

The Fed Attention Mechanism

Daniel Milo*

October 21, 2026

[Click here for the latest version](#)

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 shape information processing costs and determine forecasters' allocation of attention. Supporting this framework, I find that Fed communications shift attention toward the Fed and the topics it emphasizes, crowding out independent economic research. As forecasters shift attention toward the Fed, they rely more on its communications to form their own economic expectations. Attention to the Fed improves the accuracy of short-rate forecasts but worsens forecasts of economic outcomes and longer-term rates. I next show that asset prices react to the magnitude of Fed communication surprises, measured across forecasters, particularly surprises about the Fed's economic outlook. These findings suggest that Fed communications make the private sector better informed about near-term interest rates, but less informed about the economy and longer-term rates.

*New York University, Stern School of Business, dm4766@stern.nyu.edu. I thank Viral Acharya, Alexander Bartik, Yannis Cabossioras, Simon Gilchrist, Dan Greenwald, Arpit Gupta, Sebastian Hillenbrand, Toomas Laarits, Sydney Ludvigson, Jeffrey Meli, Alan Moreira, Holger Mueller, Laura Nicolae, Cecilia Parlatore, Thomas Philippon, Rob Richmond, Tony Saunders, Alexi Savov, Philipp Schnabl, Julia Selgrad, Bill Shu, Olivier Wang, Courtney Weigand, and Sean Wu for helpful comments and conversations. I gratefully acknowledge funding from the NYU Stern Center for Global Economy and Business.

1 Introduction

Does Fed communication make the private sector better informed? Over the past thirty years, the Federal Reserve has substantially increased the amount it communicates with markets. Initially, policymakers expected this information would help markets better anticipate Fed policy, reducing uncertainty about economic and financial outcomes [Bernanke cite]. However, the emergence of several asset pricing anomalies around Fed communications has suggested a more complex role [cites]. More recently, the new Fed Chairman Kevin Warsh has decided to communicate less with financial markets, arguing that communication shifts market attention towards the Fed and away from independent economic research [cite].

I begin by formalizing this debate with a conceptual framework where forecasters engage in costly processing of both Fed and other economic information. The model speaks to three main questions. First, how do Fed communications influence the attention of financial markets? Second, how does Fed attention impact the accuracy of economic and financial expectations? Third, does Fed attention help explain asset price reactions to FOMC announcements?

To speak to this framework I leverage the written reports of professional macroeconomic forecasters who detail the development of their quantitative macrofinancial expectations as well as the communications they expect from the Fed at each of their announcements. Specifically, I collect over 300,000 reports from professional forecasters spanning the past 26 years. I then process the resulting dataset with Large Language Models (LLMs) to extract over 500,000 quantitative macrofinancial forecasts, to construct a balanced panel of over 30,000 Fed communications expectations, and to quantify the degree of Fed attention for over 2.4 million paragraphs of text.

Through the lens of the model, I present three main findings. First, increases in Fed communications are associated with increased forecaster research about the Fed and the topics it emphasizes, leaving less independent economic research. Second, researching the Fed improves forecasts of short term rates, while worsening expectations of mid-long term rates, in part from worse economic expectations that are more similar to those of the Fed. Third, asset price reactions to Fed communications reflect this costly processing, with the magnitude and composition of Fed communication errors strongly predicting aggregate asset price reactions.

Conceptual Framework.

In the model, a representative forecaster predicts economic fundamentals and the Fed’s policy, which she believes to be related to economic fundamentals. To form her forecasts, the agent engages in costly research of both the Fed and the economy. The relative cost of Fed research is realized at the start of the model and reflects the amount that the Fed communicates. Because research about either object also informs her prediction of the other, she rationally places positive weight on both signals when forming each forecast. For example, research about economic fundamentals informs both the forecaster’s economic forecast and her forecast of the Fed.

A lower cost of Fed research induces the forecaster to optimally increase research on the Fed. This provides the forecaster with a more precise signal about the Fed, which improves the precision of both her Fed and economic forecasts, holding economic research fixed. However, with a more precise economic prior and less scope for economic precision to inform her Fed forecast, she may reduce her economic research. Specifically, the forecaster will weigh three forces when deciding whether to reduce economic research. First, the forecaster will find further economic research less informative for her Fed forecast. Second, the forecaster may prefer greater economic precision to complement her increased Fed precision. Finally, the forecaster may face an increasing marginal cost of economic research as she researches the Fed more. In the simple case with additive preferences and a fixed marginal cost, the forecaster reduces economic research to the point of accepting lower economic precision, as she substitutes cheaper Fed precision for economic precision.

Dataset

I introduce a novel dataset that directly speaks to the model. Specifically, I use the written reports of professional forecasters, which mirror the analytical effort markets engage in to process economic and Fed information, including the detailed expectations they form for Fed communications before FOMC announcements.

To process the unstructured corpus I utilize Large Language Models (LLMs) and develop novel measurement techniques. First, I quantify the share of research about the Fed for over 300,000 reports. The categorization separately identifies the subject of reasoning at the local level (over 2.4 million paragraphs of text) and the stated motivations for research at the article level. Combining these two measures allows for a more complete picture of the Fed’s influence on the private sector’s research agenda. Second, I extract the specific expectations forecasters form for each Fed announcement. Specifically, I identify all verifiable predictions forecasters make for each

announcement, aggregate them into a unified survey, and answer the survey for each forecaster using their reports. The resulting balanced panel measures the communications forecasters expected the Fed to give, which were realized, and crucially allows for cross-sectional identification of Fed communication surprises. I further leverage LLMs to pairwise compare expected communications versus actual communications. I fit a latent score that best explains these comparisons, representing a cross forecaster FOMC announcement shock.

Fed Communications Influence Forecasters’ Research Agenda

With a novel dataset that speaks to the conceptual framework, I begin the results section by asking whether Fed communications influence the research allocation of forecasters. I start by showing that forecasters devote a substantial share of their research to analyzing the Fed. On average, forecasters spend 25% of the text in their reports reasoning about the Fed. This share is highly cyclical, reaching about 40% right before FOMC announcements, and falling to near 20% thereafter. The magnitude of this peak of Fed research around announcements has increased over time, despite the Fed communicating more with financial markets.

I next ask whether the topics the Fed emphasizes in its communications influence the topics forecasters subsequently research. The typical challenge with such a test is common news shocks, such as geopolitical events, influencing both the Fed’s and forecasters’ topics simultaneously. To overcome this, I leverage the nature of the Fed’s minutes, which reflect the reasoning behind the Fed’s most recent decision with a three-week lag. I show that topics the Fed emphasizes in its minutes predict changes in topic emphasis of forecasters from before to after the minutes are released.

Fed Attention and Expectation Formation

Given that forecasters spend a substantial share of their research reasoning about the Fed, I next ask how Fed research influences the quantitative expectations of forecasters. The model intuitively predicts that increased research about the Fed will cause forecasters to realize lower Fed funds forecast errors, while economic forecast errors may increase depending on the forecaster’s preferences and research costs. The challenge with empirically testing the relationship between Fed research and forecast accuracy is that several time periods of heightened uncertainty such as the Great Financial Crisis or Pandemic periods may cause both elevated Fed research and large forecast errors. I overcome this challenge by leveraging the cross-sectional nature of my dataset, which allows me to compare forecasters with different levels of Fed

research that publish forecasts on the same day.

I find that forecasters who have recently researched the Fed more than their peers realize lower Fed funds forecasting error as well as higher GDP and inflation forecasting error. In the context of the model, this suggests that Fed research sufficiently crowds out economic research to lower economic precision, either due to increasing marginal costs of economic research or decreasing marginal benefits of economic research. I next ask how Fed research intensity predicts the accuracy of rate forecasts across the yield curve. I begin by showing that Fed research improves Fed funds forecasts over the next four quarters, with the advantage waning at five to six quarters into the future. Farther into the yield curve, forecasts of the two-year, five-year, and 10-year Treasury yields all worsen with increased Fed research, while forecasts of the 30-year are unaffected. This suggests that while Fed communications can reduce uncertainty of near term rates, they may increase uncertainty of mid to long term rates, by crowding out economic research.

Given the limited attention of forecasters, I next ask to what extent forecasters incorporate Fed information into their own economic expectations. Specifically, I examine signals that forecasters get from FOMC announcements by comparing their unique communication expectations to realized communications. In doing so, I inductively generate a questionnaire for each FOMC announcement, which yields a balanced panel of communication expectations. This allows me to observe for each announcement what specifically a forecaster expected the Fed to do and say, and whether that prediction was realized, or resulted in a hawkish or dovish surprise. To compare the overall shock that each forecaster realizes I leverage methodology developed by (shliefer et al list all names), by pairwise comparing the set of forecast realizations across forecasters. I then fit a latent score that best represents these pairwise comparisons, providing the first cross-sectional measure of FOMC announcement shocks.

I next compare these cross-sectional shocks to the quantitative forecast revisions that forecasters make around FOMC announcements. Leveraging the high frequency nature of my dataset, I compare the last forecast made before an FOMC announcement to the first forecast made following the announcement, with a maximum window of ten days. This allows for much tighter identification than existing papers that rely on monthly forecast revision windows. I find that forecasters more hawkishly surprised than their peers raise their Fed funds, GDP, and inflation forecasts relative to their

peers. This result suggests that forecasters compensate for their limited attention by learning about the economy through their signals from the Fed. Specifically, more surprised forecasters interpret a larger share of their signal from the Fed as economic information. Note that this still allows for the policy content of Fed announcements to outweigh the information content in aggregate. To show this, I replicate the results of Bauer and Swanson (2023) in my sample, which uses variation across meetings, and find a similar magnitude.

Consistent with these findings I show that around Fed announcements, forecasters' Fed funds forecasts become more accurate. Forecast accuracy changes for GDP depend on whether the Fed was initially more accurate. The forecast accuracy gap between the Fed and a forecaster strongly predicts whether a forecaster will become more accurate following the meeting. This again is consistent with forecasters' processing of Fed communications causing them to revise towards the Fed's view.

Asset Pricing Implications Having established that Fed communications influence forecasters' analytical research and quantitative macroeconomic forecasts, in the final part of the paper I explore the asset pricing implications of such a mechanism. I first show that forecasters' macroeconomic forecasts strongly explain their banks' asset price targets in the cross-section. This confirms that the Fed's influence on macroeconomic expectations extends to asset price expectations. I next interpret high-frequency asset price reactions to FOMC announcements through the communication surprises to forecasters. Specifically, I aggregate the communication surprises across forecasters creating a consensus measure of different communication dimensions for each announcement. I first test whether the quantity of information revealed predicts the magnitude of asset price reactions. I show that it strongly predicts these magnitudes; specifically, a one standard deviation move in the share of incorrect communication forecasts predicts a one standard deviation increase in S&P 500 and 10yr Treasury reaction magnitudes. This suggests markets react not only to the content, but also to the quantity of information released at Fed announcements. Next, I consider the degree of cross-sectional announcement surprises by measuring the agreement rate across forecasters in their pre-FOMC expectations. The greater the agreement among forecasters, the stronger the pass-through of short to long term rates.

Literature This paper contributes to three strands of literature. First, it contributes to a literature studying asset price reactions to FOMC announcements. This

includes papers documenting unexpected reactions for long-term yields (Hanson and Stein, 2015; Hillenbrand, 2025; ?) and equities (Lucca and Moench, 2015), as well as several papers applying structure to these asset price reactions to determine whether meaningful changes in expectations can be gleaned (Jarociński and Karadi, 2020; Cieslak and Khorrami, 2026; Bauer et al., 2026). This paper contributes to this literature by interpreting FOMC announcements through the surprise in forecasters' communication expectations.

The puzzles in the first strand of the literature have motivated researchers to explore forecaster surveys and Fed as well as news text to identify mechanisms of monetary policy transmission. Several papers have debated the existence of a "Fed Information Effect", where markets learn about the state of the economy from the Fed using surveys (Romer and Romer, 2000; Campbell et al., 2012; Nakamura and Steinsson, 2018; Bauer and Swanson, 2023a; Karnaukh and Vokata, 2022) and news sentiment (Acosta, 2026). Other papers have studied markets' beliefs about the Fed's reaction function (Bianchi et al., 2022; Bauer et al., 2024) and disagreement between the Fed and markets about the economy (Sastry, 2024; Caballero and Simsek, 2022). Another group of papers links investors' attention and information demand to announcement returns and asset price responses (Kroner, 2025; Fisher et al., 2022; Benamar et al., 2020). This paper nests these three mechanisms into a tractable model where forecasters believe the Fed view to be important but imperfect (hence a reaction function), and so rationally put weight (hence an information effect) on the Fed's view (hence an attention mechanism). This paper also introduces a novel dataset that links the text of forecasters' research with their quantitative forecasts, and provides a cross-section for identification.

Third, I contribute to an emerging literature that uses the text of market professionals to help explain their quantitative forecasts. Specifically, several papers have found the text of equity research reports helpful for explaining their quantitative forecasts and equity asset price reactions (Bastianello et al., 2025; Ke, 2024; Wang and Zhou, 2026; Asquith et al., 2005). Other studies examine text from asset managers (Andolfatto and Bastianello, 2025), and surveys of macroeconomists and households (Andre et al., 2022, 2023). I contribute to this literature by looking at professional forecasters where higher-order beliefs are important.

Organization of the paper. The paper is organized as follows: section 2 details a conceptual framework that guides the empirical tests. Section 3 introduces

the dataset explaining the institutional background and measurement techniques employed. Section 4 explains how Fed communications influence both the aggregate attention forecasters give to the Fed and the topics they emphasize. Section 5 explores the implications of Fed attention on forecast accuracy and explains the "Fed Information Effect" through limited attention. Section 6 relates the specific Fed communication surprises forecasters realize to aggregate asset price reactions across FOMC announcements. Finally, section 7 concludes with broader policy implications.

2 Conceptual Framework: The Attention Mechanism

2.1 The Forecasting Problem

A representative forecaster predicts X , the latent state of the economy, and F , the Fed's policy assessment. At date 3, F and X are both realized. Figure I presents the timeline. The forecaster perceives F to deviate from X by B , reflecting the Fed's economic measurement error and policy preferences:

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

At date 1, the forecaster engages in costly research that yields noisy signals s_F and s_X about F and X , respectively. The signals have the following properties:

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

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

$$\varepsilon_F \perp \varepsilon_X, \quad (\varepsilon_F, \varepsilon_X) \perp (X, B). \quad (4)$$

The signal errors are independent of X and B and of each other. The precisions of the signals are proportional to the amount of effort the forecaster places on researching that topic. By exerting effort a_X to collect information about the latent economic state, the forecaster improves the precision of her economic signal and reduces uncertainty about X . Similarly, by exerting effort a_F to interpret the Fed signal, the forecaster improves its precision and reduces uncertainty about F . Given this signal

structure, the posterior precisions are:

$$P_X \equiv \text{Prec}(X \mid s_X, s_F) = \underbrace{\frac{1}{\sigma_X^2}}_{\text{prior precision}} + \underbrace{a_X}_{\text{economic research}} + \underbrace{\frac{a_F}{1 + \sigma_B^2 a_F}}_{\text{Fed research}}, \quad (5)$$

$$P_F \equiv \text{Prec}(F \mid s_X, s_F) = \underbrace{\frac{1}{\sigma_X^2 + \sigma_B^2}}_{\text{prior precision}} + \underbrace{a_F}_{\text{Fed research}} + \underbrace{\left[\frac{\frac{1}{\sigma_X^2} + a_X}{1 + \sigma_B^2 \left(\frac{1}{\sigma_X^2} + a_X \right)} - \frac{1}{\sigma_X^2 + \sigma_B^2} \right]}_{\text{economic research}}. \quad (6)$$

Because $F = X + B$, each signal also contains information about the other forecasting object. Define ℓ_F as the marginal economic precision supplied by Fed research and ℓ_X as the marginal Fed precision supplied by economic research. The marginal effects of research on posterior precision at an interior allocation are

$$\frac{\partial P_X}{\partial a_X} = 1 > \frac{1}{(1 + \sigma_B^2 a_F)^2} = \frac{\partial P_X}{\partial a_F} \equiv \ell_F, \quad (7)$$

$$\frac{\partial P_F}{\partial a_F} = 1 > \frac{1}{\left[1 + \sigma_B^2 \left(\frac{1}{\sigma_X^2} + a_X \right) \right]^2} = \frac{\partial P_F}{\partial a_X} \equiv \ell_X. \quad (8)$$

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

$$\frac{\partial^2 P_X}{\partial a_F^2} = -\frac{2\sigma_B^2}{(1 + \sigma_B^2 a_F)^3} < 0, \quad (9)$$

$$\frac{\partial^2 P_F}{\partial a_X^2} = -\frac{2\sigma_B^2}{\left[1 + \sigma_B^2 \left(\frac{1}{\sigma_X^2} + a_X \right) \right]^3} < 0. \quad (10)$$

At an interior allocation, each unit of research produces more posterior precision when devoted to the forecasting object itself than when devoted to the other object. The relative return to cross-learning is governed by two forces. First, it decreases with the perceived variance of the Fed-specific policy component, σ_B^2 . When the Fed's policy assessment is perceived to track the latent economic state more closely, the Fed signal is more informative about X , and the economic signal is more informative about F . Second, cross-learning exhibits diminishing returns: as the forecaster devotes more

research to a cross-signal, each additional unit of that research produces a smaller increase in precision about the other forecasting object.

The forecaster values precision in both economic and policy forecasts. I use a constant elasticity of substitution (CES) specification:

$$U(P_X, P_F) = - (P_X^{-\sigma} + P_F^{-\sigma})^{1/\sigma}, \quad \sigma \geq 1. \quad (11)$$

For $\sigma > 1$, precision is complementary across forecasts: greater precision in one forecast increases the marginal value of precision in the other. At $\sigma = 1$, forecast loss is additive across the two forecasting objects. Related complementarities arise in the trading and information-acquisition setting of [Goldstein and Yang \(2015\)](#). The marginal benefits of precision are

$$m_X \equiv \frac{\partial U}{\partial P_X} = \frac{P_X^{-\sigma-1}}{(P_X^{-\sigma} + P_F^{-\sigma})^{(\sigma-1)/\sigma}}, \quad (12)$$

$$m_F \equiv \frac{\partial U}{\partial P_F} = \frac{P_F^{-\sigma-1}}{(P_X^{-\sigma} + P_F^{-\sigma})^{(\sigma-1)/\sigma}}. \quad (13)$$

Research is costly, with costs increasing in research efforts a_X and a_F . These costs of research are represented as:

$$C(a_X, a_F) = C_X a_X + C_F a_F + \kappa a_X a_F, \quad \kappa \geq 0. \quad (14)$$

The cross-term $\kappa a_X a_F$ allows research on one object to raise the marginal cost of researching the other. This can reflect task-switching costs ([Monsell, 2003](#); [Dessaint et al., 2024](#)) or fatigue ([Hirshleifer et al., 2018](#)).

At date 2, the forecaster reports their optimal predictions as the posterior means,

$$f_X = \mathbb{E}[X \mid s_X, s_F], \quad f_F = \mathbb{E}[F \mid s_X, s_F]. \quad (15)$$

Let $V_X = P_X^{-1}$ and $V_F = P_F^{-1}$ denote the posterior forecast variances. The forecaster chooses research at date 1 to minimize forecast loss and research costs:

$$\min_{a_X, a_F \geq 0} (V_X^\sigma + V_F^\sigma)^{1/\sigma} + C_X a_X + C_F a_F + \kappa a_X a_F. \quad (16)$$

2.2 Fed Communication

Fed communications provide information about the Fed's policy assessment. An increase in communications can therefore be represented by a reduction in the cost of Fed research, C_F . Under the interior-solution conditions in Appendix C, lower C_F strictly increases both the amount of Fed research and the precision of the Fed forecast. However, the response of economic research and precision depends on the forecaster's preferences and research costs. The first order conditions for economic research clarify this decision:

$$\underbrace{C_X + \kappa a_F}_{\text{marginal cost}} = \underbrace{m_X + \ell_X m_F}_{\text{marginal benefit}}. \quad (17)$$

An additional unit of economic research supplies one unit of economic precision and ℓ_X units of Fed precision. Its marginal utility therefore has two components: m_X , the marginal utility of its direct economic precision, and $\ell_X m_F$, the marginal utility of its Fed precision.

Three forces in equation (17) determine the response to greater Fed research. First, the precision gained about the Fed decreases, ℓ_X . Second, when $\sigma > 1$, improved Fed precision raises the marginal utility of economic precision, m_X . Third, when $\kappa > 0$, the marginal cost of economic research also rises. Thus, the response of economic research and precision depend on whether the forecaster's preferences for economic precision increase by enough to offset the increase in cost of economic research and decrease in benefit for Fed precision.

Additive Preferences: Consider first the special case of $\sigma = 1$ and $\kappa = 0$, where the forecaster has additive preferences over precision. The first-order condition can be rewritten as:

$$C_X = \frac{1 + \ell_F}{P_X^2} \implies P_X^2 = \frac{1 + \ell_F}{C_X}. \quad (18)$$

Greater Fed research reduces cross-learning, ℓ_F , which lowers economic precision P_X . In this case, the forecaster substitutes economic research with the cheaper Fed research such that economic precision falls.

Complementary precisions: With complementary preferences ($\sigma > 1$) and no cross-cost ($\kappa = 0$), cheaper Fed research can either improve or worsen economic

precision, depending on its initial cost. Improving the Fed forecast raises the value of an accurate economic forecast. This offsets the loss of economic research's usefulness for predicting the Fed. Economic precision improves following a cost reduction when:

$$\frac{dP_X^*}{dC_F} < 0 \iff \sigma - 1 > \frac{2[1 + (P_X/P_F)^\sigma]}{\sigma_B^2 \sqrt{P_X P_F (4 + \sigma_B^4 P_X P_F)}}. \quad (19)$$

The threshold degree of complementarity decreases as C_F falls. Along the optimal interior allocation, cheaper Fed research lowers P_X/P_F and raises $P_X P_F$, reducing the right-hand side of the condition. Thus, even with unchanged preferences, cost reductions can lower economic precision at higher Fed research costs and improve it at lower costs. Along an interior response, the direction can change at most once: once cheaper Fed research improves economic precision, further cost reductions continue to improve it. Appendix C establishes these results.

Cross-costs: Finally, allow $\kappa > 0$. Greater Fed research now additionally raises the marginal cost of economic research. The marginal utility of economic precision can be written as

$$m_X = \frac{1}{P_X^2} \left[1 + \left(\frac{P_X}{P_F} \right)^\sigma \right]^{-(\sigma-1)/\sigma} \longrightarrow \frac{1}{P_X^2} \quad \text{as } P_F \longrightarrow \infty. \quad (20)$$

As the precision of the Fed and economic forecasts increases, the marginal benefit of economic precision decreases. The marginal cost of economic research continues to rise at a rate of κ . Thus, at a low enough C_F economic precision will fall as C_F decreases. Figure A9 illustrates a parameter configuration with all three economic-precision response regions.

Economic research. The effect on economic research can be expressed directly in terms of the Fed research cost:

$$\frac{da_X^*}{dC_F} = \frac{dP_X^*}{dC_F} - \ell_F \frac{da_F^*}{dC_F}, \quad \ell_F > 0, \quad \frac{da_F^*}{dC_F} < 0. \quad (21)$$

If cheaper Fed research reduces economic precision, then $dP_X^*/dC_F > 0$ and hence $da_X^*/dC_F > 0$: independent economic research must fall. Economic precision can also improve while economic research falls, when $\ell_F da_F^*/dC_F < dP_X^*/dC_F < 0$. In this

case, the economic information gained from Fed research exceeds the information lost from independent research.

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. Finally, I conclude by summarizing the other complementary data used in the paper.

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, institutions managing \$9.7 trillion in U.S. fixed-income assets ranked professional forecasters in *Institutional Investor*'s annual survey (Johnson, 2016)¹. 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)².

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

¹The 2016 *Institutional Investor* All-America Fixed-Income Research Team survey included 535 institutions managing \$9.7 trillion in U.S. fixed-income assets.

²RBC 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.

happening in the market, which can be difficult for the buy side to ascertain without the sell side’s help.”³ 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.⁴

Recognizing the importance of professional forecasters, the Fed, industry, and academic researchers have utilized 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 investor expectations” (Correia Golay et al., 2013). Moreover, Diercks and Jendoubi (2023) find that expectations of institutional investors and professional forecasters to the New York Fed’s survey do not exhibit any systematic, persistent differences. Other longstanding 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, 2023a). 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. The most detailed analysis thus far is Hazell and Hobler (2025), which manually extracts fiscal forecasts for one specific event, the 2021 Georgia Senate runoff. In this paper, I collect over 300,000 reports from professional forecasters spanning 2000–2025.

I collect reports through proprietary client portals as well as two aggregator platforms, Refinitiv Eikon LSEG Workspace and Investext from Mergent Online. As Table (cite) demonstrates, 26 forecasting institutions are included with 10 retrieved

³Alexandra DeLuca, “What Happens When No One Has Experience Outside a ‘Big Fixed Income Bull Market’,” *Institutional Investor*, November 30, 2023.

⁴See 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.

directly from forecasters' research portals, 15 retrieved from LSEG, and one from Investext⁵.

This set of forecasters in my sample well reflects the set of investment banks typically in surveys of professional forecasters. Table (reference) compares the forecasters in the sample to those included in seven different surveys of professional forecasters finding on average a 31% overlap in forecasters. This reflects varied participation rates (on average forecasters in my sample contribute to four out of the seven surveys), a high overlap for sell-side banks (61%) and primary dealers (71%), and a low overlap for "other" forecasters (5%), which include 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 III shows variation in report frequency and length across time and forecasters. 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. 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.

3.3 Measurement

Forecaster reports consist of unstructured text, tables, and figures. Embedded in this data are macro-financial expectations and their rationale. Large Language Models (LLMs) provide a tractable cost effective 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 macro-financial forecasts, and (iii) FOMC predictions. Appendix sections detail each measurement process in more detail.

3.3.1 Measuring Fed Research

I begin measurement with a key input into the model, the forecaster's research allocation. Two key questions pin down the focus of research. First, what is the forecaster reasoning about in a given sentence or paragraph? Second, how did the forecaster

⁵The coverage of LSEG and Investext largely overlaps, but Investext was discontinued in August 2025 so coverage from LSEG is preferred.

explicitly motivate their research report? Concretely, I use an LLM to identify the share of reasoning about the Fed, about the economy, and that is factual reporting for each paragraph of text. The LLM then additionally identifies the motivation share that is attributed to the Fed at the article level. For both measurement tasks the LLM scores the shares of reasoning and motivation on a scale from 0 to 100, to reduce the impact of classification errors.

This measurement assumes that content and timing of reports accurately reflect the analytical process forecasters engage in. Several papers have argued these assumptions to hold for equity research analysts, who largely work in the same firms under similar regulations (bastianello, ke, hirshliefer et al). Moreover, the half life of client interest in macroeconomic research is short, incentivizing forecasters to publish their research in a timely manner.

Appendix [B.1](#) describes the measurement process and provides worked examples.

3.3.2 Quantitative Macro-Financial Forecasts

Forecaster reports detail their quantitative macrofinancial expectations across several common variables such as GDP, the Fed funds rate, and Treasury yields across tenors. While commonly used surveys provide these expectations of forecasters at a monthly or quarterly frequency, the reports allow pinpointing the precise timing of when forecasters form their expectations.

I develop a measurement pipeline to extract quantitative macro-financial forecasts from forecaster reports. First, I develop a multi-modal retrieval pipeline to allow extraction of forecasts from tables and figures, which I preserve as images, as well as text. Second, I create a detailed instruction guide that includes how to interpret the various measurement conventions for each variable, i.e. how GDP can be expressed at a quarterly annualized rate or a year-over-year rate. Finally, I include numerous validation checks throughout the pipeline such as extracting historical data from reports and comparing it to vintage data and double checking for missing forecasting horizons or variables that are typically in a forecaster’s report. Appendix section X details the specific process.

Table [III](#) summarizes the extracted forecasts. Typically forecasters publish a quantitative forecast for each variable weekly with a forecasting horizon of four quarters. I validate the extracted forecasts against forecaster-level submissions from the Reuters Economic Polls. Table [A1](#) reports the results: correlations between extracted and

poll-submitted forecasts range from 0.975 to 0.998 for GDP growth, headline CPI inflation, and the fed funds rate, based on 85,368 matched observations.

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.

As shown in Figure X, forecasters spend a considerable share of their research in the week leading up to FOMC announcements writing about the Fed. In these reports forecasters detail the specific communications they expect from the Fed. This goes beyond the direct policy decision of the meeting, including expectations for the Summary of Economic Projections, the statement, the press conference, and other policy choices, such as balance sheet policy and discount window rates. Appendix section X details an example of the types of communications that forecasters typically expect. For example, before the June 2024 FOMC announcement, Citibank detailed

- Examples of expectations.
- Read through the reports block by block.
- Extract out any expectations.
- Structure expectations by their expected portion of the announcement, their qualifier word/phrase, and their expectation (additional structure for SEP/policy decision variables).
- Then, I consolidate similar expectations to construct a survey of binary/multiple choice questions. I then query that survey on the text of forecaster expectations.
- Finally, I validate each question on the survey: was it correct or not? Was the realization more hawkish, more dovish, or neutral relative to the answer choice?
- Provides consensus measures of other objects which we don't have. Primarily we have focused on the direct policy decision, where we have fed funds futures, but the Fed rarely surprises there.
- We've extended this to include expectations of rates over the next year, but the interpretation of movements in this object is ambiguous; they may reflect information, policy, or perhaps even risk premia.

- The primary dealers survey pre-selects questions and asks them well before the announcement, while much of the research effort happens after that, referencing figure.
- Emerging data analysis of surveys of forecasters' SEP expectations, but only starts in year X and omits expectations of statements/press conferences, etc.
- The qualifiers predict the probability of a forecast proving correct.
- Forecast errors concentrated outside the direct policy decision.
- NY Fed primary dealers survey comparison.

3.4 Supplementary Data

First, 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. From the USMPD I take the high-frequency monetary policy surprise measure of [Bauer and Swanson \(2023b\)](#), hereafter MPS_t : the first principal component of 30-minute changes in Eurodollar futures rates around FOMC announcements, capturing changes in policy-rate expectations including forward guidance over a horizon of about one year. The surprise is normalized to have a unit effect on the four-quarter-ahead Eurodollar futures rate, measured in percentage points, and signed positive for hawkish surprises. Depending on the analysis, MPS_t is computed over either the statement window alone or the full announcement window (statement plus press conference); see Section 7 for the rationale. Second, FOMC announcement text (statements, SEP projections, dot plots, and press conference transcripts) comes from the Federal Reserve's public archive.

I use the Goldman Sachs U.S. Macro-Data Assessment Platform (MAP) as a daily measure of macroeconomic data surprises. MAP combines the size of a release's surprise relative to the median Bloomberg forecast with a measure of its relevance based on historical U.S. fixed-income market reactions. I use this series in part because it is constructed by one of the forecasters in my sample and therefore reflects the indicators that forecaster considers important.

Forecast Accuracy: For forecast-accuracy comparisons, I match each forecast to the realization for its target quarter. GDP growth is measured at the first available

release using real-time real GDP vintages from ALFRED, the Federal Reserve Bank of St. Louis’s vintage database. I compute annualized quarter-on-quarter growth from the target-quarter level and the preceding quarter’s level available in that same vintage, so subsequent revisions do not enter the accuracy benchmark. Headline CPI inflation is likewise measured as annualized growth in the quarterly average price index, using the first vintage containing all three months of the target quarter. Unemployment is the quarterly average available in the corresponding first complete vintage. Fed funds realizations are the target rate at quarter-end, or the midpoint of the target range when applicable. Treasury yield realizations are the last available daily constant-maturity yields in the target quarter, obtained from FRED. These definitions match the growth rates, quarterly averages, and quarter-end levels used to standardize the forecasts.

4 Fed Communications and Market Attention

In this section I show the influence of Fed communications on forecaster attention. I first show that forecasters spend a substantial amount of their written research writing about the Fed, on average 25% of reasoning. 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, and has increased over time. 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 all of their reasoning text on the Fed.

- How changed over time.
- Variation across forecasters.

I next show that the share of research about the Fed is highly influenced by Fed communications. Figure IV plots the Fed research share around FOMC announcements documenting a near doubling of Fed research intensity from 20% to 40% in the build up 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 the announcement, rather than summarization or post processing.

As shown in Panel B, the share of research about the Fed has increased over time, particularly in the build up to scheduled announcements. During this time period the Fed has substantially increased the quantity of text that it releases at Fed announcements (Figure X) rising from a couple hundred words per meeting to over 10,000 today.

- Could discuss other interpretations of the cyclical nature/change over time including increased client focus on hedge funds.

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. The challenge with empirically testing this relationship is that the Fed and markets often discuss similar topics due to common information shocks, such as geopolitical events or macroeconomic data releases. I address this concern by focusing on the minutes that the Fed releases which reflect the reasoning behind their most recent policy decision with a three-week delay.

Specifically, I test whether topic emphasis in the minutes predicts changes in topic emphasis of forecasters from before to after the minutes. To define topics, I use the set of 40 topics defined by (cite), who trained their topics on FOMC transcripts. I also show robustness to other topic measures including embeddings. I include controls for..... Table X shows the results, specifically, I find a 1pp shock to a minutes topic share predicts a ybp shift in topic emphasis across forecasters. (need to discuss that pool forecaster text together in aggregate not cross section).

Paragraph contextualizing the results.

- Fed topic influence -> fed economic research influence.

5 Fed Attention and Forecast Accuracy

In this section I explore the role of Fed attention in forecasters' quantitative macro-financial 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. I then show that forecasters supplement their limited economic precision by incorporating Fed information into their own economic expectations. Finally, I show that consensus forecaster expectations improve for the Fed funds rate and near-term Treasuries around Fed announcements, but worsen for economic expectations when the Fed had a less accurate expectation than forecasters.

5.1 Fed Attention and Economic Forecast Accuracy

I begin by asking whether attention to the Fed predicts forecast accuracy across expectations of the Fed funds rate and economic outcomes. 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 independent economic attention. This may happen due to increasing costs of economic research or highly substitutable preferences between Fed and economic precision.

There are several empirical challenges with testing the relationship between Fed attention and forecast accuracy. First, periods of heightened uncertainty, like the Great Financial Crisis or the Pandemic, may cause both increased attention on the Fed and larger realized forecasting errors, confounding variation in the time series. To address this I turn to the cross section, focusing on variation across forecasters at the same point in time. Second, forecasters may persistently differ in their characteristics, i.e. their preferences over Fed and economic precision or their skillset, which may reflect their marginal cost. To account for this, I use forecaster fixed effects, demeaning attention and forecast accuracy by forecaster means.

My primary specification regresses realized forecast error on the share of research about the Fed, controlling for total research volume as well as date and forecaster fixed effects. I measure realized forecast error as the root mean squared error over the next three quarters which averages forecasting performance for volatile quarterly variables like GDP over a set of common forecasting horizons (see section X for more details on how forecast error is measured). I next measure Fed research share as an average over the last 30 calendar days to allow for persistent effects of research on forecasting accuracy. I also show robustness to other windows in appendix table X.

Table X presents the results. Column 1 begins without controls showing a positive

relationship between Fed attention and realized forecast error for Fed funds and GDP forecasts, perhaps reflecting time-series variation in aggregate uncertainty. Controlling for date fixed effects flips the relationship for Fed funds to negative and further adding forecaster fixed effects produces a positive result for Real GDP and Headline CPI. Specifically, I find that increased Fed attention leads to lower realized error for Fed funds forecasts, and higher realized error for GDP and CPI forecasts.

Contextualizing paragraph

- This establishes limited attention in that attention allocation predicts forecast errors.
- Interpret magnitudes.
- What might drive the variation? Forecasters may differ over time in the preferences of their clients or in their abilities, i.e., they might be better at understanding certain Fed speeches more than others.

5.2 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. Typically, professional forecasters relay expectations of the economy and the Fed to a separate Treasury rates team in charge of forecasting Treasury yields across the curve. I collect these expectations for the 2yr, 5yr, 10yr, and 30yr rates and next test the relationship between the Fed attention of professional forecasters and their treasury rates team's expectations. I then continue the specification from subsection X controlling for date and forecaster fixed effects, but focus instead on forecasting error on a quarter-by-quarter basis, not an average across quarters, given the level nature of rates forecasts.

Table X presents the results on a balanced sample of forecaster x date pairs reporting each of the next four forecasting quarters, while Appendix Table Y presents the results for the next six quarters allowing the sample composition to vary. I find that attention towards the Fed predicts lower realized forecasting errors for Fed funds forecasts over the next four quarters, but these effects appear to wane at five- to six-quarter forecasting horizons. On the other hand, expectations of two-year, five-year, and ten-year Treasury yields all worsen with attention to the Fed, especially as the

forecasting horizon increases. The results are particularly concentrated in the "belly" of the curve, i.e. the five-year Treasury yield.

- What could explain this?
- Worse economic expectations?
- More focus on short term Fed trades (could test this).
- Treasury decompositions into term premium and other factors.

6 The Information Effect

The results so far provide evidence of limited attention among forecasters: their attention allocation predicts their forecast errors. I next ask whether forecasters supplement their limited economic precision by incorporating Fed information into their own economic expectations. 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 for Fed information have relied on time-series variation of consensus revisions of economic and interest rate expectations, which effectively have tested whether the information revealed by the Fed is strong enough to counteract the intended policy effect of monetary policy ([Nakamura and Steinsson, 2018](#); [Bauer and Swanson, 2023a](#)). 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 receive from the Fed and the economy.

6.1 Identifying Forecaster Level Fed Information

The challenge is that not all attributes of a meeting are created equal, and so I need some degree of judgment. Specifically, I'm going to lean on novel measurement techniques developed in the large language model literature, Shleifer et al., for example.

And I’m going to use pairwise comparisons. I’m going to compare each set of forecasters within a given meeting with the question of which forecaster was more hawkishly surprised. In this case, I find that Citibank was more hawkishly surprised when I weigh all of their expectations. And after completing all of the pairwise comparisons for this meeting, I can fit a latent score that best explains these pairwise comparisons. And I see across the forecasters for this meeting, Citibank realized a relatively hawkish surprise, and Wells Fargo realized a relatively dovish surprise.

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. To account for this, I turn to scheduled Fed announcements, for which forecasters differ in their specific expectations and thus in the signals they ultimately receive. Specifically, I develop a cross-forecaster communication surprise score that measures how hawkishly surprised each forecaster is relative to its peers at each meeting (see Appendix B.3 for construction details). 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. This allows me to compare forecasters receiving different signals from the same Fed announcement and examine whether they incorporate them primarily as policy, inducing a negative correlation between economic and rate forecast revisions, or information, inducing a positive correlation.

6.2 Testing for Fed Information in the Cross Section

Specifically, 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, limiting the amount of confounding information.⁶ Meeting fixed effects absorb common information, such as a fiscal bill or payrolls report near the 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

⁶I impose a maximum interval of 10 calendar days between the pre-announcement and post-announcement forecasts.

expectations. The preferred specification is

$$\Delta f_{it}^k = \alpha_t^k + \lambda_{e(i,t)}^k + \beta^k \tilde{S}_{it}^k + \gamma^{k'} \mathbf{Z}_{it}^k + \varepsilon_{it}^k, \quad (22)$$

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; $\tilde{S}_{it}^k$ is the communication surprise score scaled by its standard deviation in the estimation sample for outcome k . The terms α_t^k and $\lambda_{e(i,t)}^k$ are meeting and lead-economist 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.

Table VI 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 the surprise being interpreted primarily as economic information rather than policy. In my preferred specification, with meeting fixed effects, lead-economist fixed effects, and the controls described above, a one-standard-deviation increase in the surprise score is associated with a 1.0 basis point increase in Fed funds forecasts, a 1.2 basis point increase in GDP growth forecasts, and a 1.1 basis point increase in inflation forecasts.

I next benchmark these results against the canonical approach relating market-based monetary policy surprises to average forecast revisions. Appendix Table A4 presents the results. I first show that the positive GDP coefficient found at a monthly frequency, as in Nakamura and Steinsson (2018), does not persist in narrower revision windows around announcements. Adding controls adapted from Bauer and Swanson (2023a) for recent macroeconomic news and its potentially delayed incorporation yields a negative relationship between monetary policy surprises and GDP revisions. Specifically, 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, and a 0.9 basis point decline in inflation forecasts.

- At an individual level, monetary policy news is costly to process and influences the economic expectations of individuals.
- In aggregate, policy effects dominate, but information effects still persist in the cross-section.

7 Aggregate Accuracy and Asset Price Reactions

The results so far establish limited attention at the forecaster level and Fed influence on economic expectations through this limited attention. I next ask whether consensus forecasts reflect this dynamic, or if cross-sectional differences wash out. Then, I explore whether aggregate asset price reactions to Fed communications similarly reflect the dynamic.

7.1 Consensus Accuracy around Announcements

Having established in the cross section that Fed attention predicts forecast accuracy and that forecasters incorporate Fed information into their own economic expectations, I next test for aggregate effects of Fed communications on consensus expectations. I begin by exploring whether the accuracy of consensus expectations improves or deteriorates on average around FOMC announcements. In Figure X Panel A I show that on average the accuracy of Fed funds, 2yr and 5yr Treasury forecasts improves around FOMC announcements. This would suggest that economic forecasts should improve as well, incorporating the more accurate path of policy, and potentially incorporating additional economic information from the Fed. However, I find that on average the accuracy of GDP and inflation forecasts remains flat around announcements. This suggests that the share of meetings where forecasters worsen and improve their economic forecasts is fairly balanced.

I next test for heterogeneity across meetings. I begin by testing whether Fed information may determine whether forecasters improve or worsen their economic expectations around announcements. Specifically, I compare the Fed staff's forecasts from the Greenbook and Tealbook, which are finalized before the FOMC announcement, but only released to the public with a five-year delay. Still, the communications and decisions of the Fed often reflect these forecasts (cite). Concretely, I test whether the Fed's accuracy advantage pre-announcement predicts whether forecasters become more or less accurate around the meeting. In Panel B I plot the relationship. I find that when the Fed has a more accurate economic forecast, it pulls consensus towards that more accurate view, whereas when its expectations are less accurate than consensus it pulls consensus towards its less accurate view. This suggests that Fed information is strong enough to shift consensus expectations of forecasters.

7.2 Aggregate Forecaster Communication Surprises

I next test whether consensus surprises to expected Fed communications help explain the magnitude of asset price reactions. I begin by documenting the relationship between Fed attention, communication uncertainty, 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 define communication uncertainty as the disagreement rate across forecasters for each expected attribute of a Fed communication, averaged over attributes. Finally, I define communication surprise as the share of expected communications not realized, weighted by the number of forecasters expecting each communication. Together these metrics represent the attention, uncertainty, and surprise that forecasters realize on average for each announcement.

- Figure X shows a positive relationship between research share and ...
- Contextualize with literature.
- Communication surprises concentrated outside of the direct policy decision.
- Over time appendix figures.

Table X explores the relationship between the magnitude of high frequency asset price reactions to scheduled Fed announcements and Fed research share, communication disagreement, and communication surprise. For asset price reactions I focus on the magnitude of S&P 500 and 10yr Treasury reactions over the monetary event window defined by (cite Acosta et al)⁷. I begin by showing a positive relationship between Fed research share and the magnitude of risky asset price reactions, consistent with Kroner and FMS. I next show that communication disagreement explains a large share of this relationship, while monetary policy uncertainty.... Finally, communication surprise, the share of expected communications not realized, best explains the magnitude of asset price reactions. This suggests that the marginal communication surprise affects the magnitude of risky asset price reactions.

⁷Specifically, 10 minutes before the statement to 60 minutes after the press conference starts, or 20 minutes after the statement is released for meetings without a press conference.

7.3 Types of Fed Communication Surprises

Now, having validated that the communication surprises of forecasters reflect those of broader financial markets, I next ask: can I interpret different types of monetary policy surprises by the types of communication surprises that are generated? And so particularly, I'm going to divide communication surprises into those that are about the Fed's economic assessment, so their assessment of growth, inflation, labor markets, and those about policy communication: their policy decisions, their forward guidance, and so on.

About half of all realized communication surprises are about the Fed's economic assessment. What I find is that if I compare the market-based monetary policy shock, so the change in the one-year rate, roughly speaking, to changes across the yield curve, there is a positive relationship that is larger specifically for communication surprises that are driven by economic content, suggesting that when markets observe new information about the Fed's economic view, they incorporate that to a greater degree into their longer-term yield expectations.

Table IX relates Treasury yield responses to monetary policy surprises and the economic content of communication surprises.

7.4 Contextualizing Asset Pricing Results

Why might this exist in equilibrium? Contextualizing in literature about investor focus on public information and short-term investments. Fed directly controls short term rates and reveals public signals. This in equilibrium disincentives economic research that goes against the Fed.

8 Conclusion

This paper introduces a novel dataset of the written reports from professional forecasters as well as techniques to process the corpus. The dataset effectively opens the black box into how markets formed macroeconomic beliefs and asset price expectations over the past 26 years. Future research can further leverage this approach to explore the development of other expectations, such as for fiscal policy or foreign exchange rates.

In this paper I leverage the new dataset to argue that the Fed shapes economic expectations by shifting the attention allocation of market participants. Specifically, I provide evidence that Fed communications are costly to process and that increases in Fed communication induce forecasters to increase research on the Fed. Further, in the cross-section, forecasters who research the Fed more realize lower Fed funds forecasting error at the cost of higher economic and Treasury yield forecasting errors, suggesting an attention crowdout. I also provide evidence that market participants extract economic information from Fed communications through their limited attention. Finally, I demonstrate that aggregate asset prices reflect this dynamic, with asset prices reacting to the magnitude of communication surprises and the composition, particularly the share of communication surprises about the economy.

The increase in Fed transparency over the past 30 years may have come at a cost. Specifically, it may have left market participants better informed about the near term path of policy at the cost of leaving them less well informed about the economy and medium-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.
- 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,” *Working Paper*, 2023.
- 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.
- and —, “A Reassessment of Monetary Policy Surprises and High-Frequency Identification,” *NBER Macroeconomics Annual*, 2023, 37, 87–155.
- 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.
- 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*, 2020, 34 (7), 3403–3455.
- Bernanke, Ben and Michael Woodford**, “Inflation Forecasts and Monetary Policy,” *Working Paper*, 1997.
- Bianchi, Francesco, Martin Lettau, and Sydney C. Ludvigson**, “Monetary Policy and Asset Valuation,” *The Journal of Finance*, 2022, 77 (2), 967–1017.
- 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, *Spring 2012*.
- Cieslak, Anna and Paymon Khorrami**, “Monetary Shocks and Long-Run Neutrality in Asset Pricing,” *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.
- 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.
- Fisher, Adlai, Charles Martineau, and Jinfei Sheng**, “Macroeconomic Attention and Announcement Risk Premia,” *The Review of Financial Studies*, 2022, *35* (11), 5057–5093.
- Fuerbringer, Jonathan**, “Right or Wrong, Wall Street Economists Generate Ideas and Trades,” 1999.
- Golay, Ellen Correia, Steven Friedman, and Michael McMorro**, “Understanding the New York Fed’s Survey of Primary Dealers,” *Current Issues in Economics and Finance*, 2013, *19* (6), 1–8.
- Goldstein, Itay and Liyan Yang**, “Information Diversity and Complementarities in Trading and Information Acquisition,” *The Journal of Finance*, 2015, *70* (4), 1723–1765.
- 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.
- Hirshleifer, David, Yaron Levi, Ben Lourie, and Siew Hong Teoh**, “Decision Fatigue and Heuristic Analyst Forecasts,” *Working Paper*, 2018.
- Jarociński, Marek and Peter Karadi**, “Deconstructing Monetary Policy Surprises—The Role of Information Shocks,” *American Economic Journal: Macroeconomics*, 2020, *12* (2), 1–43.
- Johnson, Thomas W.**, “The 2016 All-America Fixed-Income Research Team: Key Figures,” 2016.
- Karnaikh, Nina and Petra Vokata**, “Growth forecasts and news about monetary policy,” *Journal of Financial Economics*, 2022, *146* (1), 55–70.
- Ke, Shikun**, “Analysts’ Belief Formation in Their Own Words,” *Working Paper*, 2024.
- Kroner, Niklas**, “How Markets Process Macro News: The Importance of Investor Atten-

- tion,” *Finance and Economics Discussion Series*, 2025, (2025-022), 1–1.
- Lucca, David O. and Emanuel Moench**, “The Pre-FOMC Announcement Drift,” *The Journal of Finance*, 2015, *70* (1), 329–371.
- Monsell, Stephen**, “Task switching,” *Trends in Cognitive Sciences*, 2003, *7* (3), 134–140.
- Morris, Stephen and Hyun Song Shin**, “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.
- RBC Capital Markets**, “Execution & Unbundling Post-MiFID II: A Global Multi-Asset Market Survey,” Working Paper, 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.
- Sastry, Karthik A**, “Disagreement About Monetary Policy,” *Working Paper*, 2024.
- Veldkamp, Laura L.**, “Information Markets and the Comovement of Asset Prices,” *The Review of Economic Studies*, 2006, *73* (3), 823–845.
- Wang, Chen and Kangying Zhou**, “The Macro Alibi: Subjective Risk Attribution in Analyst Scenarios,” *SSRN Working Paper*, 2026.

Tables

Table I: Report Coverage by Contributor

Contributor	Reports	Start	End	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” are the first and last years with at least 36 reports. *Direct* denotes reports received directly from the forecaster’s proprietary research portal; *LSEG* are reports received from LSEG workspace. *Investext* are reports received from Mergent Investext. The total covers the 26 listed contributors.

Table II: Sample Overlap with Professional Forecaster Surveys

Survey	Share 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

Coverage: Shares report the percentage of each survey category represented among the 26 contributors covered by this survey comparison. “Survey total” counts all participants; “Sample only” counts forecasters in this paper’s sample who do not participate in that survey. Survey participants are matched to sample forecasters by institution. Primary dealers are a subset of sell-side banks. “Other” includes independent research providers, asset managers, universities, and nonfinancial organizations.

Average: Surveys receive equal weight.

Survey dates: 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.

Roster definitions: Bloomberg requires quarterly GDP, core PCE, unemployment, and policy-rate forecasts; Reuters requires quarterly GDP, CPI, unemployment, and fed funds paths. WSJ’s Economic Forecasting Survey includes 74 respondents, excluding Boston College, Stifel, and Daiwa, which submitted no forecasts. The NY Fed row covers the 26 primary dealers on its Survey of Market Expectations mailing list; inclusion does not necessarily imply a survey response. Primary-dealer classifications use the NY Fed’s list at the date of each survey.

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

Variable	Measure	Forecasts	Value Changes	Med. Days Btwn. Updates	Med. Horizons per Report
Fed funds rate	Level	104,793	9,489	6	4
2Y Treasury	Level	49,641	9,331	7	5
5Y Treasury	Level	39,613	7,705	7	5
10Y Treasury	Level	84,738	17,136	7	4
30Y Treasury	Level	32,965	7,080	14	5
GDP growth	QoQ SAAR	87,527	15,679	7	5
Headline CPI	QoQ SAAR	26,076	9,363	29	5
Total		425,353	75,783		

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 and CPI inflation are quarter-over-quarter seasonally adjusted annualized rates (QoQ SAAR). *Value Changes* counts consecutive changes for each contributor and target quarter. *Med. Days Btwn. Updates* is the median across contributors of each contributor’s median gap between report dates carrying that variable. *Med. Horizons per Report* is the median number of distinct target quarters per contributor and report date.

Table IV: Fed Research and Forecast Accuracy

Forecast	(1)	(2)	(3)	(4)	<i>N</i>
<i>Realized RMSE over next three quarters</i>					
Fed funds	5.7** (2.6)	-1.6*** (0.6)	6.8** (3.5)	-2.1*** (0.8)	10,637
Real GDP	12.8*** (4.6)	1.8 (1.7)	17.4*** (6.3)	5.0*** (1.8)	8,199
Headline CPI	-15.9 (13.9)	4.3 (3.2)	-7.9 (16.2)	13.9*** (5.0)	2,005
Date FE	No	Yes	No	Yes	
Forecaster FE	No	No	Yes	Yes	

Notes: .

The full specification is

$$L_{it}^v = \alpha + \beta A_{it} + \delta_{t,q}^v + \mu_i^v + \text{Controls}_{it}^v + \varepsilon_{it}^v.$$

L is the realized RMSE over the next three quarters of forecasts, measured in basis points. The coefficients shown are on A , the share of research about the Fed on the forecast date and the preceding 30 days, as defined in subsection 3.3. Coefficients are reported per 15-percentage-point increase in Fed research share, approximately one standard deviation in the broader sample. The estimation sample filters for forecaster x date pairs where a forecaster produces a forecast for each of the next three quarters on that day. Controls include a log of the total analytical text produced on the forecast date and over the preceding 30 days and absolute distance from consensus. Forecast dates: 2000–2025. GDP realizations extend through 2026Q2 and CPI through 2026Q3. Missing October 2025 CPI is linearly interpolated between September and November. GDP bundles intersecting 2020Q2–2021Q4 are excluded. Standard errors in parentheses are clustered by forecaster and target bundle. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table V: Fed Research and Rate Forecast Accuracy by Horizon

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

Notes: The following equation is estimated:

$$|\widehat{r}_{it,q}^v - r_q^v| = \alpha + \beta_{v,h} A_{it} + \delta_{t,q}^{v,h} + \mu_i^{v,h} + \text{Controls}_{it}^{v,h} + \varepsilon_{it,q}^{v,h}.$$

The dependent variable is absolute forecast error at the indicated horizon, in basis points. Coefficients are reported per 15-percentage-point increase in Fed research share, approximately one standard deviation in the broader sample. Forecast dates are 2000–2025; Treasury realizations extend through 2026Q3. Research share and volume cover the forecast date and preceding 30 days. Each regression includes forecaster and date fixed effects, log analytical text volume, and absolute distance from consensus. The sample is filtered for forecaster x date pairs with all forecasts available for each of the next four quarters on that date. Standard errors in parentheses are clustered by forecaster and target quarter. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table VI: Forecaster Surprises and Forecast Revisions

Outcome	N	(1)	(2)	(3)
Fed funds rate	1,039	1.34 ^{***} (0.42)	1.45 ^{***} (0.43)	1.01 ^{***} (0.30)
GDP growth	1,031	0.66 (0.44)	1.17 ^{***} (0.41)	1.17 ^{***} (0.41)
Inflation	871	0.70 [*] (0.42)	1.15 ^{**} (0.53)	1.08 ^{**} (0.50)
Meeting fixed effects		Yes	Yes	Yes
Lead economist fixed effects		No	Yes	Yes
Forecast controls		No	No	Yes

Notes: Cell entries represent β^k for the following regression:

$$\Delta f_{it}^k = \alpha_t^k + \lambda_{e(i,t)}^k + \beta^k \tilde{S}_{it}^k + \gamma^{k'} \mathbf{Z}_{it}^k + \varepsilon_{it}^k.$$

where i indexes forecasters, t FOMC meetings, and k outcomes. The dependent variable Δf_{it}^k is the forecast revision around the FOMC announcement in basis points: the average revision for the current quarter and the next three quarters for Fed funds, and the revision in annualized growth over the next three quarters for GDP and inflation. The forecaster surprise score $\tilde{S}_{it}^k$ is scaled by its standard deviation in each outcome's estimation sample. Controls include the pre-announcement forecast level, its level 90 days prior to the announcement, and the sum of Goldman Sachs macroeconomic release surprises from the pre-announcement forecast date through the post-announcement forecast date, averaged across the included forecast horizons; GDP and inflation also include the initial Fed funds forecast and its contemporaneous revision. The maximum interval between pre-FOMC and post-FOMC forecasts is 10 calendar days. Standard errors clustered by meeting are in parentheses. Significance levels are: ^{*} $p < 0.10$, ^{**} $p < 0.05$, ^{***} $p < 0.01$.

Table VII: Summary of FOMC Communication Accuracy by

Category	Meetings	Bank-question observations	Incorrect (%)
Policy decision	207	6,347	10.3
Statement	207	14,055	27.5
Economic projections	59	18,533	42.3
Press conference	87	4,832	20.8
Total	207	43,767	30.5

Notes: The sample contains 43,767 scored forecaster-question observations from 21 forecasters across 207 scheduled FOMC meetings, February 2000–December 2025. Each observation combines one forecaster’s verifiable answers to a meeting-specific question, recovered from pre-announcement reports and verified against the realized communication (see Appendix B.3). *Incorrect* is the average fraction of verifiable answers classified as incorrect: verified-false and partially-true answers count as incorrect, verified-true answers as correct, and unverifiable answers are excluded. Fractions are first averaged within each forecaster-question observation, then equally across observations. The total error rate weights categories by their observation counts. Meetings count distinct meetings with scored observations; categories overlap across meetings, so these counts do not add to the total. *Policy decision* covers the rate decision; *Statement* covers statement language; *Economic projections* covers the Summary of Economic Projections, including the dot plot; *Press conference* covers the chair’s remarks and Q&A.

Table VIII: Fed Research and Asset Price Reaction Magnitudes

	S&P 500			10-year Treasury yield		
	(1)	(2)	(3)	(4)	(5)	(6)
Fed research	0.21*** (0.06)	0.11* (0.07)	0.05 (0.07)	0.19** (0.08)	0.12 (0.09)	0.06 (0.08)
Communication disagreement		0.26*** (0.07)	0.10 (0.08)		0.18** (0.08)	0.04 (0.09)
Communication surprise			0.25** (0.11)			0.23** (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

$$|\Delta p_t^a| = \alpha^a + \beta_R^a R_t + \beta_U^a U_t + \beta_S^a S_t + \varepsilon_t^a.$$

$|\Delta p_t^a|$ is the absolute S&P 500 return or 10-year yield change over the full announcement window (Section 3.4). 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 fraction of answers verified as incorrect. All variables are expressed in standard-deviation units within each asset's sample. Regressions include an intercept and no additional controls. The sample covers scheduled FOMC meetings in 2000–2025, excluding March 18, 2009 for Treasuries. Standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table IX: Treasury Yield Responses to Economic Share of Communication Surprises

	2-year		5-year		10-year		30-year	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
MPS	1.43*** (0.11)	1.65*** (0.07)	1.23*** (0.12)	1.44*** (0.11)	0.72*** (0.09)	0.88*** (0.10)	0.20** (0.09)	0.32*** (0.09)
Economic content	-0.26 (0.20)	-0.13 (0.15)	-0.17 (0.23)	-0.05 (0.21)	-0.25 (0.20)	-0.16 (0.18)	-0.17 (0.20)	-0.10 (0.20)
MPS $\times$ economic content		0.37*** (0.06)		0.36*** (0.07)		0.28*** (0.06)		0.20*** (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	0.38	0.43	0.06	0.10

Notes: Each column estimates

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

Odd-numbered columns omit the interaction ($\delta^{(n)} = 0$); even-numbered columns include it. Each pair uses the same observations and controls. $\Delta y_t^{(n)}$ is the Treasury yield change over the monetary event (statement and press conference). Yield changes and MPS (the U.S. Monetary Policy Event-Study Database monetary-event factor) are in basis points. Economic_t is the share of forecast-error weight concerning economic assessments; Surprise_t is the overall forecast-error rate. Both text measures are standardized. The sample is 206 scheduled FOMC meetings, February 2000–December 2025, excluding March 18, 2009 (QE1). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figures

Figure I: Timeline of the model.






Figure II: Federal Reserve Communication over Time




Notes: Panel A plots an eight-meeting rolling average of total word counts in official announcement-day communication for scheduled FOMC meetings through September 16, 2026, on a logarithmic scale. Materials include statements, implementation notes and other same-day policy documents, Summary of Economic Projections (SEP) materials, statements on longer-run goals, and press-conference transcripts. Minutes and other documents released later are excluded. Panel B separately indexes press conferences, statements, and projection materials to 100 at their respective averages over the 16 scheduled meetings from June 2024 through April 2026. The projection baseline uses the eight SEP releases in that period; projection markers appear only at SEP meetings. The vertical dashed line marks May 22, 2026, when Kevin Warsh took office as Fed Chair. Endpoint annotations compare September 2026 with each series' pre-Warsh average, rather than averaging all meetings under Warsh. Counts use whitespace-delimited extracted text, including table cells; accessible-HTML SEP releases use a consistent conversion. The September 2026 press-conference transcript is preliminary. Source: Federal Reserve Board releases.

Figure III: Report Coverage, Publication Frequency, and Length



Notes: Publication frequency is reports divided by calendar weeks in the year; annual report length is mean pages per report. Panels C and E show the median and 25th–75th percentiles across forecasters each year. Panels D and F show each forecaster’s mean across eligible years, sorted separately.


Figure IV: Fed Research Intensity Around FOMC Announcements




Notes: Percent of text written about the Fed defined in Section 3.3. Panel A averages across forecasters and time periods (2000-2025). Each panel uses scheduled announcements only. Since January 2005, minutes have been released three weeks after meetings; previously, they were released on the Thursday after the following meeting.

Figure V: Forecast Accuracy around FOMC Announcements


A. Average changes in forecast accuracy



B. Real GDP




C. Headline CPI



Notes: Negative changes indicate improved accuracy. Panel A reports mean changes and 95% confidence intervals clustered by meeting, using the last forecast in the seven days before and the first forecast in the seven days after scheduled FOMC announcements during 2000–2025; announcement-day forecasts are excluded. Fed funds, GDP, and CPI use root mean squared error (RMSE) over the next three quarters; Treasuries use absolute error three quarters ahead. GDP realizations extend through 2026Q2, and CPI and Treasury realizations through 2026Q3. October 2025 CPI is linearly interpolated between September and November. All outcomes require forecasts for each of the next three quarters at both endpoints; GDP excludes target bundles overlapping 2020Q2–Q4. Panels B and C compare next-three-quarter RMSE. Panels B and C use the latest private forecast within seven days up to the staff forecast date and the first forecast within seven days after the announcement. Fed accuracy advantage is private pre-announcement RMSE minus Greenbook/Tealbook staff RMSE. The sample covers 2000–June 2019 and requires all target quarters to end before 2020Q2. Dots show 15 equal-count bins; lines are fitted to individual observations. Samples and axis scales differ across outcomes. Both scatter axes contain initial private forecast error, so the relationships are descriptive.

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



Notes: Each panel groups 207 scheduled FOMC meetings from February 2000 through December 2025 into ten equal-count bins of pre-announcement Fed research share. Dots show bin means, with equal weight on each meeting. Fed research share is measured over the seven days leading up to each FOMC announcement. Communication uncertainty is the percentage of comparable forecaster pairs giving different answers to the same announcement question, pooling comparable pairs within each meeting. Communication surprise averages the percentage of incorrect answers within each forecaster-question observation and then across observations within a meeting.

A Data Sources and Processing Pipeline

I develop a document processing pipeline for macro research reports 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, complex tables. 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 imperative 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. This recursive approach produces a running description of what each page covers (country, topic, variable), which is then used to augment the text and image embeddings with hierarchical context that would otherwise be lost during chunking. 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 and proceed with a multi-modal pipeline that can pass raw images directly to a frontier LLM at runtime, which handles complex tables well.

Figure A1 summarizes the pipeline.

Figure A1: Document Processing 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. Text chunks and image summaries are then each prefixed with the document headline and page context, then embedded 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 A2 illustrates this structure.

Figure A2: Anatomy of a Document Chunk



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. For each extraction step, the LLM is given both the retrieved information and a set of instructions on how to process it. 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.

B.1 Measuring Fed Research

Figure [A3](#) illustrates direct reasoning about the Fed’s next policy decision. Figure [A4](#) 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.


Figure A3: From Report Text to Fed Research Share



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.



A. 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 A4: Economic Research Motivated by the Fed

B. 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.”

C. 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.”

Panels B–C: Société Générale, September 28, 2015.

$$\text{Fed research share} = \frac{\text{Fed reasoning} + \text{Fed motivation} \times \text{Economic reasoning}}{\text{All reasoning}}.$$

Notes: Reports are divided into text blocks. The measurement assigns shares of each block to Fed reasoning, economic reasoning, and factual reporting, and separately measures the article’s Fed-motivation share. Economic reasoning can therefore enter Fed research even when the passage itself concerns economic fundamentals. Block contributions are weighted by character count; factual reporting is excluded from the reasoning denominator. The quotations illustrate the categories, not binary scores assigned to every paragraph or numerical motivation shares for these examples.

B.2 Quantitative Forecast Extraction

The key variables in this study are the quantitative forward-looking macro-financial forecasts that analysts make. The central challenge in extracting this information is that macro-financial 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 is using often requires reading table footnotes, column headers, or surrounding text rather than just the number itself. To guide the LLM through these distinctions, the extraction prompt provides detailed instructions on how to read forecast tables and interpret text, including common practices analysts use when presenting forecasts. The LLM is also given 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 A5 summarizes the extraction pipeline.

Figure A5: Quantitative Forecast Extraction Pipeline



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) extracts all quantitative forecasts from each passage, outputting structured records with variable, measure, horizon, and value (for example, real GDP growth, quarter-over-quarter annualized, 2024Q3, 2.1%).

The raw extractions are then passed to a refining agent that flags potentially suspicious values and queries a more capable model (Gemini 3 Pro) for review. Three types of situations trigger a flag: first, when multiple passages from the same date yield conflicting values for the same variable and horizon; second, when a forecast exhibits a reversion pattern (moving from value A to B and back to A in a short window), suggesting a possible extraction error; and third, when a variable or measure type appears for the first time for a given contributor, since the lack of historical context makes it harder to verify the extraction. In each case, the flagged chunks are

sent to the review model along with surrounding context, and the model is asked to confirm or correct the extracted values.

I validate the extracted point forecasts by comparing them against independently recorded submissions to the Reuters Economic Polls. For GDP growth, headline CPI inflation, and the fed funds rate, I match bank-quarter pairs where the same institution submitted a forecast to the Reuters poll and also published a report from which the pipeline extracted a forecast for the same target quarter, requiring the extraction date and poll submission date to fall within ± 7 days. Table A1 reports the results: correlations between extracted and poll-submitted forecasts range from 0.975 to 0.998 across these three variables, based on 85,368 matched observations from 25 forecasters.

Table A1: Extraction Validity: Correlation with Reuters Economic Poll Submissions

Variable class	Extracted	Matched	Correlation
	Sample	Refinitiv	
Real GDP growth (QoQ SAAR)	130,791	30,800	0.977
Headline CPI inflation (YoY)	141,412	19,042	0.975
Fed funds rate (EOP)	135,000	35,526	0.998


Notes: Pearson correlation between report-extracted forecasts and Reuters (Refinitiv) Economic Poll submissions by the same bank for the same target quarter, matched to the nearest poll submission within ± 7 days of the report date. Extracted counts report forecasts from institutions present in both datasets; Matched counts the subset with an eligible poll match. QoQ SAAR denotes quarter-on-quarter growth at a seasonally adjusted annual rate; YoY denotes year-on-year growth; EOP denotes quarter-end values.

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. Returning to the source reports to answer this questionnaire places forecasts expressed in different formats and at different levels of detail on a common basis. I then compare these

expectations with the realized announcement and construct cross-sectional scores of how forecasters were surprised. Figure A6 summarizes the procedure.

Figure A6: 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. Verified outcomes and those answers enter pairwise comparisons, which are combined into cross-sectional surprise scores. Prior official releases can supply comparison baselines; current-meeting outcomes enter only after expectations are reconstructed.

Source selection and forecast identification. I define the pre-announcement window as the three weeks leading up to each meeting. Forecasters typically detail their Fed calls in the final week, but the longer window also captures earlier previews. Within this window, the RAG pipeline identifies passages likely to contain predictions for the upcoming announcement. A forecast-identification prompt reads each retrieved chunk, including its original table or figure when applicable, and returns a set of forecasts, or no forecasts if the chunk contains none. Chunks with forecasts or relevant Fed commentary are retained for later use.

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. This structure facilitates comparison within each forecaster's reports. Repeated forecasts are merged, while a later revision replaces an earlier view about the same attribute. Genuine alternatives remain distinct. The result is a consolidated account of what each forecaster expected the upcoming meeting to deliver.

Standardizing the granularity of forecasts. Even after consolidation, forecasts can differ in their level of detail. One forecaster might expect the Fed to lower its GDP projections at every horizon, while another lists the expected projection for each year separately. Treating these statements as individual observations would make comparisons depend on how each forecaster presented its call. I address this problem in two steps. First, extraction places forecasts in a common structure where possible. For example, an SEP forecast records the variable, horizon, and statistic, such as a median or central-tendency range, together with the expected value. This value can be numerical or qualitative, so an expectation of a downward revision can be retained without inventing a point forecast.

Second, the consolidated forecasts across forecasters are used to construct a unified questionnaire for the meeting. Questions define common announcement attributes and response options describe their possible values or content. Broad and detailed forecasts can thus be mapped to the same questions. Related questions are reconciled to reduce overlap while preserving distinct expectations. Both the topics and the level of detail are drawn from what forecasters discuss in their reports, allowing the questionnaire to vary across meetings. Prior official releases supply baselines when a forecast concerns a change, such as a revision to an SEP projection.

Recovering answers and confidence. The questionnaire is then answered separately for each forecaster. The model receives its previously extracted forecasts together with the retained source text and original visuals. Rereading this evidence allows the model to recover answers from either prose or tables, correct earlier interpretations, and identify expectations missed in the first pass. Answers can be explicit or supported implicitly by the source; questions without sufficient evidence remain unanswered. Silence is not an expectation of no change.

Answers retain confidence language such as “expect,” “likely,” or “could,” as well as explicit probabilities such as “roughly 30%.” A forecaster can assign probabilities to several alternatives, and these are recorded separately rather than collapsed into a single certain prediction. Confidence in an event is distinguished from words inside the predicted Fed message: “likely” in a forecast that the Fed will describe rate cuts as likely belongs to the message. If an answer fits an existing question but none of its options, equivalent write-in answers are combined into an additional option.

Verifying the announcement. A separate model call compares all options for each question with the relevant official announcement text, figures, and tables, including the press conference transcript and SEP where available. It records whether each option is true, false, partially true, or not verifiable; whether that assessment is objective or requires judgment; and whether the realization is hawkish, dovish, or neutral relative to the option. For example, a realized 75 basis point increase is hawkish relative to an expected 50 basis point increase and dovish relative to an expected 100 basis point increase. Verification does not receive forecaster identities or confidence language. Its option-level results are subsequently joined to each forecaster’s answers.

Cross-sectional surprise scores. A common questionnaire makes expectations comparable at the attribute level, but does not by itself determine how to aggregate them. Questions differ in economic importance, and forecasters do not all answer the same questions. Simply summing forecast errors would give equal weight to unequal pieces of news and make scores depend on coverage. Constructing a common weighted index would instead require weights for each question and a treatment of missing expectations. I therefore use pairwise comparisons: given the realized announcement, which of two forecasters was more hawkishly surprised? These narrower judgments are then combined into a cross-sectional ranking.

For each pair, the model receives the questionnaire answers, their confidence language, and the verified outcomes. It is instructed to focus on questions answered by both forecasters and to consider related evidence jointly. For example, an inflation-related surprise may appear both in the SEP projections and in the statement’s discussion of inflation. These need not be two independent surprises merely because they appear as separate questions. The comparison prompt directs the model to recognize such overlap rather than mechanically count questions, and to distinguish central expectations from qualified alternatives. This approach still requires judgment about the importance of the evidence, but does not require imputing a complete set of forecasts or assigning a fixed numerical weight to every question. The inputs contain neither asset prices nor subsequent forecaster revisions.

I aggregate the comparisons using a Bradley–Terry model. Let $s_{i,t}$ denote forecaster i ’s relative hawkish-surprise score at meeting t . The probability that i is judged

more hawkishly surprised than j is modeled as

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

The comparison outcome is coded as $y_{ij,t} = 1$ when i is more hawkishly surprised and 0 when j is more hawkishly surprised. A comparable judgment is coded as 1/2, contributing half a win to each forecaster. Each pair receives equal weight; the model’s “modest” versus “decisive” distinction does not affect estimation. Scores maximize the resulting log likelihood with a ridge penalty of $0.01 \sum_i s_{i,t}^2/2$, which keeps estimates finite when the comparisons perfectly separate forecasters.


Scores are fitted using the forecasters included in each meeting’s estimation sample and centered at zero within the meeting. Higher values mean more hawkish surprise relative to peers, not a larger common surprise for the meeting. For regression interpretation, I scale scores by their pooled standard deviation across the estimation sample. This is distinct from standardizing each meeting separately. The same framework produces component-specific scores for policy, economic information, and balance-sheet news. A separate comparison of which forecaster was more surprised in absolute terms produces a relative magnitude score, irrespective of direction.

Figure [A7](#) illustrates the comparison using Wells Fargo and Citi expectations for the June 12, 2024 announcement.

Figure A7: From Announcement Expectations to Relative Surprise

Question	Wells Fargo	Citi	Realized announcement
Longer-run policy-rate dot	Revised higher	Unchanged	Revised higher
Statement on inflation	Could upgrade inflation language	Upgrades inflation language	Upgrades inflation language
Powell on the policy stance	Not addressed in supplied reports	Restrictive	Restrictive

Example: Pre-announcement expectations for the June 12, 2024 FOMC meeting. Citi’s June 10 preview combines a numerical projection table, proposed statement language, and expectations for Powell’s remarks. Each is mapped to a common meeting question. A tentative claim remains distinct from a firm forecast; absence from the supplied reports is not coded as agreement with the outcome.



Notes: Each grid contains 47 questions in the same order. Colors summarize the realized announcement relative to each forecaster’s expectations. Split cells retain different scenarios within an answer; segment widths are not probabilities or estimated attention shares. The comparison uses the full answers and their context, rather than counting colored cells. Pairwise judgments across forecasters feed the Bradley–Terry scores described in Appendix B.3. The three examples above are selected questions; the grid summarizes the full questionnaire. Source: the worked Wells Fargo–Citi comparison in the presentation.

C Research Responses to Fed Communication

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

The economic-research first-order condition is equation (17). The corresponding condition for Fed research is

$$MC_F \equiv C_F + \kappa a_X = m_F + \ell_F m_X \equiv MB_F. \quad (23)$$

Fed research. Let $C_F^L < C_F^H$ and let a_F^L, a_F^H be Fed research at 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 and does not require convexity. For the strict local results in the main text, assume an interior optimum satisfying the strict second-order condition. Let H be the Hessian of $C - U$ with respect to (a_X, a_F) at that optimum, and write $D_H = H_{XX}H_{FF} - H_{XF}^2 > 0$. Differentiating the first-order conditions gives

$$\frac{da_F^*}{dC_F} = -\frac{H_{XX}}{D_H} < 0, \quad \frac{da_X^*}{dC_F} = \frac{H_{XF}}{D_H}. \quad (24)$$

Positive definiteness concerns the full objective $C - U$; the product cost $\kappa a_X a_F$ is not itself jointly convex. These derivatives describe a regular interior branch, rather than switches between distinct optima.

Fed forecast precision. Since $\partial P_F / \partial a_F = 1$ and $\partial P_F / \partial a_X = \ell_X$,

$$\frac{dP_F}{dC_F} = \frac{\ell_X H_{XF} - H_{XX}}{D_H}. \quad (25)$$

To establish the sign, write $v_X = P_X^{-1}$, $v_F = P_F^{-1}$, $q = \text{Cov}(X, F \mid s_X, s_F) > 0$, and $\delta = v_X v_F - q^2 > 0$. Gaussian updating gives $\ell_X = q^2 / v_F^2$. Let $S = v_X^\sigma + v_F^\sigma$ and $\eta = (\sigma - 1)v_X^{\sigma-2}v_F^{\sigma-2}S^{1/\sigma-2} \geq 0$. Within this proof, write $w_X = P_X^2 m_X$ and $w_F = P_F^2 m_F$. Differentiation gives

$$H_{XX} = \eta v_X^2 \delta^2 + 2w_X v_X^3 + 2w_F q^2 v_X,$$

$$H_{FF} = \eta v_F^2 \delta^2 + 2w_X q^2 v_F + 2w_F v_F^3.$$

Since $v_X > q^2/v_F$, these expressions imply $H_{XX} > \ell_X^2 H_{FF}$. Positive definiteness then implies

$$H_{XF} < \sqrt{H_{XX} H_{FF}} < \frac{H_{XX}}{\ell_X}.$$

Thus $dP_F/dC_F < 0$: cheaper Fed research strictly increases Fed forecast precision on the regular interior branch.

Economic precision and response regions. Consider a connected regular interior response to C_F . Write $b = \sigma_B^2$, $f = a_F$, and $P = P_X$. Holding f fixed, substitute $a_X = P - \sigma_X^{-2} - f/(1 + bf)$ and define

$$d = \frac{b}{1 + bf}, \quad z = \frac{P_X}{P_F} = \ell_F + dP.$$

The marginal benefit of economic research, expressed in these coordinates, is

$$M(P, f) \equiv m_X + \ell_X m_F = \frac{1 + \ell_F z^{\sigma-1}}{P^2 (1 + z^\sigma)^{(\sigma-1)/\sigma}}. \quad (26)$$

The first-order condition is $M(P, f) = C_X + \kappa f$. Since $M_P = -H_{XX} < 0$, implicit differentiation gives

$$\frac{dP}{df} = \frac{M_f - \kappa}{-M_P}, \quad (27)$$

where M_f holds economic precision fixed. At $\sigma = 1$, this condition reduces to $P^2 = (1 + \ell_F)/(C_X + \kappa f)$. Setting $\kappa = 0$ gives the rearranged form of equation (18). Because $\ell_F = (1 + bf)^{-2}$ strictly decreases with f , economic precision falls as cheaper Fed research raises f ; a positive cross-cost reinforces this result when $\sigma = 1$.

Now suppose $\sigma > 1$ and $\kappa = 0$. The CES preferences imply

$$\frac{\partial m_X}{\partial P_F} = (\sigma - 1) \frac{m_F}{P_X [1 + (P_X/P_F)^\sigma]} > 0. \quad (28)$$

Write

$$\rho \equiv P_X/P_F = z. \quad (29)$$

Substituting the CES marginal utilities and rearranging the first-order condition gives

$$P_X^2 = \frac{1}{C_X} \frac{1 + \ell_F \rho^{\sigma-1}}{(1 + \rho^\sigma)^{(\sigma-1)/\sigma}}. \quad (30)$$

Direct differentiation gives

$$M_f = b\sqrt{\ell_X} m_F \left[(\sigma - 1) \frac{1 - \ell_X \ell_F}{1 + \rho^\sigma} - 2\sqrt{\ell_X \ell_F} \right]. \quad (31)$$

The information structure implies

$$\sigma_B^4 P_X P_F = \frac{(1 - \sqrt{\ell_X \ell_F})^2}{\sqrt{\ell_X \ell_F}}, \quad \frac{2\sqrt{\ell_X \ell_F}}{1 - \ell_X \ell_F} = \frac{2}{\sigma_B^2 \sqrt{P_X P_F (4 + \sigma_B^4 P_X P_F)}}.$$

The prefactor and $-M_P$ are positive. Since $da_F^*/dC_F < 0$, these identities establish equation (19). Equivalently, $M_f > 0$ if and only if $K > 1$, with equality when $K = 1$, where

$$K(P, f) = \frac{(\sigma - 1)[z/\ell_F - \ell_F/z]}{2(1 + z^\sigma)}. \quad (32)$$

To establish how the condition changes with C_F , define

$$N(z, \ell_F) = \frac{1 + \ell_F z^{\sigma-1}}{(1 + z^\sigma)^{(\sigma-1)/\sigma}}, \quad C_X P^2 = N.$$

Since $z = \ell_F + dP > \ell_F$,

$$N_z = \frac{(\sigma - 1)z^{\sigma-2}(\ell_F - z)}{(1 + z^\sigma)^{(2\sigma-1)/\sigma}} < 0, \quad N_{\ell_F} > 0, \quad N - \ell_F N_{\ell_F} > 0.$$

Write $D = 2C_X P - dN_z > 0$. Differentiating $C_X P^2 = N$ along the optimal response gives

$$\frac{dz}{df} = -\frac{2d[d\ell_F N_{\ell_F} + C_X P(z + \ell_F)]}{D} < 0,$$

$$\frac{dP}{df} + dP = \frac{2d[N - \ell_F N_{\ell_F} - z N_z]}{D} > 0.$$

Because

$$(\ell_X \ell_F)^{-1/2} = z/\ell_F = 1 + bP(1 + bf),$$

the second inequality implies $d\sqrt{\ell_X \ell_F}/df < 0$. Since $df/dC_F < 0$, both $d\rho/dC_F$

and $d\sqrt{\ell_X\ell_F}/dC_F$ are positive. The precision-product identity above then implies $d(P_X P_F)/dC_F < 0$. The right side of equation (19) therefore increases with C_F . Equivalently,

$$\frac{d}{dC_F} \left[\frac{2[1 + (P_X/P_F)^\sigma]}{\sigma_B^2 \sqrt{P_X P_F (4 + \sigma_B^4 P_X P_F)}} \right] > 0. \quad (33)$$

Along a connected regular interior response, economic precision can change direction at most once, and only from falling to rising as C_F falls.

Dividing the two research first-order conditions gives

$$\frac{C_F}{C_X} = \rho \frac{\rho^\sigma + \sqrt{\ell_X \ell_F}}{1 + \sqrt{\ell_X \ell_F} \rho^\sigma}.$$

For $0 < \sqrt{\ell_X \ell_F} < 1$, this ratio exceeds one exactly when $\rho > 1$, so

$$C_F > C_X \iff \rho > 1, \quad C_F < C_X \iff \rho < 1. \quad (34)$$

Cross-costs and a low-cost precision decline. When $\kappa > 0$, the economic-research condition is $M(P, f) = C_X + \kappa f$. Rearranging equation (27) gives

$$\underbrace{-M_P}_{>0} \frac{dP}{df} = \underbrace{M_f}_{\text{change in marginal benefit}} - \underbrace{\kappa}_{\text{change in marginal cost}}. \quad (35)$$

Here M_f holds economic precision fixed. Thus $M_f < \kappa$ implies $dP/df < 0$ and, since $df/dC_F < 0$ at a regular interior optimum, $dP/dC_F > 0$. The marginal benefit itself need not fall along the optimum; economic precision adjusts until marginal benefit matches the higher marginal cost.

An explicit sufficient condition follows from equation (31), which remains valid when $\kappa > 0$. Define $R \equiv \sigma_X^{-2} + b^{-1}$. Since $b\sqrt{\ell_X} = 1/(R + a_X)$, for $\sigma > 1$ that equation implies

$$M_f < (\sigma - 1) \frac{m_F}{R + a_X} < (\sigma - 1) \frac{C_F + \kappa a_X}{R + a_X},$$

where the second inequality uses the Fed-research condition $C_F + \kappa a_X = \ell_F m_X + m_F$.

Consequently,

$$(\sigma - 1)C_F < \kappa[R + (2 - \sigma)a_X] \implies \frac{dP_X^*}{dC_F} > 0. \quad (36)$$

For $1 < \sigma \leq 2$, this yields the allocation-independent bound

$$0 < C_F < \frac{\kappa}{\sigma - 1} \left(\frac{1}{\sigma_X^2} + \frac{1}{\sigma_B^2} \right) \implies \frac{dP_X^*}{dC_F} > 0, \quad (37)$$

provided the optimum remains regular and interior. This bound is sufficient; it need not equal the turning point. At $\sigma = 1$, the exact expression $P^2 = (1 + \ell_F)/(C_X + \kappa f)$ already establishes the decline.

An interior low-cost decline is possible for any fixed $\sigma > 1$. To establish existence, suppose

$$\bar{P} \equiv C_X^{-1/2} > R, \quad a_0 \equiv \bar{P} - R > 0. \quad (38)$$

The inequality means that independent economic research remains worthwhile even after exhausting the information available from the Fed. For sufficiently small positive κ , a regular interior optimum then exists for C_F sufficiently close to zero, and $dP_X^*/dC_F > 0$ there.

To see this, first set $C_F = 0$ as an auxiliary limit and let κ decrease toward zero. The forecaster's objective is bounded above by $G(a_X) = -1/(R + a_X) - C_X a_X$, whose unique maximum is at a_0 . Choosing $a_X = a_0$ and $f = \kappa^{-1/3}$ approaches this upper bound. It follows that an optimum satisfies $a_X \rightarrow a_0$, $P \rightarrow \bar{P}$, $f \rightarrow \infty$, and $\kappa f \rightarrow 0$. For sufficiently small κ , the optimum is attained with both research choices positive: its value exceeds the boundary limit $-1/R$, and the cross-cost bounds f whenever a_X is bounded away from zero.

The Fed-research condition at $C_F = 0$ gives

$$\kappa a_X = \ell_F m_X + m_F, \quad \kappa f^2 \longrightarrow \frac{C_X}{b^2 a_0}.$$

The second limit uses $m_X \rightarrow C_X$ and $m_F \sim \bar{P}^{\sigma-1} f^{-(\sigma+1)}$. Meanwhile, differentiating equation (26) at fixed P gives

$$M_f \sim (\sigma - 1)P^{\sigma-2} f^{-(\sigma+1)}. \quad (39)$$

Hence $M_f/\kappa = O(\kappa^{(\sigma-1)/2}) \rightarrow 0$: the additional marginal benefit becomes smaller than the cost increment κ . The optimum is regular for sufficiently small κ , since $H_{XX} \rightarrow 2/\bar{P}^3$, $H_{FF} \sim 2C_X/(b^2 f^3)$, and $H_{XF} = O(f^{-2})$, so the Hessian is positive definite. Continuity extends the strict inequality $M_f < \kappa$ to an interval of positive C_F . This proves an interior low-cost decline without requiring perfect Fed information at any positive κ .

Economic research and response regions. Write $MB_X \equiv m_X + \ell_X m_F$ for the marginal benefit of economic research. The mixed derivative is

$$H_{XF} = \kappa + 2q^2(w_X v_X + w_F v_F) - \eta v_X v_F \delta^2 = \kappa - \frac{\partial MB_X}{\partial a_F}. \quad (40)$$

Since $D_H > 0$, a reduction in C_F raises economic research precisely when $H_{XF} < 0$. At $\sigma = 1$, $\eta = 0$, so $H_{XF} > 0$ and economic research falls. When $\sigma > 1$, the final term captures the opposing preference complementarity force. In precision notation, the same decomposition is

$$\frac{\partial MB_X}{\partial a_F} = \frac{\partial^2 U}{\partial P_X \partial P_F} \left(1 - \sqrt{\ell_X \ell_F}\right)^2 - 2q^2(P_X m_X + P_F m_F), \quad (41)$$

where

$$\frac{\partial^2 U}{\partial P_X \partial P_F} = (\sigma - 1) \frac{P_X^{-\sigma-1} P_F^{-\sigma-1}}{(P_X^{-\sigma} + P_F^{-\sigma})^{2-1/\sigma}}.$$

Using the main-text notation, the derivative holding a_X fixed can also be written as

$$\frac{\partial MB_X}{\partial a_F} = \frac{m_F}{P_X} \left[(\sigma - 1) \frac{(1 - \sqrt{\ell_X \ell_F})^2}{1 + \rho^\sigma} - 2\sqrt{\ell_X \ell_F} (1 + \rho^{-\sigma}) \right]. \quad (42)$$

At $\kappa = 0$, economic research increases when the bracketed expression is positive. The resulting research threshold divided by the precision threshold is

$$\frac{1 + \sqrt{\ell_X \ell_F}}{1 - \sqrt{\ell_X \ell_F}} (1 + \rho^{-\sigma}) > 1.$$

At fixed unequal precisions this cross-partial tends to zero as $\sigma \rightarrow \infty$, so larger σ does not universally strengthen complementarity. This fixed-allocation comparison is not a comparative static of the optimal allocation with respect to σ .

Along the cost response,

$$\frac{da_X^*}{da_F^*} = -\frac{H_{XF}}{H_{XX}} = \frac{\partial MB_X/\partial a_F - \kappa}{-\partial MB_X/\partial a_X},$$

where $H_{XX} = -\partial MB_X/\partial a_X > 0$. Thus $r = -da_X^*/da_F^* = H_{XF}/H_{XX}$ in the forecast-accuracy results. The corresponding Fed-precision response is

$$\frac{dP_F}{da_F^*} = 1 - \ell_X r > 0, \tag{43}$$

where the inequality follows from the Fed-precision result above.

D Appendix: Additional Figures

Figure A8: Research Windows Around the Minutes Release




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

Construction: Compare topic shares in research before and after the minutes week. Residualize forecaster topic changes and minutes topic emphasis using their respective controls, then average across banks within meeting and topic.

Notes: The diagram follows the weekly windows in the presentation: two research weeks between the FOMC week and minutes week, and two research weeks after the minutes week. Forty topic clusters follow Hansen, McMahon, and Prat (2018); the slide analysis covers minutes releases in 2006–2025. 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. The release lag helps distinguish the Fed’s earlier discussion from contemporaneous research, but does not by itself establish exogenous variation. This figure explains the comparison, not its estimated coefficient.


Figure A9: Three Economic-Precision Response Regions⁸



Notes: The figure plots change in economic precision for a reduced cost of Fed research depending on the initial set of Fed research for a specific set of parameters.


The illustration uses $\sigma = 2.8$, $\kappa = 13$, $C_X = 1$, held fixed throughout. The prior economic precision is $1/\sigma_X^2 = 0.005$ ($\sigma_X^2 = 200$), and the policy-wedge variance is $\sigma_B^2 = 30$. The displayed range is $0.075 \leq C_F \leq 270$; the response changes sign at $C_F \approx 53.61$ and $C_F \approx 1.261$. Both research choices are strictly positive throughout this range.

Figure A10: Fed Forecast Accuracy Advantage over Time




Notes: Each dot is the meeting-average difference between private forecaster and Fed staff annualized quarter-on-quarter real GDP growth forecast RMSE over the next three quarters, in basis points. Positive values indicate that the Fed was more accurate. Fed forecasts are from the Greenbook/Tealbook; private forecasts are the latest qualifying forecasts within seven days up to the staff forecast date. The orange line averages available meeting means over the current and previous seven scheduled meetings.

Figure A11: Timing of Forecasts Around FOMC Announcements




Notes: Counts of the last forecast before and first forecast after each scheduled FOMC announcement in 2000–2025, using the named forecasters in Table I. Counts sum the next three target quarters ($h1$ – $h3$), holding each target quarter fixed across the announcement. Each available forecaster–meeting–horizon endpoint is counted separately; unchanged forecasts are included. The horizontal axis shows calendar days. At day zero, adjacent bars distinguish reports classified as before or after the announcement; reports with uncertain same-day timing are excluded. The chart displays endpoints within ten days, before imposing the revision sample’s seven-day endpoint and ten-day total-span restrictions. For Wednesday announcements, days -5 and $+2$ are the preceding and following Fridays.

Figure A12: Timing of Text Surprises and High-Frequency Asset Prices




Notes: Timing of the text surprise construction (Section 7). FOMC predictions and their reasoning are extracted from analyst preview reports published in the seven days before each scheduled meeting and verified against the realized announcement; no text written after the announcement enters the measure. High-frequency monetary policy surprises (MPS_t) and asset returns are measured over the 30-minute statement window, 10 minutes before to 20 minutes after the release, from the USMPD (Section 3.4). The axis break compresses the gap of a few days between the preview report and the announcement.

Figure A13: Communication Uncertainty over Time



Notes: Communication uncertainty 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, first shown when eight meetings are available. Higher values indicate greater disagreement about the forthcoming communication.

Figure A14: Forecast Verification Rate by Qualifier Label



Notes: Each point is a qualifier word. The vertical axis shows the fraction of pre-FOMC text forecast statements containing that qualifier that were subsequently verified as correct; the horizontal axis shows the number of statements containing it (log scale). Qualifiers are extracted from the LLM’s classification of the analyst’s hedging language (e.g., “we expect the Fed to hold rates” is labeled *expect*); negated uses are shown as *(negated)*. Qualifiers with fewer than 30 observations are excluded. The decline from high-confidence qualifiers (*almost certainly*, *will*) to low-confidence ones (*perhaps*, *risk*) provides empirical validation that the qualifier extraction captures meaningful variation in analyst confidence.

E Appendix: Additional Tables

Table A2: Rate Forecast Accuracy by Horizon: Varying Samples

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

Notes: The following equation is estimated:

$$|\hat{r}_{it,q}^v - r_q^v| = \alpha + \beta_{v,h} A_{it} + \delta_{t,q}^{v,h} + \mu_i^{v,h} + \text{Controls}_{it}^{v,h} + \varepsilon_{it,q}^{v,h}.$$

The dependent variable is absolute forecast error at the indicated horizon, in basis points. Coefficients are reported per 15-percentage-point increase in Fed research share, approximately one standard deviation in the broader sample. Forecast dates are 2000–2025; Treasury realizations extend through 2026Q3. Research share and volume cover the forecast date and preceding 30 days. Each regression includes forecaster and date fixed effects, log analytical text volume, and absolute distance from consensus. Each horizon uses forecaster x date pairs with all forecasts available from the next quarter through the indicated horizon on that date. Sample sizes vary across horizons. Standard errors in parentheses are clustered by forecaster and target quarter. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A3: Forecast Accuracy Across Fed Attention Measures

Forecast	30-day average (1)	14-day average (2)	14-day half-life (3)	<i>N</i>
<i>Realized RMSE over next three quarters</i>				
Fed funds	-2.1*** (0.8)	-1.6*** (0.5)	-2.6*** (0.9)	10,630
Real GDP	5.0*** (1.8)	3.9*** (1.3)	6.1*** (2.2)	8,196
Headline CPI	13.9*** (5.0)	9.1*** (2.8)	16.4*** (5.1)	2,005
<i>Absolute forecast error four quarters ahead</i>				
2-year Treasury	5.7*** (2.0)	2.2* (1.3)	6.2*** (2.3)	3,325
5-year Treasury	7.9*** (2.5)	3.5** (1.7)	8.8*** (2.8)	2,649
10-year Treasury	3.0** (1.4)	2.0** (0.8)	3.4** (1.5)	5,703
30-year Treasury	-1.5 (1.8)	-0.4 (1.2)	-0.8 (2.0)	1,642
Date FE	Yes	Yes	Yes	
Forecaster FE	Yes	Yes	Yes	

Notes: Each column estimates

$$L_{it}^v = \alpha + \beta A_{it}^m + \delta_t^v + \mu_i^v + \text{Controls}_{it}^v + \varepsilon_{it}^v.$$

L is realized RMSE over the next three quarters for Fed funds, GDP, and CPI, and absolute forecast error four quarters ahead for Treasuries. A^m is Fed research share using the indicated weighting measure. Effects are in basis points per 15-percentage-point increase in Fed research share. All columns include forecaster and date fixed effects, log research volume, and distance from consensus. Samples are held fixed across columns within each outcome and require forecasts for all three quarters for macro outcomes and all four quarters for Treasuries.

Average measures include the forecast date and preceding 30 or 14 days. The half-life measure weights research from d days earlier by $2^{-d/14}$, using all available past research. Shares are ratios of weighted Fed research characters to weighted analytical characters, using the research classification defined in subsection 3.3. The research-volume control uses the forecast date and preceding 30 days in every column.

Forecast dates: 2000–2025. Realization coverage follows Table IV for GDP and CPI; Treasury realizations extend through 2026Q3. October 2025 CPI is linearly interpolated between September and November. GDP target bundles intersecting 2020Q2–2021Q4 are excluded. Standard errors in parentheses are clustered by forecaster and target bundle for macro outcomes, and by forecaster and target quarter for Treasuries. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A4: Effect of Monetary Policy on Average Forecast Revisions

	(1) Monthly	(2)	(3)
	Baseline	Ten-day window	
		Baseline	With controls
Fed funds	4.99*** (1.70)	2.35*** (0.62)	2.18*** (0.62)
<i>Observations</i>	4,136	7,261	7,261
GDP growth	3.61** (1.73)	−1.01 (0.64)	−1.60** (0.65)
<i>Observations</i>	4,089	9,495	9,495
Inflation	−0.97 (1.11)	−0.39 (0.58)	−0.88* (0.51)
<i>Observations</i>	1,705	2,933	2,933

Notes: This table mirrors the tests in [Nakamura and Steinsson \(2018\)](#) and [Bauer and Swanson \(2023a\)](#) using my forecaster-report dataset. Entries report β^k from

$$\Delta f_{it}^k = \alpha_i^k + \beta^k \widetilde{\text{MPS}}_{it} + \gamma^{k'} \mathbf{Z}_{it} + \varepsilon_{it}^k.$$

Forecast revisions are in basis points, averaged over the current and next three quarters for Fed funds and the next three quarters for annualized GDP and inflation growth. MPS is the monetary-event surprise from the U.S. Monetary Policy Event-Study Database ([Acosta et al., 2025](#)).

Column (1) approximates the frequency of the Blue Chip survey by taking each forecaster’s most recent forecast as of the second US business day of each month and holding target quarters fixed across snapshots. Columns (2)–(3) use the narrowest revision window around announcements, up to 10-day intervals, together with revisions outside of announcement windows. All columns include forecaster fixed effects. Column (3) adds controls adapted from [Bauer and Swanson \(2023a\)](#): prior 30-day forecast changes for all three outcomes; Goldman Sachs MAP macro surprise scores 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 yield spread, and the S&P GSCI spot index; a time trend; and log(window days). The sample covers 2000–2025, excluding March 2020–March 2021 and June 15, 2022 (WSJ leak). Standard errors cluster by survey interval in column (1), and by announcement date or ordinary-revision midpoint in columns (2)–(3). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A5: 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.