Fed Communication and Market Attention

Daniel Milo*

November 12, 2026

Updated frequently; latest version available here

Abstract

Does Fed communication make the private sector better informed? I measure how markets divide their attention between Fed communications and other economic information using over 300,000 reports from professional macroeconomic forecasters. I develop a model in which Fed communications can crowd in economic research by reducing the need to research the Fed, or crowd it out by making Fed research less costly. Consistent with crowding out, I find that Fed communications shift attention toward the Fed and away from independent economic research. This improves the accuracy of their Fed funds forecasts, but worsens the accuracy of their forecasts for economic outlooks and medium- to longer-term Treasury yields. Fed communications also draw forecasters' economic expectations toward the Fed's by shifting their focus towards the topics it emphasizes. I further show how this mechanism helps explain consensus expectations and asset price reactions to Fed communications. These findings suggest that Fed communications improve understanding of policy while crowding out independent analysis of the economy.

*New York University, Stern School of Business, dm4766@stern.nyu.edu. I am especially grateful to my dissertation committee, Arpit Gupta, Alexi Savov, Sydney Ludvigson, and Toomas Laarits, for their continued guidance and support. I also thank Viral Acharya, Alexander Bartik, Yannis Cabossioras, Simon Gilchrist, Dan Greenwald, Sebastian Hillenbrand, Isaac Leeson, Siddhartha Lewis-Hayre, Jeffrey Meli, Alan Moreira, Holger Mueller, Laura Nicolae, Desmond O'Shaughnessy, Cecilia Parlatore, Thomas Philippon, Indira Puri, Rob Richmond, Tony Saunders, Philipp Schnabl, Julia Selgrad, Bill Shu, Johannes Stroebel, Bruce Tuckman, Venky Venkateswaran, Olivier Wang, Courtney Weigand, Sean Wu, Andy Xia, Xuran Zeng, and Luofeng Zhou for helpful comments and conversations. I gratefully acknowledge funding from the NYU Stern Center for Global Economy and Business.

1 Introduction

The effectiveness of monetary policy depends on its ability to shape market expectations (Woodford, 2005; Blinder et al., 2008). Reflecting this belief, the Fed has substantially expanded the amount it communicates with markets over the past thirty years (Figure I, Panel A). Initially, policymakers expected this increased information to enhance their influence over long-term yields, to help markets price financial assets and to reduce economic and financial uncertainty (Appendix Table A1).[1] However, the emergence of several asset pricing anomalies around Fed communications has suggested a more complex role of Fed information.[2] For example, market participants may coordinate their beliefs around Fed communication, reducing incentives for private research (Morris and Shin, 2005).

This debate takes on renewed importance as the Fed actively reconsiders its communication framework.[3] Specifically, Kevin Warsh has reduced the amount the Fed communicates (Figure I, Panel B), arguing that doing so will shift market attention away from the Fed toward economic data.[4] While prior literature has provided evidence on the aggregate effect of Fed communication on market expectations (Nakamura and Steinsson, 2018; Jarociński and Karadi, 2020; Bauer and Swanson, 2023), evidence remains limited on whether Fed communication shifts the attention of market participants and what implications may follow. This motivates my research question: Does Fed communication make the private sector better informed about financial and economic outcomes?

I contribute evidence that Fed communications influence expectations by shaping the

[1] Ben Bernanke, "Central Bank Talk and Monetary Policy," October 7, 2004.

[2] See Hanson and Stein (2015) on the response of distant real forward rates and Lucca and Moench (2015) on excess equity returns before Fed announcements.

[3] Federal Reserve Board, Communications Task Force.

[4] Kevin Warsh, FOMC press conference, July 29, 2026, p. 2. Appendix Table A1 collects these and other perspectives on central bank communication.

attention of market participants. Specifically, I introduce a novel dataset of the written reports of professional macroeconomic forecasters, which links the analytical research effort of key market participants to their quantitative macrofinancial expectations and expectations of Fed communications. Leveraging this dataset, I show three key findings. First, forecasters spend a substantial share of their research on the Fed, and Fed communications influence this allocation. Second, forecaster attention towards the Fed predicts a trade-off in forecast accuracy, consistent with attention crowdout. Specifically, forecasters that attend more to the Fed realize lower Fed funds forecasting errors, but higher forecasting errors for economic outcomes and medium- and long-term Treasury yields. Third, the Fed shifts forecasters' economic expectations by directing their attention towards the information it emphasizes. I further examine the implications of these patterns for average forecast accuracy and show that they help explain aggregate asset price reactions to Fed communications.

I begin by framing the paper through a conceptual framework of costly information processing. Specifically, a forecaster predicts economic fundamentals and the Fed's policy by engaging in costly research that yields noisy signals about each outcome. The forecaster believes that Fed policy and economic fundamentals are positively correlated, so she weights each signal according to its precision and covariance when forming each forecast. Research about the Fed therefore informs her prediction of economic fundamentals.

Fed communication can influence a forecaster's attention and subsequent expectations in two ways. First, it may provide a free public signal about Fed policy. The forecaster then reduces her private research effort on the Fed and increases her research effort on the economy, improving the precision of both predictions. Second, Fed communication may reduce the marginal cost of Fed research, inducing the forecaster to shift her research effort towards the Fed. This shift in research effort can lead to lower Fed forecasting error at

the cost of higher economic forecasting error. These contrasting predictions motivate my empirical tests of forecasters' limited attention and the cost structure of Fed communication processing.

To measure how forecasters allocate attention, I assemble a novel dataset of their written research reports. These forecasters tend to work for investment banks where they provide expertise on the economy and the Fed to institutional investors. While surveys of these forecasters have been widely used in the academic literature (e.g., Bauer et al., 2024) and by the Fed itself, analysis of their written reports remains limited. The dataset comprises over 300,000 reports from 26 forecasting teams over 26 years. The sample composition closely mirrors that of existing surveys, such as the New York Fed Primary Dealers survey. To process this unstructured corpus, I develop measurement techniques using large language models (LLMs). In total, I extract over 400,000 quantitative macrofinancial expectations from the reports for the Fed funds rate, GDP, inflation, and different Treasury yield tenors. This allows me to observe when forecasters publish updated expectations, rather than relying on infrequent surveys. I further collect over 40,000 expectations of Fed communications and measure Fed attention across more than 2.4 million paragraphs of text.

I first show that forecasters devote a substantial share of their research to the Fed, and that Fed communications influence this allocation. Specifically, I define Fed research as the share of research text reasoning about the Fed or explicitly motivated by the Fed, which is 25% on average. This share is highly cyclical, nearly doubling from 20% to 40% in the lead-up to FOMC announcements. This reflects both a rise in the absolute volume of Fed research and a decline in the volume of independent economic research. Fed communication also influences the specific topics forecasters research, particularly the topics emphasized in FOMC minutes.

Forecaster attention towards the Fed predicts a trade-off in forecast accuracy, consistent with attention crowdout. Specifically, I estimate the relationship between forecasters' Fed research share and their forecast accuracy across macrofinancial outcomes. One challenge is that unobserved forecaster confidence may influence both their research allocation and realized forecast errors. I address this concern by comparing forecasters who make forecasts on the same date and are thus exposed to the same economic environment. Additionally, I find no statistically significant relationship between Fed attention and measures of confidence in economic forecasts. A one-standard-deviation increase in Fed research (approximately 15 percentage points) predicts 2.1 basis points lower Fed funds root mean squared error over the next three quarters. This suggests that Fed communications are costly information in that the more a forecaster processes them, the more precise their expectations become. Despite improving policy expectations, a one-standard-deviation increase in Fed research predicts 5.0 and 14.0 basis points higher forecasting errors for GDP growth and quarterly annualized inflation, respectively. This pattern is consistent with Fed research crowding out independent economic research and reducing economic forecast precision.

I next assess the implications for expectations of Treasury yields, which incorporate both expectations of the short-term rate and economic outcomes. I show that a one-standard-deviation increase in Fed research predicts higher realized forecasting error for Treasury yields, especially in the "belly" of the curve. This pattern persists after controlling for near-term Fed funds, GDP, and inflation forecast errors. I also find no evidence that the relationship strengthens when term-premium changes are larger. These results suggest a role for longer-term short-rate expectations.

Fed communications also draw forecasters' economic expectations toward the Fed's view by shifting their focus toward the information it emphasizes. This issue is central to current

debates over Fed communication and echoes the longstanding "hall of mirrors" concern discussed by Bernanke (2004).[5] Prior empirical tests of the information effect have primarily focused on average responses to Fed communications, using measures such as monthly GDP surveys (Bauer and Swanson, 2023) or daily news text sentiment (Acosta, 2026). I instead compare economic forecast revisions across forecasters who expected different communications from the same Fed announcement. I first show that a hawkish market-based monetary policy surprise (an unexpected increase in short-term interest rates) predicts downward GDP forecast revisions on average, consistent with standard macroeconomic predictions (Bauer and Swanson, 2023). However, in the cross-section, forecasters respond in the opposite direction. Within an announcement, forecasters who are more hawkishly surprised by the content of Fed communications revise their GDP growth and inflation forecasts upward relative to peers, with revisions similar in magnitude to the average response. Communications that uniquely surprise forecasters may require more independent interpretation, shifting more attention toward the Fed. For market-wide surprises, forecasters can instead rely on market prices and their peers' analysis. Consistent with this interpretation, I show that following an announcement forecasters' subsequent Fed research is more closely related to the specific communications they predicted incorrectly.

The results so far reflect relative differences across forecasters in the cross-section, but leave open the aggregate effects of Fed communications on expectations. Fed communications could improve all forecasters' accuracy while still generating these relative differences. I therefore use the higher-frequency nature of the dataset to compare forecast accuracy around Fed announcements with that outside announcement periods. Around Fed announcements, forecasters' accuracy improves on average for the Fed funds rate, with no clear correspond-

[5]Kevin Warsh: "If the aim is optimal decisionmaking, we should not crowd out views on the economy" ("In Our Time," August 28, 2026, p. 8).

ing improvement for economic outcomes, for which improvements are instead concentrated around economic data releases. Long-run evidence from the Survey of Professional Forecasters additionally shows that forecasters' expectations have become more similar to one another and to Fed staff forecasts without a proportionate improvement in forecast accuracy. Finally, I show that the communication surprises that forecasters realize help interpret asset price reactions to Fed communications. The share of Fed research, cross-forecaster disagreement on communication, and the share of incorrect expectations all help predict the magnitude of risky asset price reactions. I further show that Treasury yields react more strongly to changes in short-term rates when a larger share of forecasters' communication errors concerns the Fed's economic assessment. This provides suggestive evidence that forecasters and markets in aggregate focus their attention on the economic narratives of the Fed.

These findings suggest that Fed communications improve understanding of policy while crowding out independent analysis of the economy. The evidence focuses primarily on professional forecasters, who may respond differently from the investors they inform. Additionally, the dataset used covers the post-2000 period of extensive Fed communication, so the effects on research allocation may differ across periods or central banks. For example, clarifying a poorly understood policy framework may increase demand for economic research by helping investors connect economic news to future policy. More broadly, this paper does not assess the full implications of communication for the Fed itself, including its effects on political support and optimal policy decisions.

This paper contributes to three strands of literature. First, it contributes to a literature studying asset price reactions to Fed communications. This includes papers documenting unusual reactions around Fed announcements for long-term yields (Hanson and Stein, 2015;

Hillenbrand, 2025; Beaudry et al., 2026) and equities (Lucca and Moench, 2015; Laarits, 2025), and imposing structure on asset price reactions to determine transmission mechanisms (Jarociński and Karadi, 2020; Cieslak and Khorrami, 2026; Bauer et al., 2026; Pflueger and Rinaldi, 2022; Cieslak and Pang, 2021). This paper contributes to this literature by interpreting monetary policy shocks through the specific communication surprises realized by forecasters.

The first strand of the literature has motivated a second studying monetary policy transmission mechanisms. Papers in this literature analyze surveys of professional forecasters (Romer and Romer, 2000; Campbell et al., 2012; Nakamura and Steinsson, 2018; Bauer and Swanson, 2023; Karnaukh and Vokata, 2022; Kekre et al., 2024; Bauer et al., 2024; Sastry, 2026; Bianchi et al., 2022; Lunsford, 2020), financial news text (Fisher et al., 2022; Benamar et al., 2021; Kroner, 2025), investor letters (Coggi et al., 2026; Haddad and Muir, 2026), and the text the Fed communicates (Hansen et al., 2018; Hansen and Kazinnik, 2024; Gómez-Cram and Grotteria, 2022). Most closely related are Martinez and Sinclair (2025), who link individual forecasters' expectations of Fed projections to their subsequent economic forecast revisions, and Acosta (2026), who measures changes in aggregate economic expectations using financial news text. My dataset links the comprehensive set of expected Fed communications to the same forecasters' quantitative expectations and crucially to the research underlying them. This allows me to speak to a novel mechanism through which the Fed influences expectations by shaping the attention allocation of market participants.

Third, I contribute to an emerging literature that uses the text of other market professionals to help explain their quantitative forecasts. This includes papers about equity research analysts (Bastianello et al., 2025; Ke, 2026; Wang and Zhou, 2026; Asquith et al., 2005), asset managers (Andolfatto and Bastianello, 2025), and the Fed itself (Laarits et al.,

2026). A related emerging literature studies the text underlying economic and financial expectations of households, financial professionals, and academics through open-ended surveys and experiments (Andre et al., 2022, 2026a,b; Conrad et al., 2026; Jiang et al., 2025). I contribute to these literatures by introducing and analyzing the written reports of professional forecasters, which span the past 26 years and reflect the expectations of important market participants. These professionals also differ from others studied in that they spend a substantial share of their time forming higher-order beliefs, particularly about the Fed. As a result, I provide direct empirical evidence in a high-stakes setting of how own and higher-order belief formation interact, particularly through attention allocation.

The paper is organized as follows: Section 2 details a conceptual framework that guides the empirical tests. Section 3 introduces the dataset, institutional background, and measurement techniques. Section 4 examines how Fed communications influence forecaster attention and topic emphasis. Section 5 studies forecast accuracy, and Section 6 examines the information effect. Section 7 relates average forecaster accuracy and communication surprises to aggregate asset price reactions. Section 8 concludes.

2 Conceptual Framework

I develop a conceptual framework of costly information processing in which forecasters form predictions both about the Fed and economic fundamentals. Fed communication can either provide a free public signal or reduce the marginal cost of Fed research.

2.1 The Forecasting Problem

A forecaster predicts both the Fed's policy rate, F, and economic fundamentals, X. Together these two outcomes form the basis for variables professional forecasters are evaluated on (e.g., the Fed funds rate, GDP, and Treasury yields).[6] Creating forecasts for each outcome requires overlapping but distinct preparation. The Fed often substantially disagrees with forecasters' expectations about the economy (Caballero and Simsek, 2022), and the economic views guiding its policy decisions change over time (Laarits et al., 2026). Understanding the Fed's evolving interpretation therefore requires research beyond that required to form a forecaster's own economic forecast. The forecaster believes that F and X are positively but imperfectly correlated, reflecting that Fed policy responds to economic fundamentals. I represent this relationship with a simple functional form in which the Fed's policy rate, F, deviates from economic fundamentals, X, by a scalar B, which can represent the Fed's economic measurement error or policy preferences.

$$F = X + B, \qquad X \sim \mathcal{N}(\mu_X, \sigma_X^2), \qquad B \sim \mathcal{N}(\mu_B, \sigma_B^2), \qquad X \perp B. \tag{1}$$

To form their predictions, the forecaster engages in costly research that yields noisy signals s_F and s_X about F and X, respectively. Through this framework, greater research effort produces a more precise signal. Let a_X and a_F denote economic and Fed research efforts, respectively. The signals have the following properties:

$$s_F = F + \varepsilon_F, \qquad \varepsilon_F \sim \mathcal{N}(0, 1/a_F), \tag{2}$$

[6] For example, Bloomberg's 2008 ranking evaluated forecasts of GDP, CPI, unemployment, and the federal funds target rate, while the 2011 Lawrence R. Klein Award evaluated forecasts of GDP, CPI, unemployment, and Treasury bill rates.

$$s_X = X + \varepsilon_X, \qquad \varepsilon_X \sim \mathcal{N}(0, 1/a_X). \tag{3}$$

The signal errors are independent of X and B and of each other. This structure follows a lineage of work on attention and costly information acquisition (Simon, 1955; Gabaix, 2019; Veldkamp, 2011), with applications to financial professionals such as equity research analysts (Dessaint et al., 2024), fund managers (Kacperczyk et al., 2016), and professional forecasters (Kohlhas and Walther, 2021). Given the positive relationship between the two forecasting outcomes, the forecaster incorporates both signals for each of their predictions. Specifically, the posterior precision of economic fundamentals is:

$$P_X \equiv \text{Prec}(X \mid s_X, s_F) = \underbrace{\frac{1}{\sigma_X^2}}_{\text{prior precision}} + \underbrace{a_X}_{\text{economic research}} + \underbrace{\frac{a_F}{1 + \sigma_B^2 a_F}}_{\text{Fed research}}. \tag{4}$$

Let $P_F \equiv \text{Prec}(F \mid s_X, s_F)$ denote the posterior precision of the policy rate. Appendix C gives its expression and the cross-learning derivatives.

The extent to which the forecaster learns about an outcome from research on the other outcome, "cross-learning", depends on the amount of research the forecaster has performed and their belief about the correlation between the two outcomes. When the forecaster believes that the Fed's policy rate tracks economic fundamentals more closely, the Fed signal is more informative about X, and the economic signal is more informative about F.

The forecaster values accuracy in both economic and policy forecasts. Let $V_X = P_X^{-1}$ and $V_F = P_F^{-1}$ denote the posterior forecast variances. I use an additive forecast loss function with the parameter $\gamma \in (0, 1)$ governing the relative weight on economic forecast accuracy.

Research efforts a_X and a_F incur costs:

$$C(a_X, a_F) = C_X a_X + C_F a_F + \frac{\kappa}{2}(a_X + a_F)^2, \qquad \kappa \geq 0. \tag{5}$$

The parameter κ governs how quickly marginal research costs increase with total effort. The cost function includes the cross-term $\kappa a_X a_F$, through which research on one object raises the marginal cost of researching the other, reflecting a degree of limited attention. This cost specification mirrors the approach of Dessaint et al. (2024), who model convex research costs for equity analysts and motivate the cost of multitasking through cognitive constraints.

At date 2, the forecaster reports their optimal predictions as the posterior means, $f_X = \mathbb{E}[X \mid s_X, s_F]$ and $f_F = \mathbb{E}[F \mid s_X, s_F]$. The forecaster chooses research at date 1 to minimize forecast loss and research costs:

$$\min_{a_X, a_F \geq 0} \quad \gamma V_X + (1-\gamma) V_F + C_X a_X + C_F a_F + \frac{\kappa}{2}(a_X + a_F)^2. \tag{6}$$

Proposition 1 (Reliance on Fed information)**.** *At an interior allocation, greater Fed research and less independent economic research increase the weight placed on Fed information in the economic forecast, holding other inputs fixed:*

$$\underbrace{\frac{\partial f_X}{\partial s_F}}_{\substack{\text{weight on} \\ \text{Fed information}}} = \frac{1}{1 + \underbrace{\left(\frac{1}{\sigma_X^2} + a_X\right)}_{\substack{\text{independent economic} \\ \text{precision}}} \underbrace{\left(\sigma_B^2 + \frac{1}{a_F}\right)}_{\substack{\text{Fed signal noise} \\ \text{about } X}}}. \tag{7}$$

The forecaster also relies more on Fed information when they believe policy tracks economic fundamentals more closely.

This weight measures how responsive the economic forecast is to the Fed signal. Appendix C provides the proof.

2.2 Fed Communication

Fed communication can influence a forecaster's attention and subsequent expectations in two ways. First, it may provide a free signal about future Fed policy, as it is often expressed in the literature (Morris and Shin, 2002). A free public signal may reduce the amount of private research a forecaster performs on the Fed, potentially increasing their research on the economy. This mechanism is consistent with the views of Alan Greenspan, who argued for greater communication in 2001 saying market participants should not have to "waste effort inferring the stance of monetary policy from diffuse signals."[7] Second, Fed communication may lower the cost of obtaining a given precision about Fed policy. This reduction can induce more private research on the Fed, and potentially less on the economy.

[7] Alan Greenspan, "Transparency in Monetary Policy," October 11, 2001, remarks at the Federal Reserve Bank of St. Louis Economic Policy Conference. Source.

Appendix Figure A1 summarizes the timing of the model. The marginal cost of Fed research, C_F, is realized alongside a free public signal before the forecaster chooses research efforts a_X and a_F. The forecaster then receives the resulting noisy signals s_X and s_F and forms predictions f_X and f_F, before X and F are realized.

2.3 Communication Determining Marginal Cost

I first consider the implications of Fed communication that alters the marginal cost of Fed research, C_F.

Proposition 2 (Lower Fed research costs)**.** *At an interior optimum, a marginal reduction in the cost of Fed research increases Fed research and reduces economic research. Expected squared forecast error falls for Fed policy and rises for economic fundamentals. The weight placed on Fed information in the economic forecast increases.*

The first-order condition for economic research explains the trade-off:

$$\underbrace{C_X + \kappa(a_X + a_F)}_{\text{marginal cost}} = \frac{\overbrace{\gamma + \frac{1-\gamma}{(1+\sigma_B^2 a_F)^2}}^{\text{weighted marginal precision gain}}}{\underbrace{P_X^2}_{\text{total economic precision squared}}}. \tag{8}$$

A reduction in the marginal cost of Fed research induces the forecaster to perform more research on the Fed. Holding economic research fixed, this raises the marginal cost of economic research when cost curvature is positive and reduces its marginal benefit through cross-learning. The forecaster responds by reducing economic research enough to more than offset the economic precision supplied by additional Fed research, so total economic forecast precision falls.

2.4 Communication as a Costless Signal

I next consider Fed communication that provides a free public signal. Specifically, before choosing research effort, the forecaster observes a signal of the form

$$s_0 = F + \varepsilon_0, \qquad \varepsilon_0 \sim \mathcal{N}(0, 1/\tau), \tag{9}$$

whose error is independent of the underlying shocks and private signal errors. The precision of the public signal is $\tau > 0$. I vary this precision holding research costs and forecasting preferences fixed.

Proposition 3 (Free public information)**.** *At an interior optimum with increasing marginal research costs ($\kappa > 0$), an increase in public-signal precision reduces total private research and improves both policy and economic forecast precision. If $C_F \geq C_X$, private economic research increases, private Fed research decreases, and the share of private research devoted to the economy rises.*[8]

The public signal allows the same precision to be obtained with less private effort. With positive cost curvature, the resulting reduction in total effort lowers the marginal cost of each research activity. These lower costs make it worthwhile to purchase additional forecast precision. Appendix C.1 derives the responses.[9]

[8] This condition means that researching the Fed is at least as costly as developing one's own assessment of the economy.

[9] With linear research costs, $\kappa = 0$, public information replaces private Fed research one-for-one until private Fed research reaches zero, leaving economic research and both forecast precisions unchanged. Once private Fed research reaches zero, further public information need not improve economic precision.

2.5 Discussion and Empirical Tests

Empirically distinguishing between the two mechanisms of Fed communication depends on two questions: 1) how costly Fed communications are to process, and 2) how limited forecasters' research capacity is in practice.

Communication Channel Predictions

Prediction category	**Free public signal**	**Cost channel**
Private research allocation	Fed ↓	Fed ↑
	Economic ↑	Economic ↓
Forecast accuracy	Policy ↑	Policy ↑
	Economic ↑	Economic ↓

Notes: The columns describe higher public-signal precision and lower Fed research costs, respectively. Predictions hold for interior allocations, with $\kappa > 0$ and $C_F \geq C_X$ in the public-signal column.

Fed communications can provide information directly while also changing incentives for costly Fed research. For a given reduction in Fed research costs, the cost-channel predictions persist when the accompanying increase in public-signal precision is sufficiently small. This follows from continuity of the optimal allocation, provided it remains interior. Fed communication can therefore improve policy forecasts while worsening economic forecasts even when it also provides free public information. This trade-off parallels Dessaint et al. (2024), who show that cheaper short-term information can improve short-term earnings forecasts while worsening long-term forecasts by reallocating analysts' research.

The framework does not consider complementarity in forecast preferences, whereby a

more precise Fed forecast increases demand for economic precision (see Goldstein and Yang, 2015, for related complementarities in information acquisition). Sufficiently strong complementarity could allow both types of Fed communication to increase economic research and accuracy.

Empirically, I examine how Fed communications influence research allocation, whether attention allocation predicts a trade-off in forecast accuracy, and how forecasters incorporate Fed information into their economic expectations. I also examine changes in average policy and economic forecast accuracy around announcements and over time.

3 Data

This section introduces the primary dataset of this paper, the written reports of professional forecasters. First, I review the institutional background for professional forecasters. Next, I detail the sample construction and summary statistics of the data. Then, I explain how I extract information from reports.

3.1 Institutional Background

Professional forecasters are influential specialists in macroeconomic forecasting and Fed watching. They typically reside in sell-side banks where they can sell their expertise to a wide range of clients with more general responsibilities, such as portfolio managers. As former Chase Chief economist John Lipsky put it: "We are doing the digging, analysis and thinking that they would do if they had the time" (Fuerbringer, 1999). A large share of asset managers use the research from professional forecasters: in 2016, 535 institutions managing $9.7 trillion in U.S. fixed-income assets ranked professional forecasters in *Institutional*

Investor's annual survey (Johnson, 2016). Asset managers also find professional forecasters' research highly valuable; a 2018 survey from RBC found that a majority of FICC and multi-asset investors considered macro-research "significantly important" to their investment process (RBC Capital Markets, 2018).[10] Their economic outlooks also reach a broader audience through financial news coverage and appearances on outlets such as CNBC.[11]

Forecasters also provide a window into the broader thought process of markets. By serving a wide range of clients, they gather questions and concerns from investors and tailor their research to address them. As Jay Bacow, Morgan Stanley's co-head of securitized products research, put it: "... we talk to just about every relevant client in the market. That means that we often have a good pulse of what conversations are happening in the market, which can be difficult for the buy side to ascertain without the sell side's help."[12] Forecasters' written reports document their expectations and the analysis underlying them, making it possible to observe how they allocate analytical effort. Sell-side communication rules and firms' research policies further support the accurate representation and fair dissemination of forecasters' views.[13]

Recognizing the importance of professional forecasters, the Fed, industry, and academic researchers have used surveys of their expectations. For example, the New York Fed regularly surveys the professional forecasters at primary dealers, recognizing that their research can "inform the views of their clients" and that their expectations can serve as a "rough proxy for

[10]RBC Capital Markets (2018, p. 27) reports that 60% of Fixed-Income, Currency, and Commodities (FICC) and multi-asset respondents rated research as significantly important to their investment process, 28% as slightly important, and 12% as having no impact.

[11]See, for example, Goldman Sachs chief economist Jan Hatzius's CNBC interview, October 2, 2026.

[12]Alexandra DeLuca, "What Happens When No One Has Experience Outside a 'Big Fixed Income Bull Market'," *Institutional Investor*, November 30, 2023.

[13]See FINRA Rule 2210(d)(1) on non-misleading communications. Forecasters also apply dissemination and certification safeguards comparable to Rule 2241(g) and Regulation AC, despite their exclusions for economic commentary; e.g., Goldman Sachs, *US Economics Analyst*, March 30, 2025, pp. 10, 12.

investor expectations" (Correia Golay et al., 2013). Moreover, Diercks and Jendoubi (2023) find that expectations reported by institutional investors and professional forecasters in the New York Fed's survey do not exhibit any systematic, persistent differences. Other long-standing surveys of forecasters such as the Survey of Professional Forecasters, since 1968, and the Blue Chip Economic Indicators survey, since 1976, have supplied data for a wide range of academic papers, e.g., (Romer and Romer, 2000; Coibion and Gorodnichenko, 2012; Bauer and Swanson, 2023). More broadly, Bloomberg regularly surveys professional forecasters to form consensus expectations for economic data releases, which are then prominently displayed to investors.

3.2 Sample

Despite the importance of professional forecasters, analysis of their written reports remains limited. In this paper, I collect over 300,000 reports from professional forecasters spanning 2000–2025.[14]

I collect reports through proprietary client portals as well as two aggregator platforms, Refinitiv Eikon LSEG Workspace and Investext from Mergent Online. Table I summarizes the 26 forecasting institutions in the sample. I retrieve reports for 10 institutions directly from their research portals, for 15 from LSEG, and for one from Investext.[15]

This set of forecasters in my sample well reflects the set of investment banks typically in surveys of professional forecasters. Table II compares the forecasters in the sample to those included in seven different surveys of professional forecasters, showing that my sample includes 32% of survey participants on average across the seven surveys. This reflects varied

[14]Earlier studies analyze smaller collections of professional forecasters' reports, including Rybinski (2021), Hazell and Hobler (2025), and Ebener and McGregor (2026).

[15]The coverage of LSEG and Investext largely overlaps, but Investext was discontinued in August 2025 so coverage from LSEG is preferred.

participation rates: forecasters in my sample contribute to four out of the seven surveys on average. Overlap is high for sell-side banks (63%) and primary dealers (73%), and low for "other" forecasters (5%). The "other" category includes independent researchers, academics, non-financial corporates, and other government agencies. Overall, my sample most closely reflects the composition of the NY Fed primary dealers survey.

Figure II, Panels A and B, shows report coverage over time. The volume of reports grows as more forecasters enter the sample. This increase primarily reflects how far back their report archives extend, rather than when institutions established forecasting teams.[16] In the cross section, forecasters differ in their report frequency and average report length, with some forecasters focusing more on frequent short publications and others on infrequent lengthy publications (Appendix Figure A2).

3.3 Measurement

Forecaster reports consist of unstructured text, tables, and figures. Embedded in this data are macrofinancial expectations and their rationale. Large Language Models (LLMs) provide a tractable way to extract these expectations. In this section, I describe the measurement of three key components of forecaster reports: (i) the share of research devoted to the Fed, (ii) quantitative macrofinancial forecasts, and (iii) expected Fed communications. The appendix describes each measurement process in more detail.

3.3.1 Measuring Fed Research

I define Fed attention as the share of analytical text that directly reasons about or is motivated by the Fed. Two key questions pin down the focus of research. First, what is the

[16] For example, Nomura and Comerica contributed forecasts to the January 2000 *Blue Chip Economic Indicators* survey (p. 2), although their report archives begin in 2009 and 2018, respectively.

forecaster reasoning about in a given sentence or paragraph? Second, how did the forecaster explicitly motivate their research report? I divide each report into paragraphs of text, yielding over 2.4 million paragraphs. I then use LLMs to measure the share of text in each paragraph devoted to reasoning about the Fed, reasoning about the economy, and factual reporting. For example, forecasters often reason about the Fed, as Goldman Sachs writes: "Although this statement reinforced our belief that Fed officials would prefer to move slowly, persistent signs of strength in the economy will probably convince them to tighten again by 25 basis points at the March 21 FOMC meeting." Forecasters reason about the economy as well, such as Société Générale saying: "We expect China's activity growth to stabilise before the year-end, mostly thanks to a bigger push on credit-backed infrastructure investment."

However, measuring attention at a paragraph level alone misses Fed-motivated economic research, since forecasters often explicitly motivate their economic research with the Fed. For example, Société Générale motivates the same research on China by saying: "With the FOMC's data dependency now apparently extending beyond its shores, notably to China, we discuss our near-term outlook for Chinese data." Such research may help forecasters understand the Fed's economic assessment even when they disagree with it, as in Caballero and Simsek (2022). Tracking that assessment may require economic analysis distinct from the research forecasters would otherwise undertake to form their own economic expectations. In total, I define Fed attention as the text-weighted sum of reasoning about the Fed and the Fed-motivated share of economic reasoning, divided by all reasoning, excluding factual reporting.

The LLM identifies the motivation share that is attributed to the Fed at the article level. For both measurement tasks, the LLM assigns scores from 0 to 100, allowing paragraphs to contain multiple types of reasoning and articles to have multiple motivations. Throughout

the paper, I assess sensitivity to alternative measures, such as Fed-related keyword frequency.

I assume that the written research of forecasters reflects the timing and allocation of their underlying research effort. Related evidence comes from equity research analysts, who face similar regulations and working environments: topic emphasis in reports is predictive of their quantitative forecasts and forecast errors (Bastianello et al., 2025; Ke, 2026). Financial news articles have likewise been used to measure macroeconomic expectations at a daily frequency (Acosta, 2026). These studies motivate using the allocation of report text as a proxy for the allocation of research effort from the conceptual framework.

My measure differs from prior textual measures of Fed attention in both its data and methods. Prior literature typically measures Fed attention using the frequency of news articles containing Fed-related keywords (Boguth et al., 2019; Fisher et al., 2022; Kroner, 2025). I instead measure Fed attention in the reports of professional forecasters, who cater to institutional investors and express their own forecasts and the reasoning behind them. Further, I measure Fed attention using LLMs rather than keyword counts, allowing me to distinguish analytical reasoning from factual reporting.

Appendix B.1 describes the measurement process in more detail and provides worked examples.

3.3.2 Quantitative Macrofinancial Forecasts

Forecaster reports frequently reference their quantitative macrofinancial expectations across several common variables such as GDP, the Fed funds rate, and Treasury yields across tenors.

I develop a multimodal pipeline to extract quantitative macrofinancial forecasts from tables, figures, and text using the broader context of each report (Appendix B.2). I further provide detailed instructions that help the LLM distinguish forecasters' own expectations

from consensus forecasts and identify the units and transformations used for GDP forecasts. While commonly used surveys provide these expectations of forecasters at a monthly or quarterly frequency, the reports allow measuring changes in their published expectations at a higher frequency.

Table III summarizes the extracted forecasts from the pipeline. For most variables, forecasts are reported weekly and for the subsequent four to five quarters. I validate the extracted forecasts against forecaster-level submissions from the Reuters Economic Polls. Table A2 reports the results: correlations between extracted and poll-submitted forecasts range from 0.975 to 0.998 for GDP growth, year-over-year headline CPI inflation, the Fed funds rate, and 2-year and 10-year Treasury yields, based on 82,581 matched observations.[17]

3.3.3 Expectations of Fed Communications

I develop a second measurement pipeline that extracts the specific expectations for Fed communications for each scheduled announcement. Forecasters detail the specific communications they expect from the Fed, which span all aspects of the announcement like the press conference, statement, and economic projections the Fed releases. These expectations are typically published in the final week before the announcement (Appendix Figure A3). For example, before the June 2024 FOMC meeting, Citi forecast an unchanged longer-run dot and an upgrade to the statement's inflation language, and expected Powell to discuss the current policy stance as "restrictive".

The extraction proceeds in three stages (Appendix B.3). First, I extract the expected communications for each forecaster–meeting pair in a semi-structured way. For example, I categorize which part of the announcement the communication pertains to (e.g., statement

[17]The Reuters comparison does not cover quarterly annualized CPI inflation or 5- and 30-year Treasury forecasts.

vs press conference) and which qualifier words, if any, were used (e.g., "unlikely" vs "almost certainly"). I further fit these qualifier phrases to realized likelihood to provide probability weights (Appendix Figure A4). Second, I organize these expectations into a common questionnaire for each meeting, allowing comparisons across forecasters. Third, I query the pre-announcement reports of each forecaster to answer the questionnaire and compare each answer with the realized announcement text to determine whether the predicted communication occurred. FOMC communication text (statements, SEP projections, dot plots, press conference transcripts, and minutes) comes from the Federal Reserve's public archive.

The resulting questionnaires contain 5,752 questions across 207 FOMC meetings, with 43,767 scored forecaster–question observations (Table IV). This amounts to nearly 28 questions per meeting and approximately eight forecasters per question, reflecting differences in which communications forecasters address at each meeting. Error rates are lowest for policy decisions (10%), followed by press conferences (21%), statements (27%), and economic projections (42%; Table IV). The number of questions per meeting increases from an average of eight in 2000–2004 to 42 in 2021–2025, alongside the increase in communications from the Fed.

The results of this measurement pipeline provide a detailed and timely set of expected communications for Fed meetings. The canonical measurement of monetary policy surprises uses changes in market interest rates in narrow windows around Fed communications, reflecting the average market revision in interest rate expectations. The closest precedents capturing other aspects of Fed communication surprises are the New York Fed's Survey of Primary Dealers and the Bloomberg SEP survey, which Martinez and Sinclair (2025) use to measure surprises about the FOMC's median projections. My sample begins in 2000, compared with January 2011 for publicly available New York Fed survey results and December

2017 for Bloomberg's SEP median-projection questions.[18] I also measure expectations for statements and press conferences at meetings both with and without SEP releases. My measure captures expectations as they are published up to the day of the announcement, while surveys typically capture expectations a week or more beforehand (Appendix Figure A3). Moreover, surveys predefine a set of questions that may not match the level of granularity at which forecasters focus. For example, ahead of the September 2020 meeting, forecasters expressed differing expectations about how many FOMC participants would project rate hikes by 2023, even when they agreed that the median projection would remain at zero.

4 Fed Communications and Market Attention

In this section I show the influence of Fed communications on forecaster attention. I first show that forecasters devote a substantial share of their written research to the Fed, on average 25% of analytical text. I then show that the Fed share of research is highly cyclical, nearly doubling from 20% to 40% in the lead-up to scheduled announcements, with an accompanying increase in Fed-oriented text volume. Finally, I show that specific topics that Fed communications emphasize influence the subsequent topics that forecasters emphasize.

4.1 Fed Communication and Fed Attention

On average, forecasters spend a quarter of their analytical text on the Fed. Figure II, Panel C, shows substantial variation over time, with the annual Fed research share ranging from approximately 17% to 36%. Fed attention also appears to rise following expansions in Fed communication (Figure I, Panel A). For example, the Fed research share rose from about 20%

[18]The New York Fed began surveying primary dealers in 2002; its publicly available results cover January 2011 onward. See the New York Fed's overview of its monetary policy expectations surveys.

in 2011, when press conferences were introduced, to 30% in 2013. When press conferences were extended to every scheduled meeting in 2019, the share rose from 23% in 2018 to 32% in 2019. Institutions also differ substantially in their average research allocation: Fed research shares range from under 20% to around 60% of analytical text (Figure II, Panel D).

Figure III plots the Fed research share around FOMC announcements documenting a near doubling of Fed research intensity from 20% to 40% in the buildup to scheduled announcements and a smaller but noticeable surge around the release of minutes. Notably, the weekly share of research about the Fed peaks before each scheduled announcement suggesting the surge in research largely reflects preparation for the announcement, rather than summarization or postprocessing.

Table V relates daily research shares and volumes to FOMC announcements and minutes releases. I use reports from all 26 forecasters, comparing weekdays excluding federal holidays within each forecaster's observed report years. Regressions include forecasting institution and weekday fixed effects and indicators for proximity to GDP, payrolls, and CPI releases. I distinguish the six days before and the announcement date from the following seven days, comparing each window with other eligible dates.

Before and on announcements, Fed research share is 16.1 percentage points higher (column 1), Fed research volume rises by 112.5% of its outside-window mean (column 3), and non-Fed economic volume falls by 4.3% of its outside-window mean (column 5). These results describe a shift in published analysis toward the Fed consistent with Proposition 2, alongside an expansion in analytical output. Fed research remains elevated after announcements, while the change in non-Fed economic volume is statistically insignificant. The additional Fed research share before and on announcements increases by 2.3 percentage points per decade (column 2). Minutes releases are also associated with more Fed-oriented text.

4.2 Fed Topic Emphasis and Forecaster Topic Emphasis

I next ask whether the content of Fed communications influences the content of subsequent forecaster economic research. Specifically, I relate changes in each forecaster's discussion to the content of the minutes, controlling for the forecaster's pre-release discussion, prior-year discussion, previous change in discussion, and the content of earlier minutes. The three-week publication lag of the minutes helps mitigate a central concern with this comparison: the Fed and markets often discuss similar topics due to common information shocks, such as geopolitical events or macroeconomic data releases. The pre-release research controls capture forecasters' discussion after the meeting but before the minutes become public.

I measure content using a 40-topic model trained on FOMC transcripts and, separately, dense text embeddings, which represent the meaning of text as numerical vectors. The pre- and post-release windows contain two complete calendar weeks, as detailed in Table VI. I exclude the week of the minutes release to reduce the influence of previews and immediate summaries of the minutes.

Table VI shows that forecaster discussion shifts toward the content of the current minutes under both measures. For the topic model, a one-percentage-point increase in a topic's share of the minutes is associated with a 0.11-percentage-point increase in its share of forecaster discussion, conditional on the controls. This suggests that Fed communications further influence the specific topics forecasters research within their economic research.

The coefficients on the controls for prior research suggest short-run reversion in topic coverage. Holding prior-year coverage fixed, greater emphasis on a topic immediately before the release predicts a smaller subsequent increase or a larger decline in its share of discussion. Holding pre-release coverage fixed, greater emphasis over the prior year predicts a larger subsequent increase. Together, these patterns suggest that forecasters shift their attention

across topics from period to period while maintaining a persistent research agenda.

5 Fed Attention and Forecast Accuracy

In this section I explore the role of Fed attention in forecasters' quantitative macrofinancial expectations. I begin by showing that Fed attention predicts improved forecasts of the Fed funds rate and worse forecasts of economic outcomes. I next show that Fed attention predicts worse expectations of Treasury yields, particularly in the "belly" of the curve.

5.1 Identifying Assumptions

I relate realized forecast error to Fed research share, controlling for date and institution fixed effects, the log of total research volume, and forecast distance from consensus. Through the lens of the model, I expect attention towards the Fed to lead to increased precision of forecasts for the Fed funds rate. Forecast precision for economic outcomes, such as GDP or inflation, will depend on the extent to which Fed attention crowds out economic attention.

There are several empirical challenges in estimating the relationship between Fed attention and forecast accuracy. First, periods of heightened uncertainty, such as the Great Financial Crisis or the pandemic, may cause both increased attention to the Fed and larger realized forecast errors. I address this concern using date fixed effects, comparing forecasters facing the same aggregate information environment. Second, forecasters may persistently differ in their skills or the types of clients they serve. Institution fixed effects account for these persistent differences. The estimates use differences from each institution's usual Fed

research share, comparing forecasters on the same date. The resulting specification is:

$$L_{it}^{v} = \beta^{v} A_{it} + \delta_{t}^{v} + \mu_{i}^{v} + \gamma^{v} \log R_{it} + \theta^{v} D_{it}^{v} + \varepsilon_{it}^{v}. \tag{10}$$

Here, L_{it}^{v} is forecaster i's realized forecast error for outcome v, A_{it} is the share of research devoted to the Fed, R_{it} is total research volume, and D_{it}^{v} is the absolute distance of the forecast from consensus. The terms δ_{t}^{v} and μ_{i}^{v} denote date and institution fixed effects. I primarily measure Fed attention using the broad research share defined in Section 3.3.

For β to identify the effect of attention on forecast error, the time-varying forecaster-level differences that drive variation in attention must not affect forecast error other than through attention allocation. One concern is that forecasters may vary over time in their forecasting ability. For example, a forecaster may excel at understanding supply shocks, but struggle to decode Fed balance sheet policy. Forecasters may play to their strengths, performing more research or writing more extensively about the topics they are already more confident in. However, if confidence relative to peers drove research allocation, I would expect forecasters to anchor more closely to consensus for forecasts in which they are less confident. That is because forecasters face reputational penalties for being on the wrong side of consensus and can easily observe their peers' forecasts.[19] I test this implication in Appendix Table A3, showing that attention towards the Fed predicts more out-of-consensus Fed funds forecasts, with statistically insignificant relationships for GDP and inflation forecasts. This pattern is consistent with Fed research increasing confidence in Fed funds forecasts, with little evidence of an associated decline in confidence in economic forecasts. Variation in attention could instead reflect changes in forecasting preferences, which may stem from differences in clientele networks served.

[19] Forecasters can observe their peers' forecasts on Bloomberg, for example on the ECFC screen.

5.2 Fed Attention and Economic Forecast Accuracy

I begin by testing the relationship between Fed attention and forecast accuracy for the Fed funds rate, GDP, and inflation. I measure forecast error as the root mean squared error over the next three quarters for each variable. This averages forecasting performance for volatile quarterly change variables like GDP over a set of common forecasting horizons. For headline CPI inflation, I use directly reported quarter-over-quarter annualized forecasts of the quarterly-average price index.[20] I evaluate GDP and inflation forecasts against their first available releases, excluding subsequent data revisions, and interest rate forecasts against quarter-end realizations. I measure Fed research share over the forecast date and the preceding 30 days to allow for persistent effects of research on forecasting accuracy.

Greater Fed attention predicts lower Fed funds forecast errors but higher GDP and quarterly-inflation errors (Table VII). Figure IV illustrates these opposing relationships for Fed funds and GDP across quintiles of residual Fed research share. Specifically, a 15-percentage-point increase in Fed attention, one standard deviation, is associated with 2 basis points lower Fed funds RMSE, 5 basis points higher GDP RMSE, and 14 basis points higher quarterly-inflation RMSE. The broader sample of directly reported year-over-year inflation forecasts yields no statistically significant relationship between Fed attention and RMSE over the next three quarters.

The estimated effects are substantial relative to differences in accuracy across forecasters. They represent approximately 15%, 12%, and 33% of the average absolute difference in RMSE between peers predicting the same Fed funds, GDP, and inflation outcomes on the

[20] I focus on quarterly inflation because it measures price changes within each forecast quarter, although fewer forecasters report it. Year-over-year inflation combines price changes over four quarters, partly observed at the forecast date, and allows offsetting quarterly errors to cancel.

same date.[21] Relative to average RMSE, which also reflects unforecastable shocks, the estimates correspond to a 5% reduction for Fed funds and increases of 3% for GDP and 6% for inflation. The GDP and inflation effects are comparable in magnitude to gains documented by Bianchi et al. (2022), whose machine-learning model reduces RMSE by approximately 6% and 8%, respectively, relative to median professional forecasts. Although their horizons and inflation measure differ, this comparison suggests economically meaningful changes in forecast accuracy associated with Fed attention.

Column 1, which excludes fixed effects, shows a positive relationship between Fed attention and realized forecast error for Fed funds and GDP forecasts, combining variation across dates and institutions. During periods of heightened aggregate uncertainty, forecasters may devote more research to the Fed while also making larger forecast errors. Adding date fixed effects in column 2 compares forecasters facing the same economic environment and reverses the relationship for Fed funds.

The results show that forecasters' attention allocation predicts their forecast errors, providing evidence consistent with limited attention and the accuracy trade-off in Proposition 2. A few factors may help explain the limited attention of forecasting teams. First, team attention is limited by that of the chief economist who delegates research projects and aggregates the resulting signals. Second, forecasters serve clients with limited attention who prefer simple and short forecast explanations.[22] Both factors may raise the marginal cost of incorporating additional research into a coherent outlook that clients can readily understand.

[21] Within each estimation sample, I calculate each observation's average absolute difference in RMSE from peers predicting the same target quarters on the same date, then average across observations.

[22] Taimur Baig (DBS) notes that clients "read the first page, or they probably read the first couple of bullets." Torsten Slok (then at Deutsche Bank) explains that "the whole challenge for me really is to tell you something that is easy to understand."

5.3 Fed Attention and Treasury Forecast Accuracy

I next explore the implications of Fed attention for expectations of the Treasury yield curve, which embed both expectations of the short rate and the economy. I extract expectations for the 2-, 5-, 10-, and 30-year rates from forecasters' reports and use the same measures of Fed attention and research volume as in the preceding analysis. These expectations represent the firm's Treasury forecasts, produced either by the forecasters themselves or by a closely collaborating team specializing in Treasury markets. Comparing Treasury yield expectations with economic expectations within the same firm follows prior studies using institution-level survey forecasts (e.g., Singleton, 2021). I then continue the specification in equation (10), controlling for date and institution fixed effects, but use absolute forecast error and estimate the relationship separately at each forecast horizon.

Attention towards the Fed predicts lower realized forecasting errors for Fed funds forecasts over the next four quarters. Greater Fed attention is associated with larger forecast errors for two-year, five-year, and ten-year Treasury yields, especially at longer forecast horizons. The pattern is strongest in the "belly" of the curve: four quarters ahead, a 15-percentage-point increase in Fed attention predicts 7.8 basis points higher five-year yield forecast errors. The increase is equivalent to 21% of the typical cross-forecaster accuracy gap, compared with 15% for two-year and 8% for ten-year yields. The effects on Fed funds forecasts appear to wane at five- to six-quarter forecasting horizons. Table VIII presents the results on a balanced sample of forecaster $\times$ date pairs reporting each of the next four forecasting quarters, while Appendix Table A4 presents the results for the next six quarters allowing the sample composition to vary.

The effect of Fed attention on Treasury expectation accuracy could reflect errors about the expected path of future short-term rates or term premia. I test whether the association

between Fed attention and Treasury forecast errors varies with the size of realized changes in the expected-short-rate and term-premium components of yields. I find no evidence that this relationship strengthens when term-premium changes are larger (Appendix Table A5). The association between Fed attention and larger Treasury forecast errors remains after controlling for the institution's near-term Fed funds, GDP, and inflation forecast errors (Appendix Table A6). One potential interpretation is that the results may be driven by the longer-run economic and Fed expectations of forecasters, which are not quantitatively reported.

5.4 Robustness

I next check robustness to alternative measurements of Fed attention, control specifications, and inference procedures. Appendix Table A7 varies the window over which Fed attention is measured, showing robustness to a 14-day window and an exponentially weighted mean. Appendix Table A8 shows robustness to direct Fed reasoning, article-level Fed motivation, and Fed keyword frequency. This suggests that results are not driven by the specific definition of Fed attention used. I next show robustness to replacing institution fixed effects with chief economist fixed effects (Appendix Table A9), which allow for changes in forecasters' research as their leadership changes.

Finally, Appendix Table A10 shows robustness to alternative inference procedures. The baseline standard errors are two-way clustered by institution and target bundle (target quarter for individual-horizon regressions). With relatively few institutions and overlapping forecast horizons, conventional clustered inference may overstate precision. The main pattern remains when using cluster-based t critical values or Driscoll–Kraay standard errors with a one-year bandwidth.

6 Fed Attention and Economic Expectations

The preceding results link greater Fed attention to less accurate economic forecasts. I next examine whether Fed communications draw those expectations toward the Fed's own expectations by shifting forecasters' attention toward the information it emphasizes. Specifically, I compare economic forecast revisions across forecasters who expected different communications from the same Fed announcement. This question is central to current debates over Fed communication and echoes the longstanding "hall of mirrors" concern discussed by Bernanke (2004).[23]

Greater reliance on the Fed's views can strengthen asset price reactions to its communications (Hillenbrand, 2025), reduce the independent information the Fed can glean from market outcomes (Bernanke and Woodford, 1997; Morris and Shin, 2005), and lead to more correlated forecast errors and asset returns (Veldkamp, 2006). Prior papers have named this phenomenon the "Fed information effect" and modelled it as Fed communications reflecting the Fed's private information (Nakamura and Steinsson, 2018). Prior empirical tests have related economic forecast revisions to market-based monetary policy surprises (Nakamura and Steinsson, 2018; Bauer and Swanson, 2023), including tests using individual inflation forecasts (Friedrichs and Miller, 2024). More directly, Martinez and Sinclair (2025) relate individual economic forecast revisions to forecaster-specific surprises about SEP projections. I extend this approach to expectations about the broader content of Fed announcements, including statements and press conferences, using a sample beginning in 2000. Through my framework of limited attention, I instead argue that the Fed influences economic expectations by shifting market attention and thus the precision of the signals market participants

[23] Kevin Warsh argues that "If the aim is optimal decisionmaking, we should not crowd out views on the economy" ("In Our Time," August 28, 2026, p. 8).

receive from the Fed and the economy.

6.1 Identifying Forecaster-Level Fed Information

While I observe forecasters' attention allocation, I do not directly observe the signals they extract from their own economic research and their Fed research. I observe only their resulting forecasts, which combine both. I therefore turn to scheduled Fed announcements, for which forecasters differ in their specific expectations and thus in the signals they ultimately receive. I summarize these differences using a cross-sectional communication surprise score.

For each meeting, I use a large language model to compare pairs of forecasters' pre-announcement expectations with the realized communications and judge which forecaster was more hawkishly surprised. I then fit a Bradley–Terry model to combine these judgments into a score for each forecaster. Scores are centered within the meeting: a higher score indicates a more hawkish surprise relative to peers. I evaluate all Fed communications jointly, since each can convey information about the Fed's economic outlook. Pairwise comparisons allow the model to weigh the importance of different communications rather than simply count incorrect predictions. They also avoid requiring separately generated numerical ratings to be calibrated across prompts (Qin et al., 2024). Appendix B.3 provides construction details and an example using Citi and Wells Fargo forecasts for the June 2024 announcement.

Conceptually, a forecaster-specific communication surprise may be more costly to process than an aggregate surprise. Forecasters can rely on their peers' analysis and market-wide asset price reactions to help process common surprises, whereas individual surprises require more independent interpretation.

Forecasters focus on the specific Fed communication surprises they realize. Appendix Table A11 shows that forecasters' subsequent Fed-related reasoning is more closely related

to communications they predicted incorrectly. Incorrect expectations predict greater post-announcement similarity between a forecaster's Fed reasoning and the corresponding announcement question, equivalent to approximately 4–6% of mean pre-announcement similarity, including when controlling for pre-announcement discussion.

6.2 Testing for Fed Information in the Cross Section

I compare changes in quantitative forecasts around FOMC announcements with forecaster surprise scores. I measure revisions from the last forecast before an FOMC announcement to the first forecast afterward, with no more than ten calendar days between them, limiting the amount of confounding information (Appendix Figure A5). Meeting fixed effects absorb common forecast revisions associated with each announcement. I additionally control for differential exposure to macroeconomic releases using macroeconomic surprise scores, and for potential delayed incorporation of news using forecasters' initial and lagged expectations. This test assumes that, conditional on these controls, forecaster-specific responses to other news are unrelated to the communication-surprise score.

The preferred specification is

$$\Delta f_{it}^{k} = \alpha_{t}^{k} + \lambda_{i}^{k} + \beta^{k} \widetilde{S}_{it}^{k} + \boldsymbol{\gamma}^{k\prime} \mathbf{Z}_{it}^{k} + \varepsilon_{it}^{k}, \tag{11}$$

where i indexes forecasters, t FOMC meetings, and k Fed funds, GDP growth, or inflation. The dependent variable Δf_{it}^{k} is the forecast revision; $\widetilde{S}_{it}^{k}$ is the communication surprise score scaled by its standard deviation in the estimation sample for outcome k. The terms α_{t}^{k} and λ_{i}^{k} are meeting and institution fixed effects, respectively. The controls $\mathbf{Z}_{it}^{k}$ include the pre-announcement forecast, its level 90 days before the announcement, and macroeconomic

release surprises over the revision window. GDP and inflation specifications also include the initial Fed funds forecast; columns (3) and (5) additionally condition on its contemporaneous revision. These columns ask whether communication surprises predict economic forecast revisions beyond their association with revisions to expected policy rates.

Table IX presents the results. I find a positive relationship between forecaster surprise scores and revisions to the Fed funds rate, validating the score as a forecaster-specific Fed signal. Forecasters also increase their GDP and inflation forecasts in response to a more hawkish surprise, consistent with incorporating information from Fed communications into their economic expectations as in Proposition 1. In the specification that also controls for contemporaneous Fed funds revisions for GDP and inflation, a one-standard-deviation increase in the surprise score is associated with a 0.9 basis point increase in Fed funds forecasts, a 1.0 basis point increase in GDP growth forecasts, and a 0.8 basis point increase in inflation forecasts. The inflation response is significant at the 10% level. The estimates are additionally robust to chief economist fixed effects (Appendix Table A12). Together with the evidence on subsequent research, this pattern suggests that Fed communications shape economic expectations by directing forecasters' attention toward the information they convey.

I next benchmark these results against the canonical approach relating market-based monetary policy surprises to average forecast revisions. I use the U.S. Monetary Policy Event-Study Database (USMPD) of Acosta et al. (2025), which provides high-frequency changes in interest rates and asset prices around FOMC communication events. I use their monetary policy surprise measures, hereafter MPS_t, constructed as the first principal component of changes in money-market futures rates spanning approximately one year.[24] I measure

[24]The measures combine federal funds futures with Eurodollar futures, replaced by SOFR futures beginning in January 2022, and are normalized to have a one-percentage-point effect on the daily change in the one-year

monetary policy surprises and asset returns following the monetary-event definition of Acosta et al. (2025), including press conferences when they occur (Appendix Figure A6).[25]

In the preferred specification, market-based monetary policy surprises predict downward GDP forecast revisions (Appendix Table A13). Using my dataset at a monthly frequency and omitting macroeconomic controls, I reproduce the positive GDP relationship reported by Nakamura and Steinsson (2018). Moving to higher-frequency revision windows alone removes this positive relationship. Further adding controls adapted from Bauer and Swanson (2023) reproduces a negative relationship. In this specification, a one-standard-deviation increase in the market-based monetary policy shock is associated with a 2.2 basis point increase in Fed funds expectations, a 1.6 basis point decline in GDP growth expectations (consistent with Bauer and Swanson (2023)), and a 0.9 basis point decline in inflation forecasts.

7 Aggregate Accuracy and Asset Price Reactions

The cross-sectional results so far provide evidence consistent with an attention tradeoff across forecasters, but leave open the aggregate implications of Fed communications. I next examine average forecast accuracy around FOMC announcements and its long-term trends. I then assess whether forecasters' communication surprises help explain asset price reactions.

7.1 Average Forecaster Accuracy around Announcements

The higher frequency of forecasts in my dataset allows me to compare expectations immediately before and after individual FOMC announcements. Monthly surveys such as Blue

Treasury yield.

[25]The window begins 10 minutes before the statement and ends 60 minutes after the press conference begins, or 20 minutes after the statement when there is no press conference. Appendix Figure A6 illustrates this timing relative to the publication of forecasters' expectations.

Chip combine revisions around announcements with updates elsewhere in the month.

Fed funds forecast error falls by 2.3 basis points around FOMC announcements, while the 0.2 basis point decline in GDP forecast error is not statistically significant (Table X). These unadjusted means compare forecasts for the next three quarters, published together before and after each announcement, at most ten calendar days apart.

The regressions compare these changes with other forecast intervals, controlling for initial forecast error, other economic releases, and forecaster, year-quarter, and weekday effects. The adjusted Fed funds decline is smaller and imprecisely estimated in the full sample, becoming marginally significant when broad crisis years are excluded. GDP forecasts instead improve less around FOMC announcements than over other intervals. Excluding the broad crisis years, GDP forecasts improve more after GDP releases than over other intervals.

Appendix Table A14 shows that Fed funds forecasts are revised more frequently around FOMC announcements, while economic forecasts are more frequently revised around economic data releases. Changes in the expected policy path should also affect forecasts of economic activity and inflation. The lack of a corresponding increase in economic forecast updating around Fed announcements is consistent with either an offsetting information effect or limited attention to processing economic news around FOMC announcements.

7.2 Long-Run Trends in Forecast Disagreement and Accuracy

I next ask whether the longer-run evolution of aggregate forecasts is consistent with greater reliance on the Fed. The Survey of Professional Forecasters (SPF) extends back to the 1970s allowing for a comparison of forecasting performance from before the Fed began increasing communications in the 1990s. Figure V, Panel A, shows that the standard deviation of GDP growth forecasts across forecasters and their root mean squared distance from Fed staff fore-

casts have declined substantially over time, in part reflecting a reduction in macroeconomic uncertainty. However, Panel B shows that realized forecast RMSE divided by each of these two measures is generally higher in the later part of the sample. Some of this pattern may reflect large unanticipated aggregate shocks, such as the Great Financial Crisis. However, Bianchi et al. (2022) show that a machine learning benchmark using real-time information outperforms professional forecasters on average over 1995–2018, with particularly large gains during 2013–2018. Over a period in which the Fed became more transparent and communicated more, forecasters have therefore come to agree more closely with one another and with the Fed, without a comparable improvement in forecast accuracy.

7.3 Aggregate Forecaster Communication Surprises

I next test whether the share of incorrect predictions of Fed communications helps explain the magnitude of asset price reactions. I begin by documenting the relationship between Fed attention, communication disagreement, and communication surprises. Specifically, I measure Fed attention as the average across forecasters of the share of research about the Fed in the week leading up to an announcement. I measure communication disagreement as the share of comparable forecaster pairs that give different answers to the same announcement question (Appendix Figure A7). Finally, I define communication surprise as the probability-weighted share of incorrect expectations, averaged equally across scored forecaster–question observations (Appendix B.3). Figure VI shows that meetings preceded by a higher Fed research share also exhibit greater disagreement about the upcoming communications and higher communication error rates. Greater uncertainty about an announcement may induce more Fed research while also generating more disagreement and larger communication errors.

Table XI relates the magnitude of S&P 500 and 10-year Treasury yield reactions to

Fed research share, communication disagreement, and communication surprise over the full announcement window. Fed research share is positively associated with both reaction magnitudes when entered alone. Its coefficient falls after adding communication disagreement and becomes statistically insignificant after adding realized communication surprise. In the full specifications, a one-standard-deviation increase in the share of incorrect communication predictions is associated with increases of 0.25 and 0.24 standard deviations in the absolute equity return and Treasury yield change, respectively. Thus, forecasters' communication errors contain information about market reactions beyond pre-announcement attention and disagreement. Appendix Table A15 enters each text measure separately and controls for SEP releases, press conferences, and selected balance-sheet announcements. Communication disagreement and surprise remain positively associated with reaction magnitudes after removing the signed response to market-based monetary policy surprises.

7.4 Types of Fed Communication Surprises

I next distinguish communication surprises about the Fed's economic assessment from those about policy. For each meeting, I measure economic content as the share of communication errors concerning the Fed's assessment of economic conditions, outlook, and risks. Appendix B.3 details the question classification and probability weighting.

About half of all realized communication surprises concern the Fed's economic assessment. The positive relationship between market-based monetary policy surprises and Treasury yield changes is stronger when communication surprises contain more economic content. For the 10-year Treasury, the estimated sensitivity to monetary policy surprises is about 32% greater when economic content is one standard deviation above its mean. This suggests that news about the Fed's economic view helps explain the response of longer-term yields.

Table XII relates Treasury yield responses to monetary policy surprises and the economic content of communication surprises. Appendix Table A16 checks whether the stronger yield response reflects higher overall communication-error rates, larger monetary policy shocks, current-rate surprises, or differences between meeting types. Across these specifications, Treasury yields respond more strongly to monetary policy surprises when a larger share of communication errors concerns the Fed's economic assessment.

8 Conclusion

The Federal Reserve influences the expectations of market participants by shaping their attention allocation. I provide evidence for this mechanism using a novel dataset of the written research reports of professional macroeconomic forecasters. More broadly, the dataset links forecasters' written analysis to their quantitative expectations. This creates opportunities to study other expectations, such as those concerning fiscal policy, and other aspects of expectation formation, such as the narratives underlying them.

In this paper, I show that Fed communications shift attention toward the Fed and away from independent economic research. Greater Fed research predicts more accurate Fed funds forecasts but less accurate forecasts of economic outcomes and medium- to longer-term Treasury yields. Fed communications also draw forecasters' economic expectations toward the Fed's view by shifting their focus toward the topics it emphasizes. These dynamics are further reflected in consensus forecast accuracy trends over time and in asset price reactions to Fed communications.

These findings raise questions about how the Fed should communicate with financial markets. In particular, the Fed should consider how the detail it provides about its economic

outlook and expected policy path affects incentives for private economic research. These costs will need to be weighed against other objectives of communication, including public accountability and credibility. Overall, the findings suggest that the increase in Fed transparency over the past thirty years may have improved understanding of near-term policy at the cost of less accurate expectations about economic outcomes and longer-term interest rates.

References

Acosta, Miguel, "The Perceived Causes of Monetary Policy Surprises," *Working Paper*, 2026.

_ , Andrea Ajello, Michael Bauer, Francesca Loria, and Silvia Miranda-Agrippino, "Financial Market Effects of FOMC Communication: Evidence from a New Event-Study Database," Working Paper 2025-30, Federal Reserve Bank of San Francisco 2025.

Andolfatto, Andrea and Federico Bastianello, "From Numbers to Words: Breaking Down Institutional Beliefs," *Working Paper*, 2025.

Andre, Peter, Carlo Pizzinelli, Christopher Roth, and Johannes Wohlfart, "Subjective Models of the Macroeconomy: Evidence From Experts and Representative Samples," *Review of Economic Studies*, 2022, *89* (6), 2958–2991.

_ , Ingar Haaland, Christopher Roth, Mirko Wiederholt, and Johannes Wohlfart, "Narratives About the Macroeconomy," *The Review of Economic Studies*, 2026a. Advance online publication.

_ , Philipp Schirmer, and Johannes Wohlfart, "Mental Models of the Stock Market," *The Quarterly Journal of Economics*, 2026b, *141* (4), 2925–2978.

Asirvatham, Hemanth, Elliott Mokski, and Andrei Shleifer, "GPT as a Measurement Tool," Working Paper 34834, National Bureau of Economic Research February 2026.

Asquith, Paul, Michael B. Mikhail, and Andrea S. Au, "Information content of equity analyst reports," *Journal of Financial Economics*, 2005, *75* (2), 245–282.

Bartik, Alexander, Arpit Gupta, and Daniel Milo, "The Costs of Housing Regulation: Evidence From Generative Regulatory Measurement," *Working Paper*, 2025.

Bastianello, Francesca, Paul H Décaire, and Marius Guenzel, "Mental Models and Financial Forecasts," *Working Paper*, 2025.

Bauer, Michael D. and Eric T. Swanson, "An Alternative Explanation for the "Fed Information Effect"," *American Economic Review*, 2023, *113* (3), 664–700.

Bauer, Michael D, Carolin E Pflueger, and Adi Sunderam, "Perceptions About Monetary Policy," *The Quarterly Journal of Economics*, 2024, *139* (4), 2227–2278.

Bauer, Michael D., Maik Schmeling, and Andreas Schrimpf, "Risk Appetite and Monetary Transmission," *Working Paper*, 2026.

Beaudry, Paul, Paolo Cavallino, and Tim Willems, "Anatomy of a rise: Monetary policy and the post-Covid surge in long-term interest rates," *VoxEU*, CEPR, September 21, 2026.

Benamar, Hedi, Thierry Foucault, and Clara Vega, "Demand for Information, Uncertainty, and the Response of U.S. Treasury Securities to News," *The Review of Financial Studies*, 2021, *34* (7), 3403–3455.

Bernanke, Ben and Michael Woodford, "Inflation Forecasts and Monetary Policy," *Journal of*

Money, Credit and Banking, 1997, *29* (4), 653–684.

Bianchi, Francesco, Sydney C. Ludvigson, and Sai Ma, "Belief Distortions and Macroeconomic Fluctuations," *American Economic Review*, 2022, *112* (7), 2269–2315.

Blinder, Alan S., Michael Ehrmann, Marcel Fratzscher, Jakob De Haan, and David-Jan Jansen, "Central Bank Communication and Monetary Policy: A Survey of Theory and Evidence," *Journal of Economic Literature*, 2008, *46* (4), 910–945.

Boguth, Oliver, Vincent Grégoire, and Charles Martineau, "Shaping Expectations and Coordinating Attention: The Unintended Consequences of FOMC Press Conferences," *Journal of Financial and Quantitative Analysis*, 2019, *54* (6), 2327–2353.

Caballero, Ricardo J. and Alp Simsek, "Monetary Policy with Opinionated Markets," *American Economic Review*, 2022, *112* (7), 2353–2392.

Campbell, Jeffrey R., Charles L. Evans, Jonas D. M. Fisher, and Alejandro Justiniano, "Macroeconomic Effects of Federal Reserve Forward Guidance," *Brookings Papers on Economic Activity*, 2012, *2012* (1), 1–80.

Cieslak, Anna and Hao Pang, "Common Shocks in Stocks and Bonds," *Journal of Financial Economics*, 2021, *142* (2), 880–904.

_ **and Paymon Khorrami**, "Monetary Shocks and Long-Run Neutrality in Asset Pricing," *Working Paper*, 2026.

Coggi, Patrick, Francesco A. Franzoni, Pietro Lazzaretto, and Michael Weber, "Monetary Policy Expectations of Sophisticated Investors," *Working Paper*, 2026.

Coibion, Olivier and Yuriy Gorodnichenko, "What Can Survey Forecasts Tell Us about Information Rigidities?," *Journal of Political Economy*, 2012, *120* (1), 116–159.

Conrad, Christian, Julius Theodor Schoelkopf, Michael Weber, and Frank Brueckbauer, "Beyond the Numbers: Professional Forecasters' Narratives about Inflation and Stock Market Performance," *Working Paper*, 2026.

Correia Golay, Ellen, Steven Friedman, and Michael McMorrow, "Understanding the New York Fed's Survey of Primary Dealers," *Current Issues in Economics and Finance*, 2013, *19* (6), 1–8.

Dessaint, Olivier, Thierry Foucault, and Laurent Frésard, "Does Alternative Data Improve Financial Forecasting? The Horizon Effect," *The Journal of Finance*, 2024, *79* (3), 2237–2287.

Diercks, Anthony M. and Haitham Jendoubi, "Expectations of financial market participants," in "Handbook of Economic Expectations" 2023, pp. 385–410.

Du, Wenxin, Kristin Forbes, and Matthew Luzzetti, "Quantitative Tightening Around the Globe: What Have We Learned?," *Journal of Money, Credit and Banking*, 2026.

Ebener, Luca and Thomas McGregor, "What analysts hear and why it matters for monetary policy," Working Paper Series 3298, European Central Bank October 2026.

Fisher, Adlai, Charles Martineau, and Jinfei Sheng, "Macroeconomic Attention and Announcement Risk Premia," *The Review of Financial Studies*, 2022, *35* (11), 5057–5093.

Friedrichs, Melanie and David S. Miller, "The Current Month Fed Information Effect," FEDS Notes, Board of Governors of the Federal Reserve System November 2024.

Fuerbringer, Jonathan, "Right or Wrong, Wall Street Economists Generate Ideas and Trades," *The New York Times*, October 12, 1999.

Gabaix, Xavier, "Behavioral Inattention," *Handbook of Behavioral Economics*, 2019.

Goldstein, Itay and Liyan Yang, "Information Diversity and Complementarities in Trading and Information Acquisition," *The Journal of Finance*, 2015, *70* (4), 1723–1765.

Gómez-Cram, Roberto and Marco Grotteria, "Real-time price discovery via verbal communication: Method and application to Fedspeak," *Journal of Financial Economics*, 2022, *143* (3), 993–1025.

Haddad, Valentin and Tyler Muir, "Everything but the Interest Rate: Measuring Central Bank Interventions," *Working Paper*, 2026.

Hansen, Anne Lundgaard and Sophia Kazinnik, "Can ChatGPT Decipher Fedspeak?," *Working Paper*, 2024.

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

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

Hazell, Jonathon and Stephan Hobler, "Do Deficits Cause Inflation? A High Frequency Narrative Approach," *Working Paper*, 2025.

Hillenbrand, Sebastian, "The Fed and the Secular Decline in Interest Rates," *The Review of Financial Studies*, 2025, *38* (4), 981–1013.

Jarociński, Marek and Peter Karadi, "Deconstructing Monetary Policy Surprises—The Role of Information Shocks," *American Economic Journal: Macroeconomics*, 2020, *12* (2), 1–43.

Jiang, Zhengyang, Hongqi Liu, Cameron Peng, and Hongjun Yan, "Investor Memory and Biased Beliefs: Evidence from the Field," *The Quarterly Journal of Economics*, 2025, *140* (4), 2749–2804.

Johnson, Thomas W., "The 2016 All-America Fixed-Income Research Team: Key Figures," *Institutional Investor*, July 20 2016.

Kacperczyk, Marcin, Stijn Van Nieuwerburgh, and Laura Veldkamp, "A Rational Theory of Mutual Funds' Attention Allocation," *Econometrica*, 2016, *84* (2), 571–626.

Karnaukh, Nina and Petra Vokata, "Growth forecasts and news about monetary policy," *Jour-*

nal of Financial Economics, 2022, *146* (1), 55–70.

Ke, Shikun, "Dissecting Misreactions in Earnings Expectations: Evidence from Analyst Reports," *Working Paper*, 2026.

Kekre, Rohan, Moritz Lenel, and Federico Mainardi, "Monetary Policy, Segmentation, and the Term Structure," *Working Paper*, 2024.

Kiefer, Nicholas M. and Timothy J. Vogelsang, "A New Asymptotic Theory for Heteroskedasticity-Autocorrelation Robust Tests," *Econometric Theory*, 2005, *21* (6), 1130–1164.

Kohlhas, Alexandre N. and Ansgar Walther, "Asymmetric Attention," *American Economic Review*, 2021, *111* (9), 2879–2925.

Kroner, Niklas, "How Markets Process Macro News: The Importance of Investor Attention," *Finance and Economics Discussion Series*, 2025, (2025-022).

Laarits, Toomas, "Pre-Announcement Risk," *Working Paper*, 2025.

_ , Ben Matthies, Kaushik Vasudevan, and Will Yang, "Policy by Committee," *Working Paper*, 2026.

Lucca, David O. and Emanuel Moench, "The Pre-FOMC Announcement Drift," *The Journal of Finance*, 2015, *70* (1), 329–371.

Lunsford, Kurt G., "Policy Language and Information Effects in the Early Days of Federal Reserve Forward Guidance," *American Economic Review*, 2020, *110* (9), 2899–2934.

Martinez, Andrew B and Tara M Sinclair, "When the Fed Reveals Its Hand: The SEP and Monetary Policy Surprises," *Working Paper*, 2025.

Morris, Stephen and Hyun Song Shin, "Social Value of Public Information," *American Economic Review*, 2002, *92* (5), 1521–1534.

_ and _ , "Central Bank Transparency and the Signal Value of Prices," *Brookings Papers on Economic Activity*, 2005, *2005* (2), 1–66.

Nakamura, Emi and Jón Steinsson, "High-Frequency Identification of Monetary Non-Neutrality: The Information Effect," *The Quarterly Journal of Economics*, 2018, *133* (3), 1283–1330.

Pflueger, Carolin and Gianluca Rinaldi, "Why does the Fed move markets so much? A model of monetary policy and time-varying risk aversion," *Journal of Financial Economics*, 2022, *146* (1), 71–89.

Qin, Zhen, Rolf Jagerman, Kai Hui, Honglei Zhuang, Junru Wu, Le Yan, Jiaming Shen, Tianqi Liu, Jialu Liu, Donald Metzler, Xuanhui Wang, and Michael Bendersky, "Large Language Models are Effective Text Rankers with Pairwise Ranking Prompting," in "Findings of the Association for Computational Linguistics: NAACL 2024" Association for Computational Linguistics 2024, pp. 1504–1518.

RBC Capital Markets, "Execution & Unbundling Post-MiFID II: A Global Multi-Asset Market Survey," Survey report, RBC Capital Markets 2018.

Romer, Christina D. and David H. Romer, "Federal Reserve Information and the Behavior of Interest Rates," *American Economic Review*, 2000, *90* (3), 429–457.

Rybinski, Krzysztof, "Ranking professional forecasters by the predictive power of their narratives," *International Journal of Forecasting*, 2021, *37* (1), 186–204.

Sastry, Karthik A, "Disagreement About Monetary Policy," *American Economic Journal: Macroeconomics*, 2026, *18* (2), 486–522.

Simon, Herbert A., "A Behavioral Model of Rational Choice," *The Quarterly Journal of Economics*, 1955, *69* (1), 99–118.

Singleton, Kenneth J., "Presidential Address: How Much "Rationality" Is There in Bond-Market Risk Premiums?," *The Journal of Finance*, 2021, *76* (4), 1611–1654.

Veldkamp, Laura L., "Information Markets and the Comovement of Asset Prices," *The Review of Economic Studies*, 2006, *73* (3), 823–845.

_ , *Information Choice in Macroeconomics and Finance*, Princeton University Press, 2011.

Wang, Chen and Kangying Zhou, "The Macro Alibi: Subjective Risk Attribution in Analyst Scenarios," *SSRN Working Paper*, 2026.

Woodford, Michael, "Central Bank Communication and Policy Effectiveness," *Working Paper*, 2005.

Tables

Table I: Report Coverage by Contributor

Contributor	Reports	Start year	End year	Source
Morgan Stanley	34,060	2000	2025	LSEG
Goldman Sachs	30,148	2000	2025	Direct
BofA Global Research	28,943	2009	2025	LSEG
Deutsche Bank	28,570	2000	2025	LSEG
JPMorgan	17,833	2000	2025	Investext
Barclays	17,140	2008	2025	Direct
UBS	16,835	2000	2025	LSEG
ING	16,054	2000	2025	LSEG
Société Générale	15,164	2004	2025	LSEG
Citi	13,591	2010	2025	Direct
Nomura	13,505	2009	2025	Direct
NatWest	9,688	2000	2025	LSEG
Credit Suisse	9,509	2000	2023	LSEG
HSBC	9,447	2006	2025	LSEG
Natixis	9,283	2012	2025	LSEG
CIBC	9,122	2000	2025	LSEG
UniCredit	8,477	2007	2025	LSEG
Wells Fargo	7,474	2003	2025	Direct
BMO Capital Markets	5,138	2000	2025	LSEG
Scotiabank	4,008	2012	2025	Direct
Oxford Economics	3,606	2014	2025	LSEG
BNP Paribas	1,473	2005	2025	Direct
PNC	1,335	2015	2025	Direct
Wolfe Research	1,219	2018	2025	LSEG
Comerica	978	2006	2025	Direct
Rabobank	862	2016	2025	Direct
Total	313,462			

Notes: Counts cover 2000–2025. Start and end years are the first and last years with at least 36 reports. The total covers the 26 listed contributors.
Sources: Direct: forecasters' research portals; *LSEG*: LSEG Workspace; *Investext*: Mergent Investext.

Table II: Sample Overlap with Professional Forecaster Surveys

Survey	Survey participants in sample				Counts	
	All	Sell-side banks	Primary dealers	Other	Survey total	Sample only
Blue Chip	29%	72%	85%	3%	48	12
SPF	15%	55%	67%	3%	47	19
Bloomberg	33%	54%	62%	9%	49	10
Reuters	34%	47%	79%	5%	65	4
Consensus Econ.	30%	78%	71%	6%	27	18
WSJ	26%	74%	87%	4%	74	7
NY Fed (dealers)	62%	62%	62%	—	26	10
Average	32%	63%	73%	5%	48	11

Notes: I match the institutions in Table I to the names on each survey roster. Each percentage is the fraction of institutions on that survey's roster covered by the sample, within the indicated category. The last two columns give the survey's size and the number of sample institutions outside it. Dealers are included among sell-side banks. Independent research firms, asset managers, universities and nonfinancial organizations fall under Other. Averages are across surveys, with equal weight on each.

For Bloomberg, respondents must supply quarterly GDP, core PCE and policy-rate forecasts. For Reuters, they must supply quarterly GDP, CPI and fed funds paths. I exclude three WSJ participants who did not submit forecasts: Boston College, Stifel and Daiwa. The NY Fed list is the primary-dealer mailing list for its Survey of Market Expectations, so it can include dealers who did not respond. I use dealer status at the survey date, from the NY Fed's lists.

Sources: Survey rosters: Blue Chip, August 2023; Survey of Professional Forecasters (SPF), 2023 Q3; Bloomberg, February 2026; Reuters, 2023; Consensus Economics, July 2025; WSJ, January 2026; NY Fed, July 2026.

Table III: Summary of Extracted Quarterly Forecasts, 2000–2025

Variable	Measure	Forecasts	Value changes	Median	
				Days between reports	Horizons per report
Fed funds rate	Level	104,793	9,489	4	4
2Y Treasury	Level	49,269	9,139	7	5
5Y Treasury	Level	39,206	7,466	7	5
10Y Treasury	Level	81,875	15,752	6	4
30Y Treasury	Level	32,686	6,951	7	5
GDP growth	QoQ SAAR	87,527	15,679	5	5
Headline CPI	QoQ SAAR	26,076	9,363	7	5
Headline CPI	YoY	58,371	19,746	7	5
Total		479,803	93,585		

Notes: A *forecast* is one contributor's forecast for a target quarter on a report date. Targets include the current and later quarters. Rates are quarter-end levels in percent; GDP growth is quarter-over-quarter seasonally adjusted annualized growth (QoQ SAAR). CPI inflation is reported separately as QoQ SAAR and year-over-year growth (YoY). *Value changes* counts consecutive changes for each contributor and target quarter. I calculate calendar days between consecutive report dates for each contributor and variable. *Days between reports* is the median across all these intervals. *Horizons per report* is the median number of distinct target quarters per contributor and report date.

Table IV: Summary of FOMC Communication Accuracy by Category

Category	Forecaster–question observations	Incorrect (%)
Policy decision	6,347	9.8
Statement	14,055	27.1
Economic projections	18,533	42.2
Press conference	4,832	20.6
Total	43,767	30.3

Notes: The sample covers 207 scheduled FOMC meetings, February 2000–December 2025. Observations are scored forecaster–question pairs. *Incorrect* is the probability-weighted share of expected communications not realized (see Appendix B.3). *Policy decision* covers direct policy choices at the meeting; *Statement*, language in the statement; *Economic projections*, the Summary of Economic Projections; and *Press conference*, the chair's remarks and Q&A.

Table V: Research Allocation around FOMC Announcements

Dependent variable: Daily research share or volume.

	Fed share (pp)		Fed volume (% of baseline)		Non-Fed volume (% of baseline)	
	(1)	(2)	(3)	(4)	(5)	(6)
Before/on announcement	16.1***	16.0***	112.5***	112.4***	−4.3***	−4.5***
Days −6 to 0	(0.5)	(0.5)	(4.1)	(3.7)	(1.4)	(1.4)
After announcement	4.2***	4.2***	24.2***	23.9***	−0.4	−0.7
Days +1 to +7	(0.5)	(0.4)	(2.5)	(2.5)	(1.3)	(1.3)
Near minutes	6.6***	6.4***	38.7***	36.0***	−3.5*	−4.8**
Days ±3	(0.6)	(0.6)	(3.3)	(3.3)	(1.9)	(1.8)
Before/on × trend		2.3***		26.4***		0.0
		(0.7)		(4.8)		(2.1)
After × trend		0.1		−3.3		−4.0**
		(0.6)		(3.3)		(1.7)
Time trend	No	Yes	No	Yes	No	Yes
N	58,473	58,473	129,847	129,847	129,847	129,847
R^2	0.10	0.10	0.12	0.12	0.23	0.23

Notes: The regression specification is

$$Y_{id} = \alpha_i + \delta_{w(d)} + \beta P_d + \gamma Q_d + \mu M_d + \boldsymbol{\theta}'\mathbf{Z}_d + \tau T_d + \lambda P_d T_d + \rho Q_d T_d + \varepsilon_{id}.$$

N counts institution–days for 26 forecasters during 2000–2025, excluding weekends and federal holidays. Fed share is the percentage of analytical text about the Fed; volumes are analytical characters normalized by their pooled means outside announcement windows (Section 3.3). Here pp denotes percentage points. P_d, Q_d, and M_d denote the listed windows. Day 0 is the announcement or minutes release date; minutes windows begin in 2005. α_i, $\delta_{w(d)}$, and $\mathbf{Z}_d$ denote institution effects, weekday effects, and macro-release controls (Section 4.1). Even columns add T_d, a trend in decades centered at the sample mean year, and its interactions. Share regressions weight positive-text days by analytical characters; volume regressions include zero-output days. Parentheses contain standard errors clustered by calendar quarter; stars use t_{G-1} with G quarters. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table VI: Minutes Content and Changes in Forecaster Discussion

Dependent variable: Revision in forecaster topic share (pp) or embedding coordinate (SD) around Fed minutes releases

	(1) Topic model	(2) Dense embeddings
Current minutes	0.11***	0.07***
	(0.02)	(< 0.01)
Pre-period forecaster text	−0.83***	−0.74***
	(0.01)	(0.01)
Prior-year forecaster text	0.73***	0.37***
	(0.02)	(0.01)
Previous change in forecaster text	−0.01	0.00
	(0.01)	(< 0.01)
Minutes-history controls	Yes	Yes
Coordinate fixed effects	Yes	Yes
Forecaster–meeting pairs	2,394	2,394
Text dimensions	40	4,096
Coordinate-level observations	95,760	9,805,824

Notes: The regression specification is $\Delta x_{itk} = \alpha_k + \beta m_{tk} + \boldsymbol{\theta}'\mathbf{Z}_{itk} + \varepsilon_{itk}$. i, t, k index forecasters, releases, and coordinates; m is minutes content and $\mathbf{Z}$ history controls (Section 4.2). The minutes-release week is excluded (Figure A8). The topic model follows Hansen, McMahon and Prat (2018). Units: percentage points (pp) for topics; pooled within-coordinate standard deviations (SD) for each embedding variable. Sample: 159 equally weighted releases, 2005–2025. Parentheses: meeting-level Bartlett HAC standard errors (seven lags). Normal-reference significance: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table VII: Fed Research and Forecast Accuracy

Dependent variable: RMSE over the next three quarters (bp).

	Fed research share (15 pp)				
Forecast	(1)	(2)	(3)	(4)	N
Fed funds	5.7**	-1.6***	6.8**	-2.1***	10,637
	(2.6)	(0.6)	(3.5)	(0.8)	
GDP growth	12.8***	1.8	17.4***	5.0***	8,199
	(4.6)	(1.7)	(6.3)	(1.8)	
CPI inflation	-15.9	4.3	-7.9	14.0***	2,005
	(13.9)	(3.2)	(16.2)	(5.0)	
Date FE	No	Yes	No	Yes	
Institution FE	No	No	Yes	Yes	

Notes: The regression specification is

$$L_{it}^{v} = \alpha^{v} + \beta^{v} A_{it} + \delta_{t}^{v} + \mu_{i}^{v} + \boldsymbol{\gamma}^{v\prime} \mathbf{Z}_{it}^{v} + \varepsilon_{it}^{v}.$$

Variable definitions and controls follow equation (10) and Section 5.2. I use forecasts dated 2000–2025 with all three quarterly horizons present, holding the sample fixed across columns within each row. GDP bundles intersecting 2020Q2–2021Q4 are omitted. October 2025 CPI is interpolated between September and November. Appendix Table A10 reports alternative inference; Table A9 replaces institution with chief economist fixed effects. Parentheses report standard errors clustered by forecaster and target bundle. Stars use normal-reference p-values: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table VIII: Fed Research and Rate Forecast Accuracy by Horizon

Dependent variable: Absolute forecast error (bp).

	Fed research share (15 pp)				
	Forecast horizon (quarters ahead)				
Forecast	1	2	3	4	*N*
Fed funds	-0.7*	-1.8**	-3.0***	-3.0**	7,644
	(0.4)	(0.8)	(1.1)	(1.3)	
2-year Treasury	0.6	2.0	3.0**	5.7***	3,326
	(1.0)	(1.4)	(1.4)	(2.0)	
5-year Treasury	3.2**	5.7***	5.9***	7.8***	2,650
	(1.4)	(2.0)	(2.1)	(2.5)	
10-year Treasury	0.4	1.0	2.5*	3.0**	5,705
	(1.0)	(1.3)	(1.4)	(1.4)	
30-year Treasury	-1.8	-1.9	-1.0	-1.5	1,642
	(1.3)	(1.6)	(1.9)	(1.8)	

Notes: The regression specification is

$$|\hat{r}^{v}_{it,q} - r^{v}_{q}| = \alpha + \beta_{v,h} A_{it} + \delta^{v,h}_{t} + \mu^{v,h}_{i} + \boldsymbol{\gamma}^{v,h\prime} \mathbf{Z}^{v,h}_{it} + \varepsilon^{v,h}_{it,q}.$$

Research share, controls, and fixed effects follow column (4) of Table VII. I use forecasts dated 2000–2025 with all four quarterly horizons present, holding the sample fixed across columns within each row. Appendix Table A4 allows the sample to vary by horizon. Parentheses report standard errors clustered by forecaster and target quarter. Stars use normal-reference p-values: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table IX: Forecaster Surprises and Forecast Revisions

Dependent variable: Forecast revision around FOMC announcements (bp).

	Fed funds *Current and next three quarters*	GDP growth *Next three quarters*		CPI inflation *Next three quarters*	
	(1)	(2)	(3)	(4)	(5)
Forecaster surprise (1 SD)	0.92***	1.08**	0.98**	0.81*	0.75*
	(0.29)	(0.42)	(0.42)	(0.43)	(0.44)
Fed funds revision (bp)			0.10		0.07
			(0.13)		(0.07)
Meeting fixed effects	Yes	Yes	Yes	Yes	Yes
Institution fixed effects	Yes	Yes	Yes	Yes	Yes
Forecaster–meeting pairs	1,037	1,031	1,031	871	871

Notes: The regression specification is

$$\Delta f_{it}^{k} = \alpha_{t}^{k} + \lambda_{i}^{k} + \beta^{k} \widetilde{S}_{it}^{k} + \delta^{k} \Delta f_{it}^{FF} + \boldsymbol{\gamma}^{k\prime} \mathbf{Z}_{it}^{k} + \varepsilon_{it}^{k}.$$

i, t, k index forecasters, meetings, and outcomes. Higher surprise scores indicate more hawkish surprises relative to peers; SD is the outcome-specific sample standard deviation. Fed funds uses average quarterly revisions; GDP and inflation use revisions in annualized compounded growth. Forecast pairs span at most ten days around scheduled announcements, 2000–2025. Section 6.2 defines the controls $\mathbf{Z}$. Columns (3) and (5) add the contemporaneous Fed funds revision, itself an announcement response; $\delta = 0$ otherwise. Samples and score scaling are fixed within outcome. Appendix Table A12 replaces institution fixed effects with chief economist fixed effects. Parentheses report meeting-clustered standard errors. Stars use normal-reference p-values: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table X: Forecast Accuracy around FOMC Announcements

Dependent variable: Change in forecast RMSE (bp).

	Fed funds		GDP growth	
	Next three quarters		*Next three quarters*	
	Full	Excl. crisis years	Full	Excl. crisis years
	(1)	(2)	(3)	(4)
Unadjusted mean, FOMC pairs	-2.3***	-2.4***	-0.2	0.0
	(0.7)	(0.8)	(0.4)	(0.3)
Adjusted differences from other forecast intervals:				
FOMC announcement	-1.2	-1.4*	0.8*	0.9**
	(0.7)	(0.8)	(0.4)	(0.3)
Payroll release	-0.8*	-0.5	-1.5	-0.9
	(0.4)	(0.4)	(0.9)	(0.9)
CPI release	-0.2	-0.3	0.2	0.1
	(0.3)	(0.3)	(0.3)	(0.2)
Initial GDP release	0.6	0.4	-1.7*	-2.5***
	(0.6)	(0.6)	(0.9)	(0.8)
Institution FE	Yes	Yes	Yes	Yes
Year-quarter FE	Yes	Yes	Yes	Yes
Observations	7,450	6,402	7,899	7,014
R^2	0.11	0.11	0.08	0.07

Notes: Negative changes indicate improved accuracy. Each pair compares forecasts for the same three target quarters, published at most ten days apart. FOMC pairs straddle one scheduled announcement. The regression specification is

$$\Delta\text{RMSE}_{ij} = \alpha_i + \delta_{q(j)} + \eta_{d(j)} + \boldsymbol{\beta}'\mathbf{R}_{ij} + \boldsymbol{\gamma}'\mathbf{Z}_{ij} + \varepsilon_{ij},$$

where i indexes institutions, j forecast intervals, q quarters, and d weekdays. $\mathbf{R}$ contains the displayed release indicators. Each meeting or ordinary interval end date receives equal total weight. Controls $\mathbf{Z}$ are listed in the appendix construction details (p. 110). Forecasts date from 2000–2025. In columns (2) and (4), I drop 2001, 2007–2009, and 2020–2021. I exclude bundles overlapping 2020Q2–2021Q4 in all columns. Counts and R^2 refer to regressions. Parentheses report standard errors clustered by institution and year-quarter; stars use t with degrees of freedom equal to the smaller cluster count minus one. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table XI: Fed Research and Asset Price Reaction Magnitudes

Dependent variable: Absolute announcement-window reaction (SD).

	S&P 500 return			10-year yield change		
Regressor (1 SD)	(1)	(2)	(3)	(4)	(5)	(6)
Fed research share	0.21***	0.11*	0.04	0.19**	0.12	0.05
	(0.06)	(0.07)	(0.07)	(0.08)	(0.09)	(0.08)
Communication disagreement		0.26***	0.11		0.18**	0.04
		(0.07)	(0.08)		(0.08)	(0.09)
Communication surprise			0.25**			0.24**
			(0.11)			(0.11)
Meetings	207	207	207	206	206	206
R^2	0.04	0.10	0.13	0.03	0.06	0.09

Notes: The regression specification is

$$y_t^a = \alpha^a + \beta_R^a R_t + \beta_U^a U_t + \beta_S^a S_t + \varepsilon_t^a.$$

All variables are expressed in standard-deviation units within each asset's sample. Event windows follow Section 6.2. R_t is Fed research share in the prior seven days; U_t is disagreement across forecasters' answers to announcement questions; S_t is the probability-weighted share of incorrect expectations (Appendix B.3). The sample covers scheduled FOMC meetings in 2000–2025, excluding March 18, 2009 for Treasuries. Appendix Table A15 reports MPS and meeting-type robustness checks. Parentheses report HC3 standard errors; stars use t tests with residual degrees of freedom: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table XII: Economic Content and Treasury Yield Responses

Dependent variable: Treasury yield change around FOMC announcements (bp).

	2-year		5-year		10-year		30-year	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
MPS (bp)	1.43***	1.65***	1.23***	1.44***	0.72***	0.88***	0.20**	0.32***
	(0.11)	(0.07)	(0.12)	(0.11)	(0.09)	(0.10)	(0.09)	(0.09)
Economic share (1 SD)	−0.26	−0.14	−0.16	−0.05	−0.24	−0.15	−0.16	−0.10
	(0.19)	(0.15)	(0.23)	(0.21)	(0.19)	(0.18)	(0.20)	(0.20)
MPS × economic share		0.37***		0.36***		0.28***		0.19***
		(0.06)		(0.08)		(0.06)		(0.06)
Error-rate control	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Meetings	206	206	206	206	206	206	206	206
R^2	0.78	0.83	0.57	0.61	0.37	0.43	0.06	0.10

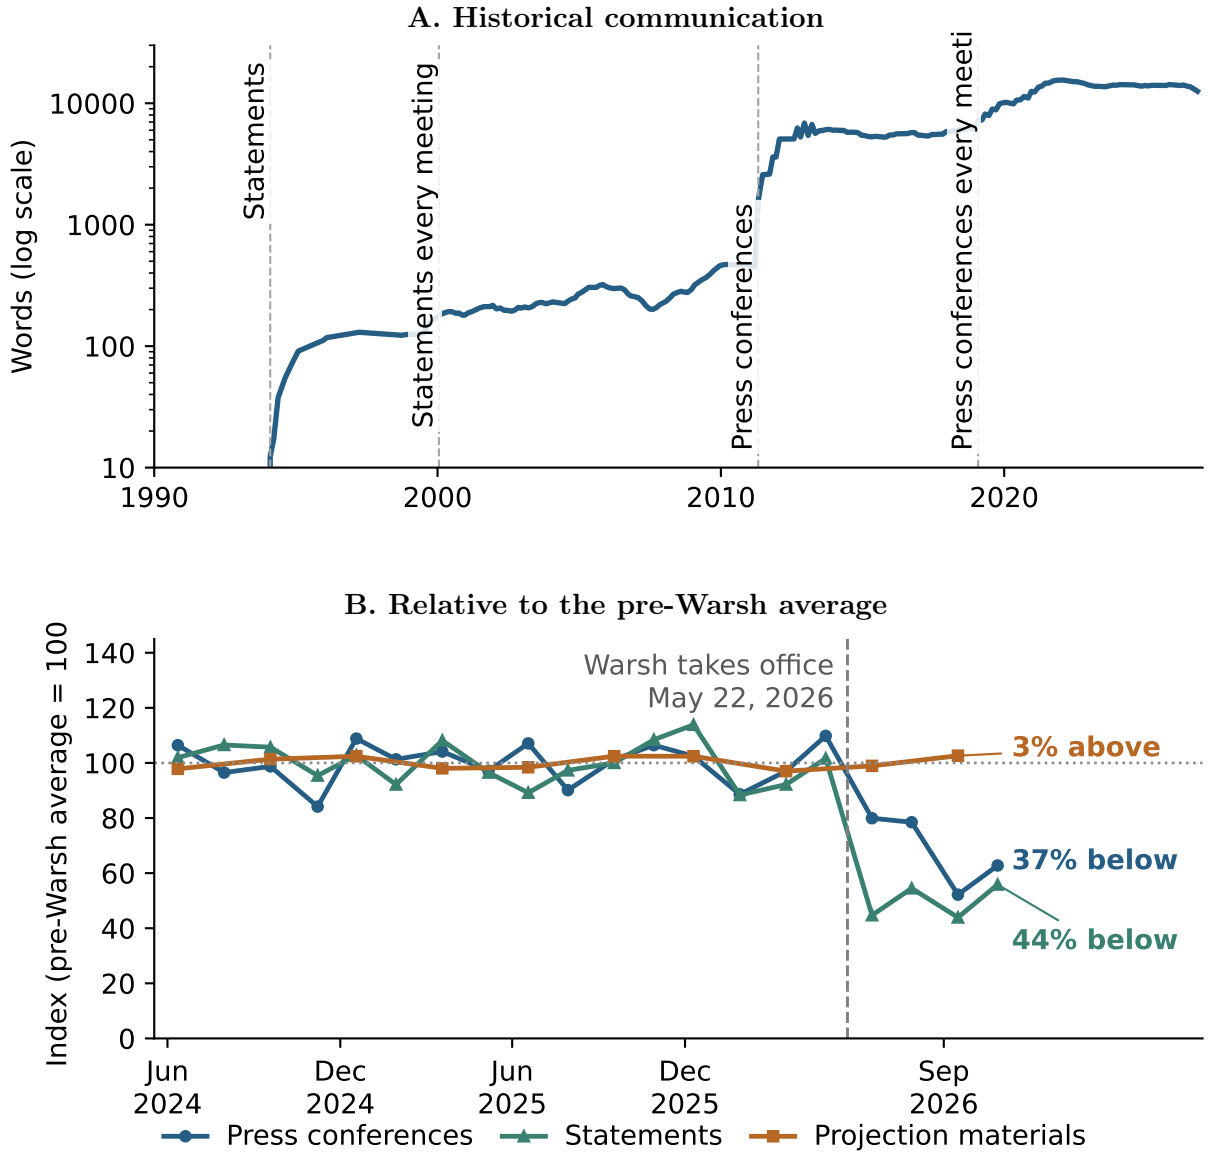
Notes: The regression specification is

$$\Delta y_t^{(n)} = \alpha^{(n)} + \beta^{(n)}\text{MPS}_t + \gamma^{(n)}E_t + \delta^{(n)}(\text{MPS}_t \times E_t) + \lambda^{(n)}S_t + \varepsilon_t^{(n)}.$$

n indexes maturity and t meetings. MPS is the monetary policy surprise factor from the U.S. Monetary Policy Event-Study Database. E is the economic share of communication errors (equation 14); S is the overall forecast-error rate. Both text measures are standardized (SD denotes standard deviation). Odd-numbered columns omit the interaction; samples and other controls are unchanged. Outcomes cover the statement and press conference windows (Section 6.2). The sample covers 206 scheduled meetings, 2000–2025, excluding March 18, 2009. Appendix Table A16 reports robustness checks. Parentheses report HC3 standard errors; stars use t tests with residual degrees of freedom: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

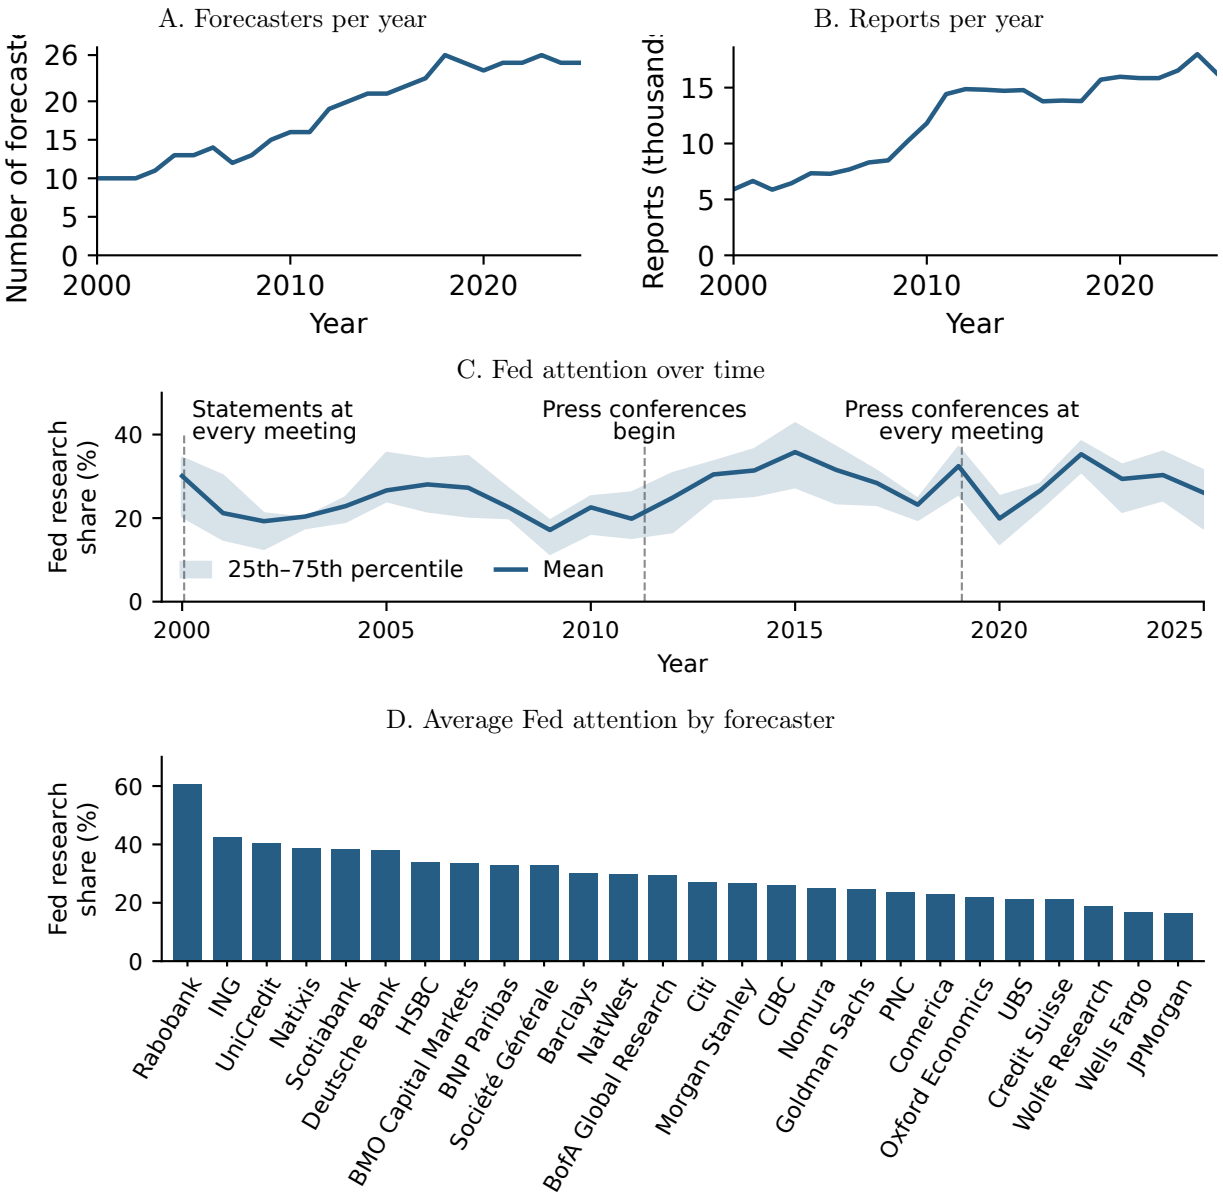
Figures

Figure I: Federal Reserve Communication over Time



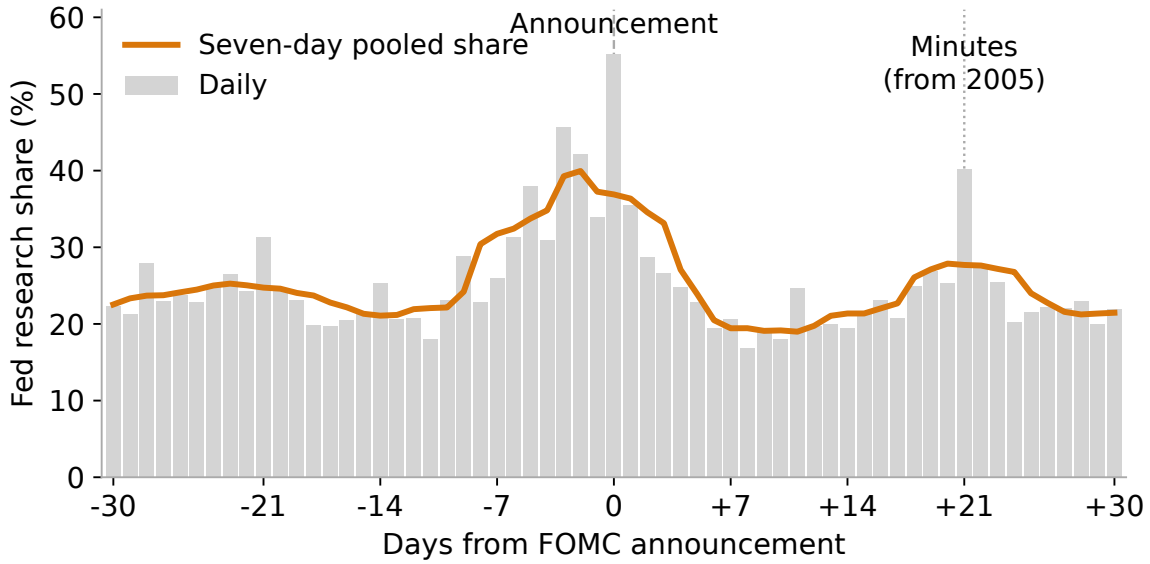
Notes: Panel A shows an eight-meeting average of word counts in official announcement-day communication through October 28, 2026 (log scale). Panel B indexes each series to its June 2024–April 2026 average (16 meetings; eight SEP releases). Endpoint labels compare October 2026 statements and press conferences and September 2026 projections with these averages. The October 2026 transcript is preliminary. Source: Federal Reserve Board releases.

Figure II: Report Coverage and Fed Attention



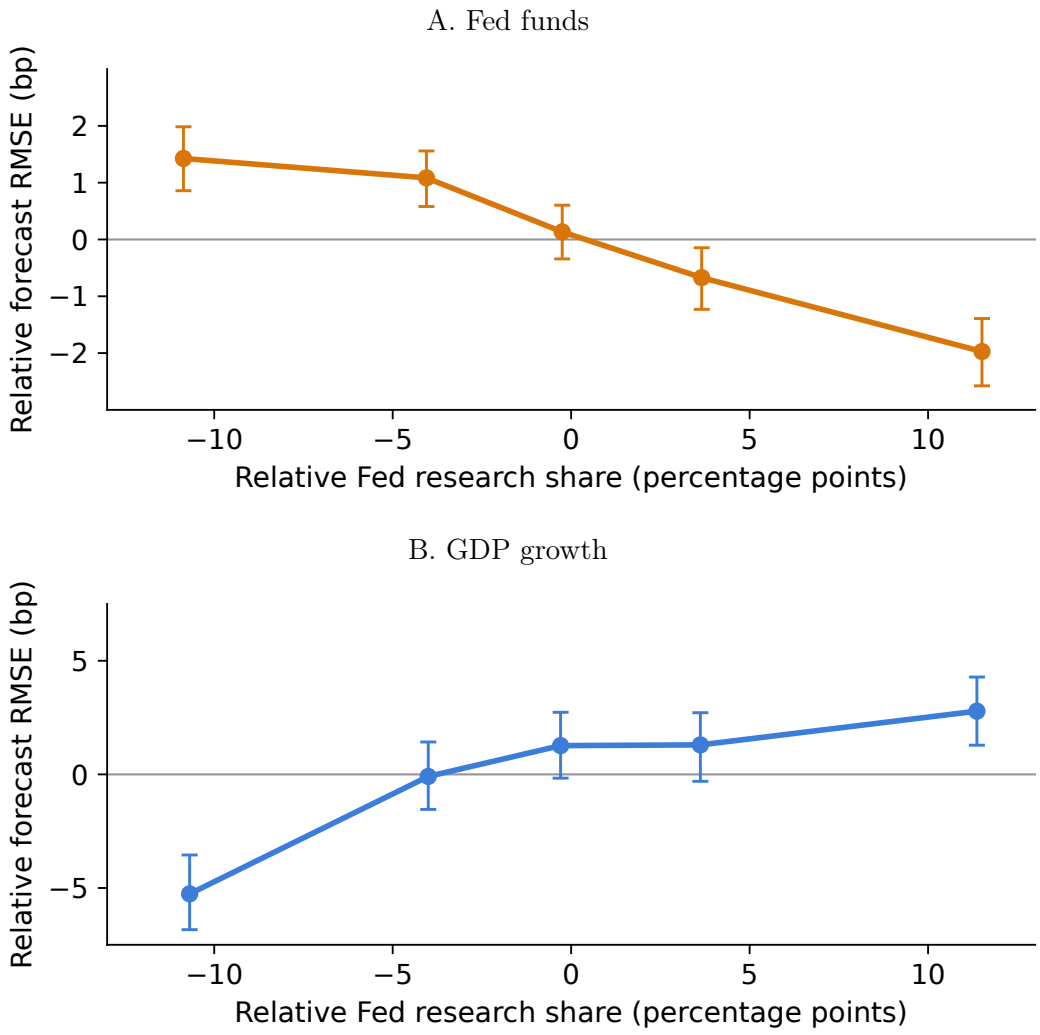
Notes: The sample covers 26 forecasters during 2000–2025. Panel A counts forecasters with at least 36 reports in a year to exclude years with only isolated reports. Fed attention is the broad Fed-related share of analytical text. Panel C shows the mean and 25th–75th percentiles of forecaster Fed research shares, weighted by annual analytical text volume. Shading represents dispersion across forecasters. Panel D shows unadjusted means of annual shares.

Figure III: Fed Research Intensity Around FOMC Announcements



Notes: Percent of analytical text devoted to the Fed, defined in Section 3.3. Gray bars show daily shares and the line shows the centered seven-day pooled share across scheduled announcements in 2000–2025. Shares are weighted by analytical text volume across forecasters and meetings. Since January 2005, minutes have been released three weeks after meetings; previously, they were released on the Thursday after the following meeting.

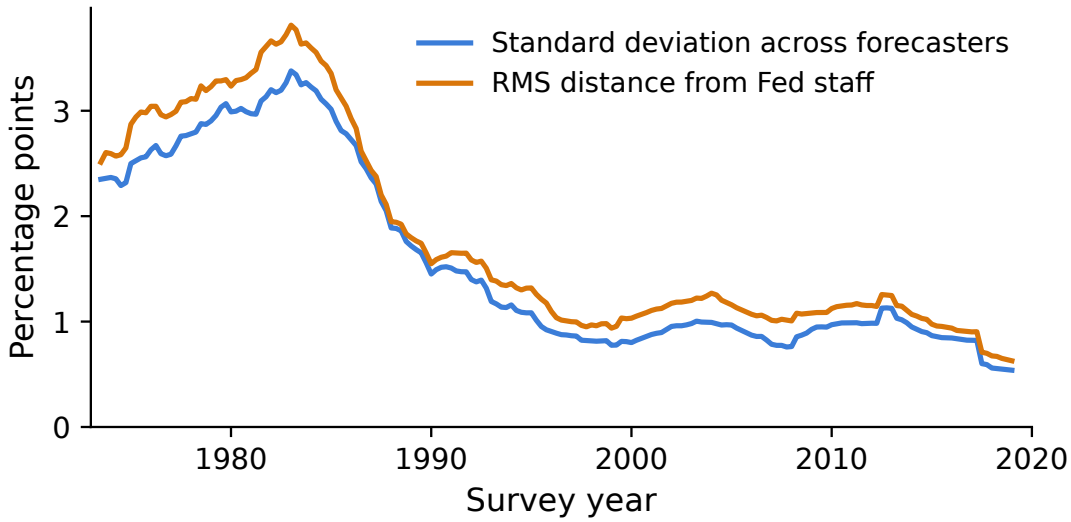
Figure IV: Fed Research and Forecast Accuracy across Quintiles



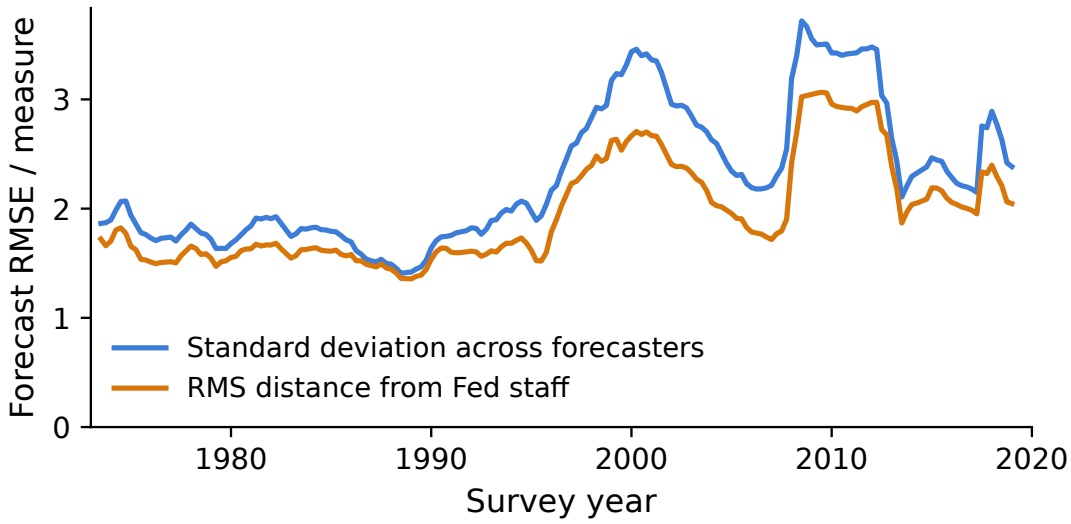
Notes: Points show mean residual Fed research share and residual forecast RMSE within five quintiles of residual research share. RMSE covers forecasts for the next three quarters. GDP growth is quarter-over-quarter growth in real GDP, seasonally adjusted and annualized. Samples and controls follow column (4) of Table VII. Error bars show 95% pointwise confidence intervals from 1,000 bootstrap samples of publication dates, holding control adjustments and quintile assignments fixed. Table VII and Appendix Table A10 report regression inference using two-way clustered standard errors and alternative procedures.

Figure V: Historical Forecast Disagreement and Accuracy

A. Forecast disagreement



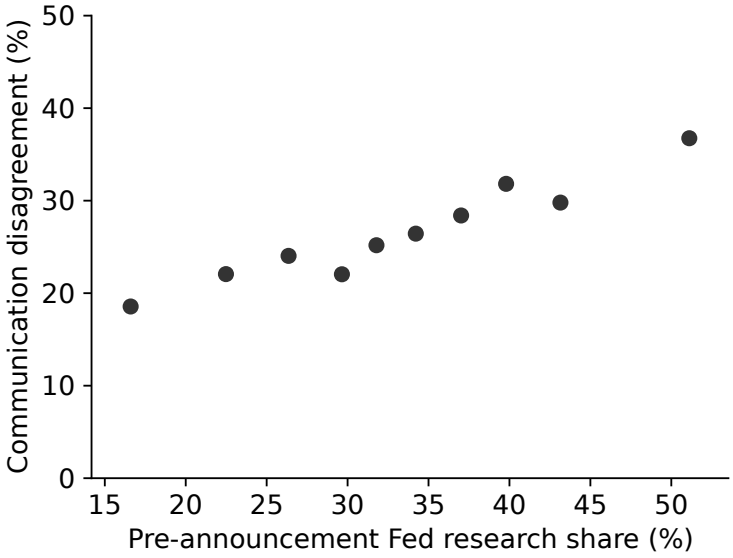
B. Forecast error relative to disagreement



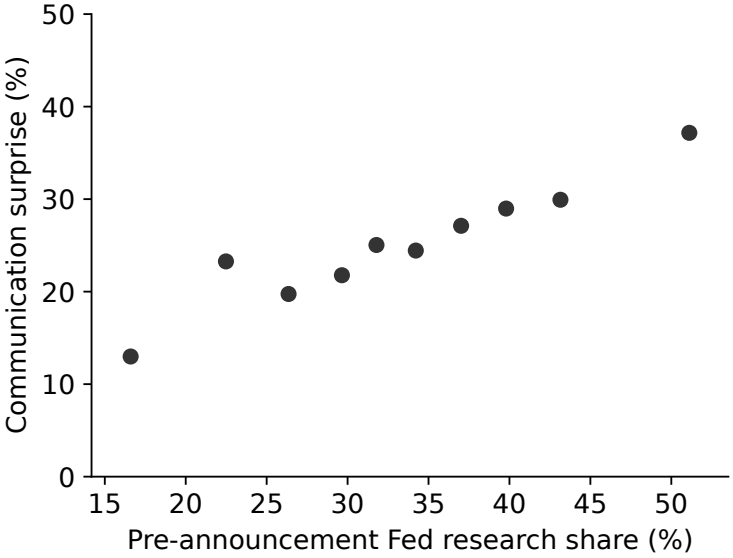
Notes: Each point summarizes annualized quarterly real GNP/GDP growth forecasts one to three quarters ahead over a trailing five-year window. The sample covers 1973–2019. Panel A shows the standard deviation across forecasters and the root mean squared distance from Fed staff forecasts. Panel B divides individual realized forecast RMSE by each measure. Sources: Philadelphia Fed Survey of Professional Forecasters (SPF), Greenbook staff forecasts, and revised outcomes from FRED.

Figure VI: Pre-Announcement Fed Research, Communication Disagreement, and Surprise

A. Communication disagreement



B. Communication surprise



Notes: Dots are means for ten equal-count bins of Fed research share in the seven days before each announcement. The sample comprises 207 scheduled FOMC meetings (February 2000–December 2025), equally weighted. Disagreement is the percentage of comparable forecaster pairs that give different answers to the same question; surprise is the probability-weighted percentage of incorrect expectations. See Appendix B.3 for construction.

A Data Sources and Processing Pipeline

I develop a document processing pipeline for macro research reports that combines page-level context summaries, parsed text, and original images of tables and figures. These components are indexed for retrieval during extraction. The pipeline is guided by three principles. First, the pipeline must allow easy querying of the most relevant pieces of each report, consistently across different data sources. Second, it must accurately capture the information in each report, particularly complex tables and figures. Third, it must provide enough context that a language model can interpret what a chunk of information means without seeing the entire document it came from. This pipeline, combined with the extraction instructions described in Section B, allows the LLM to query reports across the dataset and have enough context to interpret them.

A.1 Document Processing

Existing vision-language model (VLM) PDF parsers struggle with two aspects of macro research reports. First, document hierarchy: the ordering and nesting of headers within a document is frequently lost during parsing. Without this information, it can be ambiguous which country or topic a given table or piece of text is referring to. Second, tables in macro research reports can be quite complex, often exceeding 20 columns or 20 rows, and properly interpreting the forecasts they contain requires carefully reading footnotes, headers, and row labels. Getting the formatting of these tables correct is important for downstream extraction accuracy.

To address the document hierarchy problem, I generate page-level context summaries before parsing. A vision-language model processes each page sequentially, feeding the prior

page's summary forward as context for the next. I then attach these summaries to each chunk to provide the LLM context such as which country, topic, or time period the chunk belongs to. The same model also classifies each report by geography and topical focus to filter the corpus to U.S. and global economics reports.

I then use PaddleOCR-VL, a VLM fine-tuned for PDF parsing, to perform layout detection on each page. The model identifies and transcribes text blocks, and separately identifies tables and figures. Rather than relying on the parser's text-based transcription of tables and figures, I retain them as images. At runtime, the multimodal pipeline can pass these images directly to a frontier LLM, which handles complex tables well.

Figure A9 summarizes the pipeline.

To make the retained table and figure images searchable, I pass each image through a lightweight vision-language model (Qwen3-VL-8B) together with its page context to produce a textual summary. I prefix each text chunk and image summary with the document headline and page context, then embed it using Qwen3-Embedding-8B.

A.2 Retrieval-Augmented Generation

The resulting data from the pipeline described in Section A.1 is a set of chunks (either text, table, or figure) with their associated page context and embedding vectors. These chunks are queryable by their metadata: contributor, date, geography, topic, and report headline. Figure A10 illustrates this structure.

This structure turns the corpus into a searchable database. To extract information about a specific variable or topic, I first apply hard filters on the metadata (e.g., date range, contributor, geography) and then semantically search via the embedding vectors for the most relevant chunks. This retrieval-augmented generation (RAG) approach provides a

deterministic, computationally tractable way to identify the relevant passages from hundreds of thousands of documents without processing each one in full (Bartik et al., 2025).

At extraction time, the LLM sees the underlying chunk content (the text itself, or the image of the table or figure), the page context for that chunk, and the document headline and date. The contributor, geography, and topic fields serve as hard filters for retrieval but are not passed to the LLM. Section B describes each extraction task in detail.

B LLM-Based Extraction

Section A describes the process of constructing easily queryable, context-aware chunks from raw PDF reports. This section describes the extraction pipelines that operate on those chunks: how the LLM is instructed to process the retrieved information for each task. Section 3.3 introduces these measures and their role in the paper.

B.1 Measuring Fed Research

I measure Fed research in three stages. First, I standardize reports by removing boilerplate language and splitting reports into substantive articles. Second, I determine the extent to which forecasters explicitly motivated their article with the Fed. Third, I determine the share of reasoning in each paragraph of text within an article about the Fed. Together, this procedure produces a standardized measure of Fed research.

Throughout this procedure I use Qwen3.8-27B-FP8 as the LLM. Fractional scores allow mixed content to contribute to more than one category.

I begin by standardizing the formatting of articles across forecasters. Forecasters differ in the amount of boilerplate/regulatory language included in their reports as well as in report length. For example, some forecasters prefer to publish shorter reports, while others publish longer reports that aggregate collections of shorter articles. I provide the LLM with ordered text blocks from the PDF parser described in Appendix A.1 and task it with grouping text blocks into substantive articles and dropping boilerplate text. I exclude all figures and tables for this procedure. Forecasters may also republish articles. To accommodate this, I exclude repeated text when aggregating research shares. The resulting articles provide a common unit for measuring motivation across different report formats.

I provide the LLM with each substantive article's text and ask it to assess, on a scale from 0–100, the extent to which the forecaster explicitly motivated the article by the Fed. An article qualifies as Fed-motivated when its central objective is to understand the Fed or when the source explicitly connects its economic investigation to a Fed action or communication. I also measure the share of reasoning about the Fed for each block of text identified by the PDF parser. The LLM allocates 100 points across direct Fed reasoning, economic reasoning, and factual reporting.

The following example illustrates why economic reasoning can also count as Fed research: the report explicitly motivates its analysis of China by the FOMC's attention to foreign developments. These passages illustrate the classification categories without assigning numerical scores to the examples.

> **Economic reasoning**
>
> "We expect China's activity growth to stabilise before the year-end, mostly thanks to a bigger push on credit-backed infrastructure investment."
>
> **Fed motivation for that research**
>
> "With the FOMC's data dependency now apparently extending beyond its shores, notably to China, we discuss our near-term outlook for Chinese data."
>
> Société Générale, September 28, 2015.

I construct Fed attention as the character-weighted share of reasoning directed towards the Fed. Let c_p be the character count of retained block p, and let $r_{F,p}$, $r_{X,p}$, and f_p be its Fed-reasoning, economic-reasoning, and factual shares, divided by 100 so that $r_{F,p}+r_{X,p}+f_p = 1$. Let m_a be the similarly normalized Fed-motivation share of article a. For forecaster i and research window w, let $\mathcal{A}_{iw}$ contain the articles published in that window and $\mathcal{P}_a$ the retained

blocks in article a. Fed attention is

$$A_{iw} = 100 \times \frac{\sum_{a \in \mathcal{A}_{iw}} \sum_{p \in \mathcal{P}_a} c_p \left(r_{F,p} + m_a r_{X,p} \right)}{\sum_{a \in \mathcal{A}_{iw}} \sum_{p \in \mathcal{P}_a} c_p \left(r_{F,p} + r_{X,p} \right)}. \tag{12}$$

The numerator includes direct Fed reasoning and the Fed-motivated share of economic reasoning. The denominator includes all reasoning and excludes factual reporting. I sum these quantities before taking their ratio, both for daily measures and for the research windows specified in each analysis. Windows with no analytical text have an undefined share. Failed extractions are treated as missing measurements rather than zero Fed attention.

Figure A11 illustrates direct reasoning about the Fed's next policy decision.

B.2 Quantitative Forecast Extraction

For each contributor and each calendar date, the pipeline filters to reports classified as U.S. or global economics and retrieves the most relevant passages via RAG. A lightweight model (Gemini 3 Flash) then extracts all quantitative forecasts from each passage along with their associated transformations and respective horizons. For example, an extracted forecast can be of the form real GDP growth quarter over quarter annualized for 2024q3 of 2.1%. The central challenge in extracting this information is that macrofinancial forecasts can be presented in a variety of formats. GDP, for example, can be reported in levels or growth rates; growth rates can be quarter-over-quarter, quarter-over-quarter annualized, or year-over-year; and values can be real or nominal. Identifying which transformation an analyst uses requires guiding the LLM in how to process the table footnotes, column headers, or surrounding text for context. Specifically, the extraction prompt provides detailed instructions on how to read forecast tables and interpret text, including common practices analysts use when presenting forecasts. I also give the LLM access to ALFRED vintage macroeconomic data (the historical values that would have been available at the time of the report), transformed into a variety of common formats. Analysts frequently include historical data for the past few quarters in their forecast tables, and when the LLM can match these historical values to a specific ALFRED vintage transform, it gains confidence in which measure is being discussed.

Figure A12 summarizes the extraction pipeline.

To improve accuracy, I flag cases for review, including conflicting patterns across passages from the same date, forecasts that oscillate between two values in a short window (A $\rightarrow$ B $\rightarrow$ A), and a forecaster's first forecast for a specific target. In each case a reviewer agent (Gemini 3 Pro) reviews the flagged forecasts with additional context of surrounding passages and the passages from which prior and subsequent forecasts are extracted.

I validate the extracted point forecasts by comparing them against independently recorded submissions to the Reuters Economic Polls. Specifically, I match forecasts for GDP, year-over-year headline CPI inflation, the Fed funds rate, and the 2-year as well as 10-year Treasury yields. The comparison does not cover quarterly annualized CPI inflation or 5- and 30-year Treasury forecasts. I match each report forecast to the nearest poll submission by the same institution for the same target quarter within seven days before or after the report date. In Table A2 I report the correlations between poll submissions and my extractions, finding correlations ranging from 0.975 to 0.998 across the variables. In total these performance metrics are based on 82,581 matched observations.

B.3 FOMC Forecast Extraction

Ahead of each FOMC announcement, forecasters form a *Fed call*: a detailed expectation of the policy decisions, quantitative projections, and qualitative statements the announcement will contain, together with their reasoning. I recover these calls from research reports, identify the announcement attributes forecasters discuss, and organize their expectations into a common questionnaire for each meeting. I then compare these expectations with the realized announcement to construct communication errors and cross-sectional surprise scores. Figure A13 summarizes the procedure.

Source selection and forecast identification. I allow expectations from reports published up to three weeks before each meeting, prioritizing more recent expectations when expectations conflict. Forecasters typically detail their Fed calls in the final week before the Fed announcement (Appendix Figure A3), after they are surveyed by, for example, the NY Fed. Throughout the three-week window, I use RAG to identify passages likely to contain

predictions for the upcoming announcement. I then extract all expectations from each passage and denote which passages had at least some relevant context for the forecaster's Fed call. For example, Citi's June 10, 2024 preview discussed the longer-run policy-rate dot, the statement's inflation language, and Powell's description of the policy stance ahead of the June 12 announcement.

Organizing each forecaster's expectations. The extracted forecasts are categorized by the part of the announcement they concern, such as the statement, SEP, policy decision, or press conference. I use an LLM to compare all expectations within the same category (e.g., statement expectations or SEP inflation expectations) across a forecaster's reports to identify duplicates and superseded expectations. The result is a consolidated list of expected communications from each forecaster.

Standardizing the granularity of expected communications. To standardize expectations and allow for cross-forecaster comparisons, I next provide an LLM the set of expected communications across forecasters within each category (i.e. statement expectations or all SEP inflation expectations). The LLM organizes the expectations expressed in these pre-announcement reports into a common questionnaire. The expected communications differ across forecasters in their degree of granularity. For example, while some forecasters may broadly comment on the tone they expect from the press conference (such predictions are excluded given the subjectivity in verifying tone), others detail specific phrases or topics they expect. The LLM then answers the questionnaire separately for each forecaster. For each question, the LLM returns to the forecaster's source passages to identify the answer supported by the report. Questions without sufficient supporting evidence remain unanswered. When available, answers retain confidence language such as "expect," "likely," or "could," as

well as explicit probabilities such as "roughly 30%." For this announcement, Citi expected an unchanged longer-run dot, while Wells Fargo expected it to be revised higher. Citi also expected Powell to describe the policy stance as "restrictive," a question not addressed in Wells Fargo's supplied reports.

Verifying the announcement. I provide the LLM with each questionnaire question and the realized announcement content, asking it to identify the corresponding communication value. Some expectations cannot be verified from the published material. Before the June 2024 meeting, for example, Barclays expected St. Louis Fed President Alberto Musalem to submit a longer-run rate projection above 2.55%. This prediction cannot be verified from the released projections because the SEP dots are anonymous.

Probability-weighted communication errors. When stated, I use the confidence qualifier attached to each retained answer to weight alternatives within a forecaster–question observation. Unambiguous numerical point probabilities enter directly, i.e. "30% chance of X". For other confidence labels, such as "likely", I estimate probabilities in sample from event realization rates. Qualifiers with more than 30 observations receive their empirical realization rate. For less common qualifiers, I use a similarity-weighted average of the rates for the three nearest well-observed qualifiers in text-embedding space. I exclude "expect," the most common qualifier, as a neighbor. Estimated probabilities are bounded between 0.01 and 0.99. Figure A4 plots the realization rates for common qualifiers; "expect" is the default when no qualifier is stated.

Let $\mathcal{A}_{iqt}$ contain forecaster i's verifiable retained answers to question q at meeting t, p_{iqat} denote the probability weight of answer a, and d_{qat} indicate that it is incorrect. The

probability-weighted error is

$$e_{iqt} = \frac{\sum_{a \in \mathcal{A}_{iqt}} p_{iqat} d_{qat}}{\sum_{a \in \mathcal{A}_{iqt}} p_{iqat}}. \tag{13}$$

I average e_{iqt} equally across questions for forecaster-level measures and across forecaster–question observations for meeting-level measures. Probability weighting within each question gives less weight to low-probability alternatives while retaining a binary error for single-answer questions. The measure remains conditional on which expectations forecasters report.

Economic content of communication surprises. I classify each question by its object using Gemini 3.8 Flash, providing only the question text and explicit classification rules. Policy decisions include current rate decisions and announced balance-sheet actions. Policy guidance includes future policy intentions, reaction-function language, policy stance, and policy-rate projections, including the dots. Economic assessments include economic conditions, outlook, and risks, as well as SEP economic projections. These rules apply across statements, SEPs, and press conferences; questions combining objects are classified by their principal requested outcome.

Let e_{iqt} be the probability-weighted error in equation (13), using the same verification rule. Let $\mathcal{O}_t$ contain the scored forecaster–question observations, and let E_{qt} indicate that question q concerns an economic assessment. The economic share is

$$S_t^{\text{econ}} = \frac{\sum_{(i,q) \in \mathcal{O}_t} E_{qt} e_{iqt}}{\sum_{(i,q) \in \mathcal{O}_t} e_{iqt}}. \tag{14}$$

Each forecaster–question observation receives equal weight, so questions covered by more forecasters receive greater aggregate weight. In the regressions, Economic_t is this share standardized to mean zero and unit sample standard deviation over the estimation sample.

For example, consider a meeting for which forecasters make 20 scored predictions, each with a single answer. If ten predictions are incorrect and six of those concern the Fed's economic assessment, then the overall communication error rate is 50% and the economic share is 60%.

Cross-sectional surprise scores. A unique aspect of this dataset is that I can compare the communication surprises of different forecasters facing the same announcement. I summarize these differences with a score for how hawkishly surprised each forecaster was relative to their peers. Doing so requires a degree of judgment because expected communications vary in their importance, the magnitude of their errors, and which forecasters address them. I therefore compare forecasters pairwise, asking which was more hawkishly surprised, and then fit a score to these comparisons. This approach follows Qin et al. (2024) and Asirvatham et al. (2026). For each pair, Gemini 3.7 Flash receives both forecasters' questionnaire answers, their confidence language, and the verified outcomes. The LLM is instructed to focus on questions that both forecasters answer and to return an inconclusive answer when the two responses are not comparable. In the June 2024 example, the longer-run dot was revised higher, making Wells Fargo's prediction correct and Citi's incorrect. Citi was therefore more hawkishly surprised than Wells Fargo on this dimension.

I fit the pairwise judgments using a Bradley–Terry model. Each forecaster receives a score $s_{i,t}$ at meeting t. A higher score means that the forecaster was more hawkishly surprised relative to peers. The probability that forecaster i is judged more hawkishly surprised than forecaster j is

$$p_{ij,t} = \frac{\exp(s_{i,t})}{\exp(s_{i,t}) + \exp(s_{j,t})}.$$

The result of each comparison is coded $y_{ij,t} = 1$ when forecaster i is more hawkishly surprised and zero otherwise. Inconclusive comparisons are coded as $y_{ij,t} = 1/2$. I estimate the scores

by maximizing the log likelihood with a ridge penalty of $0.01 \sum_i s_{i,t}^2/2$ to keep scores finite.

C Research Responses to Fed Communication

Posterior precision and cross-learning. The posterior precision of the policy rate is

$$P_F \equiv \text{Prec}(F \mid s_X, s_F) = \underbrace{\frac{1}{\sigma_X^2 + \sigma_B^2}}_{\text{prior precision}} + \underbrace{a_F}_{\text{Fed research}} + \underbrace{\left[\frac{\frac{1}{\sigma_X^2} + a_X}{1 + \sigma_B^2\left(\frac{1}{\sigma_X^2} + a_X\right)} - \frac{1}{\sigma_X^2 + \sigma_B^2}\right]}_{\text{economic research}}. \tag{15}$$

The marginal effects of research on posterior precision at an interior allocation are

$$\frac{\partial P_X}{\partial a_X} = 1 > \frac{1}{(1 + \sigma_B^2 a_F)^2} = \frac{\partial P_X}{\partial a_F}, \tag{16}$$

$$\frac{\partial P_F}{\partial a_F} = 1 > \frac{1}{\left[1 + \sigma_B^2\left(\frac{1}{\sigma_X^2} + a_X\right)\right]^2} = \frac{\partial P_F}{\partial a_X}. \tag{17}$$

The cross-signal returns decline with research on the cross-signal:

$$\frac{\partial^2 P_X}{\partial a_F^2} = -\frac{2\sigma_B^2}{(1 + \sigma_B^2 a_F)^3} < 0, \tag{18}$$

$$\frac{\partial^2 P_F}{\partial a_X^2} = -\frac{2\sigma_B^2}{\left[1 + \sigma_B^2\left(\frac{1}{\sigma_X^2} + a_X\right)\right]^3} < 0. \tag{19}$$

Proof of Proposition 1. The adjusted Fed signal $s_F - \mu_B$ is a signal of X with independent error variance $\sigma_B^2 + 1/a_F$. Gaussian updating gives

$$f_X = \frac{\mu_X/\sigma_X^2 + a_X s_X + \frac{a_F}{1+\sigma_B^2 a_F}(s_F - \mu_B)}{P_X}.$$

Differentiating with respect to s_F and substituting equation (4) gives equation (7). Its denominator increases with a_X and σ_B^2 and decreases with a_F, establishing the stated comparative statics. □

Optimal research. Throughout, $C_X, C_F > 0$, $\kappa \geq 0$, $0 < \gamma < 1$, and the prior variances are positive. A finite optimum exists: forecast loss is nonnegative, whereas research costs grow without bound as $a_X + a_F \to \infty$. The objective is continuous on the nonnegative research domain, using the limiting posterior precisions when either research choice is zero.

Write $v_X = P_X^{-1}$, $v_F = P_F^{-1}$, and $q = \mathrm{Cov}(X, F \mid s_X, s_F) > 0$. Let h be the Hessian of forecast loss alone. Gaussian updating gives

$$\begin{aligned} h_{XX} &= 2v_X\left[\gamma v_X^2 + (1-\gamma)q^2\right], \\ h_{FF} &= 2v_F\left[\gamma q^2 + (1-\gamma)v_F^2\right], \\ h_{XF} &= 2q^2\left[\gamma v_X + (1-\gamma)v_F\right]. \end{aligned} \tag{20}$$

The posterior covariance matrix is positive definite, so $\Delta \equiv v_X v_F - q^2 > 0$. These expressions imply that h is positive definite. The Hessian of the full objective is $H = h + \kappa \mathbf{1}\mathbf{1}'$, where $\mathbf{1} = (1,1)'$. It is therefore positive definite, the objective is strictly convex, and the optimum is unique. Write $D_H = H_{XX}H_{FF} - H_{XF}^2 > 0$.

Expanding equation (8) gives

$$C_X + \kappa(a_X + a_F) = \frac{\gamma + \frac{1-\gamma}{(1+\sigma_B^2 a_F)^2}}{\left[a_X + \frac{1}{\sigma_X^2} + \frac{a_F}{1+\sigma_B^2 a_F}\right]^2}.$$

The corresponding condition for Fed research is

$$\begin{aligned} C_F + \kappa(a_X + a_F) &= \frac{1-\gamma}{P_F^2} + \frac{\gamma}{P_X^2(1+\sigma_B^2 a_F)^2} \\ &= \frac{(1-\gamma) + \frac{\gamma}{\left[1+\sigma_B^2\left(\frac{1}{\sigma_X^2}+a_X\right)\right]^2}}{P_F^2}. \end{aligned} \tag{21}$$

The second equality follows from $P_X^{-2}(1+\sigma_B^2 a_F)^{-2} = q^2$ and $P_F^{-2}[1+\sigma_B^2(1/\sigma_X^2 + a_X)]^{-2} = q^2$.

Proof of Proposition 2.

Research allocation. Let $C_F^L < C_F^H$ and let (a_X^L, a_F^L) and (a_X^H, a_F^H) be the respective global optima. Adding the two optimality inequalities gives $(C_F^H - C_F^L)(a_F^L - a_F^H) \geq 0$, hence $a_F^L \geq a_F^H$. This weak monotonicity result allows corners. At an interior optimum, differentiating the first-order conditions gives

$$\frac{da_F^*}{dC_F} = -\frac{H_{XX}}{D_H} < 0, \qquad \frac{da_X^*}{dC_F} = \frac{H_{XF}}{D_H} > 0. \tag{22}$$

The second sign follows from

$$H_{XF} = 2q^2\left[\gamma v_X + (1-\gamma)v_F\right] + \kappa > 0. \tag{23}$$

Thus cheaper Fed research increases Fed research and reduces economic research.

Forecast precision. Rearranging the two research conditions gives

$$\begin{aligned} P_X^2 &= \frac{\gamma + \frac{1-\gamma}{(1+\sigma_B^2 a_F)^2}}{C_X + \kappa(a_X + a_F)}, \\ P_F^2 &= \frac{(1-\gamma) + \frac{\gamma}{\left[1+\sigma_B^2\left(\frac{1}{\sigma_X^2}+a_X\right)\right]^2}}{C_F + \kappa(a_X + a_F)}. \end{aligned} \tag{24}$$

Using $\partial P_X/\partial a_F = q^2/v_X^2$ and $\partial P_F/\partial a_X = q^2/v_F^2$, the precision responses are

$$\frac{dP_X^*}{dC_F} = \frac{H_{XF} - \frac{q^2}{v_X^2}H_{XX}}{D_H} > 0, \qquad \frac{dP_F^*}{dC_F} = \frac{\frac{q^2}{v_F^2}H_{XF} - H_{XX}}{D_H} < 0. \tag{25}$$

The signs follow from the strictly positive expressions

$$\begin{aligned} H_{XF} - \frac{q^2}{v_X^2}H_{XX} &= \frac{2(1-\gamma)q^2\Delta}{v_X} + \kappa(1 - \frac{q^2}{v_X^2}), \\ H_{XX} - \frac{q^2}{v_F^2}H_{XF} &= 2\Delta\left[\frac{\gamma v_X(v_X v_F + q^2)}{v_F^2} + \frac{(1-\gamma)q^2}{v_F}\right] + \kappa(1 - \frac{q^2}{v_F^2}). \end{aligned}$$

Cheaper Fed research therefore worsens economic forecast accuracy while improving policy forecast accuracy at an interior optimum. This result holds with additive costs, $\kappa = 0$, and with positive cost curvature. Once economic research reaches zero, further Fed research improves economic precision through cross-learning.

The weight on the Fed signal in the economic forecast is $[a_F/(1 + \sigma_B^2 a_F)]/P_X$. At an interior optimum, a marginal reduction in C_F increases the numerator and reduces the denominator, so this weight rises. □

C.1 Free Public Information

Cost-saving comparison. Denote optimal private research without a public signal by (a_X^0, a_F^0). For public precision $\tau < a_F^0$, the forecaster could instead choose $(a_X^0, a_F^0 - \tau)$, preserving the original precision of both forecasts. With positive cost curvature, $\kappa > 0$, the marginal cost of each research activity is lower at this comparison allocation:

$$\underbrace{C_i + \kappa(a_X^0 + a_F^0 - \tau)}_{\text{with a public signal}} < \underbrace{C_i + \kappa(a_X^0 + a_F^0)}_{\text{without a public signal}}, \qquad i \in \{X, F\}. \tag{26}$$

Proof of Proposition 3. Let $z = a_F + \tau$ denote total direct Fed-signal precision. The posterior variances depend on the public signal's precision but not its realization. The forecaster's problem is

$$\min_{a_X \geq 0,\ z \geq \tau} \quad \gamma V_X(a_X, z) + (1-\gamma) V_F(a_X, z) + C_X a_X + C_F(z - \tau) + \frac{\kappa}{2}(a_X + z - \tau)^2. \tag{27}$$

At $\kappa = 0$, changing τ adds only the constant $-C_F\tau$ and changes the lower bound on z. Until that bound binds, the unique optimum in (a_X, z) is unchanged, so $da_F^*/d\tau = -1$ and $da_X^*/d\tau = 0$.

For $\kappa > 0$, use the same Hessian expressions as above, now in (a_X, z), with posterior variances and covariance conditional on all three signals. Write h_{zz} and h_{Xz} for the corresponding forecast-loss derivatives. At an interior optimum, differentiating the first-order conditions gives

$$\begin{aligned} \frac{da_X^*}{d\tau} &= \frac{\kappa(h_{zz} - h_{Xz})}{D_H}, & \frac{dz^*}{d\tau} &= \frac{\kappa(h_{XX} - h_{Xz})}{D_H}, \\ \frac{da_F^*}{d\tau} &= \frac{dz^*}{d\tau} - 1, & \frac{d(a_X^* + a_F^*)}{d\tau} &= -\frac{\det h}{D_H} < 0. \end{aligned} \tag{28}$$

Thus total private effort falls. The individual signal-precision responses depend on the differences between own and cross curvature of forecast loss.

Research allocation. The two research conditions imply

$$C_F - C_X = (1-\gamma)(v_F^2 - q^2) - \gamma(v_X^2 - q^2).$$

Consequently,

$$\begin{aligned} h_{XX} - h_{Xz} &= 2(1-\gamma)v_F\Delta - 2v_X(C_F - C_X), \\ h_{zz} - h_{Xz} &= 2\gamma v_X\Delta + 2v_F(C_F - C_X). \end{aligned} \tag{29}$$

If $C_F \geq C_X$, the second expression is positive, so economic research rises. Since total private effort falls, private Fed research decreases. The share of private research devoted to the economy therefore rises. These allocation signs need not hold when $C_F < C_X$.

Forecast precision. Both posterior precisions rise at an interior optimum for arbitrary positive C_X, C_F when $\kappa > 0$. To show this, write $\delta_X = h_{XX} - h_{Xz}$ and $\delta_z = h_{zz} - h_{Xz}$. Then

$$\frac{dP_X^*}{d\tau} = \frac{\kappa(\delta_z + \frac{q^2}{v_X^2}\delta_X)}{D_H} > 0, \qquad \frac{dP_F^*}{d\tau} = \frac{\kappa(\delta_X + \frac{q^2}{v_F^2}\delta_z)}{D_H} > 0. \tag{30}$$

The signs follow from

$$\begin{aligned} \delta_z + \frac{q^2}{v_X^2}\delta_X &= \frac{2\Delta}{v_X^2}\left\{\gamma v_X q^2 + (1-\gamma)\left[v_X(v_F^2 - q^2) + v_F q^2\right]\right\} > 0, \\ \delta_X + \frac{q^2}{v_F^2}\delta_z &= \frac{2\Delta}{v_F^2}\left\{(1-\gamma)v_F q^2 + \gamma\left[v_F(v_X^2 - q^2) + v_X q^2\right]\right\} > 0. \end{aligned}$$

Here $v_X > q$ and $v_F > q$ at an interior allocation. Thus both forecast errors have lower conditional expected squares, even when the forecaster does not increase both direct signal precisions. □

These comparative statics vary public precision holding research costs and preferences fixed. At a given public precision, differences in private research costs or forecasting preferences can instead generate the cross-sectional policy–economic accuracy tradeoff.

Both channels operating together. At an interior optimum, the positive-definite Hessian implies, by the implicit function theorem, that optimal research varies continuously with research costs and public-signal precision. Forecast precisions and signal weights are also continuous functions of these choices. Fix a reduction in Fed research costs for which the allocation remains interior. Holding public precision fixed, the reduction increases Fed research, policy forecast precision, and reliance on Fed information, while reducing economic research and economic forecast precision. These strict changes persist for a sufficiently small accompanying increase in public-signal precision.

No private Fed research. If private Fed research is zero, $z = \tau$. While economic research remains positive and this constraint continues to bind,

$$\begin{aligned} \frac{da_X^*}{d\tau} &= -\frac{h_{Xz}}{h_{XX}+\kappa} < 0, \\ \frac{dP_X^*}{d\tau} &= \frac{q^2}{v_X^2} - \frac{h_{Xz}}{h_{XX}+\kappa}, \\ \frac{dP_F^*}{d\tau} &= 1 - \frac{q^2}{v_F^2}\frac{h_{Xz}}{h_{XX}+\kappa} > 0. \end{aligned} \tag{31}$$

The economic-precision response can have either sign. The policy-precision response is positive by the expression for $H_{XX} - \frac{q^2}{v_F^2} H_{XF}$ above, with the additional nonnegative term $\kappa \frac{q^2}{v_F^2}$. If both private research choices are zero, further public precision improves both forecasts directly while both constraints remain binding.

D Appendix: Additional Figures

Figure A1: Timeline of the model.

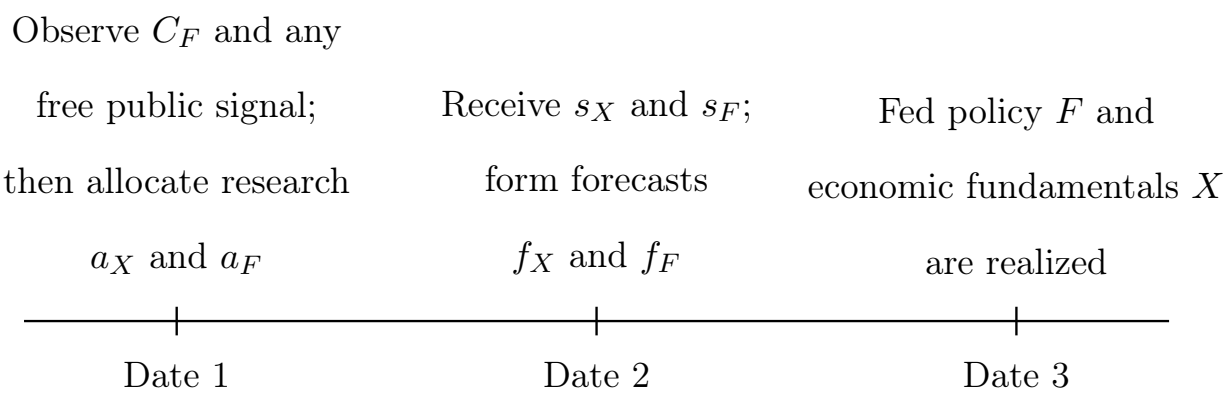
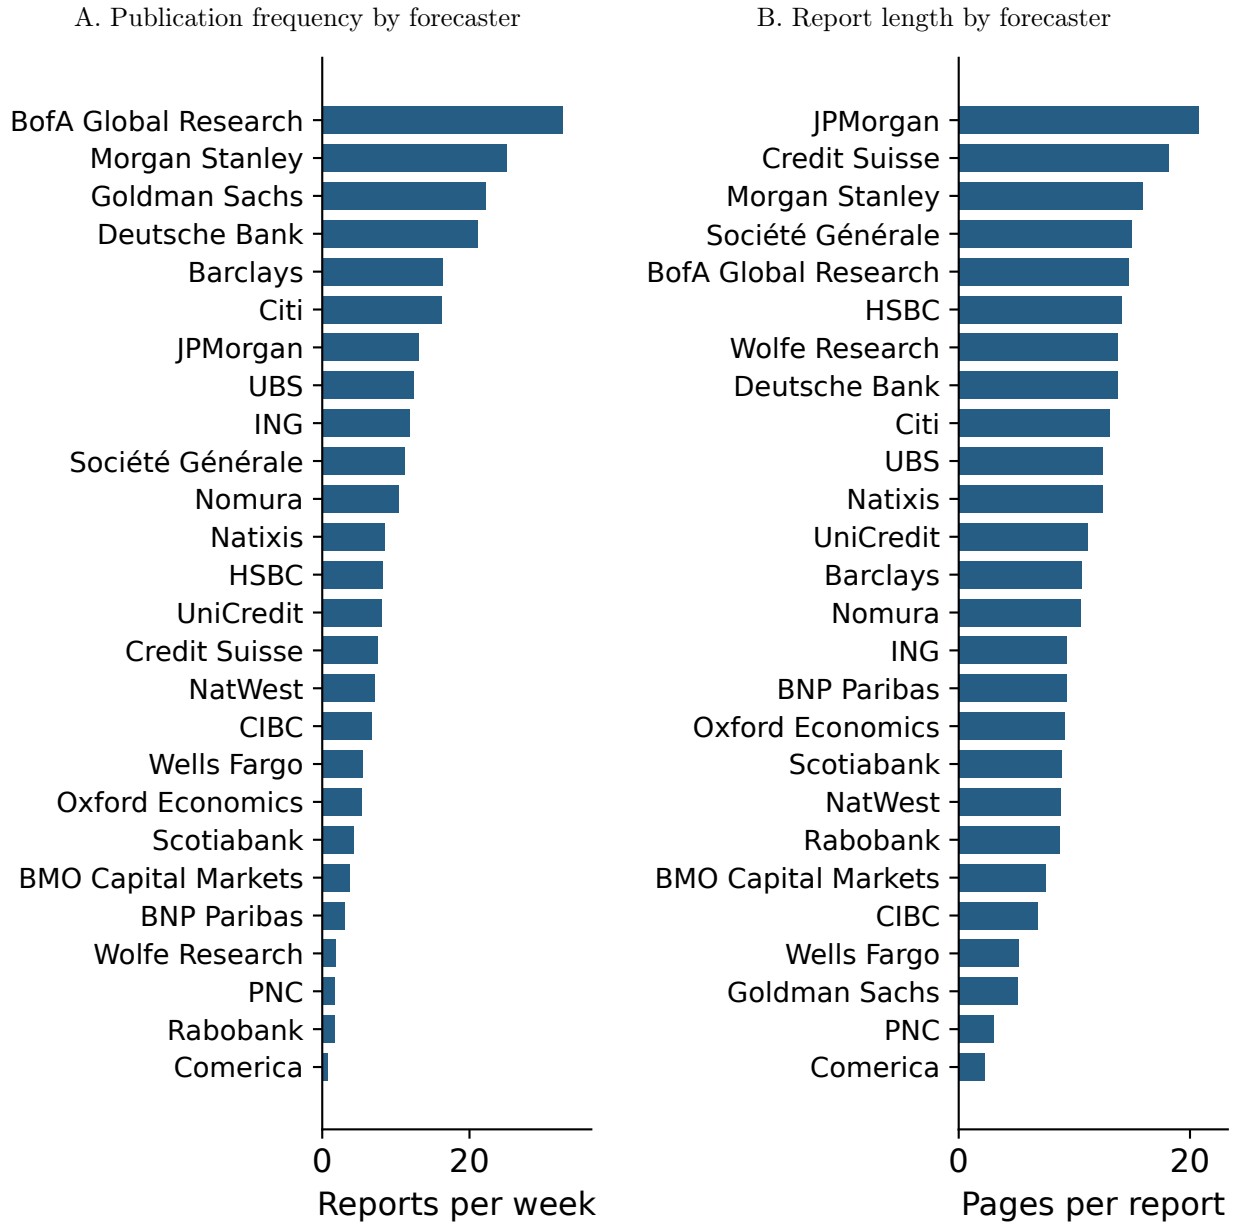
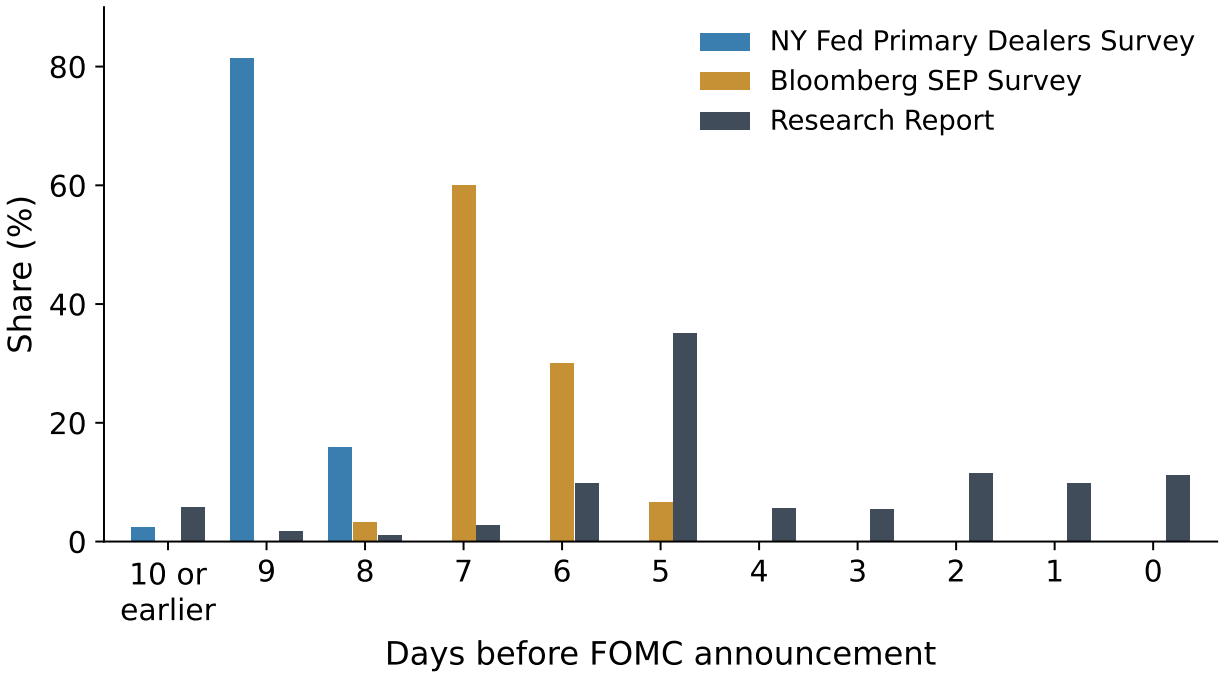


Figure A2: Publication Frequency and Report Length



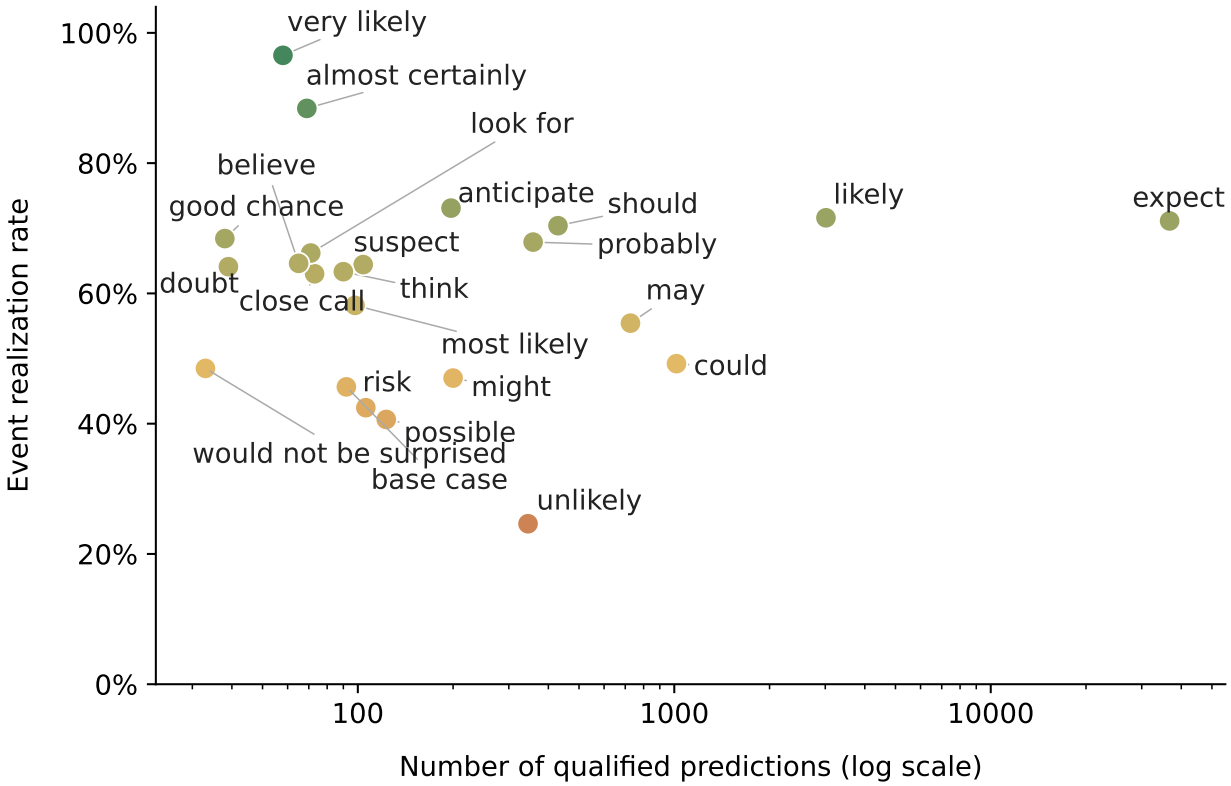
Notes: The sample covers the 26 forecasters in Figure II during 2000–2025. Bars average each forecaster's annual reports per calendar week and mean pages per report over all available years, sorted separately.

Figure A3: Fed Announcement Survey Deadlines vs Publication Date of Expectations



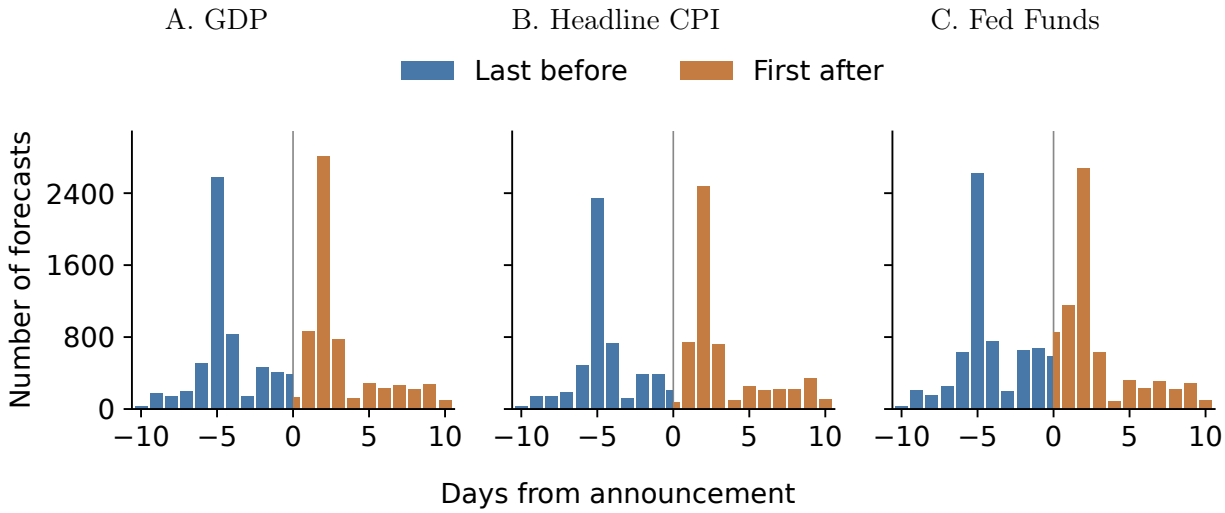
Notes: This figure plots the distribution of days before a scheduled FOMC announcement that expectations were published on, for research reports, or surveyed on, for the NY Fed Primary Dealers Survey and Bloomberg SEP Survey. The NY Fed survey spans 2011-2025, including the successor Survey of Market Expectations in 2025, and reflects the survey due date. The Bloomberg survey contains the due date of the surveys from December 2017 through December 2025.

Figure A4: Confidence Qualifiers and Communication Realization Rates



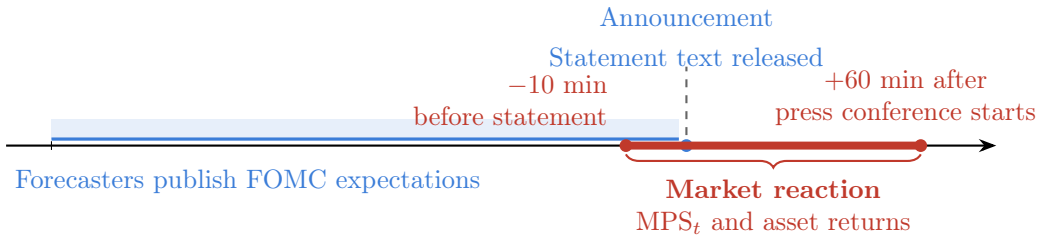
Notes: Each point shows a qualifier's in-sample realization rate and number of verifiable qualified predictions across 207 scheduled FOMC meetings, February 2000–December 2025. Alternatives sharing a qualifier within a forecaster–question observation form one prediction, realized if any alternative is true or partially true. Only qualifiers with more than 30 observations are shown. These rates supply the probability weights in equation (13); unambiguous stated numerical probabilities enter directly. The qualifier "expect" is used when no qualifier is stated. Appendix B.3 describes the treatment of rare qualifiers and unverifiable events.

Figure A5: Timing of Forecasts Around FOMC Announcements



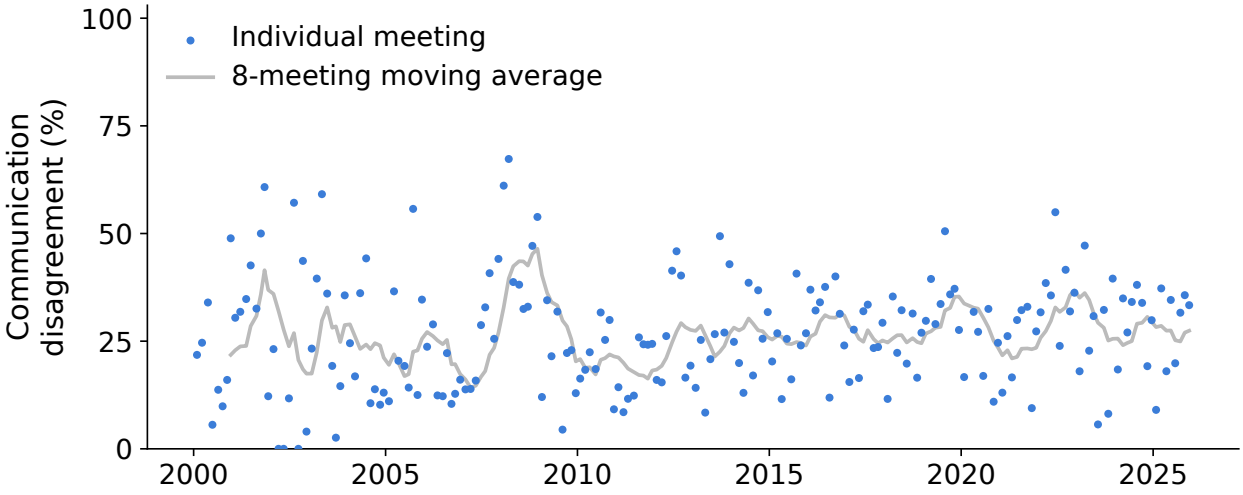
Notes: The bars count the last forecast before and first forecast after each scheduled FOMC announcement in 2000–2025, for the institutions in Table I. In each panel, a forecaster can contribute up to three forecasts on each side, one for each of the next three target quarters. Reports on the announcement day are classified as before or after the announcement. The chart includes forecasts within ten days on either side; the revision regressions retain pairs no more than ten days apart. For Wednesday announcements, days −5 and +2 are Fridays.

Figure A6: Timing of Text Surprises and High-Frequency Asset Prices



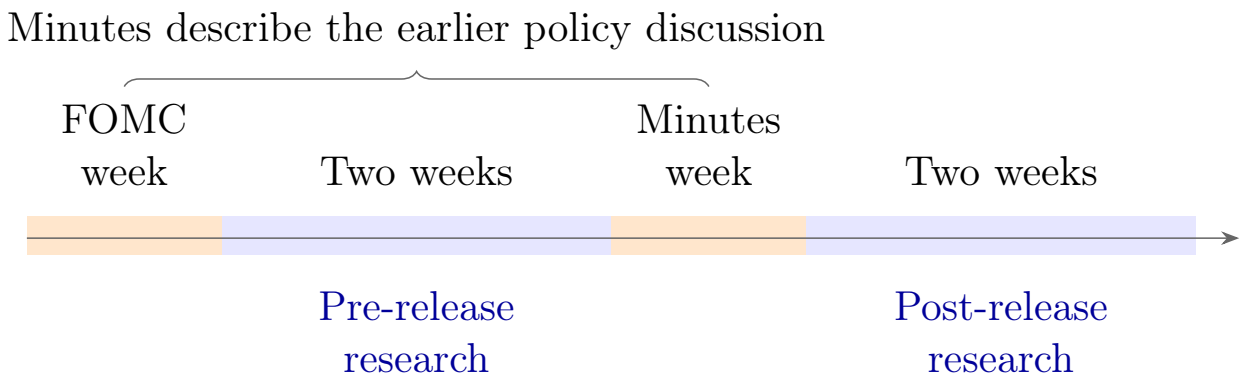
Notes: Forecasts and their reasoning come from preview reports published in the three weeks before each meeting (Appendix B.3). I use the latest forecast when reports disagree, then compare it with the announcement. MPS and asset returns use the USMPD window shown above (Section 6.2). Without a press conference, the window ends 20 minutes after the statement. Not to scale.

Figure A7: Communication Disagreement over Time



Notes: Communication disagreement is the share of forecaster pairs giving different answers to the same Fed announcement question, expressed as a percentage. It equals one minus agreement, with questions weighted by the number of forecaster-pair comparisons. Each dot represents one of 207 scheduled FOMC meetings from February 2000 through December 2025. The questions cover announcement communications, including press conferences; their construction is described in Appendix B.3. The gray line is a trailing eight-meeting average. Higher values indicate greater disagreement about the forthcoming communication.

Figure A8: Research Windows Around the Minutes Release



Comparison: Does unexpected topic emphasis in the minutes predict a change in the topics forecasters research after the release?

Construction: Compare topic shares in research before and after the minutes week, using the regression specification and controls in Table VI.

Notes: The diagram illustrates the research windows in Table VI: two research weeks between the FOMC week and minutes week, and two research weeks after the minutes week. Forty topic clusters follow Hansen et al. (2018); the sample covers minutes releases in 2005–2025. Controls include pre-release forecaster topic shares, own prior-year coverage and earlier coverage changes, and the previous and prior-eight average minutes topic shares.

Figure A9: Document Processing Pipeline

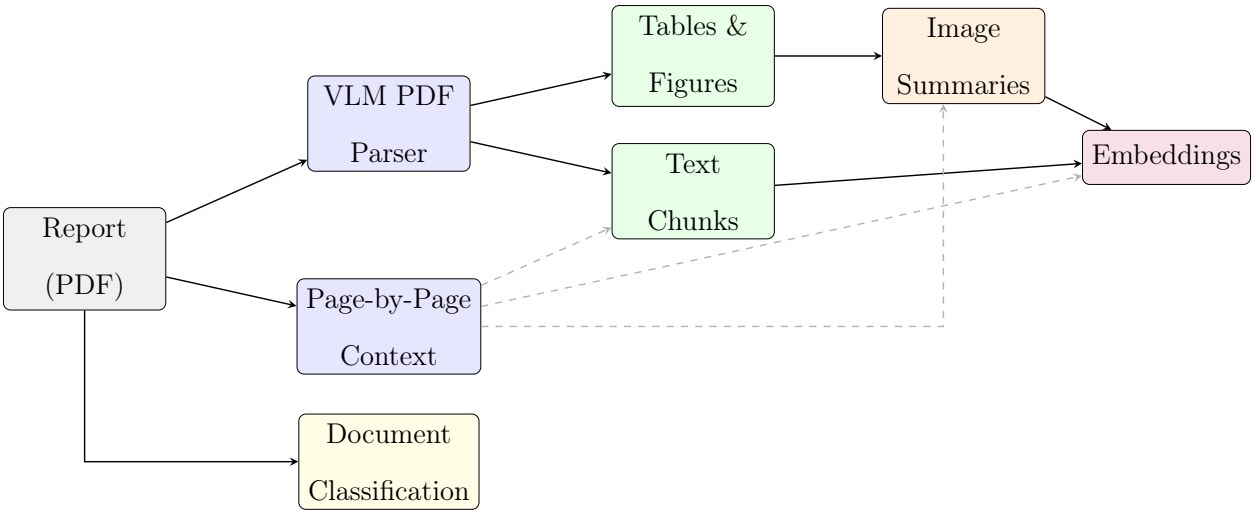


Figure A10: Anatomy of a Document Chunk

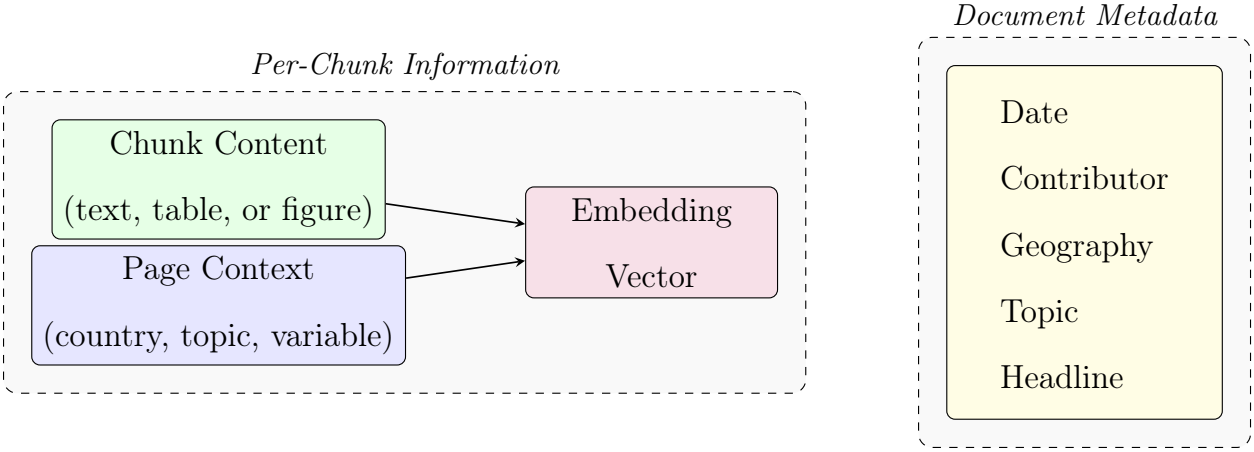


Figure A11: From Report Text to Fed Research Share

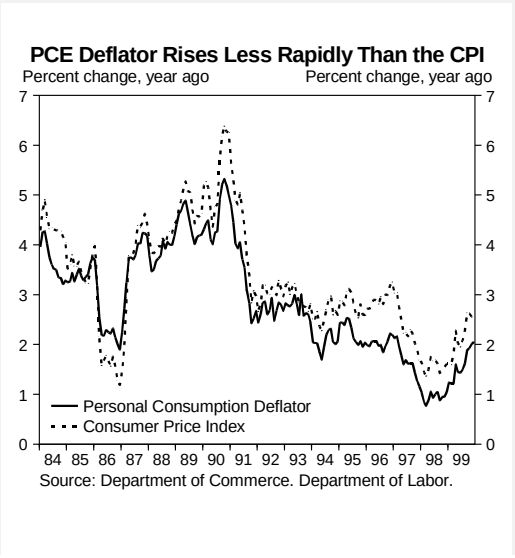
Goldman Sachs

GS Global Economics Website
Economics Research from the GS Financial Workbench℠ at https://www.gs.com

U.S. Economics Analyst

Inflation: It's the Direction that Matters

- The only surprise in the Federal Open Market Committee's decision to boost interest rates by 25 basis points and to voice concern about inflation risks was the moderate tone of its statement given recent labor cost acceleration.
- Although this statement reinforced our belief that Fed officials would prefer to move slowly, persistent signs of strength in the economy will probably convince them to tighten again by 25 basis points at the March 21 FOMC meeting.
- In divining future policy moves, it makes little difference whether market participants follow changes in Chairman Greenspan's preferred measure of consumer inflation—the deflator for personal consumption expenditures—or changes in the CPI. This is because the Fed responds to actual or expected *changes* in inflation. For that purpose the two tell almost the same story.



Reasoning about the Fed

"Although this statement reinforced our belief that Fed officials would prefer to move slowly, persistent signs of strength in the economy will probably convince them to tighten again by 25 basis points at the March 21 FOMC meeting."

Goldman Sachs, *US Economics*, February 4, 2000. The shaded paragraph reasons about the Fed's next policy decision.

Figure A12: Quantitative Forecast Extraction Pipeline

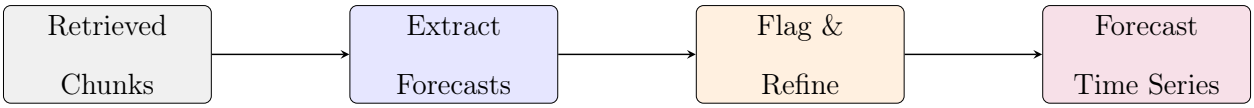
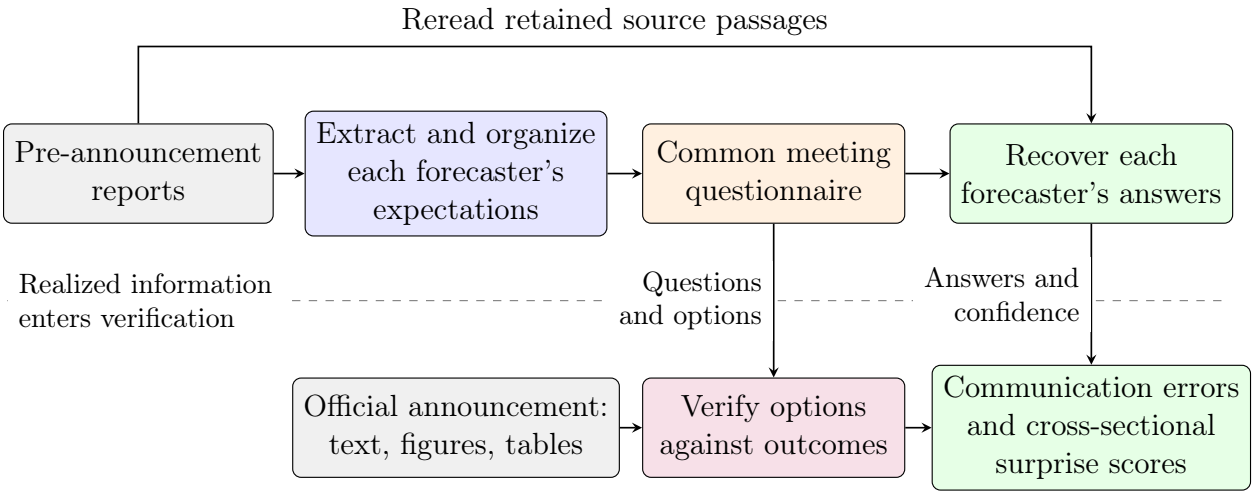


Figure A13: FOMC Forecast Extraction and Surprise Scoring



Notes: Questions are derived from forecasters' reports separately for each meeting. Their source evidence is reread to recover answers and confidence language. Comparing answers with verified outcomes gives communication errors, which are aggregated into an overall error rate and an economic share of errors.

E Appendix: Additional Tables

Table A1: Selected Views on Central Bank Communication

Speaker and date	Quotation
Ben Bernanke[a] October 7, 2004	"When the monetary policy committee regularly provides information about its objectives, economic outlook, and policy plans ... markets will price financial assets more efficiently ... the policy committee attains increased influence over the most economically relevant long-term yields, reduced financial and economic uncertainty, and, in all probability, better economic outcomes."
Kevin Warsh[b] July 29, 2026	"In the intermeeting period, market attention centered on real data and real economic developments ... Market participants are learning to play the ball, not the referee—and market prices will continue to respond in the direction and magnitude they see fit ... the central bank need not always and everywhere be the center of attention."
Stephen Morris and Hyun Song Shin[c] 2005	"Public pronouncements can thus crowd out private opinions, and the market may cease to function as a way of aggregating and revealing diverse, private judgments about the world in the way that Hayek envisaged."
Stephen Poloz[d] May 25, 2020	"When market participants pay attention only to the words we say about future policy and begin to discount the flow of actual economic data, they set up a one-way bet in markets. Market prices cease acting as a check for our own projections."
Donald Kohn[e] January 9, 2005	"In any case, the risks of herding, of overreaction, of too little scope for private assessments of economic developments to show through, would seem to be high for central bank talk about policy interest rates."
Tommaso Padoa-Schioppa[f] July 15, 2004	"The danger that I see is that the market becomes lazy. And, in a way, if it is spoiled to the point that it is told everything in advance, it relies not on its own analysis, but on the analysis of the central bank."

Notes: Selected excerpts; ellipses indicate omissions.

Sources: [a] Federal Reserve speech, *Central Bank Talk and Monetary Policy.* [b] FOMC press conference transcript. [c] *Central Bank Transparency and the Signal Value of Prices*, p. 19. [d] Bank of Canada, *Monetary Policy in Unknowable Times*, p. 11. [e] Federal Reserve speech, *Central Bank Communication*, Policy Inclination section. [f] *Wall Street Journal Europe* interview, July 15, 2004.

Table A2: Extraction Validity: Correlation with Reuters Economic Poll Submissions

Variable	Extracted forecasts	Matched forecasts	Correlation
Fed funds rate (Level)	99,925	35,261	0.998
2Y Treasury (Level)	46,134	947	0.995
10Y Treasury (Level)	76,258	1,916	0.987
GDP growth (QoQ SAAR)	82,608	25,930	0.978
Headline CPI (YoY)	55,372	18,527	0.975

Notes: I match each extracted forecast to the nearest Reuters (Refinitiv) Economic Poll submission from the same bank for the same target quarter, within seven days of the report date. The last column reports the Pearson correlation across these matched pairs. Counts refer to bank–report-date–target-quarter observations. Extracted forecasts come from banks covered by each Reuters series; matched forecasts are the subset with a poll match. I use reports dated 2000–2025, excluding 2020. Variable definitions follow Table III.

Table A3: Fed Research and Distance from Consensus

Dependent variable: Mean absolute distance from consensus (bp).

	Fed funds	GDP growth	CPI (QoQ)	CPI (YoY)
		Next three quarters		
	(1)	(2)	(3)	(4)
Fed research share (15 pp)	1.1**	-0.9	-2.4	1.1
	(0.5)	(1.3)	(2.5)	(0.9)
Date FE	Yes	Yes	Yes	Yes
Institution FE	Yes	Yes	Yes	Yes
Log research volume	Yes	Yes	Yes	Yes
Observations	10,637	8,199	2,005	5,880

Notes: The regression specification is

$$D_{it} = \mu_i + \delta_t + \beta A_{it} + \boldsymbol{\gamma}'\mathbf{Z}_{it} + \varepsilon_{it}.$$

Research share, fixed effects, and controls follow column (4) of Table VII, excluding distance from consensus from the controls. I keep the Fed funds and GDP samples from that table. The two CPI columns use directly reported quarterly annualized (QoQ) and year-over-year (YoY) forecasts, with all three horizons available. Consensus is the median of each other forecaster's latest forecast for the same target quarter, published on the report date or in the preceding 14 days. Observations are forecaster–date pairs from 2000–2025. Parentheses report standard errors clustered by forecaster and target bundle. Stars use normal-reference p-values: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A4: Rate Forecast Accuracy by Horizon: Varying Samples

Dependent variable: Absolute forecast error (bp).

	Fed research share (15 pp)					
	Forecast horizon (quarters ahead)					
Forecast	1	2	3	4	5	6
Fed funds	-1.0**	-2.0**	-2.9**	-3.0**	-1.8	0.8
	(0.4)	(0.9)	(1.1)	(1.3)	(1.5)	(2.1)
N	15,868	12,309	10,637	7,644	5,062	2,910
2-year Treasury	-0.6	0.3	1.0	5.7***	10.8***	10.8***
	(1.0)	(1.0)	(1.4)	(2.0)	(2.6)	(3.6)
N	6,761	6,421	5,814	3,326	1,855	973
5-year Treasury	1.5	3.0**	3.2**	7.8***	11.5***	11.0***
	(1.0)	(1.3)	(1.6)	(2.5)	(3.6)	(4.1)
N	5,093	4,959	4,646	2,650	1,551	850
10-year Treasury	0.6	1.0	1.7*	3.0**	4.2*	5.4*
	(0.8)	(1.1)	(1.0)	(1.4)	(2.3)	(3.1)
N	12,482	10,825	9,471	5,705	3,221	1,700
30-year Treasury	-0.1	0.9	0.7	-1.5	-4.7**	-2.3
	(0.8)	(1.3)	(1.5)	(1.8)	(1.9)	(5.2)
N	4,051	3,924	3,605	1,642	708	288

Notes: The regression specification is

$$|\hat{r}^{v}_{it,q} - r^{v}_{q}| = \alpha + \beta_{v,h} A_{it} + \delta^{v,h}_{t} + \mu^{v,h}_{i} + \text{Controls}^{v,h}_{it} + \varepsilon^{v,h}_{it,q}.$$

Research share, controls, and fixed effects follow Table VIII. The forecasts were published in 2000–2025. For each column, I require forecasts for all quarters from the next quarter through that horizon, from the same forecaster on the same date. I report the number of forecaster–date observations below each estimate. Table VIII instead uses the same sample for all four horizons. Parentheses report standard errors clustered by forecaster and target quarter. Stars use normal-reference p-values: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A5: Treasury Forecast Accuracy: Yield-Component Interactions

Dependent variable: Absolute Treasury forecast error (bp).

	2-year		5-year		10-year	
	Four quarters ahead		*Four quarters ahead*		*Four quarters ahead*	
	(1)	(2)	(3)	(4)	(5)	(6)
Fed research share (15 pp)	5.65***	6.19***	7.80***	8.50***	2.96**	3.02**
	(1.95)	(2.08)	(2.53)	(2.42)	(1.43)	(1.46)
× term-premium change (pp)		2.11		-3.12		1.18
		(8.23)		(7.10)		(1.91)
× expected-rate change (pp)		-2.48		-4.12**		-1.86
		(1.71)		(1.97)		(2.28)
Yield-component interactions	No	Yes	No	Yes	No	Yes
Forecaster–date observations	3,326	3,326	2,650	2,650	5,705	5,705

Notes: The regression specification is

$$|\hat{r}_{it,q} - r_q| = \mu_i + \delta_{tq} + \beta A_{it} + \theta A_{it} P_{tq} + \kappa A_{it} R_{tq} + \boldsymbol{\gamma}' \mathbf{Z}_{it} + \varepsilon_{it,q}.$$

A denotes Fed research share; controls $\mathbf{Z}$ and sample selection follow Table VIII. P is the absolute change in the Adrian–Crump–Moench term premium; R is the absolute change in the component reflecting expected average short rates. I measure both for the same maturity, from the close before publication to target quarter-end, in pp. Each change is centered at its median across unique date–target pairs; date–target effects absorb their levels. The sample is 2000–2025, with the same observations in each column pair. Standard errors in parentheses are clustered by institution and target quarter. Stars use normal-reference p-values: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A6: Treasury Forecast Accuracy: Near-Term Error Controls

Dependent variable: Absolute Treasury forecast error (bp).

	2-year		5-year		10-year	
	Four quarters ahead		*Four quarters ahead*		*Four quarters ahead*	
	Baseline	Controls	Baseline	Controls	Baseline	Controls
	(1)	(2)	(3)	(4)	(5)	(6)
Fed research share (15 pp)	7.0***	8.8***	7.6**	8.5***	4.7**	5.1**
	(2.6)	(2.2)	(3.1)	(2.8)	(2.1)	(2.1)
Fed funds RMSE (pp)		105.0***		78.1***		48.8***
		(12.5)		(13.6)		(8.0)
GDP RMSE (pp)		1.1		1.2		0.5
		(0.7)		(0.8)		(1.5)
CPI forecast error (SD)		5.8		4.2		2.1
		(5.3)		(6.6)		(5.2)
Observations	2,335	2,335	1,924	1,924	3,791	3,791

Notes: The regression specification is

$$|\hat{r}_{it,q} - r_q| = \mu_i + \delta_{tq} + \beta A_{it} + \boldsymbol{\theta}'\mathbf{E}_{it,q} + \boldsymbol{\gamma}'\mathbf{Z}_{it} + \varepsilon_{it,q}.$$

Fed research share A, baseline controls $\mathbf{Z}$, fixed effects, and standard errors follow Table VIII. The added controls $\mathbf{E}$ are the forecaster's errors for Fed funds, GDP, and CPI. For Fed funds and GDP, I use RMSE over the next four quarters. For CPI, I use either quarterly RMSE or the target-quarter year-over-year absolute error, taking the more recent forecast and using year-over-year when both are from the same date. I standardize CPI errors separately for the two formats. The macro forecasts must be from the same institution and no more than 14 days old. Each column pair uses the same forecaster–date observations from 2000–2025. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A7: Forecast Accuracy: Attention Windows and Weighting

Dependent variable: Forecast error (bp).

Forecast	Fed research share (15 pp)			N
	30-day average	14-day average	14-day half-life	
	(1)	(2)	(3)	
A. RMSE over the next three quarters				
Fed funds	-2.1***	-1.6***	-2.6***	10,630
	(0.8)	(0.5)	(0.9)	
GDP growth	5.0***	3.9***	6.1***	8,196
	(1.8)	(1.3)	(2.2)	
Quarterly CPI	14.0***	9.1***	16.4***	2,005
	(5.0)	(2.8)	(5.1)	
B. Absolute error four quarters ahead				
2-year Treasury	5.7***	2.2*	6.2***	3,325
	(2.0)	(1.3)	(2.3)	
5-year Treasury	7.9***	3.5**	8.8***	2,649
	(2.5)	(1.7)	(2.8)	
10-year Treasury	3.0**	2.0**	3.4**	5,703
	(1.4)	(0.8)	(1.5)	
30-year Treasury	-1.5	-0.4	-0.8	1,642
	(1.8)	(1.2)	(2.0)	
Date FE	Yes	Yes	Yes	
Institution FE	Yes	Yes	Yes	

Notes: The regression specification is

$$L_{it} = \mu_i + \delta_t + \beta A_{it} + \boldsymbol{\gamma}'\mathbf{Z}_{it} + \varepsilon_{it}.$$

Specifications and standard errors (in parentheses) follow column (4) of Table VII for Panel A and Table VIII for Panel B. N counts forecaster–date observations; within each row, I use the same observations in all three columns. Forecasts date from 2000–2025. The average columns pool Fed research and total analytical characters over the forecast date and preceding 30 or 14 days. For the half-life measure, I weight both Fed research and total analytical characters from d days earlier by $2^{-d/14}$, using all available past research (Section 3.3). The research-volume control uses the forecast date and preceding 30 days in all columns. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A8: Forecast Accuracy and Alternative Measures of Fed Attention

Dependent variable: Forecast error (bp).

Forecast	Fed attention (1 SD)			N
	Direct Fed reasoning (1)	Fed motivation (2)	Keywords (3)	
A. RMSE over the next three quarters				
Fed funds	-1.7**	-1.4**	-2.7***	10,637
	(0.7)	(0.6)	(1.0)	
GDP growth	6.6***	2.4**	4.8**	8,199
	(2.1)	(1.2)	(2.4)	
Quarterly CPI	11.0***	7.3*	7.8**	2,005
	(3.8)	(4.1)	(3.4)	
B. Absolute error four quarters ahead				
2-year Treasury	4.7**	4.2***	3.4**	3,326
	(2.1)	(1.6)	(1.7)	
5-year Treasury	6.1**	5.9***	8.2***	2,650
	(2.9)	(2.1)	(3.0)	
10-year Treasury	2.3	2.2**	3.8**	5,705
	(2.0)	(1.1)	(1.6)	
30-year Treasury	-2.7	-1.0	2.2	1,642
	(1.7)	(1.6)	(2.7)	
Date FE	Yes	Yes	Yes	
Institution FE	Yes	Yes	Yes	

Notes: The regression specification is

$$L_{it} = \mu_i + \delta_t + \beta A_{it} + \boldsymbol{\gamma}'\mathbf{Z}_{it} + \varepsilon_{it}.$$

Specifications and standard errors (in parentheses) follow column (4) of Table VII for Panel A and Table VIII for Panel B. Coefficients are per sample standard deviation of each measure. N counts forecaster–date pairs from 2000–2025, held fixed across columns within each row. Measures cover the forecast date and preceding 30 days. Column (1) uses direct Fed reasoning as a share of analytical text. Column (2) averages the share of each article motivated by the Fed, weighted by characters. Column (3) uses the character share of paragraphs mentioning Fed, Federal Reserve, FOMC, Federal Open Market Committee, Greenspan, Bernanke, Yellen, or Powell. Columns (2) and (3) include factual prose. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A9: Chief Economist Fixed Effects: Forecast Accuracy

Dependent variable: Forecast error (bp).

Forecast	Fed research share (15 pp)		N
	Institution FE	Chief economist FE	
	(1)	(2)	
A. RMSE over the next three quarters			
Fed funds	-2.1***	-1.9***	10,637
	(0.8)	(0.7)	
GDP growth	5.0***	3.5**	8,199
	(1.8)	(1.4)	
Quarterly CPI	14.0***	12.9***	2,005
	(5.0)	(4.4)	
B. Absolute error four quarters ahead			
Fed funds	-3.0**	-2.6**	7,644
	(1.3)	(1.1)	
2-year Treasury	5.7***	4.6***	3,326
	(2.0)	(1.8)	
5-year Treasury	7.8***	6.5***	2,650
	(2.5)	(2.3)	
10-year Treasury	3.0**	2.4**	5,705
	(1.4)	(1.0)	
30-year Treasury	-1.5	0.0	1,642
	(1.8)	(2.4)	

Notes: The regression specification is

$$L_{it} = \mu_{g(i,t)} + \delta_t + \beta A_{it} + \boldsymbol{\gamma}'\mathbf{Z}_{it} + \varepsilon_{it},$$

where $g(i,t)$ indexes institutions in column (1) and institution–chief economist pairs in column (2). I assign chief economists using the forecast publication date. Research share, controls, date fixed effects, and standard errors (in parentheses) follow column (4) of Table VII for Panel A and Table VIII for Panel B. N counts forecaster–date pairs from 2000–2025; each row uses the same observations in both columns. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A10: Fed Research and Forecast Accuracy: Alternative Inference

Dependent variable: Forecast error (bp).

	Fed research share (15 pp)			
	Estimate	95% confidence interval		
		Clustered	Clustered	Driscoll–Kraay
		Normal reference	*t reference*	*One-year bandwidth*
A. RMSE over the next three quarters				
Fed funds	-2.11	[-3.69, -0.52]	[-3.79, -0.42]	[-4.13, -0.09]
GDP	4.97	[1.34, 8.59]	[1.11, 8.83]	[0.27, 9.67]
Quarterly CPI inflation	13.95	[4.25, 23.66]	[3.40, 24.51]	[3.43, 24.48]
B. Absolute error four quarters ahead				
Fed funds	-3.00	[-5.46, -0.54]	[-5.61, -0.38]	[-6.11, 0.12]
2-year Treasury	5.65	[1.83, 9.47]	[1.55, 9.75]	[1.51, 9.79]
5-year Treasury	7.80	[2.84, 12.76]	[2.46, 13.13]	[3.32, 12.27]
10-year Treasury	2.96	[0.15, 5.77]	[-0.03, 5.95]	[0.28, 5.64]
30-year Treasury	-1.46	[-4.90, 1.98]	[-5.18, 2.26]	[-4.93, 2.02]

Notes: The regression specification is

$$L_{it} = \mu_i + \delta_t + \beta A_{it} + \boldsymbol{\gamma}'\mathbf{Z}_{it} + \varepsilon_{it}.$$

Panel A uses column (4) of Table VII. Panel B uses the four-quarter-ahead estimates in Table VIII. I keep the samples and regressions from those tables and compare their 95% confidence intervals. The forecasts span 2000–2025. For clustered inference, I cluster by institution and target bundle in Panel A, and by institution and target quarter in Panel B. The t intervals use degrees of freedom equal to the smaller cluster count minus one. For Driscoll–Kraay, I aggregate scores by day and use a 365-calendar-day Bartlett bandwidth, with fixed-b critical values (Kiefer and Vogelsang, 2005).

Table A11: Incorrect Expectations and Post-Announcement Fed Reasoning

Dependent variable: Discussion similarity (% of pre-meeting mean).

Announcement day after release and following day

Equal weight per:	Observation		Forecaster–meeting	
	(1)	(2)	(3)	(4)
Incorrect share (0 to 1)	3.72***	3.62***	5.90***	5.98***
	(1.15)	(0.98)	(1.86)	(1.62)
Pre-meeting discussion	No	Yes	No	Yes
Forecaster–meeting FE	Yes	Yes	Yes	Yes
Meeting–question FE	Yes	Yes	Yes	Yes
Observations	9,472	9,472	9,472	9,472

Notes: The regression specification is

$$D^{\text{post}}_{imq} = \alpha_{im} + \delta_{mq} + \beta E_{imq} + \gamma D^{\text{pre}}_{imq} + \varepsilon_{imq},$$

where i indexes forecasters, m meetings, and q communication questions. E is the share of a forecaster's answers to a question that were incorrect. D measures sentence-level TF-IDF cosine similarity to the question, weighted by sentence word counts and paragraph-level direct Fed reasoning and divided by total retained report words. Pre-meeting discussion covers days -7 to -1. Coefficients and standard errors are multiplied by $100/\overline{D}^{\text{pre}}$, using the same pre-meeting mean in all columns.

The sample covers 2000–2025. Each observation has an assessable expectation and report text in both windows. Columns (3)–(4) give each forecaster–meeting equal total weight. Parentheses report standard errors clustered by forecaster and meeting; stars use t with degrees of freedom equal to the smaller cluster count minus one. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A12: Chief Economist Fixed Effects: Forecast Revisions

Dependent variable: Forecast revision around FOMC announcements (bp).

	Fed funds *Current and next three quarters*	GDP growth *Next three quarters*		CPI inflation *Next three quarters*	
	(1)	(2)	(3)	(4)	(5)
Forecaster surprise (1 SD)	1.01***	1.30***	1.22***	1.20**	1.14**
	(0.30)	(0.41)	(0.41)	(0.49)	(0.49)
Fed funds revision (bp)			0.08		0.07
			(0.13)		(0.07)
Meeting fixed effects	Yes	Yes	Yes	Yes	Yes
Chief economist fixed effects	Yes	Yes	Yes	Yes	Yes
Forecaster–meeting pairs	1,037	1,031	1,031	871	871

Notes: This table repeats Table IX with chief economist fixed effects in place of institution fixed effects. The regression specification is

$$\Delta f_{it}^{k} = \alpha_{t}^{k} + \lambda_{c(i,t)}^{k} + \beta^{k}\widetilde{S}_{it}^{k} + \delta^{k}\Delta f_{it}^{FF} + \boldsymbol{\gamma}^{k\prime}\mathbf{Z}_{it}^{k} + \varepsilon_{it}^{k}.$$

Chief economist assignments, $c(i,t)$, follow Table A9. All other variables and controls are defined in Table IX. Standard errors are clustered by meeting. Stars use normal-reference p-values: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A13: Effect of Monetary Policy on Average Forecast Revisions

Dependent variable: Average forecast revision (bp).

	Monetary policy surprise (1 SD)		
	Monthly	Ten-day window	
Forecast	(1)	(2)	(3)
Fed funds	4.99***	2.37***	2.19***
Current and next three quarters	(1.70)	(0.63)	(0.62)
N	4,129	7,013	7,013
GDP growth	3.62**	−1.02	−1.63**
Next three quarters	(1.73)	(0.64)	(0.65)
N	4,081	9,477	9,477
CPI inflation	−0.97	−0.39	−0.88*
Next three quarters	(1.11)	(0.58)	(0.51)
N	1,700	2,930	2,930
Economic controls	No	No	Yes

Notes: The regression specification is

$$\Delta f_{it}^{k} = \alpha_i^k + \theta^k I_{it}^{\text{FOMC}} + \beta^k \widetilde{\text{MPS}}_{it}^{k} + \boldsymbol{\gamma}^{k\prime}\mathbf{Z}_{it} + \varepsilon_{it}^{k}.$$

All columns include institution fixed effects and an FOMC-interval indicator. The appendix construction details (p. 110) describe forecast matching, MPS scaling, and controls. Forecast dates are 2000–2025, excluding March 2020–March 2021 and June 15, 2022 (pre-announcement policy news). N counts forecaster–revision intervals. Standard errors cluster by survey interval in column (1), and by announcement date or ordinary-revision midpoint in columns (2)–(3). $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Construction Details for Announcement Tests

Forecast revisions (Table A13). Column (1) uses each forecaster's latest forecast as of the second US business day of each month, holding target quarters fixed. Columns (2)–(3) use the closest forecasts bracketing each announcement, at most ten days apart, and revisions outside announcement windows. GDP and CPI growth forecasts are annualized. For monthly revisions, I sum scheduled MPS between the survey cutoffs. Announcement revisions in columns (2)–(3) receive that meeting's MPS. Other intervals receive zero. I divide the shock by its sample standard deviation across survey intervals containing announcements in column (1), or meetings in columns (2)–(3), within each outcome's estimation sample. June 15, 2022 is excluded because pre-announcement news could affect revisions without entering the announcement-window shock. Column (3) follows Bauer and Swanson (2023) in controlling for prior forecast revisions and economic news: prior 30-day forecast changes for all three outcomes; Goldman Sachs MAP surprises over the revision window and the 21 days before the initial forecast; six-month core CPI acceleration; 13-week changes in the S&P 500, the 10-year minus 3-month Treasury spread, and the S&P GSCI spot index; a time trend; and log(window days).

Forecast accuracy (Table X). I exclude announcement-day reports and intervals containing unscheduled or multiple announcements. Year-quarter effects use the meeting date for FOMC intervals and the second publication date otherwise; weekday effects use the second publication date. Release indicators mark events strictly between publications. Controls include initial RMSE, interval sums of signed and absolute macro surprises, log interval length, forecast ages, market movements before the initial publication, and holiday controls at the second publication (Table A14). Unadjusted means do not require these controls.

Meeting types (Table A15). SEP and press-conference indicators use USMPD; announcement-time SEP releases begin April 27, 2011. The balance-sheet dates are March 18, 2009; August 10, September 21, and November 3, 2010; September 21, 2011; June 20, September 13, and December 12, 2012; June 19, September 18, and December 18, 2013; March 20, 2019; January 26 and March 16, 2022. These combine the five dates in Hanson and Stein (2015, footnote 7), subsequent QE and taper announcements, and scheduled QT announcements in Du et al. (2026). The Treasury sample excludes March 18, 2009.

Table A14: Forecast Updating around FOMC and Economic Releases

Dependent variable: Daily forecast revision probability (pp).

	Fed funds		GDP growth		CPI inflation	
	Next three quarters		*Next three quarters*		*Next three quarters*	
	Baseline	Controls	Baseline	Controls	Baseline	Controls
	(1)	(2)	(3)	(4)	(5)	(6)
Days 0 to 6 after a release						
FOMC announcement	0.74***	0.76***	0.23	-0.02	-0.14	-0.01
	(0.22)	(0.23)	(0.19)	(0.18)	(0.19)	(0.31)
Payroll release		0.72***		1.08**		2.52**
		(0.20)		(0.41)		(1.04)
CPI release		0.33*		0.18		3.43*
		(0.17)		(0.19)		(1.75)
Initial GDP release		-0.62***		1.06***		-0.66
		(0.20)		(0.37)		(0.66)
Institution FE	No	Yes	No	Yes	No	Yes
Year-quarter FE	No	Yes	No	Yes	No	Yes
Observations	112,046	112,046	105,251	105,251	41,524	41,524
R^2	0.00	0.03	0.00	0.04	0.00	0.11
Base daily revision rate (%)	1.56		2.17		4.24	
	(0.18)		(0.13)		(0.78)	

Notes: Each observation is an institution–calendar day. U_{it} is 100 for any numerical change in the next-three-quarter forecast path and zero otherwise, including days without a report. Sample years and target exclusions follow Table X, columns (1) and (3). CPI is quarterly annualized inflation. The regression specification is

$$U_{it} = \alpha_i + \delta_{q(t)} + \eta_{d(t)} + \boldsymbol{\beta}'\mathbf{R}_t + \boldsymbol{\gamma}'\mathbf{Z}_{it} + \varepsilon_{it}.$$

R marks days 0–6 after the displayed releases. Columns (1), (3), and (5) include only a constant and the FOMC indicator. The remaining columns add institution, year-quarter, and weekday effects, and controls **Z**: forecast ages and revision history; signed and absolute Goldman Sachs macro surprises; prior 7- and 30-day signed and absolute S&P 500 returns and two- and ten-year yield changes; and holidays within three days. Each pair uses the same observations. The bottom row reports the daily revision rate outside FOMC windows. Parentheses report standard errors clustered by institution and year-quarter; stars use t with degrees of freedom equal to the smaller cluster count minus one. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A15: Asset Price Reaction Magnitudes: MPS and Meeting-Type Controls

Dependent variable: Absolute announcement-window reaction (SD).

	Absolute reaction			Absolute residual reaction		
Text measure (1 SD)	(1)	(2)	(3)	(4)	(5)	(6)
A. S&P 500						
Fed research	0.10			−0.02		
	(0.07)			(0.07)		
Communication disagreement		0.26***			0.21***	
		(0.06)			(0.07)	
Communication surprise			0.31***			0.21***
			(0.08)			(0.08)
Meeting-type indicators	Yes	Yes	Yes	Yes	Yes	Yes
Meetings	207	207	207	207	207	207
R^2	0.12	0.18	0.19	0.10	0.14	0.13
B. 10-year Treasury yield						
Fed research	0.07			0.04		
	(0.08)			(0.09)		
Communication disagreement		0.16**			0.23***	
		(0.07)			(0.07)	
Communication surprise			0.21**			0.19*
			(0.08)			(0.10)
Meeting-type indicators	Yes	Yes	Yes	Yes	Yes	Yes
Meetings	206	206	206	206	206	206
R^2	0.10	0.12	0.12	0.12	0.17	0.15

Notes: The regression specification is

$$y_t = \alpha + \beta X_t + \boldsymbol{\delta}'\mathbf{Z}_t + \varepsilon_t, \qquad r_t = a + b\,\mathrm{MPS}_t + u_t.$$

Each column enters one text measure X_t from Table XI. $\mathbf{Z}_t$ contains indicators for an SEP released with the announcement, a press conference, and a selected balance-sheet announcement (appendix construction details, p. 110). Columns (1)–(3) use absolute asset price reactions, while columns (4)–(6) use absolute reactions after residualizing for the monetary policy shock (Section 6.2). Samples and event windows match Table XI.

$^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A16: Economic Content and Treasury Yield Responses: Robustness

A. Two- and five-year Treasuries

Dependent variable: Treasury yield change around FOMC announcements (bp).

Reported coefficient: MPS (bp) × economic share (1 SD).

Maturity	(1)	(2)	(3)	(4)	(5)
2-year	0.37***	0.35***	0.24**	0.21***	0.20**
	(0.06)	(0.06)	(0.09)	(0.08)	(0.09)
R^2	0.83	0.84	0.88	0.86	0.89
5-year	0.36***	0.35***	0.27***	0.17*	0.20*
	(0.08)	(0.07)	(0.10)	(0.09)	(0.11)
R^2	0.61	0.61	0.68	0.65	0.69
MPS × total error rate	No	Yes	Yes	Yes	Yes
Shock-size/current-rate controls	No	No	Yes	No	Yes
Meeting-type controls	No	No	No	Yes	Yes
Meetings	206	206	206	206	206

Notes: The regression specification is

$$\Delta y_t = \alpha + \beta \mathrm{MPS}_t + \gamma E_t + \delta(\mathrm{MPS}_t \times E_t) + \lambda S_t + \boldsymbol{\theta}'\mathbf{Z}_t + \varepsilon_t.$$

The table reports δ. Outcomes, E, and S follow Table XII. $\mathbf{Z}$ contains the indicated controls. Shock-size/current-rate controls are MPS|MPS|, MP1, and MP1×E; MP1 is the USMPD current-rate surprise. Meeting-type controls include the indicators in Table A15 and their interactions with MPS. The sample is fixed at 206 scheduled meetings, 2000–2025, excluding March 18, 2009. Parentheses report HC3 standard errors; stars use t tests with residual degrees of freedom: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A16: Economic Content and Treasury Yield Responses: Robustness (continued)

B. Ten- and thirty-year Treasuries

Dependent variable: Treasury yield change around FOMC announcements (bp).

Reported coefficient: MPS (bp) × economic share (1 SD).

Maturity	(1)	(2)	(3)	(4)	(5)
10-year	0.28***	0.27***	0.20**	0.18*	0.19*
	(0.06)	(0.07)	(0.10)	(0.10)	(0.11)
R^2	0.43	0.43	0.49	0.45	0.50
30-year	0.19***	0.19**	0.16	0.16*	0.19
	(0.06)	(0.07)	(0.10)	(0.10)	(0.12)
R^2	0.10	0.10	0.13	0.12	0.15
MPS × total error rate	No	Yes	Yes	Yes	Yes
Shock-size/current-rate controls	No	No	Yes	No	Yes
Meeting-type controls	No	No	No	Yes	Yes
Meetings	206	206	206	206	206

Notes: The regression specification is

$$\Delta y_t = \alpha + \beta \mathrm{MPS}_t + \gamma E_t + \delta(\mathrm{MPS}_t \times E_t) + \lambda S_t + \boldsymbol{\theta}'\mathbf{Z}_t + \varepsilon_t.$$

The table reports δ. Outcomes, E, and S follow Table XII. $\mathbf{Z}$ contains the indicated controls. Shock-size/current-rate controls are MPS|MPS|, MP1, and MP1×E; MP1 is the USMPD current-rate surprise. Meeting-type controls include the indicators in Table A15 and their interactions with MPS. The sample is fixed at 206 scheduled meetings, 2000–2025, excluding March 18, 2009. Parentheses report HC3 standard errors; stars use t tests with residual degrees of freedom: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.