - \rho(1-\Lambda_t)} + \sigma_R \varepsilon_t^R, \tag{7}$$

where $y_t^{obs} = \hat{R}_t - \phi\hat{\Pi}_t$. Kalman filter recursions are used to obtain the private sector agent's date-t estimate of the state variables in period t, given the laws of motion for the state variables, $\hat{r}_t^*$ and λ_t, and the observation equation.[6] Solving the model forward, up to a first-order approximation, the solution for inflation and the nominal interest rate are

$$\hat{\Pi}_t = \frac{\hat{r}_t^*}{\phi}\left[\frac{\phi-\rho}{\phi-\rho(1-\Lambda_t)} + \frac{\rho}{\phi-\rho(1-\Lambda_{t|t})}\right] - \frac{\sigma_R\varepsilon_t^R}{\phi}, \tag{8}$$

and

$$\hat{R}_t = \frac{\phi+\rho\Lambda_{t|t}}{\phi-\rho(1-\Lambda_{t|t})}\hat{r}_t^*, \tag{9}$$

respectively, where $\Lambda_{t|t}$ denotes the logistic function evaluated at $\lambda_{t|t}$. From (9), the private sector's h-step ahead central forecast for the interest rate is

$$\mathbb{E}_t^{PS}\left[\hat{R}_{t+h}\right] = \frac{\phi+\rho\Lambda_{t|t}}{\phi-\rho(1-\Lambda_{t|t})}\rho^h\hat{r}_t^*, \tag{10}$$

which depends on the estimated $\lambda_{t|t}$. The now-cast error $\Lambda_t - \Lambda_{t|t}$ affects interest rate expectations into the infinite future, acting as a wedge between full information and imperfect information an distorting the sequence of expected interest rates.

2.1 Simulation

We simulate 600 periods of data from our model with a burn-in of 100 periods to illustrate how central bank communicating their reaction function can affect interest rate expectations. We assume the following parameterisation for illustrative purposes: $\phi = 1.5$, $\rho = 0.9$, $\sigma_R = 0.01$, $\sigma_\lambda = 0.02$, $\sigma_r = 0.03$.

Panel A in Figure 1 shows the true and estimated responsiveness coefficient. The es-

[5]This abstracts from the central bank's expectations depending on the information set of the private sector, resulting in a simpler fixed-point solution. The model nests the FIRE equilibrium as a special case if either $\sigma_\lambda = 0$ or $\sigma_R = 0$.

[6]Details of the belief updating equations with the extended Kalman filter are in Appendix C.

Figure 1: Tracking Problem and the Effect of Imperfect Information on Expected Interest Rates

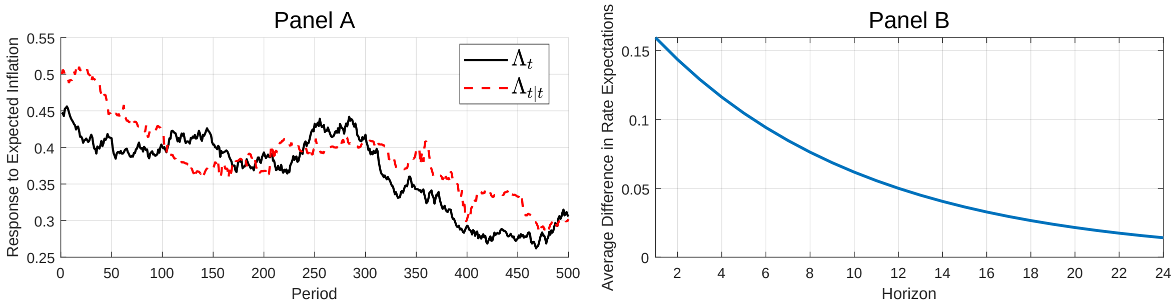

Notes: Panel A plots the true value of Λ_t and the filtered estimate $\Lambda_{t|t}$ from a 600-period simulation, after dropping the first 100 periods. Panel B plots the average absolute percentage difference in nominal interest rate expectations between two cases at each point in the simulation: one that uses the estimated $\lambda_{t|t}$, and another that uses the true λ_t. This is calculated as $100 \times |\mathbb{E}_t[R_{t+h} \mid \lambda_t] - \mathbb{E}_t[R_{t+h} \mid \lambda_{t|t}]|$. This illustrates the average percentage difference in interest rate expectations that would be closed by releasing a precise signal of λ_t.

timated responsiveness generally moves with the true responsiveness, as the Kalman filter ensures $\lambda_t - \lambda_{t|t}$ has a zero mean. Over finite periods of time where the true responsiveness may trend upwards or downwards, the estimated responsiveness can persistently lag behind. For example, over a finite period of time when λ_t is trending downwards, successive revealing signals of λ_t can result in repeated dovish surprises, which is clear from equation (10).

Next we consider the effect of this source of imperfect information on the private sector's expectations of interest rates. At each point in the simulation, we forecast the interest rate 24 periods ahead under two different information sets. One forecast is conditional on the private sector's filtered estimate of $\lambda_{t|t}$, and the other on the true λ_t. We think of the forecast under the true value of λ_t as reflecting a scenario when the central bank perfectly communicates their reaction function. Panel B in Figure 1 plots the average absolute difference in interest rate expectations between these two cases, calculated as the average of $100 \times |\mathbb{E}_t[R_{t+h} \mid \lambda_t] - \mathbb{E}_t[R_{t+h} \mid \lambda_{t|t}]|$ at each $h = 1, ..., 24$. This plots how interest expectations would be revised from a perfectly revealing communication of the reaction function. This effect is strongest for shorter-term forecasts, as ρ^h diminishes the effect of the forecast error $\lambda_t - \lambda_{t|t}$ as h increases.

The central bank communication we focus on in our empirical work is the FOMC minutes, which we interpret as containing signals for the central bank's outlook and reaction function. In practice, the FOMC minutes is not perfectly revealing, but in our model, we assume that the signal noise of the communication is zero. Modifying this to instead be a second noisy signal, slightly complicates the problem without changing the implications. Therefore, the high-frequency market response around the release of the minutes reflects this news. For the Federal Reserve, there are multiple communications before the release of the minutes, such as press conferences and intermeeting speeches. We do not consider sequences of communications within a period in our model.

In practice, central banks communicate about a broad range of topics and each topic may have different implications for market participants over time. This multidimensional nature of central bank communication creates additional layers of complexity in the information transmission process. Market participants must simultaneously process signals about multiple economic variables, potentially with varying degrees of forward-looking orientation and different tones. Our topic-level regressions should therefore be interpreted as a reduced-form decomposition of that generic surprise into observable thematic components, not as a structural estimate of separate policy-rule parameters. While our main hypothesis is still focused on inflation expectations, this broader approach allows us to better understand the different ways in which central bank communication shapes market expectations across the entire yield curve and across different asset classes.

3 Data & Methodology

3.1 Text and Financial Market Data

The text corpus is the FOMC meeting minutes from February 1997 to March 2023, which we obtained from the Fed Board of Governors website. We focus on the FOMC minutes because it is a highly detailed document that outlines the most recent policy meeting and views of the FOMC. Moreover, unlike with the FOMC announcement, the market response to the minutes is entirely due to changes in expectations as it does not accompany a policy change.

We obtain 43,152 sentences from the FOMC minutes over our sample, which we clean with commonly used pre-processing techniques (e.g. stop-word removal and lowercasing), resulting in a vocabulary of 4,049 words.[7] We then use a Structural Topic Model (STM) (Roberts et al., 2013) to measure the different topics covered in the FOMC minutes over time. The financial market data we use are the first four Eurodollar (ED) futures, Treasury

[7] Appendix B contains details of the text pre-processing.

futures (2, 5, and 10 year), and Standard and Poors (S&P500) futures.[8]

We use high-frequency changes in financial market data to identify the overall information content of the minutes. We use a time window of 60 minutes (10 minutes before the release and 50 minutes after) as the minutes are about 3 000 to 5 000 words long each. This aims to strike a balance between the information in the minutes being absorbed into market prices, and before other news events that may potentially contaminating the market response.[9] The minutes are released several weeks after the initial policy meeting, and the response is measured over a tight time window, which means leftover information from past communications is unlikely to have any effect on this market response. For example, if Treasury yields increased around the release of the minutes, we would infer the minutes contained 'hawkish news' about future monetary policy, or 'dovish news' if yields fell.[10] An example of this is illustrated in Figure 2 which shows the response around the release of the FOMC minutes of 1st January, 2009, where the price of 2-, 5- and 10-year Treasuries fell, indicating an increase in expected interest rates along the yield curve.

Figure 2: Market Response to an FOMC minutes Release

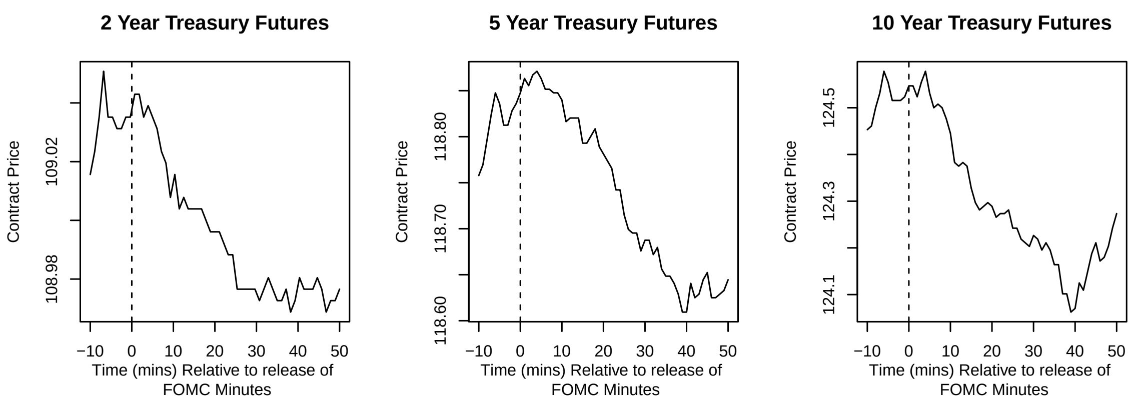

Notes: This figure shows the change in the price of 2-, 5- and 10-year Treasury futures in a 60 minute window around the release of the FOMC minutes on January 6th 2009 at 2pm, Eastern Standard Time.

[8]Our ED and Treasury Futures data go to 2023-04-12, but S&P500 data only goes to 2021-10-13.

[9]The literature on monetary policy surprises often uses a one-hour window when studying the effect of the FOMC minutes on financial markets (Swanson and Jayawickrema, 2023). Our results are robust to both a 30-minute window and 60-minute window, shown in Appendix D.

[10]Referring to the empirically observed market response as 'hawkish' or 'dovish' refers the directional change in interest rate expectations. Whereas in our model, structural hawkishness is the responsiveness to expected inflation, and an upward revision in interest rate expectations could occur from either a higher λ_t when $r_t^* < 0$, or a lower λ_t when $r_t^* > 0$.

3.2 Methodology

In this section, we describe the empirical methodology used to quantify and analyse the impact of forward-looking central bank communication on financial markets. First, we apply a Structural Topic Model (STM) (Roberts et al., 2016) to extract the different topics discussed in the FOMC minutes. Second, we measure forward-looking language within each sentence using tense-based linguistic features. Third, we classify the tone of each sentence in the minutes as either hawkish, dovish or neutral. Finally, we combine these data to construct tense oriented hawkish-dovish topic shares for each FOMC minutes to see their effect on market expectations in an event study analysis.

3.2.1 Measuring Tense

We construct an indicator for whether a sentence contains forward-looking language with a computational linguistic tool and a dictionary approach. The computational linguistic tool we use is the Tense-Mood-Voice (TMV) classification tool of Ramm et al. (2017).[11] TMV is a rule-based linguistic classification of verb phrases to detect the temporal orientation of a sentence. Table 1 illustrates the process for how TMV determines the tense of three sentences. TMV first detects the verbal complex of the sentence, then, through a series of linguistic rules, determines the tense. The tense column classification has the following interpretation: "past" denotes a verb anchored in past time (e.g., projected), "condI" (conditional I) marks modal constructions typically with would expressing hypothetical or projected situations (e.g., would slow), and "future" indicates future reference, usually formed with will + verb (e.g., will decline).

[11] This classification method is discussed in the context of evaluating the text of policy documents in Byrne et al. (2023*a*) and applied to aspects of central bank communications in Byrne et al. (2023*b*).

Table 1: Verbal Complex and Tense Annotation

Sentence	Verbal Complex	Tense
Exports from some U.S. manufacturing industries **will decline** due to softer foreign markets and import competition.	will decline	futureI
The staff **projected** that the expansion of economic activity **would slow** for a time to a pace appreciably below the estimated growth of the economy's potential and then **would pick up** to a rate more in line with that potential.	projected	past
	would slow	condI
	would pick up	condI
With regard to the outlook for inflation, members **referred** to widespread indications of increasingly tight labor markets and to statistical and anecdotal reports of faster increases in labor compensation.	referred	past

Notes: This table illustrates the TMV annotation process for three sentences found in FOMC meeting minutes. Verbal complexes are first identified within each sentence, after which their tense is assigned by applying a set of rule-based linguistic patterns. Where a tense classification of a verbal complex is made, this is reported in the final column, but tense is not always assigned to every verbal complex.

The "futureI" classification appears in direct forward-looking language, such as the first sentence of Table 1. It is more common for the FOMC minutes to discuss potential economic developments or policy responses, which would likely recieve a tense classification of "condI", such as sentence two in Table 1. We classify a sentence as forward-looking if any verbal complex in a sentence has a "futureI", "futureII" or "condI" tense label applied by TMV. We further refine the future tense classification with a second-layer dictionary approach, where if a sentence contains a word in our dictionary, we also classify it as forward-looking. This second-layer is often redundant, but since the minutes are written to refer to a meeting that took place a few weeks previously, the past tense is often used when discussing the future, which can be missed by TMV alone. For example, the third sentence in Table 1 has the verbal complex "referred" classified with a tense of "past", despite the sentence discussing the economic outlook. TMV does not classify ths sentence as forward-looking, but our second-layer dictionary does. Our final forward-looking classification of a sentence is the union of the TMV and dictionary layer classifications.

We evaluate the performance of our classifier by comparing its output with 200 randomly selected sentences that were manually labelled as forward-looking or not.[12] We report some commonly used statistics in applied NLP to evaluate the accuracy of a classification task. Precision refers to the proportion of sentences classified as forward-looking that were correctly identified, scoring 73%. Recall measures the proportion of all truly forward-looking sentences that were successfully detected, scoring 77%. The harmonic mean of precision and recall is called the F1 score, which is 0.75. An F1 score of 0.75 reflects good performance for a binary text classification, giving us confidence in our measure of forward-looking language.

Figure 3: Tense and Tone of FOMC Minutes

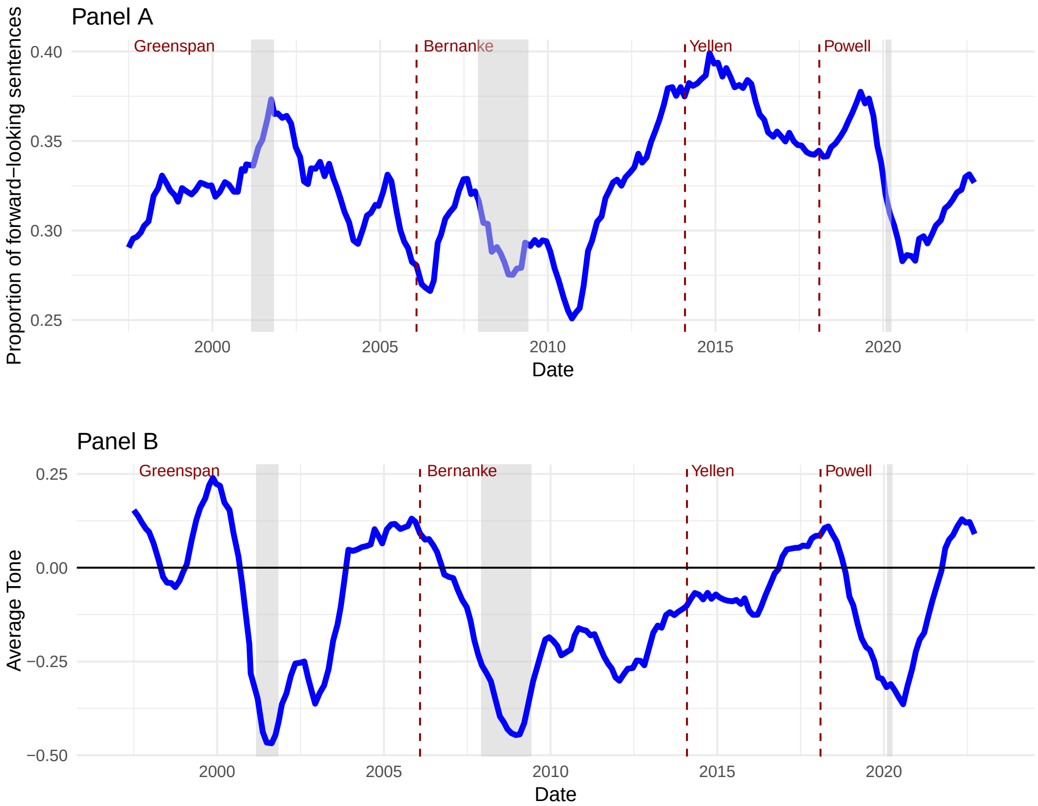

Notes: Panel A plots the proportion of forward-looking sentences in each minutes. Panel B plots the average hawkish-dovish tone of sentences in each FOMC meeting minutes over time. Changes in the Federal Reserve Chair are indicated with dashed vertical lines, and shaded NBER recession bands. Both series are smoothed using an 8-period centred moving average, corresponding to approximately 52 weeks.

Panel A in Figure 3 plots the proportion of forward-looking sentences contained in each FOMC minutes over time from our classification. These changes reflect the strategic choice of forward-looking communication. The FOMC officially began using forward guidance in its post-meeting statements at the turn of the millennium, but this shows forward-looking

[12] Ramm et al. (2017) report an 82% accuracy for tense classification with TMV on general text data. In our particular context of technical central bank text, the classification task is more complicated.

language had been meaningfully present in communications before. After the financial crisis and from the early 2010s, the proportion of forward-looking sentences increased considerably, as there was heightened focus on when interest rates would start to 'lift-off' from the zero-lower bound. At the onset of the COVID-19 pandemic, forward-looking language use fell sharply to its lowest levels in a decade, potentially reflecting the severity of the pandemic and increased emphasis on the present versus the future conditions.

3.2.2 Measuring Tone

To measure the tone of sentences in the minutes, we use a fine-tuned RoBERTa model from Shah et al. (2025). They fine-tune a large language model with sentences of minutes from central banks around the world, manually classified sentences as either "hawkish", "dovish" or "neutral", abstracting from how hawkish or dovish a sentence is. We numerically define the tone of each sentence as -1 for dovish, 0 for neutral and 1 for hawkish. Their model results in an additional classification of "not related to monetary policy", which detects sentences such as procedural text or lists of names; we drop all 2,197 sentences with this classification.

Panel B in Figure 3 plots the average tone of each FOMC minutes over time. A more positive tone indicates a more contractionary stance, whereas a negative one indicates a more expansionary stance. As expected from the tone of the minutes, they tend to be more dovish (negative tone) on average during economic downturns, such as the dot-com bubble, the Global Financial Crisis and the COVID-19 pandemic.

3.2.3 Measuring Topics

We use a Structural Topic Model (Roberts et al., 2013) to construct topic shares. The algorithm models the frequency of words by first defining a data generation process for each document[13], then determines the most likely parameter values to explain the word frequencies. In this framework, a topic constitutes a mixture of words, where each word has a probability of belonging to a topic, which in other words means that topics are clusters of words that commonly co-occur across documents. Each document represents a mixture of topics, meaning that a single document can belong to multiple topics. This ensures that

[13]Standard computational linguistics terminology refers to "documents" as the fundamental unit of observation being fed to the algorithm. In our analysis, we consider individual sentences as documents.

the sum of the topic proportions across all topics for a document is one, and the sum of the word probabilities for a given topic is one.

An advantage of STM over topic models such as Latent Semantic Analysis (LSA) (Landauer and Dumais, 1997) or Latent Dirichlet Allocation (LDA) (Blei et al., 2003) is its ability to use variation from additional variables when generating topics. These variables provide additional structure in the formation of the topics that might inform the language used.

Specifically, STM incorporates "prevalence" covariates that model how the proportion of each topic varies across documents as a function of observed metadata (e.g., time period, author demographic, or document type). Unlike LDA, which assumes topic proportions are drawn from a symmetric Dirichlet prior without any dependence on document characteristics, STM uses a logistic-normal prior where the mean of the topic proportion distribution is a linear function of covariates. This allows the model to explain why some topics appear more frequently in certain contexts and to quantify the effect of specific variables on topic prevalence.

Formally, the generative process for each document d with vocabulary size V follows these steps:

1. Draw the share of a document belonging to each topic from a conditional distribution based on document covariates X_d:

$$\vec{\theta}_d \mid X_d\gamma, \Sigma \sim \text{LogisticNormal}\,(\mu = X_d\gamma, \Sigma) \tag{11}$$

   where X_d is a 1-by-p vector of prevalence covariates, γ is a p-by-$(K-1)$ matrix of coefficients and Σ is a $(K-1)$-by-$(K-1)$ covariance matrix. The covariates dated at the time of the FOMC meeting that we use are: the output gap (defined as the 100 times the log difference between actual and potential GDP from the Congressional Budget Office), the unemployment rate, the PCE inflation rate, the effective federal funds rate, and Fed Chair fixed effects.

2. For each topic k, form the global distribution over words (shared across all documents) using:

$$\beta_k \propto \exp(m + \kappa_k^t) \tag{12}$$

where m and κ_k^t are V-length vectors containing one entry per word in the vocabulary. These components represent the baseline word distribution m and the topic-specific deviation κ_k^t.

3. For each word in the document ($n \in \{1, \ldots, N_d\}$):

   - Draw word's topic assignment:

$$z_{d,n} \mid \vec{\theta}_d \sim \text{Multinomial}(\vec{\theta}_d) \tag{13}$$

   - Draw observed word from that topic:

$$w_{d,n} \mid z_{d,n}, \beta_{k=z_{d,n}} \sim \text{Multinomial}(\beta_{k=z_{d,n}}) \tag{14}$$

We initialize the model with LDA.[14] The number of topics, K, is chosen to provides a coherent yet concise representation of the corpus. We select K by estimating a set of models over $K = 2, ..., 60$ topics and evaluate them based on their semantic coherence and exclusivity. Semantic coherence (Mimno et al., 2011) measures how often a topic's most probable words co-occur together. The criterion is maximized by fewer topics dominated by very common words (Roberts et al., 2014). Conversely, exclusivity uses the FREX metric (Airoldi and Bischof, 2016), which measures the weighted harmonic mean of a word's exclusivity and frequency ranks. This metric favours more topics, identifying words frequently used when discussing a topic yet relatively unique to it.

[14] See Roberts et al. (2013) for details and a comparison of different approaches and Roberts et al. (2016) for more details on the estimation process.

Figure 4: Number of Topics Frontier

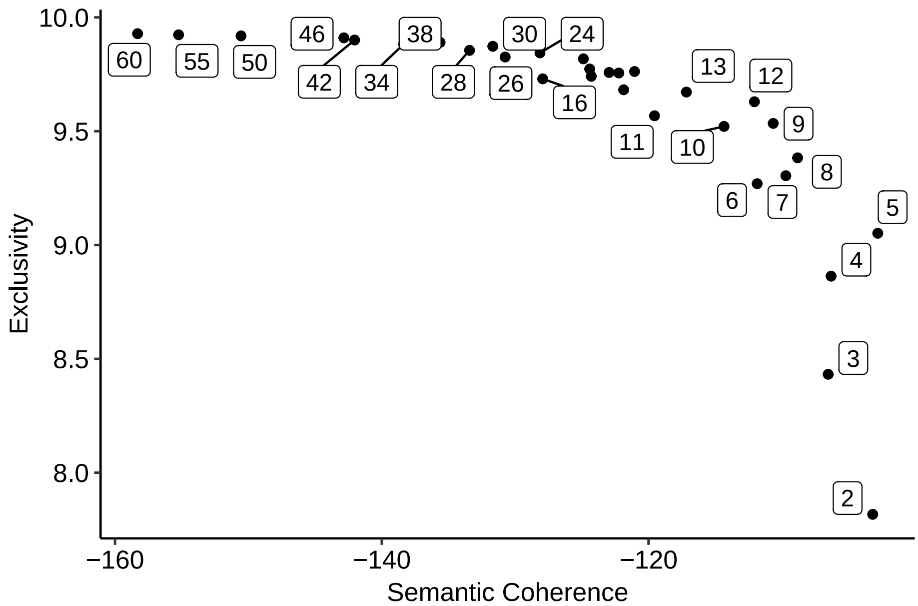

Notes: This figure presents the exclusivity and semantic coherence score for different estimated STMs, where the number of topics used in the estimation, K labels each point. The figure shows an approximate "frontier" for a trade-off between topic exclusivity and semantic coherence over different values of K.

Figure 4 plots the coherence and exclusivity score pairs for each K, which traces out a trade-off between semantic coherence and exclusivity. A model with a higher score in both coherence and exclusivity would dominate a model with lower scores in both measure. This analysis helps to narrow the choice set of K and informs judgement on how many topics to include. We choose to prioritise semantic coherence, resulting in a lower number of topics that are more interpretable as our interest lies in whether the FOMC discusses a certain broad topic rather than a specific aspect of it, and select the 12-topic model. A further reason for us prioritising a lower value of K is that more topics reduce the degrees of freedom in our already small sample size downstream regressions. We manually label each of the 12 topics based on its most representative words and sentences. Word clouds are a way to visualise the high-probability words of each topic, and are reported in Tables A1 and A2. Figure 5 shows word clouds for our 12-topic model with the 30 most likely words in each topic and our assigned label underneath each cloud.

Figure 5: Topic Word Clouds

Spending | Credit | Interest Rate | Inflation

Open. Mkt. Operations | Uncertainty | Supply Chain | Fin. Stability

Construction | Growth | Ind. Prod. | Fin. Mkts.

Notes: Word clouds for the twelve topics identified through textual analysis of FOMC meeting minutes. Larger words indicate higher frequency within each topic. Topics reflect key themes in Federal Reserve communication, including spending, growth, policy, markets, and global risks.

Figure 6 plots smoothed estimates over time for our 12 topics. The figure shows some topics have become increasingly mentioned in the FOMC minutes over time, such as uncertainty or interest rates, whereas other have become less prominent such as uncertainty or industrial production. These trends highlight broad changes in the topics of discussion in the FOMC that shape the policy stance.

Figure 6: Temporal Evolution of Topics

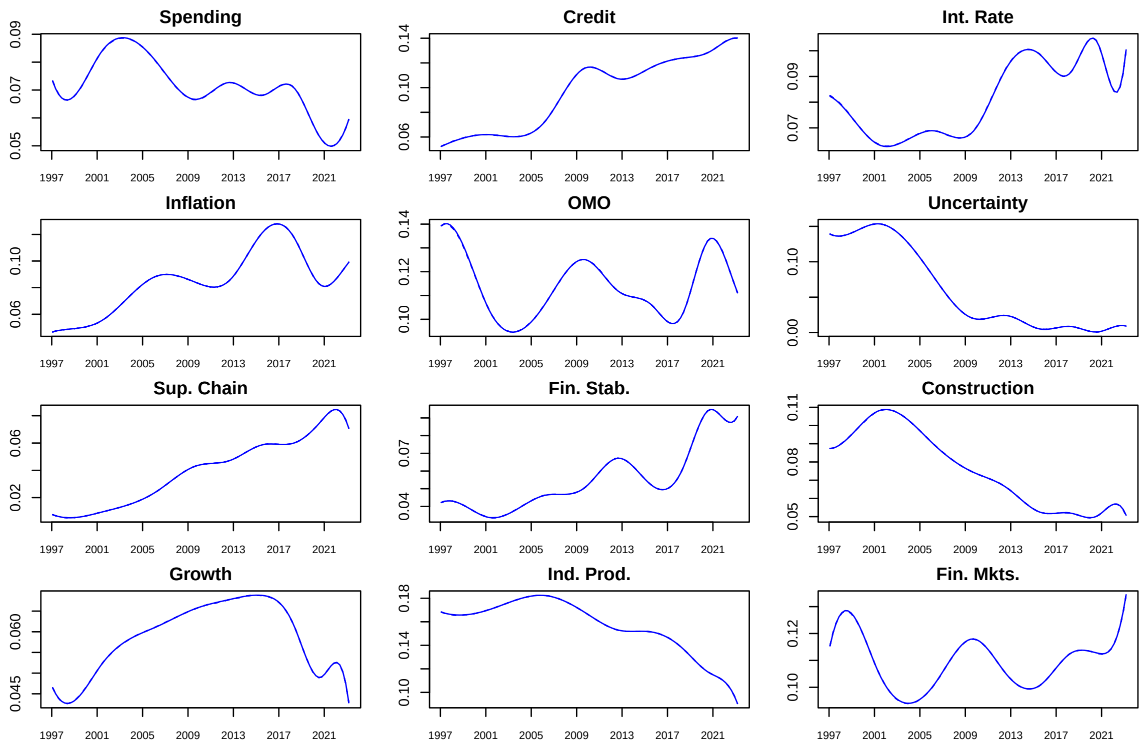

Notes: This figure shows the smoothed estimated topic shares over time from the STM. The vertical axis represents the share of the FOMC minutes at each date (horizontal axis) that was about that particular topic.

3.2.4 Combining Tense, Tone and Topic

Given our measures of tense, tone and topic, we measure the textual information of each FOMC minutes with future-tone weighted topic shares. Topic modelling does not guarantee an expectations topic for a dimension of interest, like inflation, will be obtained. Our approach takes those sentences that are forward-looking and then infer the topic shares of those sentences as forward-looking topic shares.[15].

Formally, let $\hat{\theta}_{m,j,k}$ denote the estimated share of topic k in sentence j in FOMC minutes m, and the dovish-neutral-hawkish tone score is given by $T_{m,j} \in \{-1, 0, 1\}$. Then, the "tone weighted share" of topic k in sentence j for FOMC minutes m is

$$TWS_{m,j,k} = T_{m,j} \cdot \hat{\theta}_{m,j,k}. \tag{15}$$

We denote the future-tense indicator for sentence j in FOMC minutes m with $F_{m,j}$ and for

[15]A smaller number of topics in our STM that results in more general topics helps to more accurately select the forward-looking content from these shares, by reducing dispersion of the content of interest among a larger set of topics. On the other hand, a too small number risks compressing different topics together.

each sentence, we also calculate "future tone weighted shares" of each topic,

$$FTWS_{m,j,k} = F_{m,j} \cdot T_{m,j} \cdot \hat{\theta}_{m,j,k}. \tag{16}$$

We average both $TWS_{m,j,k}$ and $FTWS_{m,j,k}$ at the FOMC minutes level, m to obtain $FTWS_{m,k}$ and $TWS_{m,k}$ for each topic, which lie in the interval $[-1, 1]$. A value of $FTWS_{m,k}$ closer to one corresponds to a more hawkish forward-looking communication of topic k in FOMC minutes m. Similarly, a value closer to -1 reflects a more dovish forward-looking communication about topic k. We average our sentence-level measure to the meeting level as we can only observe how financial markets reacted to the release of the minutes as a whole. For interpretability in our regressions, we standardise the topic shares such that a one-unit increase in $FTWS_{m,k}$ corresponds to a one-standard deviation increase.

Tense weighting allows us to measure the extent to which each topic appears in a forward-looking context, following Byrne et al. (2023*a*). Adding tone weighting allows us to capture the overall information about the policy stance within each forward-looking topic. If the FOMC minutes communicate hawkish information about future inflation, we ask whether this information surprises financial market participants. Topics that systematically surprise investors over certain periods indicate dimensions of policy that the FOMC did not fully convey to markets prior to the release of the minutes. In the context of our model, and related to Bauer and Swanson (2023), if the responsiveness of certain dimensions of monetary policy has trended over some periods, rational market participants may be repeatedly surprised by successive communications as they gradually learn the true policy rule parameter.[16] The series for $FTWS_{m,k}$ and $TWS_{m,k}$ are shown in Appendix E.

3.3 Event Study Analysis

We use an event study analysis to identify the effect of forward-looking central communications on financial markets. Our aim is twofold: to investigate whether certain dimensions of the FOMC minutes have been informative to financial markets, and to determine whether forward-looking information has an effect beyond that of contemporaneous or backward-

[16]A related, but more general, argument is made in Farmer et al. (2024), who show that imperfect information about the structure of the data-generating process can lead to persistent forecast errors.

looking information in the communication. We therefore estimate regressions of the form

$$\Delta y_m = \alpha + \sum_{k=1}^{12} \beta_k \cdot TWS_{m,k} + \sum_{k=1}^{12} \gamma_k \cdot FTWS_{m,k} + \varepsilon_m, \tag{17}$$

where Δy_m represents the high frequency market response of a financial market variable such as interest rates or stock prices around the release of FOMC minutes m.[17] The market responds to the release of the minutes as a whole, so we include the decomposition of the minutes into all topics estimated in Section 3.2.3 as explanatory variables. Our main coefficients of interests are the γ_k, which capture whether the forward-looking information in the minutes contain news beyond the tone-weighted shares. The dependent variables, Δy_m are the high-frequency market responses in a 60 minute window around the FOMC minutes release.

We study the effects of the FOMC minutes on the S&P500 index, and yields of two-year, five-year and ten-year Treasury futures. In addition, we look at the response to interest rate expectations over a horizon of a year through Eurodollar futures contracts that expire in the current-quarter (ED1), 1-quarter ahead (ED2), 2-quarters ahead (ED3) and 3-quarters ahead (ED4). Responses to Eurodollar futures and Treasuries are measured in percentage points (yields). For the S&P500, we measure the market response in percentage changes. Figure D1 shows histograms showing the range for these market responses to the release of the FOMC minutes. Typical responses (2 standard deviations) for the ED responses are around one to two basis points, around one basis point for the Treasury yields, and 0.1% for the SP500 response.

With topic modelling, the obtained topic shares are estimates of the 'true' topic share. Estimating the share of latent topics is not an "end in itself, but rather a first step in a larger econometric strategy" (Battaglia et al., 2024). Using the mean of $\hat{\theta}_{m,j,k}$ in a downstream econometric estimation does not take into account the uncertainty of this constructed regressor. We account for this measurement error in the topic shares by sampling from the posterior distribution of $\hat{\theta}_{m,j,k}$ at the sentence level. For each draw, d, we obtain $\hat{\theta}^{(d)}_{m,j,k}$ and construct the sentence level variables $TWS^{(d)}_{m,j,k}$ and $FTWS^{(d)}_{m,j,k}$, averaged at the document-level to estimate equation (17). We collect the the point estimates, $\{\hat{\alpha}^{(d)}, \hat{\beta}^{(d)}, \hat{\gamma}^{(d)}\}$ over each draw and construct quantiles for these estimates around the mean value of each coefficient

[17]Event studies of this form a popular method for studying how information affects market expectations since seminal papers such as Kuttner (2001) or Gurkaynak et al. (2005).

in our main result tables.[18]

4 Results

4.1 Asset Price Responses to the FOMC Minutes

Table 2 shows the effect on 2-, 5- and 10- year Treasury yields to the future-tone-weighted topic shares over our full sample, February 1997 to March 2023. The results highlight some specific dimensions of forward-looking communication have been informative to financial markets. In particular, more hawkish future inflation language is associated with investors revising upward their beliefs of future interest rates, with a diminishing effect along the yield curve, as predicted in our model. The effect is strongest on 2-year yields, with a one-standard deviation increasing the future-tone-weighted inflation topic raising the two-year yield 0.8 basis points. The effect diminishes to half a basis point at 10-year yields. The positive coefficients imply that markets tended to under-predict the hawkish future inflation content of the minutes on average over the sample. The forward-looking open market operations (OMO) content of the minutes had a stronger effect at longer maturity, 5- and 10- year Treasuries than shorter-term interest rate expectations. As this topic relates to asset purchases, it is plausible that news surrounding the unconventional large-scale asset purchases over this period, which primarily focused on longer-maturity bonds, came across through the minutes. From one FOMC minutes to the next, the variation in the language is typically not large (as seen in Figure 3), but over time or in periods of rapid change, our results show a small but not negligible effect of these learning dynamics on market prices.

Our theoretical framework showed that imperfect information about the forward-looking term in the policy function leads to communications generating surprises in interest rate expectations. An upward revision in expected nominal rates is consistent with learning about a more reactionary central bank when the natural rate is negative, or a less reactionary central bank when the natural rate is positive. Yet, our reduced-form NLP framework measures hawkishness or dovishness as information indicative of a tighter or looser nominal policy stance, with the high-frequency market responses factoring in any news regarding the state of the economy. The empirical notion of hawkishness therefore captures both effects, confirming that markets update their beliefs in response to the varied signals in the minutes.

[18]The STM posterior is non-conjugate and asymmetric in general which motivates our choice of using quantiles for constructing confidence intervals and reporting significance level for our statistics of interest.

Table 2: Asset Responses to Forward-Looking Communications

	2Y Treasuries	5Y Treasuries	10Y Treasuries	S&P 500
Inflation X F	0.008***	0.006***	0.005***	0.025
	[0.003, 0.012]	[0.002, 0.011]	[0.002, 0.009]	[-0.074, 0.128]
Spending X F	0.003	0.002	0.002	-0.102
	[-0.004, 0.010]	[-0.005, 0.008]	[-0.003, 0.007]	[-0.242, 0.031]
Credit X F	0.001	0.001	0	-0.064
	[-0.003, 0.004]	[-0.003, 0.004]	[-0.003, 0.003]	[-0.172, 0.039]
Int. Rate X F	-0.005	-0.006*	-0.007**	0.116*
	[-0.012, 0.002]	[-0.013, 0.001]	[-0.013, -0.001]	[-0.042, 0.277]
OMO X F	0.004*	0.006***	0.005***	0.025
	[-0.000, 0.009]	[0.002, 0.010]	[0.002, 0.009]	[-0.079, 0.128]
Uncertainty X F	-0.002	0.006	0.001	-0.135
	[-0.012, 0.008]	[-0.003, 0.015]	[-0.006, 0.009]	[-0.305, 0.033]
Supply Chain X F	-0.001	-0.003	-0.003	0.089
	[-0.009, 0.007]	[-0.012, 0.006]	[-0.010, 0.004]	[-0.155, 0.331]
Fin. Stability X F	0.005	0.007	0.007*	0.128
	[-0.006, 0.017]	[-0.004, 0.018]	[-0.002, 0.017]	[-0.157, 0.417]
Construction X F	0	-0.001	0	0.05
	[-0.007, 0.006]	[-0.007, 0.005]	[-0.005, 0.005]	[-0.067, 0.168]
Growth X F	-0.001	-0.002	-0.001	0.06
	[-0.009, 0.008]	[-0.010, 0.006]	[-0.008, 0.005]	[-0.120, 0.238]
Ind. Prod X F	-0.002	0	-0.001	0.03
	[-0.007, 0.003]	[-0.005, 0.004]	[-0.005, 0.003]	[-0.070, 0.131]
Fin. Mkts. X F	-0.006**	-0.004*	-0.003*	0.045
	[-0.011, -0.001]	[-0.009, 0.000]	[-0.007, 0.000]	[-0.067, 0.158]
TWS Included?	Yes	Yes	Yes	Yes
Observations	210	210	210	198
R-squared	0.043	0.059	0.058	0.177

Notes: This table shows the estimated relationship between Treasury futures yields and the forward-looking tone-weighted topic shares; the tone-weighted topic shares are included as controls but not reported here. The S&P 500 response is measured in percentage point changes, and the response for Treasuries is measured in percentage points. Simulated quantiles are used to calculate the 95% confidence intervals in square brackets, and significance stars reflect * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

The response to the future interest rate topic is negative along the yield curve on average, with a similar magnitude as the inflation topic. This implies that markets had been overly hawkish with respect to the overall path of interest rates up until the release of the FOMC minutes, as information in the minutes regarding future interest rates had the effect of investors revising downwards their beliefs of future interest rates along the yield curve. The effect is strongest at the front-end of the yield curve in terms of magnitude and significance, indicating less news regarding longer-term interest rates on average. In the context of our learning model, this would suggest that periods with the most hawkish language were also periods where the market's ex-ante overestimation of the rate path was at its most extreme. As a result, higher levels of realized hawkishness correspond to larger dovish expectation shortfalls, resulting in the negative coefficient estimates.

Table 3 shows the market response for Eurodollar futures contracts up to one-year ahead.

We also include MPS which is the first principal component of the change in the first four Eurodollar futures, reflecting the overall change in short-term interest rates up to one-year ahead. We scale MPS such that a one-unit increase corresponds to a one-percentage point change in the ED4 rate on average in sample. The results are similar to the Treasury response, with forward-looking inflation having the same qualitative implications for belief revisions of investors, however increasing with (relatively) longer horizons.

Table 3: Asset Responses to Forward-Looking Communications

	ED1	ED2	ED3	ED4	MPS
Inflation X F	0	0.003*	0.004*	0.004*	0.003*
	[-0.002, 0.002]	[-0.001, 0.006]	[-0.000, 0.009]	[-0.001, 0.009]	[-0.001, 0.007]
Spending X F	-0.001	-0.001	0	0	-0.001
	[-0.004, 0.001]	[-0.006, 0.005]	[-0.006, 0.007]	[-0.007, 0.007]	[-0.007, 0.005]
Credit X F	0.002**	0.003*	0.001	0.002	0.003*
	[0.000, 0.003]	[-0.000, 0.006]	[-0.002, 0.005]	[-0.002, 0.005]	[-0.001, 0.006]
Int. Rate X F	-0.002	-0.006**	-0.007*	-0.011***	-0.007**
	[-0.005, 0.001]	[-0.012, -0.000]	[-0.014, 0.000]	[-0.018, -0.003]	[-0.013, -0.001]
OMO X F	0.002*	0.002	0.003*	0.005**	0.003*
	[-0.000, 0.003]	[-0.002, 0.005]	[-0.001, 0.008]	[0.000, 0.009]	[-0.001, 0.007]
Uncertainty X F	0	0.007*	0.004	0.005	0.005
	[-0.004, 0.004]	[-0.001, 0.015]	[-0.006, 0.014]	[-0.005, 0.016]	[-0.004, 0.013]
Supply Chain X F	0	0.003	0.001	0.002	0.002
	[-0.003, 0.004]	[-0.004, 0.010]	[-0.007, 0.009]	[-0.006, 0.011]	[-0.005, 0.009]
Fin. Stability X F	0.002	0.002	0.004	0.004	0.004
	[-0.002, 0.007]	[-0.006, 0.011]	[-0.007, 0.015]	[-0.008, 0.016]	[-0.006, 0.014]
Construction X F	0	-0.003	-0.001	0.001	-0.001
	[-0.003, 0.002]	[-0.008, 0.002]	[-0.007, 0.005]	[-0.006, 0.007]	[-0.007, 0.004]
Growth X F	0	-0.005	-0.003	-0.003	-0.003
	[-0.003, 0.004]	[-0.011, 0.002]	[-0.011, 0.005]	[-0.012, 0.006]	[-0.010, 0.004]
Ind. Prod X F	-0.001	0	0	0	0
	[-0.002, 0.001]	[-0.004, 0.004]	[-0.005, 0.004]	[-0.005, 0.005]	[-0.005, 0.004]
Fin. Mkts. X F	-0.002*	-0.004**	-0.004*	-0.005**	-0.004**
	[-0.004, 0.000]	[-0.008, -0.000]	[-0.008, 0.000]	[-0.010, -0.000]	[-0.008, -0.000]
TWS Included?	Yes	Yes	Yes	Yes	Yes
Observations	210	210	210	210	210
R-squared	0.097	0.074	0.047	0.078	0.067

Notes: This table presents estimated responses of Eurodollar futures to topic-specific communication variables; the tone-weighted topic shares are included as controls but not reported here. MPS is the first principal component of the change in the first four Eurodollar futures contracts, measured in ED4 units. Simulated 95% confidence intervals reported in square brackets, and significance stars are calculated from the simulated quantiles intervals reflecting, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

One concern relates to the large number of covariates that are likely correlated in this context, and a fairly small sample size. To address this, we estimate an additional specification using the full topic model but restricting attention to the three topics that are typically the most relevant for central banks: growth, inflation, and interest rates. The results from this parsimonious specification are qualitatively and quantitatively similar to Tables 2 and 3, implying that the findings are not spuriously driven by the dimensionality of the topic space.

4.2 Responses Over Time

Motivated from the learning dynamics of our theoretical framework, in finite samples, markets would likely react more strongly to new information when there is greater uncertainty in the current state estimate.[19] In this section, we ask how much the market responses have changed over time by estimating rolling window regressions for 125 events at a time, using the simulation method before with 1 000 draws for each window.[20] We estimate (17) with all topics starting in February 1997, with the first window ending in August 2012.

Figure 7: Temporal Evolution of Future Topic Coefficients Across Yields

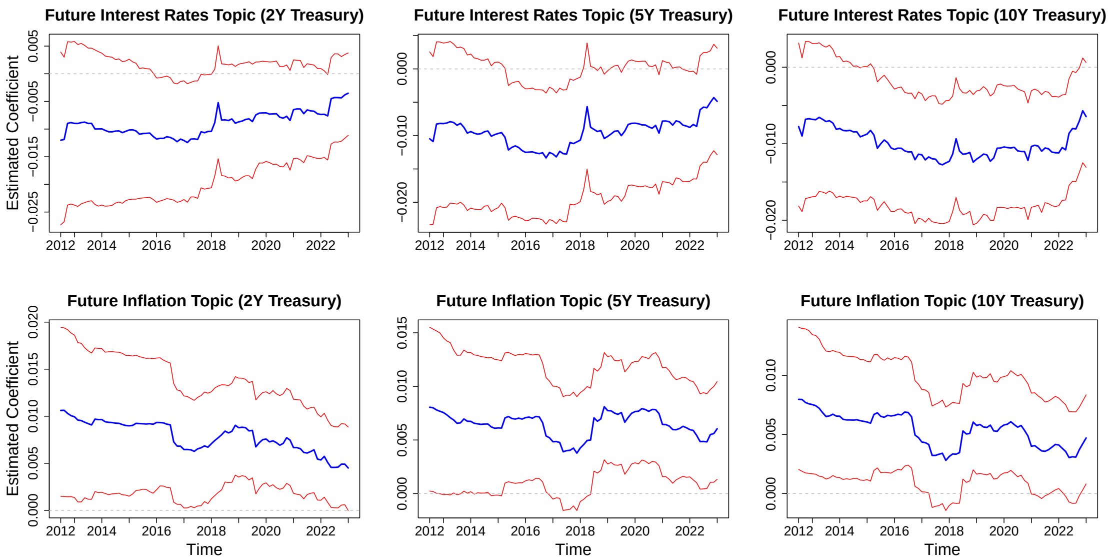

Notes: This figure presents select coefficient estimates from rolling window estimations using 125 events in each window. Each row corresponds to a different future-topic, and the columns to different asset responses. The vertical axis represents the estimated coefficient, and the bands reflect 90% confidence intervals obtained with 1 000 draws for each window.

Figure 7 plots these dynamic coefficient estimates. The first row shows the estimated coefficients for the future interest rate topic, and the second row shows the estimated coefficients for the future inflation topic; the first, second, and third columns correspond to the 2-, 5-, and 10-year Treasury yield responses, respectively. The first row demonstrates that the dovish market reactions to the future interest rate topic became more pronounced over

[19]When the variance of the state estimate is low, the market response to new information (the Kalman gain) would also be low.

[20]The use of 1 000 draws per window instead of 10 000 because of the additional computational cost of the rolling window regressions.

the zero lower bound period, particularly at the longer end of the yield curve. These results indicate that during this period, markets repeatedly overestimated how soon the Fed would 'lift off' interest rates, maintaining an overly hawkish prior belief. Consequently, when the minutes dedicated forward-looking language to the interest rate path, the release systematically adjusted the market's premature lift-off expectations downward. This finite-sample learning dynamic perfectly mirrors the broad mechanism illustrated in our theoretical model, confirming that finite-sample trends in the policy function can persistently surprise markets in a single direction.

The second row shows that the market response to hawkish future inflation has remained relatively stable over time along the yield curve. There is a slight downward movement toward the end of the sample during the pandemic, a period when market focus likely shifted to concerns about nearer-term inflation. Nonetheless, the persistently positive coefficients demonstrate that forward-looking inflation language has reliably and repeatedly informed financial market pricing throughout the sample.

5 Conclusions

In this paper, we show that financial markets respond to forward-looking language in Federal Open Market Committee (FOMC) communications. By combining structural topic modelling with tone and tense analysis, we document that the information conveyed in the FOMC minutes varies substantially across topics and temporal orientations. In particular, we find that markets revise upward their expectations of future interest rates when the minutes contain more hawkish forward-looking discussions of inflation.

We also find that forward-looking discussions of interest rates were, on average, more dovish than markets expected during the zero lower bound period. More generally, our results show that certain dimensions of central bank communication have systematically surprised markets over time. We interpret these findings through an imperfect information model in which market participants learn about the central bank's responsiveness to its economic outlook in a Bayesian manner. In this framework, communication affects expectations because it reveals information about the central bank's reaction function that is not directly observable to the private sector.

Our findings suggest that the temporal dimension of monetary policy communication is a key component of market interpretation, operating alongside the topical content of the message itself. More broadly, the results highlight the importance of studying not only what central banks communicate, but also how they frame that communication over time.

References

Airoldi, E. M. and Bischof, J. M. (2016), 'Improving and evaluating topic models and other models of text', *Journal of the American Statistical Association* **111**(516), 1381–1403.

Battaglia, L., Christensen, T., Hansen, S. and Sacher, S. (2024), 'Inference for regression with variables generated from unstructured data', *arXiv preprint arXiv:2402.15585* .

Bauer, M. D. and Swanson, E. T. (2022), A reassessment of monetary policy surprises and high-frequency identification, Technical report, National Bureau of Economic Research.

Bauer, M. D. and Swanson, E. T. (2023), 'A reassessment of monetary policy surprises and high-frequency identification', *NBER Macroeconomics Annual* **37**(1), 87–155.

Bennani, H., Fanta, N., Gertler, P. and Horvath, R. (2020), 'Does central bank communication signal future monetary policy in a (post)-crisis era? the case of the ecb', *Journal of International Money and Finance* **104**, 102167.

Blanchard, O. J., L'Huillier, J.-P. and Lorenzoni, G. (2013), 'News, noise, and fluctuations: An empirical exploration', *American Economic Review* **103**(7), 3045–70.

Blei, D. M., Ng, A. Y. and Jordan, M. I. (2003), 'Latent dirichlet allocation', *Journal of machine Learning research* **3**(Jan), 993–1022.

Blinder, A. S. (2018), Through a crystal ball darkly: The future of monetary policy communication, *in* 'AEA Papers and Proceedings', Vol. 108, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, pp. 567–571.

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

Byrne, D., Goodhead, R., McMahon, M. and Parle, C. (2023*a*), The central bank crystal ball: Temporal information in monetary policy communication, Technical report, Central Bank of Ireland.

Byrne, D., Goodhead, R., McMahon, M. and Parle, C. (2023*b*), Measuring the temporal dimension of text: An application to policymaker speeches, Research technical papers, Central Bank of Ireland.

Clarida, R., Gali, J. and Gertler, M. (1999), 'The science of monetary policy: a new keynesian perspective', *Journal of economic literature* **37**(4), 1661–1707.

Clarida, R., Gali, J. and Gertler, M. (2000), 'Monetary policy rules and macroeconomic stability: evidence and some theory', *The Quarterly journal of economics* **115**(1), 147–180.

Del Negro, M. and Eusepi, S. (2011), 'Fitting observed inflation expectations', *Journal of Economic Dynamics and Control* **35**(12), 2105–2131.
URL: *https://ideas.repec.org/a/eee/dyncon/v35y2011i12p2105-2131.html*

Farmer, L. E., Nakamura, E. and Steinsson, J. (2024), 'Learning about the long run', *Journal of Political Economy* **132**(10), 3334—-3377.

Gurkaynak, R. S., Sack, B. and Swanson, E. T. (2005), 'Do actions speak louder than words? the response of asset prices to monetary policy actions and statements'.

Handlan, A. (2022), 'Text shocks and monetary surprises: Text analysis of fomc statements with machine learning'.

Hansen, S. and McMahon, M. (2016), 'Shocking language: Understanding the macroeconomic effects of central bank communication', *Journal of International Economics* **99**, S114–S133.

Hansen, S., McMahon, M. and Prat, A. (2018), 'Transparency and Deliberation Within the FOMC: A Computational Linguistics Approach*', *The Quarterly Journal of Economics* **133**(2), 801–870.

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.

Landauer, T. K. and Dumais, S. T. (1997), 'A solution to plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge.', *Psychological review* **104**(2), 211.

Larsen, V. H., Thorsrud, L. A. and Zhulanova, J. (2021), 'News-driven inflation expectations and information rigidities', *Journal of Monetary Economics* **117**, 507–520.

Leombroni, M., Vedolin, A., Venter, G. and Whelan, P. (2021), 'Central bank communication and the yield curve', *Journal of Financial Economics* **141**(3), 860–880.

Levine, P., Pearlman, J., Perendia, G. and Yang, B. (2012), 'Endogenous Persistence in an estimated DSGE Model Under Imperfect Information', *Economic Journal* **122**(565), 1287–1312.

Mcauliffe, J. and Blei, D. (2007), 'Supervised topic models', *Advances in neural information processing systems* **20**.

Mimno, D., Wallach, H., Talley, E., Leenders, M. and McCallum, A. (2011), Optimizing semantic coherence in topic models, *in* 'Proceedings of the 2011 conference on empirical methods in natural language processing', pp. 262–272.

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

Nakamura, E. and Steinsson, J. (2018), 'High-Frequency Identification of Monetary Non-Neutrality: The Information Effect*', *The Quarterly Journal of Economics* **133**(3), 1283–1330.
URL: *https://academic.oup.com/qje/article/133/3/1283/4828341*

Pearlman, J., Currie, D. and Levine, P. (1986), 'Rational expectations models with partial information', *Economic Modelling* **3**(2), 90–105.

Ramm, A., Loáiciga, S., Friedrich, A. and Fraser, A. (2017), Annotating tense, mood and voice for english, french and german, *in* 'Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (ACL 2017), System Demonstrations', Association for Computational Linguistics, pp. 1–6.

Roberts, M. E., Stewart, B. M. and Airoldi, E. M. (2016), 'A model of text for experimentation in the social sciences', *Journal of the American Statistical Association* **111**(515), 988–1003.

Roberts, M. E., Stewart, B. M., Tingley, D. and Airoldi, E. M. (2013), The structural topic model and applied social science, *in* 'Advances in neural information processing systems workshop on topic models: computation, application, and evaluation', Vol. 4, Harrahs and Harveys, Lake Tahoe, pp. 1–20.

Roberts, M. E., Stewart, B. M., Tingley, D., Lucas, C., Leder-Luis, J., Gadarian, S. K., Albertson, B. and Rand, D. G. (2014), 'Structural topic models for open-ended survey responses', *American journal of political science* **58**(4), 1064–1082.

Shah, A., Sukhani, S. and Pardawala, H. (2025), 'Words that unite the world: A unified framework for deciphering global central bank communications'.

Svensson, L. E. (1997), 'Inflation forecast targeting: Implementing and monitoring inflation targets', *European Economic Review* **41**(6), 1111–1146.

Swanson, E. T. and Jayawickrema, V. (2023), 'Speeches by the fed chair are more important than fomc announcements: An improved high-frequency measure of us monetary policy shocks', *Unpublished Manuscript* .

Wilcox, K. T., Jacobucci, R., Zhang, Z. and Ammerman, B. A. (2023), 'Supervised latent dirichlet allocation with covariates: A bayesian structural and measurement model of text and covariates.', *Psychological Methods* .

Appendix

A Topic Words and Sentences

Appendix Table A1: Representative Words and Sentences

Topic	**Representative Words**	**Representative Sentences**
Spending	capital household job orders	Forward-looking indicators of business equipment spending–such as increases in new and unfilled capital goods orders, along with upbeat [...] Recent indicators of business equipment spending–such as rising new orders of nondefense capital goods excluding aircraft [...] Forward-looking indicators of business equipment spending–such as a rising backlog of unfilled orders for nondefense capital goods [...]
Credit	credit loans changed survey	Credit remained broadly available to higher-score borrowers seeking conforming mortgages but tightened further, from already tight levels [...] In the April Senior Loan Officer Opinion Survey on Bank Lending Practices (SLOOS), banks reported having tightened their C&I lending [...] The Senior Loan Officer Opinion Survey on Bank Lending Practices conducted in April indicated that banks had tightened lending [...]
Int. Rate	funds target path stance	The Committee agreed to maintain the target range for the federal funds rate at 0 to 1/4 percent and to reaffirm in its postmeeting statement [...] A technical adjustment of the IOER rate to a level 5 basis points below the top of the target range could keep the effective federal funds rate [...] The Committee agreed to maintain the target range for the federal funds rate at 0 to 1/4 percent and to reaffirm in the statement [...]
Inflation	expectations energy measures core	With longer-term inflation expectations assumed to remain stable, resource slack projected to diminish slowly, and changes in commodity [...] The staff continued to project that total PCE inflation would remain near the Committee's 2 percent objective over the medium term [...] The staff's forecast of total PCE price inflation for this year was revised down somewhat, reflecting slightly lower projected paths for consumer [...]
OMO	securities open operations agency	In order to ensure the effective conduct of open market operations, while assisting in the provision of short-term investments for foreign and international accounts maintained at the Federal Reserve Bank of New York and accounts [...]
Uncertainty	stimulus restraint prospective moderation	A great deal of uncertainty surrounded the behavior of productivity growth going forward, but some further pickup, and the associated ability [...] While the members believed that some slowing in the expansion was inevitable, they felt that substantial uncertainty surrounded [...] More fundamentally, however, the members believed that current growth in aggregate demand, should it persist, would continue to exceed [...]

Notes: This table reports six labelled topics from the Structural Topic Model (STM) estimation from the FOMC meeting minutes. For each topic, the seven most representative words are displayed alongside the most representative sentences for that topic. Topic labels are assigned based on the representative words and sentences.

Appendix Table A2: Representative Words and Sentences

Topic	**Representative Words**	**Representative Sentences**
Supply Chain	observed couple pandemic several	Participants noted that supply chain bottlenecks and labor shortages had continued to limit businesses' ability to meet strong demand [...] While several participants pointed to signs of incremental improvement in supply chains, a few others remarked that business contacts [...] Some of these participants also suggested that labor supply constraints were the main impediments to further improvement in labor market [...]
Financial Stability	risks stability asset downside	A few participants stressed that the activities of some nonbank financial institutions presented vulnerabilities to the financial system [...] The staff provided presentations covering the efficacy of the Federal Reserve's asset purchases, the effects of the purchases [...] Participants recognized that the invasion of Ukraine by Russia was causing tremendous human and economic hardship for the Ukrainian people.
Construction	housing homes house country	At the same time, nonresidential construction activity appeared to have been well maintained in many parts of the country [...] With mortgage rates at their recent reduced levels and incomes continuing to rise, the cash flow affordability of home ownership was exceptionally favorable. There were widespread anecdotal and other reports of high levels of home sales and few reports of faltering housing demand.
Growth	activity suggested contacts expand	In their discussion of the economic situation and the outlook, meeting participants viewed the information received since the Committee [...] In their discussion of the economic situation and outlook, meeting participants regarded the information received during the intermeeting [...]
Industrial Production	quarter real production second	The U.S. international trade deficit narrowed in May as the value of exports of goods and services rose slightly and the value of imports declined [...] Real personal consumption expenditures grew in June and in the second quarter as a whole, largely reflecting a brisk rise in purchases [...] Real state and local government purchases, however, appeared to increase; the payrolls of these governments expanded briskly in September [...]
Financial Markets	intermeeting yields currencies economies	In foreign exchange markets, the trade-weighted value of the dollar declined substantially over the intermeeting period in relation [...] In foreign exchange markets, the trade-weighted value of the dollar rose somewhat further over the intermeeting period in relation [...] In foreign exchange markets, the trade-weighted value of the dollar has depreciated slightly over the period in relation to other major currencies [...]

Notes: This table reports six labelled topics from the Structural Topic Model (STM) estimation from the FOMC meeting minutes. For each topic, the seven most representative words are displayed alongside the most representative sentences for that topic. Topic labels are assigned based on the representative words and sentences.

B Future-Oriented Sentence Classification and Cleaning Method

We use a two-stage NLP pipeline to identify forward-looking language in FOMC meeting minutes. The method combines grammatical tense detection with keyword-based lexical matching in order to capture both explicitly marked future constructions and more general forward-looking statements.

First, we clean the raw meeting transcripts to remove noise such as URLs, HTML tags, and non-printable characters, and then segment them into individual sentences.

We identify forward-looking sentences in two stages:

1. **TMV-based tense detection**: We first parse each sentence using the Mate dependency parser, after which we apply TMV to recover verb tense and mood information from the parsed output. This stage identifies sentences with explicit future-marking grammatical constructions. In particular, we classify a sentence as forward-looking if TMV identifies verbal forms corresponding to `condI`, `futureI`, or `futureII`. This allows us to capture sentences whose forward-looking content is expressed syntactically rather than through specific keywords.

2. **Keyword-based matching**: In a second step, we compare sentences against a predefined dictionary of future-oriented terms. If a sentence contains any token or phrase from this set, we classify it as forward-looking. This lexical layer wants to capture forward-looking content that may not appear in an explicit future tense construction but nevertheless conveys expectations, plans, projections, or anticipated developments.

We based the choice of keywords on repeated validation exercises and manual review of forward-looking sentences. We converged on the following set of terms for the lexical matching layer:

```
expect, prospect, outlook, view, judgment, belief, believe, gauge, will,
shall, forthcoming, forecast, anticipate, predict, ahead, forward, commit,
intend, aim, anticipation, scenario, path, future, plan, project, trajectory,
direction, objective, goal, trend, outcome, propose, speculate, judgement,
pledge, foresee, forsee, expectation, projection, likely to, projected
to.
```

We ultimately classify a sentence as forward-looking if it is identified by either stage of the procedure. The final output is therefore an indicator variable equal to one if the sentence is classified as forward-looking, and zero otherwise.

C Extended Kalman Filter Recursions

This section describes the tracking problem for $(\lambda_t,\ \hat{r}_t^*)$. We apply the extended Kalman filter to our non-linear set-up as is standard in the literature. First, denote $\zeta_t = (\lambda_t, \hat{r}_t^*)'$,

$\epsilon_t = (\varepsilon_t^\lambda, \varepsilon_t^r)'$, and $y_t^{obs} = (\lambda_t^{obs}, r_t^{obs})'$. The transition equation and the observation equations can be written as

$$\begin{aligned}\zeta_t &= F\zeta_{t-1} + Q\epsilon_t \\ y_t^{\text{obs}} &= h(\zeta_t) + D\varepsilon_t^R\end{aligned}$$

where

$$F = \begin{bmatrix}\rho_\lambda & 0 \\ 0 & \rho_r\end{bmatrix}, \; QQ' = \begin{bmatrix}\sigma_\lambda^2 & 0 \\ 0 & \sigma_r^2\end{bmatrix}, \; h(\zeta_t) = \begin{bmatrix}\frac{\rho \hat{r}_t^*}{(\phi-\rho)\cdot\exp\{-\lambda_t\}-\phi} \\ \hat{r}_t^*\end{bmatrix}, \; DD' = \begin{bmatrix}\sigma_R^2 & 0 \\ 0 & 0\end{bmatrix}.$$

The extended Kalman filter handles the non-linear observation equation by taking a first-order Taylor expansion at each point in time. The Jacobian matrix

$$H_t = \left.\frac{\partial h(\zeta_t)}{\partial \zeta_t}\right|_{\zeta_t=\zeta_{t|t-1}}$$

is a matrix of constants when evaluated at the previous period's state forecast. Kalman filter recursions are then calculated in the steps:

Update step:

$$\begin{aligned}\zeta_{t|t} &= \zeta_{t|t-1} + P_{t|t-1}H_t'(H_tP_{t|t-1}H_t' + DD')^{-1}(y_t^{\text{obs}} - h(\zeta_{t|t-1})) \\ P_{t|t} &= P_{t|t-1} - P_{t|t-1}H_t'(H_tP_{t|t-1}H_t' + DD')^{-1}H_tP_{t|t-1}'\end{aligned}$$

Forecast step:

$$\begin{aligned}\zeta_{t+1|t} &= F\zeta_{t|t} \\ P_{t+1|t} &= FP_{t|t}F' + QQ'\end{aligned}$$

where $P_{t|t}$ is the posterior covariance matrix of the state estimate at time t. The initial condition for $P_{1|0}$ is set to QQ', and $\zeta_{1|0}$ is set to zero.

The left-hand side of the observation equation is $(\hat{R}_t - \phi\hat{\Pi}_t, \; \hat{r}_t^*)'$. The endogenous solution of both $\hat{R}_t$ and $\hat{\Pi}_t$ are a function of the now-cast $\lambda_{t|t}$ obtained with the Kalman filter, which in-turn depends on the state variables and shocks. Both the now-cast and endogenous co-state variables are determined by this period's shocks given the value of the state variables in the previous period. The filtering problem introduces an additional state variable to the overall model for the now-cast of the state through the Kalman filter recursions. Given a sequence of shocks and an initial condition, the recursions and transition equation trace out the path for the endogenous co-state variables in the model.

D Robustness Checks

Appendix Table D1: Asset Responses to Forward-Looking Communications: Feb 1997 to Sept 2007

	2Y Treasuries	5Y Treasuries	10Y Treasuries	S&P 500
Inflation X F	0.010*	0.006*	0.007*	-0.051
	[-0.002, 0.021]	[-0.003, 0.016]	[-0.001, 0.014]	[-0.204, 0.104]
Spending X F	0.009	0.005	0.005	-0.002
	[-0.007, 0.024]	[-0.008, 0.018]	[-0.005, 0.015]	[-0.197, 0.193]
Credit X F	-0.001	-0.002	-0.003	-0.075
	[-0.015, 0.015]	[-0.014, 0.011]	[-0.013, 0.007]	[-0.280, 0.128]
Int. Rate X F	-0.016	-0.014	-0.010	-0.143
	[-0.038, 0.007]	[-0.032, 0.004]	[-0.024, 0.004]	[-0.429, 0.137]
OMO X F	0.001	0.001	0	-0.094
	[-0.017, 0.020]	[-0.014, 0.017]	[-0.012, 0.012]	[-0.324, 0.135]
Uncertainty X F	-0.006	0.003	-0.004	-0.063
	[-0.029, 0.015]	[-0.015, 0.021]	[-0.018, 0.010]	[-0.336, 0.220]
Supply Chain X F	-0.006	-0.008	-0.004	-0.097
	[-0.049, 0.037]	[-0.044, 0.029]	[-0.032, 0.026]	[-0.649, 0.448]
Fin. Stability X F	0.02	0.018	0.013	0.108
	[-0.021, 0.061]	[-0.017, 0.051]	[-0.014, 0.040]	[-0.389, 0.612]
Construction X F	0	0.001	0.002	0.061
	[-0.012, 0.012]	[-0.009, 0.011]	[-0.005, 0.010]	[-0.090, 0.211]
Growth X F	-0.001	-0.004	-0.002	-0.035
	[-0.019, 0.017]	[-0.019, 0.011]	[-0.014, 0.010]	[-0.271, 0.192]
Ind. Prod X F	-0.003	0	-0.002	-0.004
	[-0.014, 0.009]	[-0.009, 0.010]	[-0.009, 0.005]	[-0.146, 0.142]
Fin. Mkts. X F	-0.007	-0.003	-0.002	0.183**
	[-0.020, 0.005]	[-0.013, 0.007]	[-0.010, 0.006]	[0.025, 0.343]
TWS Included?	Yes	Yes	Yes	Yes
Observations	85	85	85	85
R-squared	0.225	0.279	0.33	0.343

Notes: This table shows estimated responses of Treasury futures yields and the S&P500 index for the pre-financial crisis period. All regressions include the tone-weighted topic shares as covariates, which are not reported here. Treasury responses are measured in percentage points; the S&P500 response is measured in percent changes. Simulated 95% confidence intervals are shown in square brackets, and significance stars are based on the simulated confidence intervals: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Appendix Table D2: Eurodollar Responses to Forward-Looking Communications: Feb 1997 to Sept 2007

	ED1	ED2	ED3	ED4	MPS
Inflation X F	0.001	0.003	0.008*	0.006	0.005
	[-0.004, 0.005]	[-0.006, 0.012]	[-0.003, 0.019]	[-0.005, 0.017]	[-0.005, 0.015]
Spending X F	0	0.004	0.005	0.004	0.004
	[-0.006, 0.007]	[-0.008, 0.015]	[-0.010, 0.019]	[-0.011, 0.018]	[-0.009, 0.016]
Credit X F	-0.002	-0.002	-0.002	-0.001	-0.003
	[-0.009, 0.004]	[-0.015, 0.009]	[-0.017, 0.012]	[-0.015, 0.014]	[-0.016, 0.010]
Int. Rate X F	-0.004	-0.017*	-0.021**	-0.029***	-0.020**
	[-0.013, 0.005]	[-0.034, 0.000]	[-0.042, -0.000]	[-0.051, -0.008]	[-0.038, -0.001]
OMO X F	0.003	0	0.002	0.002	0.002
	[-0.005, 0.011]	[-0.014, 0.014]	[-0.015, 0.020]	[-0.015, 0.020]	[-0.013, 0.018]
Uncertainty X F	-0.002	0.007	0.003	0.004	0.003
	[-0.011, 0.007]	[-0.009, 0.023]	[-0.017, 0.024]	[-0.016, 0.025]	[-0.015, 0.021]
Supply Chain X F	0.005	0.007	-0.005	0.002	0.004
	[-0.012, 0.022]	[-0.025, 0.040]	[-0.044, 0.035]	[-0.039, 0.043]	[-0.032, 0.039]
Fin. Stability X F	0.006	0.009	0.015	0.015	0.013
	[-0.011, 0.022]	[-0.022, 0.039]	[-0.023, 0.053]	[-0.024, 0.053]	[-0.020, 0.047]
Construction X F	0	-0.003	-0.001	0.001	-0.001
	[-0.005, 0.005]	[-0.012, 0.007]	[-0.013, 0.010]	[-0.011, 0.012]	[-0.011, 0.009]
Growth X F	-0.001	-0.009	-0.004	-0.006	-0.006
	[-0.009, 0.006]	[-0.023, 0.005]	[-0.022, 0.012]	[-0.023, 0.012]	[-0.021, 0.009]
Ind. Prod X F	0	0.002	0.001	0.002	0.001
	[-0.005, 0.005]	[-0.007, 0.011]	[-0.010, 0.012]	[-0.009, 0.013]	[-0.008, 0.011]
Fin. Mkts. X F	-0.002	-0.005	-0.004	-0.005	-0.005
	[-0.007, 0.003]	[-0.014, 0.005]	[-0.016, 0.008]	[-0.016, 0.007]	[-0.015, 0.006]
TWS Included?	Yes	Yes	Yes	Yes	Yes
Observations	85	85	85	85	85
R-squared	0.218	0.267	0.247	0.304	0.258

Notes: This table presents estimated responses of Eurodollar futures to topic-specific communication variables over the pre-financial crisis period. MPS is the first principal component of the change in the first four Eurodollar futures contracts, measured in ED4 units. All regressions include the tone-weighted topic shares as covariates, which are not reported here. Simulated 95% confidence intervals are shown in square brackets. Significance: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Appendix Table D3: Asset Responses to Forward-Looking Communications: Feb 1997 to Jan 2020

	2Y Treasuries	5Y Treasuries	10Y Treasuries	S&P 500
Inflation X F	0.008***	0.006***	0.005**	0.002
	[0.003, 0.013]	[0.002, 0.011]	[0.001, 0.009]	[-0.087, 0.092]
Spending X F	0.003	0.002	0.002	-0.053
	[-0.004, 0.011]	[-0.005, 0.009]	[-0.004, 0.008]	[-0.180, 0.072]
Credit X F	-0.001	-0.002	-0.002	-0.050
	[-0.005, 0.004]	[-0.006, 0.003]	[-0.005, 0.002]	[-0.147, 0.049]
Int. Rate X F	-0.006	-0.007	-0.008**	-0.005
	[-0.014, 0.003]	[-0.014, 0.002]	[-0.015, -0.001]	[-0.156, 0.143]
OMO X F	0.003*	0.006**	0.005**	-0.006
	[-0.002, 0.009]	[0.001, 0.011]	[0.001, 0.009]	[-0.102, 0.092]
Uncertainty X F	-0.002	0.005	0.001	-0.124
	[-0.013, 0.009]	[-0.004, 0.015]	[-0.007, 0.008]	[-0.285, 0.035]
Supply Chain X F	0	-0.002	-0.001	0.011
	[-0.011, 0.010]	[-0.012, 0.009]	[-0.010, 0.008]	[-0.211, 0.234]
Fin. Stability X F	0.007	0.008	0.009*	0.121
	[-0.009, 0.021]	[-0.006, 0.023]	[-0.004, 0.021]	[-0.152, 0.395]
Construction X F	0	0	0	0.034
	[-0.007, 0.007]	[-0.006, 0.006]	[-0.005, 0.005]	[-0.076, 0.144]
Growth X F	-0.002	-0.004	-0.002	0.031
	[-0.011, 0.007]	[-0.012, 0.005]	[-0.009, 0.005]	[-0.139, 0.198]
Ind. Prod X F	-0.001	0.001	-0.001	0.046
	[-0.006, 0.004]	[-0.004, 0.005]	[-0.004, 0.003]	[-0.049, 0.139]
Fin. Mkts. X F	-0.006**	-0.003	-0.003	0.037
	[-0.011, -0.000]	[-0.008, 0.002]	[-0.007, 0.001]	[-0.065, 0.144]
TWS Included?	Yes	Yes	Yes	Yes
Observations	184	184	184	184
R-squared	0.11	0.124	0.127	0.121

Notes: This table reports estimated responses of Treasury yields and the S&P500 index to topic-specific communication dimensions over the pre-COVID-19 pandemic period. All regressions include the tone-weighted topic shares as covariates, which are not reported here. Treasury responses are measured in percentage points; the S&P500 response is measured in percent changes. Simulated 95% confidence intervals are shown in square brackets. Significance: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Appendix Table D4: Eurodollar Responses to Forward-Looking Communications: Feb 1997 to Jan 2020

	ED1	ED2	ED3	ED4	MPS
Inflation X F	0	0.002	0.005*	0.005*	0.003*
	[-0.002, 0.002]	[-0.002, 0.006]	[-0.000, 0.009]	[-0.000, 0.010]	[-0.001, 0.007]
Spending X F	-0.001	0	0	0	-0.001
	[-0.004, 0.002]	[-0.006, 0.005]	[-0.007, 0.008]	[-0.007, 0.008]	[-0.007, 0.006]
Credit X F	0.001	0.001	0	0.001	0.001
	[-0.001, 0.003]	[-0.002, 0.005]	[-0.004, 0.005]	[-0.004, 0.005]	[-0.003, 0.005]
Int. Rate X F	-0.002	-0.007**	-0.008*	-0.012***	-0.008**
	[-0.005, 0.002]	[-0.013, -0.000]	[-0.016, 0.000]	[-0.021, -0.004]	[-0.015, -0.001]
OMO X F	0.001*	0.001	0.003	0.004*	0.003*
	[-0.001, 0.004]	[-0.003, 0.006]	[-0.002, 0.008]	[-0.001, 0.010]	[-0.001, 0.008]
Uncertainty X F	0	0.008*	0.005	0.006	0.005
	[-0.005, 0.004]	[-0.001, 0.016]	[-0.006, 0.015]	[-0.005, 0.017]	[-0.004, 0.014]
Supply Chain X F	0.001	0.004	0.002	0.002	0.003
	[-0.003, 0.006]	[-0.004, 0.013]	[-0.008, 0.012]	[-0.009, 0.013]	[-0.006, 0.012]
Fin. Stability X F	0.003	0.003	0.005	0.005	0.005
	[-0.003, 0.009]	[-0.008, 0.014]	[-0.009, 0.019]	[-0.010, 0.020]	[-0.008, 0.017]
Construction X F	0	-0.003	-0.001	0	-0.001
	[-0.003, 0.003]	[-0.008, 0.003]	[-0.008, 0.005]	[-0.006, 0.007]	[-0.007, 0.005]
Growth X F	0	-0.006	-0.004	-0.004	-0.004
	[-0.004, 0.003]	[-0.013, 0.001]	[-0.013, 0.005]	[-0.013, 0.005]	[-0.012, 0.004]
Ind. Prod X F	-0.001	0	0.001	0.001	0
	[-0.003, 0.001]	[-0.004, 0.004]	[-0.004, 0.006]	[-0.004, 0.006]	[-0.004, 0.005]
Fin. Mkts. X F	-0.001	-0.003	-0.003	-0.005	-0.004
	[-0.003, 0.001]	[-0.008, 0.001]	[-0.009, 0.002]	[-0.010, 0.001]	[-0.008, 0.001]
TWS Included?	Yes	Yes	Yes	Yes	Yes
Observations	184	184	184	184	184
R-squared	0.134	0.134	0.108	0.138	0.123

Notes: This table presents estimated responses of Eurodollar futures to topic-specific communication variables over the pre-COVID-19 pandemic period. All regressions include the tone-weighted topic shares as covariates, which are not reported here. MPS is the first principal component of the change in the first four Eurodollar futures contracts, measured in ED4 units. Simulated 95% confidence intervals are shown in square brackets. Significance: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Appendix Table D5: Asset Responses to Forward-Looking Communications: 30 minute window

	2Y Treasuries	5Y Treasuries	10Y Treasuries	S&P 500
Inflation X F	0.005***	0.004**	0.003**	0.013
	[0.001, 0.009]	[0.000, 0.008]	[0.000, 0.006]	[-0.068, 0.097]
Spending X F	0	-0.001	-0.001	-0.116**
	[-0.006, 0.006]	[-0.007, 0.004]	[-0.006, 0.003]	[-0.228, -0.008]
Credit X F	0.003*	0.003*	0.002*	-0.024
	[-0.000, 0.005]	[-0.000, 0.005]	[-0.001, 0.004]	[-0.111, 0.062]
Int. Rate X F	-0.004	-0.004	-0.004	0.156**
	[-0.010, 0.002]	[-0.010, 0.002]	[-0.009, 0.001]	[0.019, 0.291]
OMO X F	0.004*	0.004**	0.003*	0.008
	[-0.000, 0.007]	[0.000, 0.007]	[-0.000, 0.006]	[-0.075, 0.095]
Uncertainty X F	0.003	0.009**	0.004*	-0.043
	[-0.006, 0.011]	[0.001, 0.016]	[-0.002, 0.010]	[-0.187, 0.100]
Supply Chain X F	-0.001	-0.003	-0.004	0.081
	[-0.008, 0.006]	[-0.010, 0.004]	[-0.010, 0.003]	[-0.125, 0.286]
Fin. Stability X F	0.005	0.007*	0.007*	0.062
	[-0.005, 0.014]	[-0.002, 0.016]	[-0.001, 0.015]	[-0.179, 0.306]
Construction X F	-0.001	-0.001	0	0.074*
	[-0.006, 0.004]	[-0.006, 0.003]	[-0.004, 0.004]	[-0.022, 0.169]
Growth X F	-0.001	-0.002	0	0.082
	[-0.008, 0.006]	[-0.009, 0.004]	[-0.006, 0.005]	[-0.062, 0.228]
Ind. Prod X F	-0.003	0	-0.001	-0.01
	[-0.007, 0.001]	[-0.004, 0.003]	[-0.004, 0.002]	[-0.089, 0.074]
Fin. Mkts. X F	-0.005**	-0.004**	-0.003**	0.036
	[-0.009, -0.001]	[-0.008, -0.000]	[-0.006, -0.000]	[-0.054, 0.127]
TWS Included?	Yes	Yes	Yes	Yes
Observations	210	210	210	198
R-squared	0.107	0.117	0.11	0.286

Notes: This table reports estimated responses of Treasury yields and the S&P500 index to topic-specific communication dimensions using 30 minute window (10 before, 20 after). All regressions include the tone-weighted topic shares as covariates, which are not reported here. Treasury responses are measured in percentage points; the S&P500 response is measured in percent changes. Simulated 95% confidence intervals are shown in square brackets. Significance: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Appendix Table D6: Eurodollar Responses to Forward-Looking Communications: 30 minute window

	ED1	ED2	ED3	ED4	MPS
Inflation X F	0	0.001	0.003*	0.002	0.002
	[-0.002, 0.002]	[-0.002, 0.004]	[-0.001, 0.007]	[-0.002, 0.006]	[-0.002, 0.005]
Spending X F	-0.002	-0.002	-0.002	-0.004	-0.003
	[-0.004, 0.001]	[-0.007, 0.003]	[-0.007, 0.004]	[-0.009, 0.002]	[-0.008, 0.002]
Credit X F	0.002***	0.004***	0.004**	0.004**	0.004***
	[0.001, 0.004]	[0.001, 0.007]	[0.001, 0.006]	[0.000, 0.007]	[0.001, 0.007]
Int. Rate X F	-0.001	-0.003	-0.006**	-0.009***	-0.005*
	[-0.004, 0.002]	[-0.009, 0.002]	[-0.012, -0.000]	[-0.015, -0.002]	[-0.011, 0.001]
OMO X F	0.001	0.001	0.002*	0.003*	0.002
	[-0.001, 0.003]	[-0.002, 0.004]	[-0.001, 0.006]	[-0.001, 0.007]	[-0.001, 0.005]
Uncertainty X F	0.001	0.010***	0.007*	0.009**	0.008*
	[-0.003, 0.005]	[0.003, 0.017]	[-0.001, 0.015]	[0.001, 0.018]	[-0.000, 0.015]
Supply Chain X F	0	0.001	0.001	0.001	0.001
	[-0.003, 0.004]	[-0.005, 0.007]	[-0.006, 0.007]	[-0.007, 0.009]	[-0.006, 0.007]
Fin. Stability X F	0.002	0.003	0.005	0.006	0.005
	[-0.002, 0.007]	[-0.005, 0.011]	[-0.004, 0.015]	[-0.004, 0.016]	[-0.004, 0.014]
Construction X F	-0.001	-0.003	-0.002	0	-0.002
	[-0.003, 0.002]	[-0.007, 0.002]	[-0.007, 0.003]	[-0.005, 0.006]	[-0.006, 0.003]
Growth X F	0	-0.004	-0.003	-0.003	-0.003
	[-0.004, 0.003]	[-0.010, 0.002]	[-0.010, 0.004]	[-0.010, 0.004]	[-0.010, 0.003]
Ind. Prod X F	0	0	0	0	0
	[-0.002, 0.002]	[-0.003, 0.004]	[-0.004, 0.004]	[-0.004, 0.004]	[-0.004, 0.003]
Fin. Mkts. X F	-0.002*	-0.004**	-0.004*	-0.005**	-0.004**
	[-0.004, 0.000]	[-0.007, -0.001]	[-0.008, 0.000]	[-0.009, -0.001]	[-0.008, -0.001]
TWS Included?	Yes	Yes	Yes	Yes	Yes
Observations	210	210	210	210	210
R-squared	0.156	0.144	0.115	0.138	0.142

Notes: This table presents estimated responses of Eurodollar futures to topic-specific communication variables using 30 minute window (10 before, 20 after). All regressions include the tone-weighted topic shares as covariates, which are not reported here. Simulated 95% confidence intervals are shown in square brackets. Significance: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Appendix Table D7: Asset Responses to Forward-Looking Communications: 50 minute window

	2Y Treasuries	5Y Treasuries	10Y Treasuries	S&P 500
Inflation X F	0.005***	0.004**	0.003**	-0.011
	[0.001, 0.010]	[0.000, 0.008]	[0.000, 0.007]	[-0.098, 0.079]
Spending X F	0.001	-0.001	-0.001	-0.114*
	[-0.005, 0.008]	[-0.007, 0.005]	[-0.006, 0.004]	[-0.233, 0.006]
Credit X F	0.002	0.001	0.001	-0.039
	[-0.001, 0.005]	[-0.002, 0.004]	[-0.002, 0.003]	[-0.134, 0.052]
Int. Rate X F	-0.005	-0.006**	-0.005**	0.153**
	[-0.012, 0.001]	[-0.013, -0.000]	[-0.011, -0.000]	[0.007, 0.299]
OMO X F	0.004*	0.004**	0.004**	0.008
	[-0.000, 0.008]	[0.000, 0.008]	[0.000, 0.007]	[-0.082, 0.101]
Uncertainty X F	0.002	0.008*	0.004	-0.056
	[-0.008, 0.011]	[-0.000, 0.016]	[-0.003, 0.011]	[-0.213, 0.102]
Supply Chain X F	0	-0.001	-0.001	0.098
	[-0.008, 0.007]	[-0.008, 0.007]	[-0.008, 0.005]	[-0.118, 0.313]
Fin. Stability X F	0.004	0.006	0.006*	0.087
	[-0.006, 0.015]	[-0.004, 0.016]	[-0.003, 0.014]	[-0.170, 0.345]
Construction X F	-0.001	-0.001	-0.001	0.017
	[-0.007, 0.004]	[-0.007, 0.004]	[-0.005, 0.004]	[-0.087, 0.121]
Growth X F	0	-0.002	-0.001	0.056
	[-0.008, 0.007]	[-0.009, 0.005]	[-0.007, 0.005]	[-0.104, 0.219]
Ind. Prod X F	-0.003	0	0	0.02
	[-0.007, 0.002]	[-0.004, 0.004]	[-0.004, 0.003]	[-0.066, 0.110]
Fin. Mkts. X F	-0.005**	-0.003	-0.002	0.04
	[-0.009, -0.001]	[-0.007, 0.001]	[-0.006, 0.001]	[-0.059, 0.141]
Observations	210	210	210	198
R-squared	0.09	0.099	0.105	0.267

Notes: This table reports estimated responses of Treasury yields and the S&P500 index to topic-specific communication dimensions using 50 minute window (10 before, 40 after). All regressions include the tone-weighted topic shares as covariates, which are not reported here. Treasury responses are measured in percentage points; the S&P500 response is measured in percent changes. Simulated 95% confidence intervals are shown in square brackets. Significance: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Appendix Table D8: Eurodollar Responses to Forward-Looking Communications: 50 minute window

	ED1	ED2	ED3	ED4	MPS
Inflation X F	0	0.002	0.003*	0.003*	0.002
	[-0.002, 0.002]	[-0.002, 0.005]	[-0.001, 0.007]	[-0.001, 0.007]	[-0.002, 0.006]
Spending X F	-0.001	-0.001	-0.001	-0.002	-0.002
	[-0.004, 0.001]	[-0.007, 0.004]	[-0.008, 0.005]	[-0.009, 0.004]	[-0.008, 0.003]
Credit X F	0.002**	0.003**	0.002	0.002	0.003*
	[0.000, 0.003]	[0.000, 0.006]	[-0.001, 0.005]	[-0.002, 0.006]	[-0.000, 0.006]
Int. Rate X F	-0.001	-0.005*	-0.007**	-0.011***	-0.007**
	[-0.004, 0.002]	[-0.011, 0.000]	[-0.014, -0.000]	[-0.018, -0.004]	[-0.013, -0.001]
OMO X F	0.001*	0.001	0.002	0.003*	0.002
	[-0.001, 0.003]	[-0.002, 0.004]	[-0.002, 0.006]	[-0.001, 0.007]	[-0.001, 0.006]
Uncertainty X F	0	0.008*	0.006*	0.006*	0.005
	[-0.004, 0.004]	[-0.000, 0.015]	[-0.003, 0.015]	[-0.003, 0.016]	[-0.003, 0.013]
Supply Chain X F	0	0.003	0.002	0.004	0.002
	[-0.004, 0.004]	[-0.003, 0.010]	[-0.005, 0.010]	[-0.004, 0.012]	[-0.005, 0.010]
Fin. Stability X F	0.002	0.002	0.004	0.004	0.004
	[-0.003, 0.006]	[-0.007, 0.011]	[-0.006, 0.014]	[-0.006, 0.015]	[-0.006, 0.013]
Construction X F	0	-0.003	-0.002	0.001	-0.001
	[-0.003, 0.002]	[-0.007, 0.002]	[-0.008, 0.003]	[-0.005, 0.007]	[-0.006, 0.004]
Growth X F	0	-0.005	-0.004	-0.004	-0.003
	[-0.003, 0.004]	[-0.011, 0.002]	[-0.011, 0.004]	[-0.011, 0.004]	[-0.010, 0.004]
Ind. Prod X F	0	0	0	-0.001	-0.001
	[-0.002, 0.002]	[-0.004, 0.004]	[-0.005, 0.004]	[-0.005, 0.004]	[-0.004, 0.003]
Fin. Mkts. X F	-0.002*	-0.004**	-0.003	-0.005**	-0.004**
	[-0.004, 0.000]	[-0.007, -0.000]	[-0.007, 0.001]	[-0.009, -0.000]	[-0.008, -0.000]
TWS Included?	Yes	Yes	Yes	Yes	Yes
Observations	210	210	210	210	210
R-squared	0.117	0.122	0.096	0.123	0.115

Notes: This table presents estimated responses of Eurodollar futures to topic-specific communication variables using 50 minute window (10 before, 40 after). All regressions include the tone-weighted topic shares as covariates, which are not reported here. Simulated 95% confidence intervals are shown in square brackets. Significance: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Appendix Table D9: Asset Responses to Forward-Looking Communications- 3 topics

	2Y Treasuries	5Y Treasuries	10Y Treasuries	S&P 500
Inflation X F	0.006***	0.006***	0.004***	0.026
	[0.002, 0.009]	[0.003, 0.009]	[0.002, 0.007]	[-0.044, 0.099]
Int. Rate X F	-0.003	-0.001	-0.002	0.220***
	[-0.008, 0.002]	[-0.006, 0.003]	[-0.006, 0.002]	[0.110, 0.328]
Growth X F	-0.002	-0.001	-0.001	0.052
	[-0.006, 0.002]	[-0.004, 0.003]	[-0.004, 0.002]	[-0.043, 0.142]
TWS Included?	Yes	Yes	Yes	Yes
Observations	210	210	210	198
R-squared	0.021	0.026	0.022	0.073

Notes: This table shows the estimated relationship between Treasury futures yields and the forward-looking tone-weighted topic shares using the three main topics for central banks (inflation, interest rate, and growth); the tone-weighted topic shares are included as controls but not reported here. The S&P 500 response is measured in percentage point changes, and the response for Treasuries is measured in percentage points. Simulated quantiles are used to calculate the 95% confidence intervals in square brackets, and significance stars reflect * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Appendix Table D10: Asset Responses to Forward-Looking Communications

	ED1	ED2	ED3	ED4	MPS
Inflation X F	0	0.004***	0.004**	0.004**	0.004**
	[-0.001, 0.002]	[0.001, 0.007]	[0.001, 0.007]	[0.001, 0.007]	[0.001, 0.006]
Int. Rate X F	0	-0.003*	-0.004	-0.007***	-0.004*
	[-0.003, 0.002]	[-0.007, 0.001]	[-0.008, 0.001]	[-0.012, -0.002]	[-0.008, 0.000]
Growth X F	-0.001	-0.004**	-0.002	-0.001	-0.003
	[-0.003, 0.001]	[-0.007, -0.001]	[-0.006, 0.002]	[-0.005, 0.004]	[-0.006, 0.001]
TWS Included?	Yes	Yes	Yes	Yes	Yes
Observations	210	210	210	210	210
R-squared	0.097	0.074	0.047	0.078	0.067

Notes: This table presents estimated responses of Eurodollar futures to topic-specific communication variables using the three main topics for central banks (inflation, interest rate, and growth); the tone-weighted topic shares are included as controls but not reported here. MPS is the first principal component of the change in the first four Eurodollar futures contracts, measured in ED4 units. Simulated 95% confidence intervals reported in square brackets, and significance stars are calculated from the simulated quantiles intervals reflecting, * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

E Summary Statistics

Appendix Figure D1: Histograms of Market Responses to the FOMC Minutes

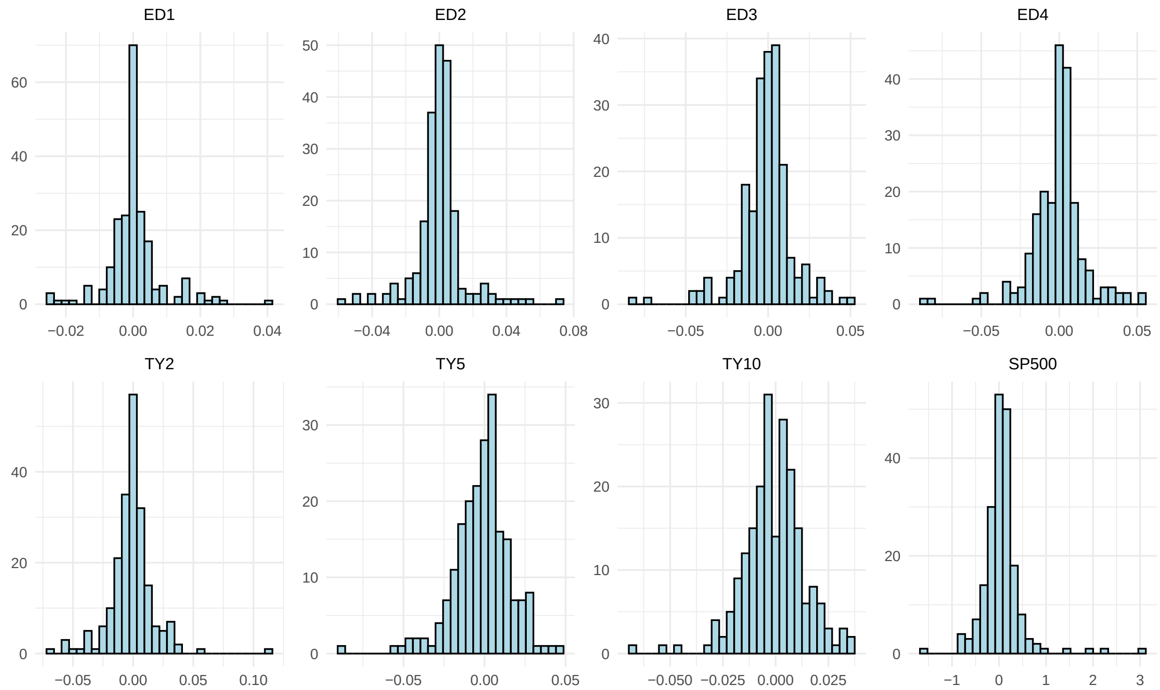

Notes: This figure plots histograms for the high-frequency financial market response to the release of the FOMC minutes over our whole sample, 1997 to 2023. ED1 to ED4 are the first four Eurodollar futures responses measured in yield changes. The TY2 to TY10 contracts are change in the yields for Treasury bill futures. The SP500 response is measured in percent changes in the SP500 futures contract.

Appendix Figure D2: Future Tone-Weighted Topic Shares

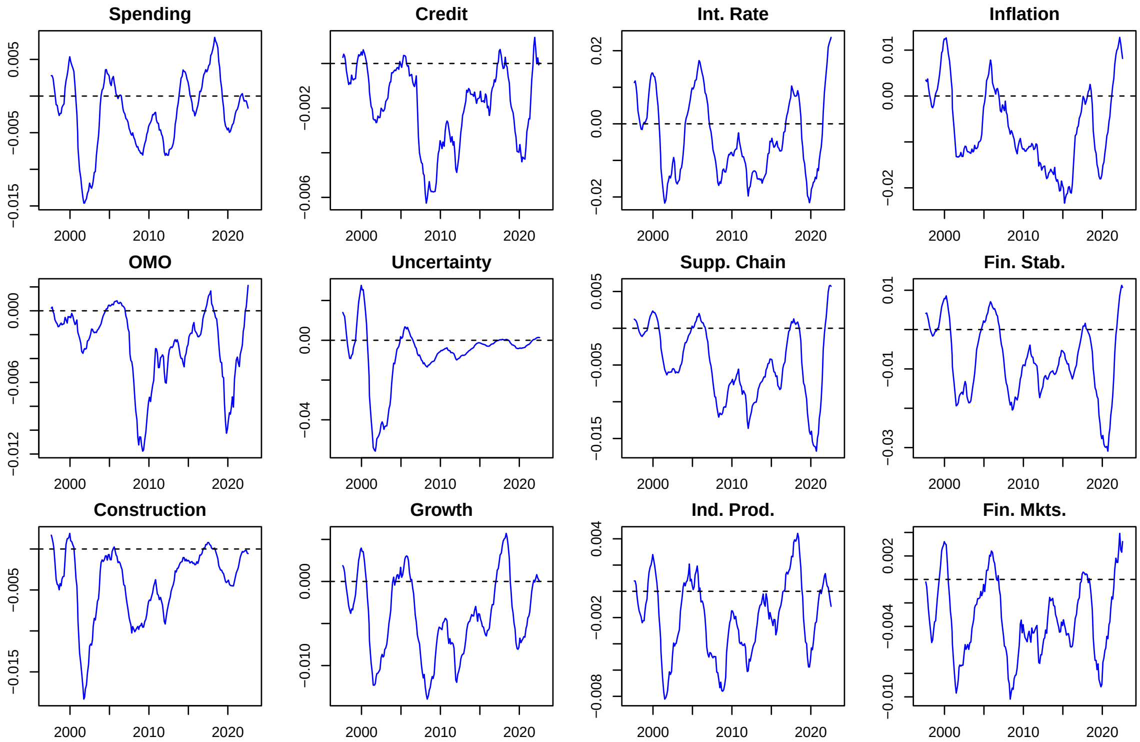

Notes: This figure plots the time series for $FTWS_{m,k}$ for the $K = 1, ..., 12$ topics using the expected value of the posterior for $\theta_{m,k,j}$ at the sentence level multiplied by the future-tense indicator and sentiment score, averaged at the document level, and then smoothed over time.

Appendix Figure D3: Tone-Weighted Topic Shares

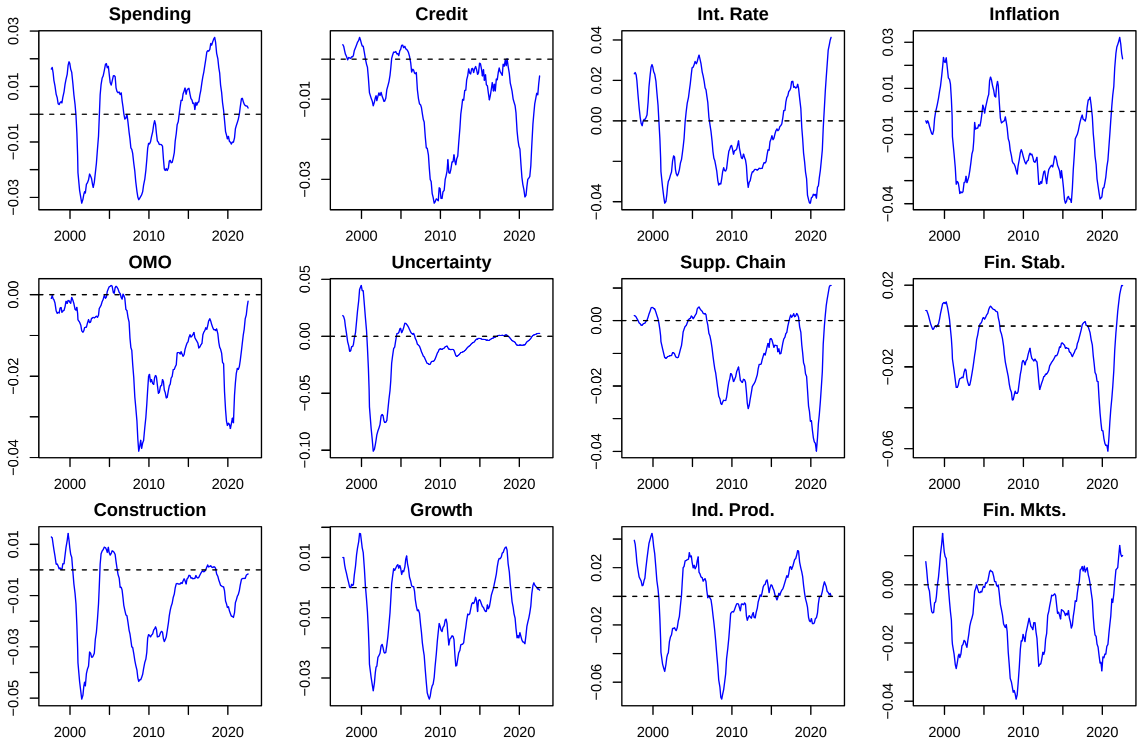

Notes: This figure plots the time series for $TWS_{m,k}$ for the $K = 1, ..., 12$ topics using the expected value of the posterior for $\theta_{m,k,j}$ at the sentence level multiplied by their sentiment score, averaged at the document level, and then smoothed over time.

Appendix Figure D4: Future Weighted Topic Shares

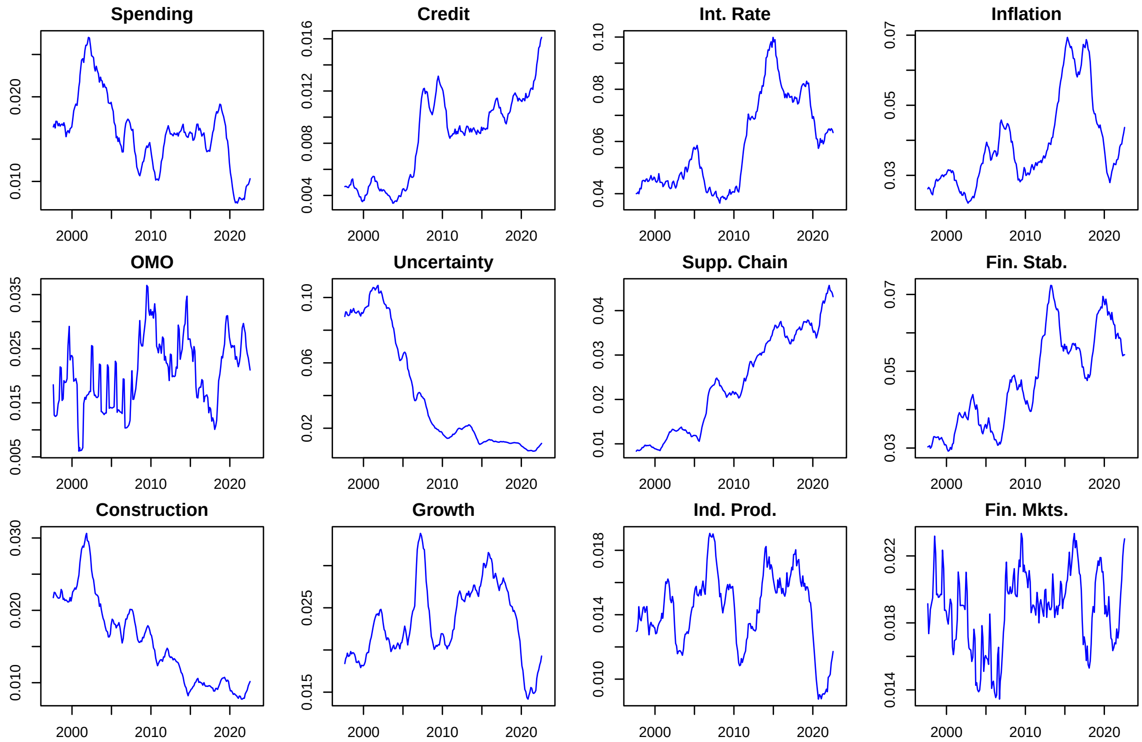

Notes: This figure plots the time series for future-weighted topic shares, which are not used in the main analysis, defined as $F_{m,j} \cdot \theta_{m,j,k}$, for the $K = 1, ..., 12$ topics using the expected value of the posterior for $\theta_{m,k,j}$ at the sentence level, weighted by the future-tense indicator, averaged at the document level, and then smoothed over time.