Monetary Policy by Committee

Toomas Laarits, Ben Matthies, Kaushik Vasudevan, and Will Yang*

March 20, 2025

Abstract

We study how groups make decisions using over forty years of meeting transcripts from the Federal Open Market Committee. We document that committee members employ distinct models of the macroeconomy, leading to persistent differences in policy preferences. On average, the committee tilts its decision towards the views of members whose models better align with recent data. These tilts induce monetary policy surprises. Guided by our findings, we construct a theoretical framework of group decision-making under model uncertainty in which committee members interpret the same data through different models. The committee may outperform its best member if members are sufficiently open to each others' interpretations, but must balance this benefit against the risk of excess sensitivity to recent data.

*Toomas Laarits is with New York University, email: tlaaris@stern.nyu.edu. Ben Matthies is with the University of Notre Dame, email: bmatthie@nd.edu. Kaushik Vasudevan is with Purdue University, email: kvasude@purdue.edu. Will Yang is with Yale University, email: will.j.yang@yale.edu. For helpful comments and conversations, we thank Nick Barberis, Chris Clayton, Zhi Da, Bill English, Paul Fontanier, Peter Hansen, Sid Kankanala, Aleksei Oskolkov, Stefano Pegoraro, John Shim, and Alp Simsek, as well as seminar participants at Yale University and the University of Notre Dame. A substantial portion of this work was completed while Vasudevan was visiting Yale University, whose hospitality he gratefully acknowledges.

1 Introduction

Group decision-making shapes economic outcomes across many contexts. While a long-standing literature following Condorcet (1785) emphasizes that groups can improve decisions by pooling diverse information, decisions are often made in groups even when information-sharing gains may be limited. For example, the vast majority of central banks make monetary policy decisions by committee, yet, as Blinder (2008) argues, "you can mostly ignore [differential information] in the monetary policy context because virtually all the data that matter are common knowledge." Consistent with this view, Romer and Romer (2008) find that the average of policymakers' forecasts adds little predictive value beyond those of central bank staff. This raises the questions of why the committee structure persists in such settings and how it influences policy decisions.

In this paper, we explore the group decision-making process by examining over forty years of transcripts from meetings of the Federal Open Market Committee (FOMC), which sets monetary policy in the United States. We document that members of the committee maintain distinct and persistent "mental models," defined as frameworks used to interpret economic data and form expectations. Analyzing how the committee aggregates its members' perspectives, we find that its policy decisions give greater weight to members whose models best explain recent economic conditions. Such tilting by the committee induces monetary policy surprises. Guided by these findings, we develop a theoretical framework of group decision-making under model uncertainty in which members interpret the same data through different models. Our framework suggests that group decision-making may enhance decision quality if members are sufficiently willing to consider others' views, but the committee must balance this potential benefit against the risk of excess sensitivity to incoming data.

Empirical Results. To conduct our study, we construct a novel data set that captures the policy preferences and their drivers for each FOMC member at each meeting from March 1976 to December 2019. For each meeting, we analyze the full text of every statement made by each committee member to systematically measure how they weigh different macroeconomic and financial variables—inflation, growth, employment, credit market conditions, and the stock market—in their policy arguments. Using a large language model, we translate these statements into member-by-meeting scores that capture both the importance of each variable (magnitude) and its policy implications (sign, where negative indicates dovish and positive indicates hawkish positions). In a limited subsample, members of the FOMC report their individual forecasts of inflation, growth, and unemployment as part of the Summary of Economic Projections (SEP). Our policy preference measures correlate strongly with these self-reported data, while providing a longer time-series and greater granularity.

With the data set in hand, we examine committee members' preferences across meetings and over time. We find that committee members exhibit stable policy preferences, and moreover, tend to be consistent in the variables they voice as driving those preferences. Member

fixed effects explain 20% to 30% of the variation in individual policy preference measures, capturing persistent differences in how members respond to economic conditions. For instance, some members consistently favor hawkish policy due to concerns about inflation, while others consistently favor dovish policy due to concerns about growth and unemployment. Moreover, these policy preference measures exhibit a strong factor structure: the first principal component explains 63% of the variation in member-by-meeting level observations across factors and 84% of the variation in member fixed effects. That is, committee members who are relatively more hawkish about inflation also tend to be relatively more hawkish about growth and unemployment. These patterns suggest that committee members operate with fixed mental models of the macroeconomy, which they use to filter and interpret economic data and form forecasts across variables. Heterogeneity in these mental models leads some members to consistently produce higher or lower forecasts across variables, which in turn gives rise to the strong fixed effects observed in policy preferences.

In order to establish the drivers of members' policy preferences, we link these preferences with the interpretations of economic data shared in committee meetings. We find that a specific set of frameworks explains substantial variation in policy preferences. We develop taxonomies of economic frameworks across inflation, growth, labor, and financial domains and then use a large language model to analyze how each committee member's meeting statements comport with these frameworks. Specifically, we identify whether a particular framework is present in a given member's speeches—for example, whether they discuss the unemployment versus inflation tradeoff embodied in the Phillips curve, or cost-push shocks—as well as the *parameters* that members assign to these frameworks—for example, the *slope* of the Phillips curve, or the *degree of pass-through* of cost-push shocks to prices. We find that policy preferences for each variable can largely be summarized by a small number of framework/-parameter pairs, with about five such pairings matching the explanatory power of fixed effects for members' policy preferences. Hawkish beliefs and stances frequently coincide with perceptions of a steep Phillips curve and high pass-through of cost-push shocks, while dovish ones are associated with discussions of persistent demand shocks and stressed financial market assessments. These findings highlight a strong link between policy preferences and the economic frameworks members share to interpret macroeconomic data.

In our first main empirical finding, we document how the FOMC aggregates the policy preferences of its members. We find that the committee systematically tilts its decision towards the views of members whose models better explain recent economic conditions. To do so, we first construct measures of each committee member's alignment with the committee's policy decision, based on their statements during the policy discussion section of the meetings. We then represent each member's mental model as a vector of their policy preferences in the previous meeting, and assess its alignment with incoming economic data, measured using quantitative indicators (Federal Reserve staff forecasts and financial data) and textual

analysis of staff assessments. Members whose models better match the recent data are more aligned with the committee's final decision. This pattern holds both across members within meetings and within members over time. The evidence suggests that the committee gives more weight to the perspectives that are better aligned with recent economic conditions.

We identify two channels that potentially drive this pattern. First, committee members whose models better match economic conditions speak more during the policy discussions, suggesting they exercise greater influence through deliberation. Second, the chair of the committee consistently shows strong alignment with the final decision, suggesting they may help steer the committee towards models that better fit the recent data.

In our second main empirical result, we demonstrate that the changes to the FOMC policy that arise from committee deliberations, and form the basis of our committee member alignment measure, have strong predictive power for policy surprises measured using asset market data. We measure monetary policy surprises as the high-frequency response of interest rate futures contracts in the 30-minute period following the FOMC's policy announcement, following the extant literature (Gürkaynak, Sack and Swanson, 2005; Nakamura and Steinsson, 2018; Bauer and Swanson, 2023*a*,*b*). The hawkish or dovish tenor of the changes to policy is strongly related to monetary policy surprises. This result indicates that the committee's tendency to tilt towards the views of the members that are most aligned with recent data is not just a reflection of internal deliberation, but also a key driver of financial market responses. This result further suggests that financial market participants may not internalize the shifting influence of different members on the committee's decision.

Theoretical Framework. Our empirical findings suggest a new perspective on group decision-making. Rather than being solely about aggregating their members' diverse information, committees also serve as a way to bring together members with different models to interpret the same information. We formalize this insight in a theoretical framework of group decision-making under model uncertainty and consider the associated tradeoffs.

In our framework, committee members observe the same data drawn from the unknown, true data-generating process. Each committee member interprets this data using a model that is partially misspecified and distinct from those of others. The committee's goal is to agree upon a model to use in making its policy decision. Each member advocates for a mixture of the committee members' models, trading off (1) a desire for a final model that matches the data well and (2) a cost for supporting a model that deviates from their own. The committee settles on a model as a Nash equilibrium. We evaluate the committee's performance based on how closely its chosen model relates to the unknown true model in expectation.

If the data are sufficiently informative, members' models are diversifying, and members are sufficiently willing to consider one another's perspectives, the committee outperforms its best member in expectation. There is always an optimal combination of members' models that outperforms even the best member's model. However, the committee does not select

this optimal combination because the data only provide a noisy signal of the optimal combination, and because individual members are stubborn in advocating for their own models. Hence, the committee only improves performance under the given conditions. Intriguingly, some stubbornness on the part of committee members may help improve the committee's performance, by weighing against excess sensitivity in the committee's choice of model when the data are noisy. Our framework illustrates a novel tradeoff associated with the committee structure—it brings a potential benefit of combining different members' models to better explain the data, but the potential risk of excess sensitivity to incoming data.

We also consider an extension where members face costs for dissent, motivated by empirical evidence that transparency at the FOMC—in the form of publicly releasing the committee's meeting transcripts—may have increased conformity (Meade and Stasavage, 2008; Hansen, McMahon and Prat, 2018). We find that the cost of dissent has ambiguous effects. Higher costs can promote beneficial model aggregation by encouraging members to consider alternatives. However, conformity pressure can generate multiple equilibria, creating unpredictability in committee decisions and risking convergence on poor models.

Overall, the theoretical results indicate that committee decision-making can be a useful tool under the right circumstances—committees may improve overall decision-making performance even in the absence of differential information. These results also carry practical implications for the design of committees. The FOMC's deliberation process, which combines a roundtable discussion of economic conditions followed by a policy discussion where potential disagreements are aired and suggested changes to the policy directive are made, provides clear structure that facilitates the sharing and combining of members' interpretations of the data. However, other features like increased transparency may have mixed effects by altering members' willingness to express diverse views.

Literature Review. Our paper contributes to the literature on group decision-making. Previous work has largely taken an informational perspective—that group decision-making may improve decision quality by aggregating diverse information (Gerling et al. (2005) and Sibert (2006) provide literature reviews). In contrast, we document that individual committee members often operate with different models to interpret the same data, indicating that information aggregation is not a sufficient explanation for the function of committees, consistent with the ideas voiced in Blinder (1999, 2008). We argue that a key function of committees is in bringing together members with different models, taking a related approach regarding the sharing of models to Schwartzstein and Sunderam (2021, 2024). Given that committee members' models may be misspecified, our work relates to a small literature in behavioral economics considering how group decision-making affects decision quality in light of behavioral biases (Charness and Sutter, 2012; Enke, Graeber and Oprea, 2023; Barahona et al., 2024).

Our work is also related to a long tradition in monetary policy and macroeconomics that considers how to make decisions in the face of model uncertainty. Prior work has presented

numerous approaches, for example, conservatism (Brainard, 1967), learning (Sargent, 1999; Cogley and Sargent, 2005; Primiceri, 2006), model averaging (Blinder, 1999), and robust control (Hansen and Sargent, 2008). Our work suggests that decision-making by committee is another way that policymakers address model uncertainty, and evaluates tradeoffs associated with the committee structure as a tool to deal with model uncertainty.

Additionally, our paper is related to recent work considering how the FOMC's policy decision is affected by the composition of its voters, in terms of geographic region (Bobrov, Kamdar and Ulate, 2024; Fos, Tamburelli and Xu, 2024; Fos and Xu, 2025) and policy preferences (Malmendier, Nagel and Yan, 2021; Bordo and Istrefi, 2023; Hack, Istrefi and Meier, 2023). Relative to this work, we document that the committee aggregates its members' preferences by tilting its decision towards the models of committee members that are better aligned with the data. In light of this behavior, our evidence is consistent with the findings of Hack, Istrefi and Meier (2025) that the FOMC appears to use a time-varying policy rule in making its policy decision. We link the committee's tilting behavior with monetary policy surprises measured from asset prices.[1] Accordingly, our results provide evidence that the committee structure of the FOMC may contribute to market participants' estimates of its reaction function deviating from its observed policy stance (Bauer and Swanson, 2023*a*,*b*; Bauer, Pflueger and Sunderam, 2024). This evidence complements arguments that the committee form of monetary policy-making may introduce complexities in clearly communicating the central bank's policy stance (Blinder, 2007, 2008; Vissing-Jorgensen, 2021).

More broadly, our paper relates to empirical work studying central bankers' policy preferences and their determinants (Belden, 1989; Chappell Jr, McGregor and Vermilyea, 2004; Schonhardt-Bailey, 2013; Hansen, McMahon and Rivera, 2014; Istrefi, 2019), and to work systematically analyzing FOMC transcripts (Hansen, McMahon and Prat, 2018; Cieslak and Vissing-Jorgensen, 2021; Shapiro and Wilson, 2022; Romer and Romer, 2023; Cieslak and McMahon, 2023; Cieslak et al., 2023; Bordo, Istrefi and Martínez, 2024). Relative to this literature, we construct a set of granular measures of individual committee members' policy preferences for each meeting across various dimensions over a long sample, and we explicitly tie these preferences in with the macroeconomic frameworks employed by committee members to interpret data. In this regard, our paper is also related to a small but growing literature that seeks to identify the mental models used by economic agents in macroeconomic and financial settings (Andre et al., 2022; Andre, Schirmer and Wohlfart, 2023; Bybee, 2023; Binetti, Nuzzi and Stantcheva, 2024; Ke, 2024; Bastianello, Décaire and Guenzel, 2024). Importantly, in addition to characterizing committee members' mental models, we also tie such models into the FOMC's policy decisions.

[1] This result relates to Handlan (2022) and Gáti and Handlan (2023), who link the text of the FOMC statement to monetary policy surprises.

2 Background and Data

To proceed with our empirical analysis, we construct several measures of FOMC committee members' policy preferences from FOMC transcripts. We summarize relevant institutional details and data sources and outline the construction of our measures.

2.1 Background

Our sample consists of regularly scheduled FOMC meetings from March 1976 through December 2019. The FOMC meets eight times per year to formulate monetary policy. The committee consists of 12 voting members: seven Governors of the Federal Reserve Board in Washington, DC, including a Chairperson; the President of the New York Fed; and four members selected on a rotating basis from the presidents of the eleven other regional Federal Reserve banks. In addition to the voting members, the non-voting Regional Fed presidents and numerous Federal Reserve staff members attend and actively participate in the meetings.

Since 1993, transcripts of the FOMC meetings have been released to the public on a five-year lag, and are available from the Federal Reserve website. These transcripts provide nearly complete accounts of every FOMC meeting from 1976 onwards, including verbatim statements and discussions by committee members.[2] We discuss the processing of the transcript data in Appendix B.1. In addition to the transcripts, the Federal Reserve also publishes materials associated with the meeting, which we use as part of our analysis. In particular, prior to each meeting, the Federal Reserve Board staff prepare Tealbook A (formerly the Greenbook), which provides discussion and forecasts of macroeconomic and financial conditions. The report is available on the Federal Reserve Board website and the forecasts contained within the reports are published on the Philadelphia Federal Reserve website. The staff of the Federal Reserve Board also prepares and distributes Tealbook B (formerly the Bluebook) prior to the meeting, which we also obtain from the Federal Reserve website. The Tealbook B includes a set of monetary policy alternatives and associated background, which serve as a baseline for the policy discussion and decision made by the committee. The committee is not restricted to choosing one of the alternatives and often combines elements from the different alternatives in their final decision.

FOMC meetings have two parts that are related to the committee's monetary policy decision: an economic and financial situation discussion and a monetary policy discussion.[3] Following Hansen, McMahon and Prat (2018), we focus our analysis on these two parts of the transcripts, and in particular on the statements by Federal Reserve governors and regional

[2] Prior to 1976, summaries of the meetings are provided as 'Memoranda of Discussion.'.

[3] In addition to these two parts, FOMC meetings typically feature presentations about open market operations, and often also include reports and discussions of special topics, such as long-run policy implementation frameworks and balance sheet normalization principles.

Federal Reserve presidents in these sections. The discussion of the economic and financial situation is led off by a presentation by members of the Federal Reserve Board staff, which is followed by questions, prepared statements by members of the committee, and a general discussion of the economic and financial situation. The subsequent monetary policy discussion is led off by the Federal Reserve Board's director of monetary affairs presenting the monetary policy alternatives, which is followed by a discussion among committee members about their policy preferences. Following this discussion, the committee votes on the proposed policy directive. Beginning in 1994, the committee began announcing the outcome of its meeting at the end of the meeting via a public statement; a substantial amount of the discussion in the monetary policy segment of the meeting following this date is with regard to the exact wording of this statement.

2.2 Extracting Information from Transcripts with Large Language Models

To proceed with our analysis, we generate a number of our empirical measures of interest, discussed in detail below, by analyzing subsections of the transcript texts using the 20241022 release of the Anthropic Claude 3.5 Sonnet large language model ("the LLM"). In generating our prompts to extract information from the transcript texts, we particularly make use of 'chain-of-thought' prompting, wherein the LLM is prompted to engage in various stages of intermediate reasoning before delivering a final answer to a query.[4] Work on optimizing the performance of large language models has found chain-of-thought prompting to significantly improve performance in complex reasoning tasks (Wei et al., 2022), which, in turn, has led it to serve as the basis for the training of the most recent generation of complex reasoning models that rival PhD level performance on math and science problems (as of this writing).[5] We discuss our empirical measures and prompts below, and present the full text of the prompts in Appendix B.2.

[4] Dunn et al. (2024) investigate the ability of various large language models in identifying topics discussed in FOMC meeting minutes. They find strong performance from the 20240620 release of the Claude 3.5 Sonnet model. To generate chain-of-thought prompts, we make particular use of the Anthropic prompt improver tool, discussed at https://www.anthropic.com/news/prompt-improver.

[5] E.g., OpenAI's o1 and o3 models, see https://openai.com/index/learning-to-reason-with-llms/.

Date	Speaker	Variable, Score	Quotes
Feb-88	Alan Greenspan	Inflation, +1	"[due to] "fears of inflation ... we have to make certain that, if we choose to be a little accommodative in the short run, we don't send the message that that's our longer-term view"
Feb-88	Alan Greenspan	Credit, -2	"if there is a problem out there, it's in the financial area... the markets need more time for healing–meaning they need some general softness in the short end of the money market."
Feb-88	Alan Greenspan	Stocks, -2	"I'm basically concerned... that if we were to indicate that we were tightening, the shock to the markets I think would break the stock market and create some real problems."
Aug-08	Richard Fisher	Inflation, +3	"Inflation expectations are, I think, becoming unmoored, and I believe they are, if not already adrift, at risk of drifting... The percentage of small business owners citing inflation as 'the number one problem they face,' rose to 20 percent in June, the highest rating since 1982."
Aug-08	Richard Fisher	Growth, -1	"Nationally I expect continued anemia on the growth side. We are, as you know, not at the extreme but at the low end of expectations for economic growth. I would expect zero growth as we approach year-end."
Jun-17	Neel Kashkari	Inflation, -3	"Core 12-month PCE inflation of 1.5 percent, in my view, is a significant shortfall below our 2 percent target."
Jun-17	Neel Kashkari	Labor, -1	"The labor market is clearly tighter than it was. But as you've indicated, this is not showing up in nominal wage growth."

TABLE 1: MEASURING POLICY PREFERENCES: EXAMPLES WITH LLM-IDENTIFIED QUOTES

Note: The table presents examples of the policy preference measures estimated for each speaker in each meeting for inflation, growth, labor markets, credit markets, and the stock market. For each meeting, members are assigned scores from -3 to +3 for each variable, with -3 indicating that a variable strongly contributes to a preference for more dovish policy and +3 indicating that a variable strongly contributes to a preference for more hawkish policy. The scores are constructed by providing an LLM with a committee member's speeches and prompting it to provide scores for each variable as well as direct quotes that support its assessments.

2.2.1 Measuring Committee Members' Policy Preferences

For each meeting in our sample, we construct measures of each individual committee member's policy preferences pertaining to each of several macroeconomic and financial market dimensions: inflation, growth/output, employment, credit market conditions, and the stock market. The result, for a given speaker in a given meeting, is a score from -3 to 3 for each of the five dimensions, where the magnitude of the score indicates the importance of the variable for their policy preference, and the sign indicates whether the variable contributes dovishly (-) or hawkishly (+), meaning supporting more accommodative or more restrictive policy. The measures capture a combination of the relative weight the member gives to a variable in their policy preferences (e.g., a coefficient in a forward-looking Taylor (1993)-style rule), as well as the member's forecast of that variable.

To construct the measures, we take all of a member's statements during the economic/-financial situation and policy discussion sections of a meeting, pass them into the LLM, and prompt the LLM to provide a score from -3 to 3 for that dimension, where -3 indicates that the variable contributes strongly to a member's preferences in the direction of a bigger rate cut (or smaller hike), 0 indicates the variable does not contribute to their policy preference, and +3 indicates that the variable contributes strongly in the direction of a bigger rate increase (or smaller cut). To ensure chain-of-thought reasoning in formulating these scores, we also provide the LLM with the set of monetary policy alternatives discussed during the meeting (extracted from the Tealbook B/Bluebook), and prompt it to formulate which alternative was preferred by the speaker; we also prompt the LLM to justify its assigned score using direct quotes from the speaker. The LLM provides its initial reasoning in text form before making a final determination of its output.

Table 1 provides examples of scores across variables and the quotes identified by the LLM to justify the scores. The first three rows present scores and justifying quotes for chairman Alan Greenspan from February 1988, in a meeting where he preferred slightly accommodative policy due to financial market concerns, but was also concerned about credibility with respect to inflation. The committee kept rates at the same level, slightly eased reserve targets, and set conservative targets on money supply expansion, with the goal of easing financial conditions while keeping inflation in check. The next two rows present scores and justifying quotes for Richard Fisher, the president of the Federal Reserve Bank of Dallas, from the FOMC meeting in August 2008. The committee kept the target Fed funds rate at 2% during the meeting, balancing concerns about poor credit market conditions and a poor growth outlook with inflationary concerns, while Fisher dissented, favoring a higher rate due to concerns about inflation. The last two rows present scores and justifying quotes for Neel Kashkari in June 2017; the committee moved to raise the target Fed funds rate by 25 bps, with Kashkari dissenting in advocacy of not moving the target interest rate due to concerns of low inflation.

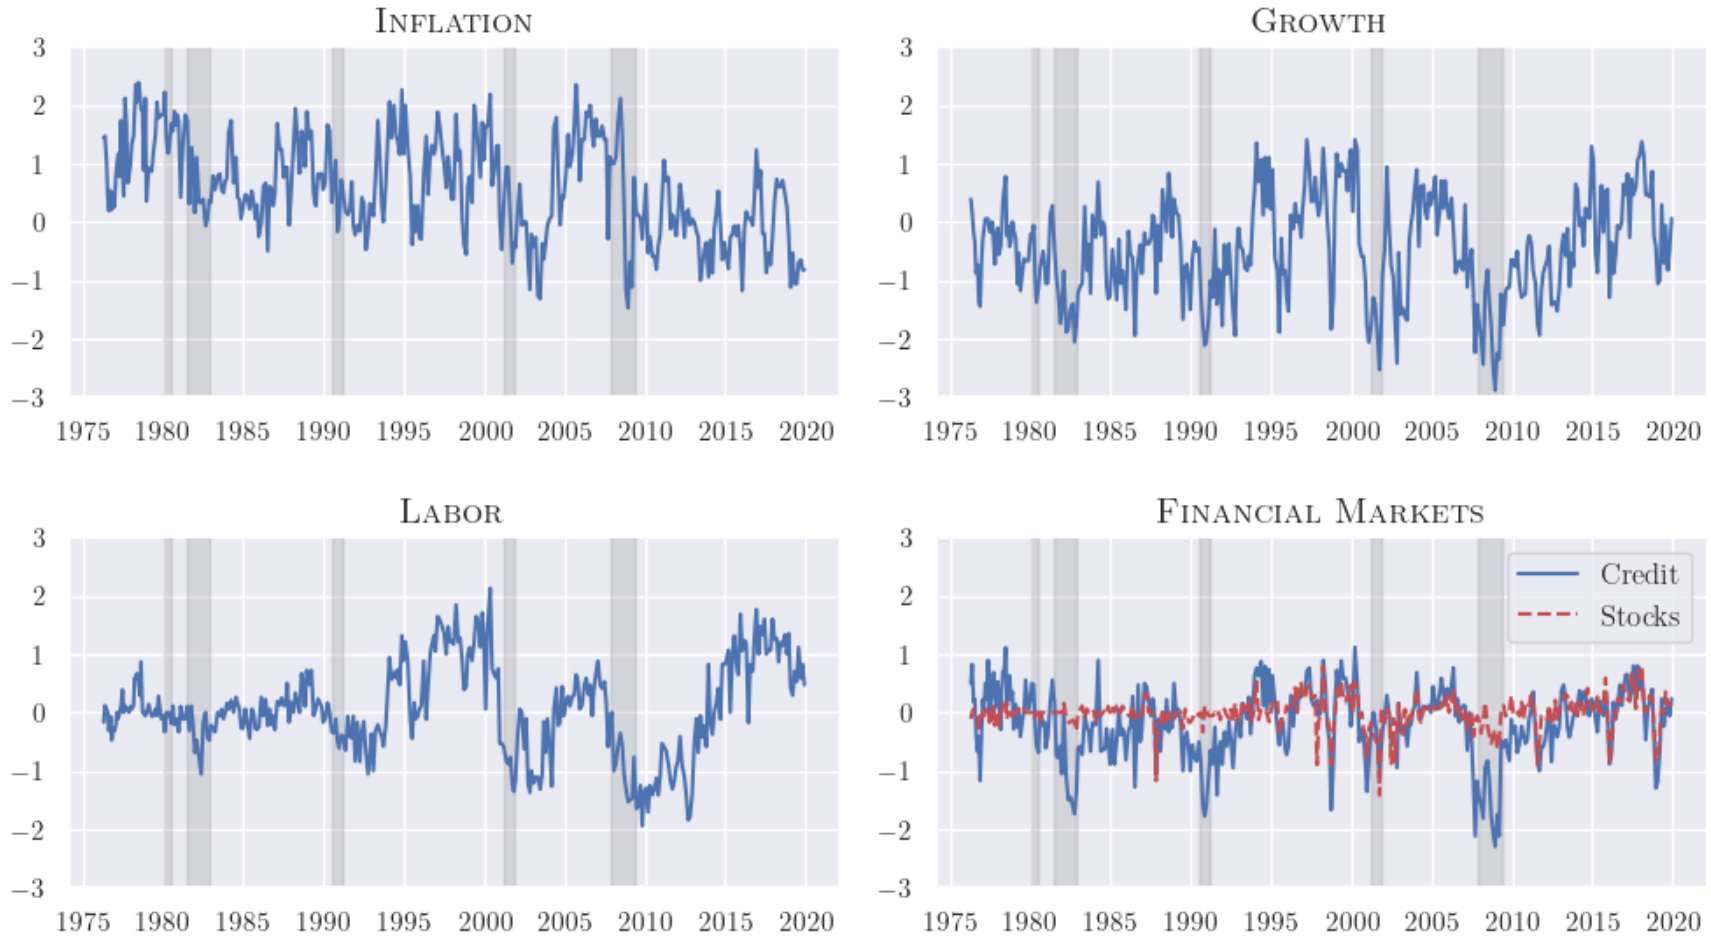


FIGURE 1: COMMITTEE MEMBERS' POLICY PREFERENCES OVER TIME

Note: The figure plots the average of committee members along the inflation, growth, labor, credit, and stock market dimensions for each FOMC meeting. For each meeting, members are assigned scores from -3 to +3 for each variable, with -3 indicating that a variable contributes strongly to a preference for more dovish policy and +3 indicating that a variable contributes to a preference for more hawkish policy. The dark grey bands indicate NBER recession periods.

Figure 1 plots the average of committee members' estimated policy preferences in each meeting for each variable. The measures display a number of intuitive properties, showing cyclical behavior and attaining troughs during NBER recessions (identified with dark grey bands). Concerns about inflation, all else equal, push towards more hawkish preferences; these are elevated in the 1970s and late 1990s, and are low in the years following the Great Financial Crisis. On the financial side, concerns about credit markets display cyclical properties, with elevated concerns during the Great Financial Crisis. The stock market tends not to be an important variable in the early part of the sample, with the exception of the stock market crash of 1987. Starting in the mid-1990s, the stock market increases in importance, with some hawkishness arising in the late 1990s (corresponding with Alan Greenspan's concerns about irrational exuberance, expressed in December 1996), and dovish considerations in periods of poor stock market performance in subsequent years corresponding with the "Fed Put" (Cieslak and Vissing-Jorgensen, 2021).[6]

We also validate policy preference measures' ability to capture cross-sectional information about committee members' policy stances by comparing them with committee members' forecasts of inflation, GDP growth, and unemployment. These forecasts are available from the Summary of Economic Projections, which has been conducted prior to four FOMC meetings per year since October 2007. Figure 2 sorts committee members into deciles based on their cross-sectionally demeaned GDP, inflation, and unemployment forecasts for the year of the meeting, and plots the average of these forecasts versus the average cross-sectionally demeaned transcript-measured growth-, inflation-, and labor-driven policy preference measures within each decile. The figure indicates that the policy preferences display a strong cross-sectional relationship with committee members' forecasts. This relationship is despite the fact that the comparison is not an apples-to-apples comparison; our policy preference measures combine *forecasts* of variables with the *weight* those variables have in committee members' policy preferences. Nevertheless, we observe a strong relationship between such preferences and forecasts, indicating the importance of committee members' beliefs in driving cross-sectional variation in their policy stances.

[6]The importance of financial variables for policy preferences may also be broadly consistent with some form of 'financial conditions targeting' by the Fed, e.g., as studied by Caballero, Caravello and Simsek (2024).

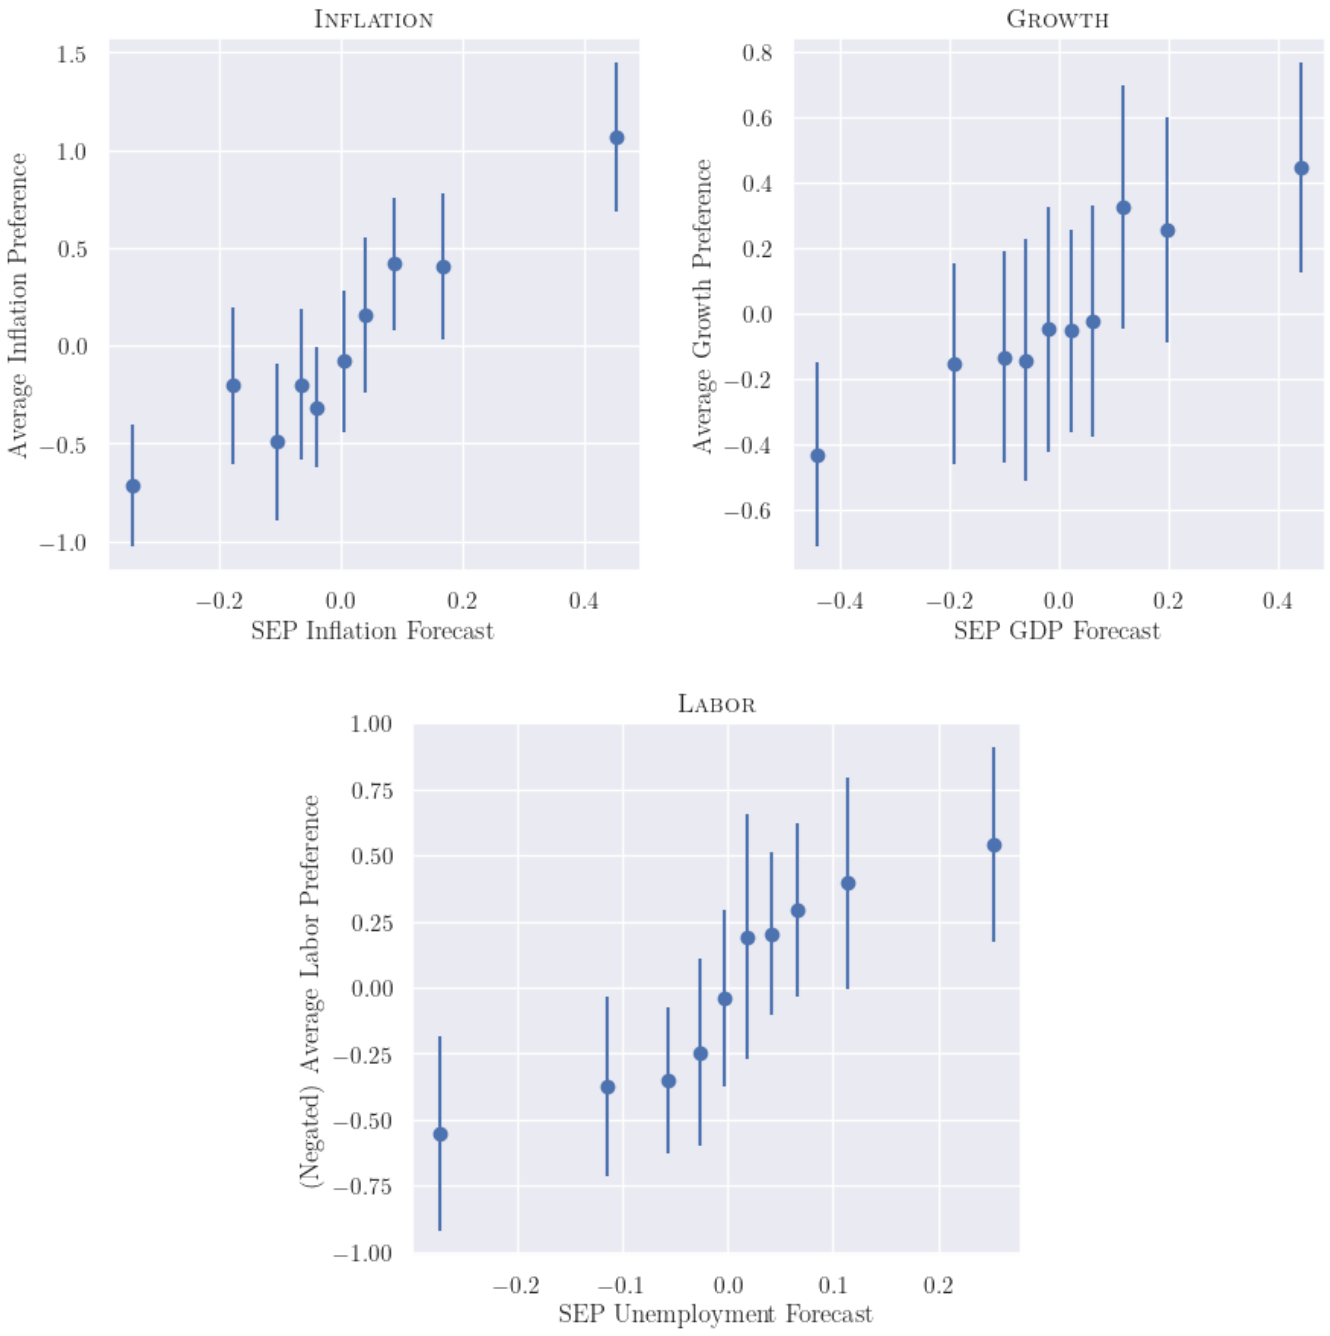


FIGURE 2: COMMITTEE MEMBERS' POLICY PREFERENCES AND SEP FORECASTS

Note: The figure sorts committee members into deciles based on their cross-sectionally demeaned GDP, inflation, and unemployment forecasts reported for the year of the meeting, and plots the average of these forecasts versus the average cross-sectionally demeaned transcript-measured growth, inflation, and labor-based policy preference within each decile. Forecasts are from the Summary of Economic Projections, conducted prior to four meetings per year since October 2007. The figure also plots plus and minus two standard error bars for each decile.

2.2.2 Alignment with the Policy Decision

For each meeting, we also construct measures of each member's alignment with the committee's policy decision. The result is a score for each speaker in each meeting ranging from -3 to 3, where 3 corresponds to perfect alignment with the decision made by the committee, and -3 corresponds to complete opposition to the decision made by the committee. Alignment with the policy decision captures both alignment with the target interest rate settled on by the committee, as well as with other decisions, for example, targets for monetary aggregates, an important aspect of policy decisions early in our sample; unconventional monetary policies, such as quantitative easing; and the inclusion or exclusion of specific words or phrases in the committee's official policy statement (e.g., as part of forward guidance). The latter two types of policies are important in the later part of our sample.

To construct our measures of alignment, we first construct the set of changes to the policy directive that occurred during each meeting. To do so for a given meeting, we provide the LLM with the policy discussion section of the meeting, and the policy alternatives from Tealbook B (or the Bluebook). We prompt the LLM to identify all of the changes to the policy directive that were adopted by the committee, including those presented in the policy alternatives, as well as those proposed and adopted in an *ad hoc* manner during the meeting. We also prompt the LLM to evaluate the hawkish and dovish tenor of each of the changes by giving each a score from -3 (highly accommodative/dovish) to +3 (highly restrictive/hawkish); we make use of these scores when studying monetary policy surprises in our main analysis.

For a given member in a given meeting, we run an additional prompt, providing the LLM with the policy discussion, the policy alternatives, and the identified changes to the policy directive for that meeting. We prompt the LLM to provide a score from -3 to 3 to capture each committee member's alignment with each of the proposed changes to the directive, as well as a score from -3 to 3 to capture the speaker's overall alignment with the policy decision of the committee. Using the resulting output, we construct two measures of policy alignment: *DirectiveAlignment*, which is simply the overall alignment score provided by the LLM for the committee member in the meeting; and *ChangeAlignment*, which is the average alignment with each of the changes to the directive made during the meeting. The two measures have a correlation of 0.91 across observations in our sample.

Table 2 provides examples to illustrate the construction of these variables. Panel A presents the policy changes and associated scores from the FOMC meeting in September 2011. The LLM identifies four changes during the meeting: a maturity-extension program to sell short-term securities and buy long-term securities, a change in reinvestment policy for MBS and agency debt, forward guidance about low interest rates, and modified statement language about inflation and global markets. The LLM identifies the first change as having a score of -2, representing substantial accommodation due to its intended effect on portfolio rebalancing, while identifying the other changes as more modestly dovish, with scores of -1.

PANEL A: EXAMPLE POLICY DIRECTIVE CHANGES

Date	Change	Score	Justifications
Sep-11	Implementation of $400 billion maturity extension program (MEP) to purchase longer-term Treasury securities and sell shorter-term securities.	-2	Represents significant accommodation through portfolio rebalancing, though less aggressive than QE due to neutral impact on balance sheet size.
Sep-11	Change in reinvestment policy to direct agency debt and MBS principal payments into agency MBS rather than Treasury securities.	-1	Support for mortgage market while maintaining existing balance sheet size. More modest than MEP but signals continued support for housing sector.
Sep-11	Maintained forward guidance on federal funds rate staying exceptionally low through mid-2013.	-1	Maintained rather than strengthened previous stance, hence moderate classification.
Sep-11	Modified statement language regarding inflation to "appears to have moderated" and added reference to "strains in global financial markets."	-1	Signals Committee's assessment of reduced inflation risks and acknowledgment of downside risks, providing implicit dovish bias while maintaining flexibility.

PANEL B: EXAMPLE SPEAKER ALIGNMENT SCORES

Date	Change	Speaker	Alignment	Justification
Sep-11	Maturity Extension	Charles Plosser	-3	He argued that current actions would be ineffective and
Sep-11	MBS Reinvestment	Charles Plosser	-2	potentially harmful, expressing particular concern about
Sep-11	Forward Guidance	Charles Plosser	-2	inflation risks and the limitations of monetary policy tools. His
Sep-11	Statement Language	Charles Plosser	-1	acceptance of some technical aspects of the statement language
Sep-11	*DirectiveAlignment*	Charles Plosser	-2	prevents a -3 rating.
Sep-11	Maturity Extension	Ben Bernanke	3	He explicitly recommended and supported all implemented
Sep-11	MBS Reinvestment	Ben Bernanke	3	changes. He clearly articulated the rationale for the maturity
Sep-11	Forward Guidance	Ben Bernanke	3	extension program, supported the MBS reinvestment strategy,
Sep-11	Statement Language	Ben Bernanke	3	and agreed with the statement language modifications.
Sep-11	*DirectiveAlignment*	Ben Bernanke	3	

TABLE 2: POLICY CHANGES AND POLICY ALIGNMENT: EXAMPLES WITH LLM JUSTIFICATIONS

Note: Panel A of the table displays examples of identified changes to the policy directive and the associated scores from -3 (highly accommodative/dovish) to +3 (highly restrictive/hawkish) from the September 2011 meeting. Panel B of the table displays examples of committee members' measured alignment with the proposed changes, which are averaged to construct the *ChangeAlignment* measure for each speaker. Panel B also displays the *DirectiveAlignment* measure, which captures a member's overall alignment with the decision. Scores of -3 indicate complete opposition, and scores of +3 indicate complete alignment. All measures are constructed by prompting an LLM after providing the policy discussion section of the FOMC meetings and the Tealbook B (Bluebook). The Justification columns report paraphrased output from the LLM justifying the given scores.

	DirectiveAlignment						
	-3.0	-2.0	-1.0	0.0	1.0	2.0	3.0
Average *ChangeAlignment*	-2.17	-0.90	-0.23	0.16	1.00	2.00	2.80
Percentage dissenting votes	90.0%	85.5%	32.2%	12.5%	5.6%	0.4%	0.0%
Number of voters	10	220	143	16	303	1381	1735

TABLE 3: DECISION ALIGNMENT AND DISSENT

Note: The table groups each voting committee member during an FOMC meeting based on their *DirectiveAlignment* score, where *DirectiveAlignment* ranges from -3 to 3, with -3 indicating strong opposition to the policy decision made by the committee and +3 indicates perfect agreement with the policy made by the committee. For each Alignment score, the table reports the average *ChangeAlignment*, where *ChangeAlignment* is the average alignment (from -3 to 3) of speakers for each of successful proposed changes to the policy directive identified during the meeting; and the percentage of voters with that *DirectiveAlignment* score that dissented. The table also reports the total number of voting FOMC members with each of the *DirectiveAlignment* scores in our sample.

Panel B of Table 2 shows examples of two committee members' alignment with the policy changes made in the September 2011 meeting. Charles Plosser, the President of the Federal Reserve Bank of Philadelphia, cast a dissenting vote for the committee's decision in that meeting. The LLM assigns him a score of -3 for the maturity extension program, a score of -1 for statement language changes, and scores of -2 for the remaining changes, with an overall *DirectiveAlignment* score of -2. The justification provided reflects Plosser's strong disagreement with the efficacy of the tools being used by the committee, but also indicates a modulated score due to his agreement with some of the language changes in the statement. In contrast, for the chairman Ben Bernanke, the LLM assigns scores of +3 for all of the individual changes and overall, citing his explicit advocacy and support for the changes.

We compare our measures of policy alignment with dissenting votes for the committee's policy decision.[7] We obtain voting data for FOMC meetings from Thornton, Wheelock et al. (2014). Table 3 groups each voting committee member in an FOMC meeting based on their *DirectiveAlignment* score, and reports the average *ChangeAlignment* score, percentage of dissents, and number of voters corresponding with each *DirectiveAlignment* score. The table reveals a strong relationship between the *DirectiveAlignment*, *ChangeAlignment*, and dissents. Of the voting members of the committee with *DirectiveAlignment* scores of -3.0, -2.0 and -1.0, indicating disagreement with the committee's decision, 90%, 85.5%, and 32.2% of them dissent, respectively. Meanwhile, only a negligible proportion of voters with positive *DirectiveAlignment* scores dissent in their voting.

[7]Work studying dissenting behavior includes Belden (1989); Chappell Jr, McGregor and Vermilyea (2004); Meade and Sheets (2005); Riboni and Ruge-Murcia (2014); Bobrov, Kamdar and Ulate (2024); Fos and Xu (2025).

As noted in prior work, dissents are uncommon, accounting for only 6.9% of votes in our sample. This tendency against dissenting is exemplified by a remark from Alan Greenspan during the policy discussion of the February 1994 FOMC meeting:

> "I would request that, if we can, we act unanimously. It is a very potent message out in the various communities with which we deal if we stand together. If we are going to get a split in the vote, I think it will create a problem for us..."

While this remark was made in the context of a meeting during which the committee wanted to hike rates and the chairman was concerned about financial market turmoil—a meeting where voting consensus may have been particularly important—Greenspan's statement underscores the strategic importance of unanimity in policy decisions, reflecting a broader institutional preference for consensus due to concerns about how the committee's policy will be received, e.g., by financial markets (Madeira and Madeira, 2019). Additionally, career concerns may also play an important role in the lack of observed dissent in voting (Meade and Stasavage, 2008). Our *DirectiveAlignment* and *ChangeAlignment* measures provide the advantage of capturing alignment in addition to misalignment with the policy decision, and of providing more granular measurement than voting dissent.

3 Committee Members' Mental Models and Preferences

With our main empirical measures in hand, we next turn to analyzing them. We document that committee members appear to operate with largely fixed and distinct mental models—ways of interpreting economic data and forming forecasts.

3.1 Committee Members' Policy Preferences

We begin by examining committee members' preferences. Table 4 reports results from regressions of the meeting-by-member policy preference measures for inflation, growth, labor, credit, and stock markets on meeting fixed effects and committee member fixed effects. Unsurprisingly, meeting fixed effects play an important role in explaining committee members' policy preferences, with R^2 values ranging from 0.31 to 0.36 for macroeconomic variables, and ranging from 0.17 to 0.20 for the financial market variables. These capture the role that macroeconomic fluctuations play in driving a common component of members' policy preferences (e.g., a tendency for all members of the committee to be more accommodative in periods of low growth and high unemployment).

However, committee member fixed effects also play an important role, with R^2 values ranging from 0.20 to 0.31 for macroeconomic variables, and from 0.07 to 0.22 for financial market variables. When taking the average across measures, meeting fixed effects explain 30% of preference variation and member fixed effects explain 29%. This finding echoes evidence

	Regression R^2					
	Inflation	Growth	Labor	Credit	Stocks	Average
Meeting FE	0.31	0.35	0.36	0.20	0.17	0.30
Member FE	0.31	0.20	0.23	0.22	0.07	0.29
Meeting + Member FE	0.53	0.51	0.53	0.40	0.24	0.56
N	6369	6369	6369	6369	6369	6369

TABLE 4: COMMITTEE MEMBERS' POLICY PREFERENCES: FIXED EFFECTS

Note: The table reports R^2 values from regressions of the inflation, growth, labor, credit, and stock market policy preference measures, and the average across measures, on meeting fixed effects and FOMC committee member fixed effects. The policy preference measures range from -3 to 3, with the signs of the measures indicating whether the variable induces a hawkish or dovish preference, and the magnitude of the measures capturing their importance for the committee member.

from other settings—such as in the stock market (Giglio et al., 2021)—regarding the importance of individual fixed effects for beliefs. The importance of member fixed effects indicates that committee members exhibit persistent policy preferences across various dimensions. For example, one committee member may consistently display more hawkish preferences due to concerns about inflation, while another may persistently exhibit more dovish preferences due to concerns about growth. Jointly, meeting and member fixed effects explain over half of the variation in members' macro variable-driven policy preferences, 24% to 40% of the variation in their financial variable-driven policy preferences, and 56% of the variation in their overall policy preferences.

Next, we explore how members' preferences are related across different variables. Table 5 reports the principal components of the policy preference measures. The first three columns of the table report principal components estimated from member-by-meeting observations. The first principal component loads positively on each of the variables, and explains 63% of the variance. These positive loadings are consistent with the first principal component representing a hawk-versus-dove dimension, where committee members and meetings that are comparatively hawkish on one dimension (e.g., inflation) also tend to be hawkish on other dimensions (e.g., growth and unemployment). The second and third principal components capture differential loadings on various variables, such as hawkishness based on inflation versus dovishness based on growth and credit market conditions (the second principal component).

The last three columns of the table report principal components from a matrix of member-level fixed effects for each variable, estimated after controlling for meeting fixed effects.[8]

[8]By simultaneously estimating both member and meeting fixed effects, our member fixed effects also account for

Comparison of Principal Components						
	Full Data			Member FE		
Principal Component Loadings	PC1	PC2	PC3	PC1	PC2	PC3
Growth	0.55	-0.40	0.13	0.52	-0.07	-0.56
Inflation	0.53	0.82	-0.20	0.55	0.25	0.76
Labor	0.42	-0.03	0.77	0.37	0.69	-0.32
Credit	0.47	-0.40	-0.59	0.52	-0.68	0.02
Stocks	0.12	-0.10	-0.09	0.12	-0.04	-0.12
Proportion of Variance Explained	0.63	0.14	0.12	0.84	0.07	0.04

TABLE 5: PRINCIPAL COMPONENTS OF COMMITTEE MEMBERS' POLICY PREFERENCES

Note: The table reports a comparison of the first three principal components from two different analyses of the inflation, growth, labor, credit, and stock market policy preference measures. The first three columns (Full Data) show the results where observations are FOMC committee member by meeting preferences. The last three columns (Member FE) displays results where observations are committee member fixed effects for each variable, estimated from regressions of the preference measures for the variables on member and meeting fixed effects.

Once again, the first principal component loads positively on all variables, capturing a hawk-versus-dove dimension and explaining 84% of the variation. Because these principal components are estimated after controlling for meeting fixed effects, the evidence indicates that a committee member who tends to be persistently hawkish about inflation also tends to be persistently hawkish about growth, employment, credit market conditions, and the stock market. The subsequent principal components capture differential loadings on variables, such as dovishness based on credit conditions versus hawkishness based on labor market conditions and inflation (the second principal component).

The principal component analysis reveals that most of the variation in the policy preference measures is explained by a hawk-versus-dove dimension. Without controlling for meeting fixed effects, committee members' policy preferences reflect the dominant macroeconomic or financial variables at the time; for example, committee members may be more concerned with inflation during inflationary periods or credit conditions more during periods of financial turmoil. However, after controlling for meeting fixed effects, the lion's share of the cross-sectional variation in policy preferences is explained by heterogeneity across members along the hawk-versus-dove dimension.[9]

composition effects on the committee when estimating an individual committee member's preference tilt. i.e., the committee might include more hawkishly oriented members in some periods and more dovishly oriented members in other periods. The separation of composition effects via regression fixed effects is due to the fact that committee members' tenures on the committee are staggered.

[9]Related, Herbst and Winkler (2021), Scanlan (2023), and Wu (2023) document evidence of a factor structure in macroeconomic forecasters' predictions across different variables.

The dominance of a single hawk–dove factor in explaining cross-sectional variation in committee members' policy preferences as related to different variables, and the close link between those preferences and committee members' macroeconomic forecasts (Figure 2) are consistent with the idea that different committee members employ distinct mental models of the macroeconomy. In particular, the evidence suggests that a given committee member may filter information and generate forecasts for different variables using a unified mental model; the use of different mental models by different members leads to persistently higher or lower forecasts across variables, which in turn gives rise to strong fixed effects in policy preferences.

3.2 Policy Preferences and Underlying Frameworks

In this section we establish a connection between committee members' policy preferences and the economic frameworks they share to interpret economic data during the committee's deliberations. We find that committee members' policy preferences are well explained by a small number of economic frameworks, with about five frameworks having the same explanatory power as fixed effects for each variable that we study.

We begin by constructing taxonomies of economic frameworks within different themes—inflation, growth, labor, and financial conditions—identifying the key parameters that distinguish different frameworks. For instance, in the inflation theme, we distinguish between Phillips Curve models emphasizing economic slack, monetarist models focused on money growth, cost-push models centered on supply shocks, and expectations-centric models. To map members' statements to these frameworks, we prompt the LLM to evaluate the statements of each member during each meeting against diagnostic criteria of the relevant taxonomy.[10] The LLM assigns probability scores indicating the presence of each framework in a given member's statements and estimates key model-specific parameter values when sufficient evidence exists. This procedure generates a vector of framework-parameter pairs for each member for each meeting, capturing the interpretations that each member expressed during the meeting. For example, we document whether an individual discusses the unemployment-inflation tradeoff in terms of a steep or flat Phillips Curve, inflation expectations as well-anchored or adaptive, and cost-push shocks as transitory or persistent.

[10] Section B.5 in the Appendix provides the complete taxonomies and details on the measure construction.

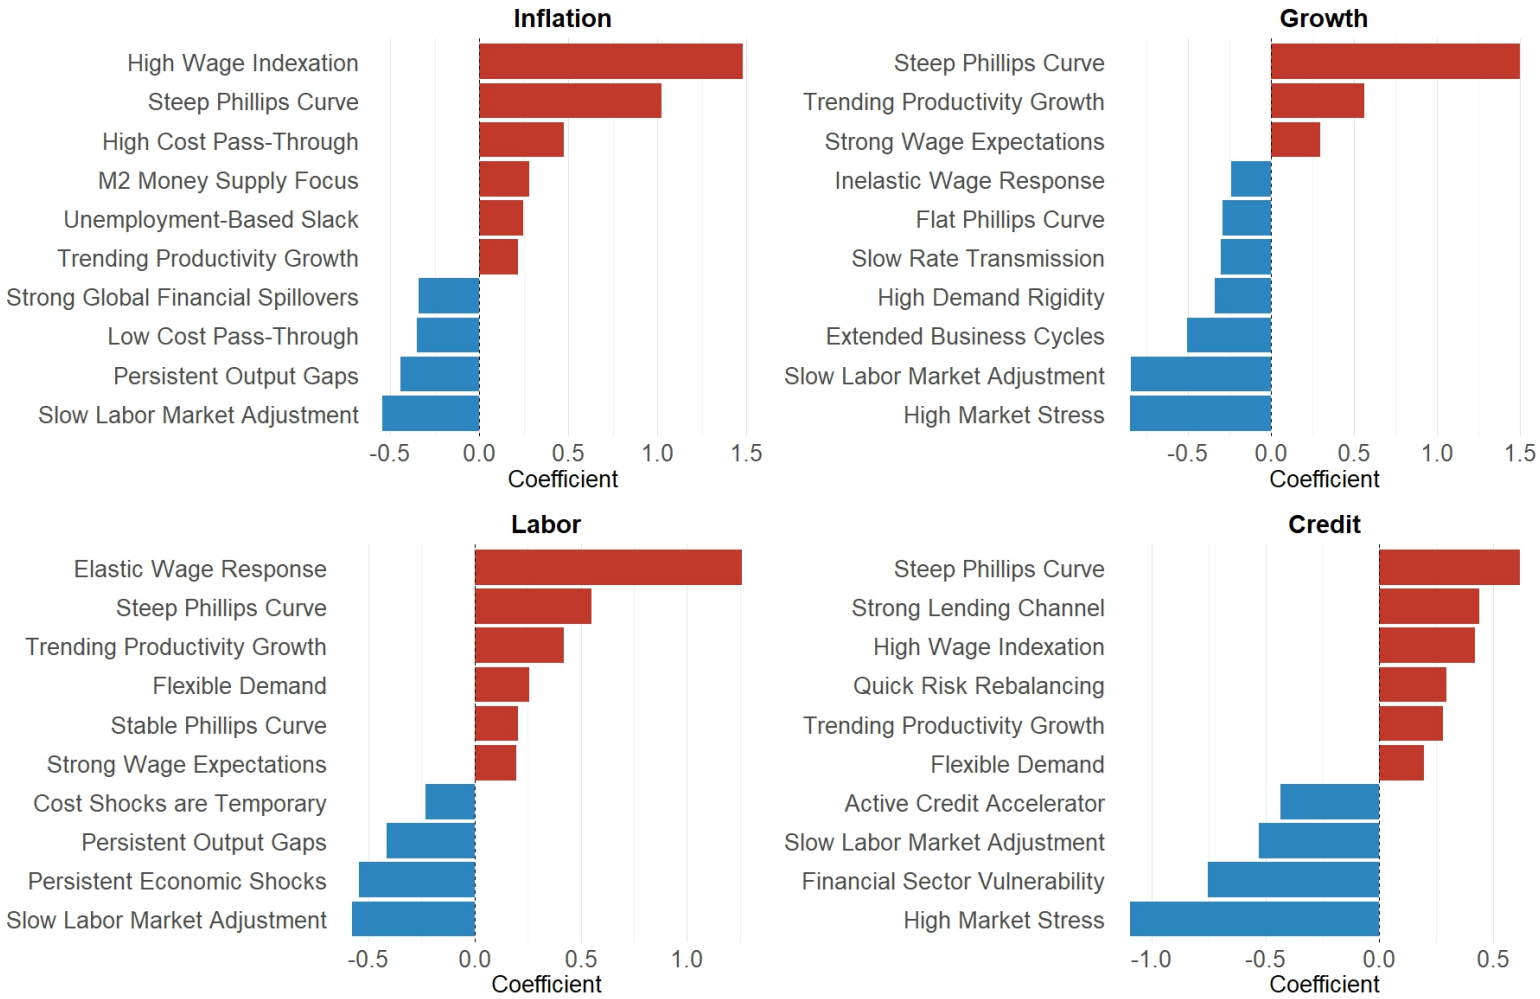


FIGURE 3: POLICY PREFERENCES AND ECONOMIC FRAMEWORKS

Note: The figure presents the relationship between committee members' economic frameworks and their policy preferences across four dimensions. Each panel corresponds to a different policy domain: inflation, growth, labor markets, and credit conditions. The horizontal bars represent the coefficients from our forward selection procedure, indicating the association between specific economic model parameters and policy positions. Red (positive) coefficients indicate parameters associated with more hawkish policy positions, while blue (negative) coefficients are linked to more dovish stances. The policy preference scale ranges from -3 (most dovish) to 3 (most hawkish).

We examine how these interpretations of economic data map to policy preferences. We use a feature selection approach to isolate the interpretations most strongly associated with members' policy preferences. We focus on parameter estimates with at least 30 occurrences to ensure reliable inference. We construct an 0/1 indicator variable for each framework-parameter pair, indicating the presence of that framework-parameter pair in each member's statements during each meeting. Then, for each policy dimension (inflation, growth, labor markets, and credit), we begin with an intercept-only model and iteratively incorporate the framework-parameter pair that produces the largest improvement in Bayesian Information Criterion (BIC). We conclude when the maximum improvement from adding an additional framework-parameter pair is lower than 2 units.

Figure 3 presents the relationships between the presence of different frameworks in a committee members' statements in a given meeting and their measured policy preferences across four variables: inflation, growth, labor markets, and credit conditions. For each variable, we plot the estimated coefficients from our forward selection procedure, where each coefficient represents the relationship between a specific parameter of an economic framework and the hawkish-dovish policy preference scale. We focus on the ten most influential parameters from the selection process. Positive coefficients (shown in red) indicate parameters associated with more hawkish policy positions, while negative coefficients (shown in blue) reflect parameters linked to more dovish stances. The parameters are ordered by coefficient magnitude to highlight the strength of these relationships.[11]

For inflation policy preferences (panel A), we find that hawkish preferences are associated with views that wages are strongly indexed to inflation, discussion of a steep Phillips Curve slope, and a high pass through of cost-push shocks. In contrast, dovish inflation policy preferences are associated with interpretations such as a persistent output gap. The growth preference results (panel B) show parallels with inflation preferences. Interpreting economic data through a steep Phillips Curve framework is associated with hawkish growth policy preferences. Discussion of high financial market stress is associated with more dovish growth preferences. Labor market preferences (panel C) show a similar pattern but with a strong relationship between hawkish preferences and views that labor markets can withstand tighter policy when wages respond elastically to economic conditions. Finally, credit condition preferences (panel D) are most influenced by financial market functioning assessments, with discussion about stressed markets strongly associated with dovish credit policy preferences. The coherent cross-domain patterns—particularly for Phillips Curve beliefs—suggest a link between the common component hawk-dove structure of policy preferences and the use of particular frameworks.

Table 6 provides the detailed progression of R^2 values as parameters are sequentially

[11] Section B.6 in the Appendix provides a more complete visual overview of these frameworks, parameters, and key values we identify, and the links to our policy preference scores.

Parameters	Inflation	Growth	Labor	Credit
1	0.15	0.13	0.18	0.100
2	0.20	0.21	0.23	0.14
3	0.23	0.26	0.25	0.16
4	0.25	0.29	0.27	0.18
5	0.26	0.31	0.29	0.20
10	0.30	0.34	0.32	0.23
Max	0.33	0.37	0.34	0.24
Member FE	0.31	0.20	0.23	0.22
Max No.	29	21	21	18

TABLE 6: INCREMENTAL EXPLANATORY POWER OF ECONOMIC FRAMEWORK PARAMETERS

Note: This table reports the cumulative R^2 values as economic framework parameters are sequentially added to models predicting policy preferences. Each column represents a different policy dimension. The Max No. row displays the maximum number of parameters meeting the BIC criteria. Max row displays the corresponding R^2. The Member FE row displays the R^2 from regressions with member fixed effects alone.

added to each model. It indicates that the use of a small number of frameworks has similar explanatory power as member fixed effects for understanding committee members' policy preferences.

Across variables, with about five parameters, framework-parameters pairs have similar explanatory power as member fixed effects. For inflation preferences, economic frameworks explain approximately 26% of the variation by the time nine parameters are included, and 30% when ten parameters are included. This compares reasonably to the 31% explained by member fixed effects alone. The first parameter (high wage indexation) alone explains 15% of variation, with steep Phillips Curve slope contributing an additional 5%. For growth preferences, five parameters explain about 31% of the variation, exceeding the 20% explained by member fixed effects. Labor market and Credit preferences show similar patterns, with five parameters explaining approximately 29% and 20% of variation compared to 23% and 22% for member fixed effects respectively. These results indicate a relationship between policy preferences and economic models used to interpret macroeconomic data. Across all policy dimensions, the use of a small set of frameworks in deliberation explains approximately the same or greater variation than member fixed effects.

4 Monetary Policy by Committee

Given the evidence that individual members use different mental models to interpret incoming economic data and formulate their policy preferences, we next analyze how the

FOMC aggregates these models in making its policy decision and connect this aggregation to monetary policy surprises measured using asset prices.

4.1 Anecdotal Evidence

In conducting our analysis, we take guidance from descriptions of the process of deliberation in the FOMC by its members, as presented in Schonhardt-Bailey (2013). Discussing this process, Alan Blinder, the former vice chairman of the Federal Reserve noted that

> "You would make a point, maybe forcefully if necessary, at one meeting in the hope that it would sink in and by the next meeting... or the meeting after that, you are starting to persuade people."

Cathey Minehan, the former president of the Federal Reserve Bank of Boston indicated that

> "I viewed the FOMC as a process with discussion evolving over a series of meetings... views expressed in one meeting could resonate and influence the discussion and decisions of subsequent meetings."

Donald Kohn, the former secretary and vice chair of the Federal Reserve, described how

> "It's... fair to say that in many meetings I got food for thought that influenced how I was looking at the data, [and it] made me go back and rethink issues, or look at the data in a different way... That was influential over the following inter-meeting period as I went about thinking what to do next time."

Taking these descriptions into account, we hone in on the hypothesis that committee members whose mental models—as measured by their statements in the prior meeting—match well with the incoming data will have more influence over the policy decision in a meeting.

4.2 The Committee Decision and Aggregating Models

To proceed, we represent each committee member's mental model in period t as a vector of their policy preferences along each dimension, i.e., as

$$m_{i,t} = \begin{pmatrix} \gamma_{i,t}^{\text{inflation}} \\ \gamma_{i,t}^{\text{growth}} \\ \gamma_{i,t}^{\text{labor}} \\ \gamma_{i,t}^{\text{credit}} \\ \gamma_{i,t}^{\text{stocks}} \end{pmatrix}, \tag{1}$$

where $\gamma_{i,t}^{k}$ is a measure of committee member i's policy preference in meeting t for variable k.

We then estimate how well committee member i's model matches incoming economic data based on the inner product of their mental model with a vector of the associated incoming

economic data. We define an economic data vector in period t as the vector

$$d_t = \begin{pmatrix} inflation_t \\ growth_t \\ labor_t \\ credit_t \\ stocks_t \end{pmatrix}, \tag{2}$$

where each row corresponds with a measure of the given variable in period t, signed such that more positive measurement should correspond with more hawkish policy. We define our primary variable of interest as

$$m_{i,t-1}(data_t) = m_{i,t-1}^T d_t, \tag{3}$$

i.e., as the inner product of the economic data vector with a member's mental model, as measured given their statements in the previous meeting. More positive values of $m_{i,t-1}(data_t)$ indicate better alignment with recent economic data. The choice to use lagged mental models is motivated by FOMC members' descriptions of influence on the committee as arising when members' interpretations as stated in a meeting carry over and influence the tenor of the discussion in subsequent meetings. We also conduct additional analysis where we measure the $m(data)$ variables using the fixed effect for member i on variable k, estimated over the full sample after controlling for meeting fixed effects.

To operationalize our $m(data)$ variables, we construct two measures of recent economic data, d_t: one using quantitative data and one using textual data. For our first measure, d_t^{quant}, we measure $growth_t$ and $labor_t$ using the Tealbook/Greenbook forecasts of real GDP growth and (the negative) of the unemployment rate for the quarter following t. We measure $inflation_t$ using forecasts of the following quarter's headline CPI growth (from 1983 onwards) and GDP deflator growth (pre-1983).[12] The Tealbook is distributed to FOMC members days before the FOMC meeting, reflecting up-to-date forecasts of the variables of interest by the Federal Reserve Board of Governors. It has the advantages of reflecting real-time information about the economy at the time of the meeting and of being part of the shared information set of committee members that serves as a baseline for economic discussion. We measure $stock_t$ and $credit_t$ as the 1-month stock market return (from Kenneth French's website) and the (negative of the) BAA–$10Y$ credit spread (from the St. Louis Federal Reserve website). We express all measures as expanding window z-scores, to ensure similarly scaled variables

[12]The sample of Greenbook forecasts for headline CPI begin with the October 1979 meeting, whereas the history of forecasts of the GDP deflator (and real GDP/unemployment) goes back to 1967. We use the GDP deflator forecasts for an additional 3 years after the sample of headline CPI forecasts begins due to excess volatility of the expanding window z-score in small samples. Our results are similar simply using GDP deflator forecasts throughout the full sample.

across dimensions.

For our second measure, d_t^{text}, we construct a score from -3 to 3 for each of inflation, growth, labor markets, credit markets, and the stock market, by inputting text of the Tealbook A (or Greenbook) distributed prior to a meeting into an LLM and prompting the LLM to evaluate whether the variable supports tighter or looser monetary policy; a score of -3 corresponds with a given variable providing strong support for an interest rate cut and +3 corresponds with a given variable providing strong support for an interest rate hike. We present the full text of the prompts used to generate this measure in Appendix B.2.5. This text-based measure has the additional advantage of incorporating the various metrics discussed in the Tealbook A in formulating its score for each variable, for example focusing on business and consumer credit across segments when considering credit market conditions.[13]

With our measures of incoming economic data in hand, we construct $m_{i,t-1}(data_t^{text})$ and $m_{i,t-1}(data_t^{quant})$ for each committee member i in period t as the inner products of the member's mental model from $t-1$, $m_{i,t-1}$, with d_t^{quant} and d_t^{text}, respectively. We then estimate meeting-by-member level regressions in which the dependent variables capture different aspects of a member's alignment with the committee's final decision—namely, *DirectiveAlignment*, *ChangeAlignment*, and a binary indicator, *Dissent* (equal to 1 if the member dissented and 0 otherwise)—using the $m(data)$ measures as regressors.

Our regressions include meeting fixed effects and member fixed effects. With meeting fixed effects, the regressions assess whether, within a given meeting, members whose mental models are more closely aligned with the recent economic data are also more aligned with the committee's decision. With member fixed effects, the regressions examine whether the committee's policy decision tends to align more closely with a member's preferences during periods when that member's mental model is better matched to recent data— controlling for the possibility that, on average, the committee's decision may be more in line with the preferences of members whose models are better aligned with the data on average.

With the exception of *Dissent*, all variables in the regressions are standardized to have zero mean and unit standard deviation. Consequently, the regression coefficients can be loosely interpreted as the standard deviation change in a member's alignment with the committee's decision corresponding to a one standard deviation change in the match between their model and the data. For *Dissent*, the coefficients represent the percentage point change in the probability of dissent associated with a one standard deviation change in how well the member's model matches the data.

Table 7 reports the regression results. With *DirectiveAlignment* and *ChangeAlignment* as the dependent variables, focusing on the regressions with meeting fixed effects (but without member fixed effects), the regression coefficients range from 0.16 to 0.25 and t-statistics

[13] Aruoba and Drechsel (2024) show that the Tealbook text contains decision-relevant information, and use it to construct a new measure of monetary policy shocks.

ranging from 3.44 to 4.31, depending upon the specification. These results indicate that in a given meeting, the decision made by the committee is more aligned with committee members whose models are well matched with the economic data. When including member fixed effects, the regression coefficients decline, ranging from 0.09 to 0.13, with t-statistics ranging from 2.35 to 3.59. The decline in coefficients suggests that part of the relationship between the $m(data)$ and $Alignment$ variables is captured by committee members whose mental models persistently match the macroeconomic data well tending to be better aligned with the committee's policy decision on average; but a substantial portion of the relationship comes from an individual member of the committee being better aligned with the committee's decision in periods in which their mental model better matches the data.

With *Dissent* as the dependent variable, coefficients range from -0.018 to -0.061, indicating that a one standard deviation decrease in $m(data)$ corresponds to a 1.8% to 6.1% increase in dissent probability. Given the unconditional dissent probability of 6.9% in our sample, our results indicate that a substantial portion of voting dissents arises from committee members whose mental models are misaligned with the recent data.

Table 8 reports results of the specifications in which we measure committee members' mental models as member fixed effects from regressions of policy preferences on member fixed effects and meeting fixed effects. Across specifications, the effect sizes in the regressions are of a similar magnitude to those in Table 7. That is, the evidence suggests that the match between macroeconomic data and a time-invariant component of a committee member's mental models explains nearly the full relationship between how well their model matches the recent macroeconomic data and their alignment with the committee's decision.

In Appendix B.3, we report variations of the results from Table 7. Table B.1 reports regression results including member type fixed effects, controlling for whether a committee member is a regional Fed president, a governor, or the chairperson of the committee, motivated by work emphasizing differential behavior between governors and Fed presidents in dissenting (Belden, 1989; Chappell Jr, McGregor and Vermilyea, 2004). These fixed effects do not substantially influence the results. Tables B.2 and B.3 report regression results in which the $m(data)$ variables are constructed using only the traditional macroeconomic variables and only the financial variables, respectively. The results reveal that both the macroeconomic and financial variables help explain which members are aligned with the committee's decision. Table B.4 reports regression results with *DirectiveAlignment* and *ChangeAlignment* as dependent variables and the *Dissent* variable as a control. The alignment variables have a substantial loading on *Dissent*, consistent with dissents being highly related to alignment. The results also indicate that the alignment variables capture additional information about committee members' policy preferences over and above dissent.

	y =*DirectiveAlignment*				y =*ChangeAlignment*				*y=Dissent*			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
$m_{i,t-1}(data_t^{quant})$	0.25	0.11			0.23	0.13			-0.06	-0.03		
	(3.66)	(2.40)			(3.88)	(2.94)			(3.68)	(2.39)		
$m_{i,t-1}(data_t^{text})$			0.16	0.09			0.19	0.13			-0.03	-0.02
			(3.44)	(2.35)			(4.31)	(3.59)			(2.35)	(1.62)
R^2	0.22	0.46	0.21	0.46	0.24	0.46	0.24	0.46	0.15	0.30	0.13	0.30
N	3659	3659	3659	3659	3659	3659	3659	3659	3659	3659	3659	3659
Meeting Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Member Fixed Effect	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes

TABLE 7: MENTAL MODELS AND POLICY ALIGNMENT

Note: The table reports results from regressions of member i's alignment with the committee's policy decision in period t on $m_{i,t-1}(data_t)$ variables, which capture how well committee member i's mental model matches the macroeconomic data in period t. The $m_{i,t-1}(data_t)$ variables are constructed as the inner product of a vector committee member i's preferences across macroeconomic and financial variables expressed in the period $t-1$ meeting with a vector of macroeconomic and financial data from period t. The *quant* superscripts indicate that the macroeconomic data are macroeconomic forecasts from the Greenbook/Tealbook A and the financial data are trailing 1-month stock market returns and the $BAA-10Y$ credit spread. The *text* superscripts indicate that the macroeconomic and financial data are measured using textual analysis of the Greenbook/Tealbook A. The dependent variables are *DirectiveAlignment* and *ChangeAlignment*—scores from -3 (complete opposition) to 3 (perfect alignment) constructed from the policy discussion sections of FOMC transcripts—and a 0/1 indicator variable if the member i dissented in the meeting in period t. The sample consists of voting members during FOMC meetings. Standard errors are clustered by meeting and committee member. t-statistics are reported in parentheses.

	y =*DirectiveAlignment*				y =*ChangeAlignment*				*y=Dissent*			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
$m_{i,FE}(data_t^{quant})$	0.23	0.14			0.21	0.17			-0.05	-0.03		
	(3.10)	(2.91)			(3.21)	(3.54)			(3.04)	(1.69)		
$m_{i,FE}(data_t^{text})$			0.17	0.12			0.20	0.17			-0.03	-0.02
			(2.94)	(2.89)			(3.91)	(4.38)			(2.00)	(1.95)
R^2	0.23	0.46	0.21	0.46	0.24	0.46	0.24	0.46	0.15	0.30	0.13	0.30
N	3659	3659	3659	3659	3659	3659	3659	3659	3659	3659	3659	3659
Meeting Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Member Fixed Effect	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes

TABLE 8: FIXED EFFECTS, MENTAL MODELS, AND POLICY ALIGNMENT

Note: The table reports results from regressions of member i's alignment with the committee's policy decision in period t on $m_{i,FE}(data_t)$ variables, which capture how well committee member i's mental model matches the macroeconomic data in period t. The $m_{i,FE}(data_t)$ variables are constructed as the inner product of a vector committee member i's mental model, measured as their fixed effects in regressions of preferences across macroeconomic and financial variables on meeting and member fixed effects, with a vector of recent macroeconomic data. The *quant* superscripts indicate that the macroeconomic data are macroeconomic forecasts from the Greenbook/Tealbook A and the financial data are trailing 1-month stock market returns and the $BAA - 10Y$ credit spread. The *text* superscripts indicate that the macroeconomic and financial data are measured using textual analysis of the Greenbook/Tealbook A. The dependent variables are *DirectiveAlignment* and *ChangeAlignment*—scores from -3 (complete opposition) to 3 (perfect alignment) constructed from the policy discussion sections of FOMC transcripts—and a $0/1$ indicator variable if the member i dissented in the meeting in period t. The sample consists of voting members during FOMC meetings. Standard errors are clustered by meeting and committee member. t-statistics are reported in parentheses.

Taken together, the evidence that committee members whose mental models better match recent macroeconomic data are better aligned with the policy decision made by the committee indicates that the committee tilts its decision towards those members' preferences.

4.3 Preference Aggregation Mechanisms

We next explore the mechanisms driving the committee's tilt towards members whose mental models are better aligned with recent macroeconomic data. We present evidence that such tilt may be driven in part by committee members with mental models that tend to match the data well speaking more on average. We also document evidence that the chairperson of the committee likely plays an important role in this aggregation.

We examine the speech patterns of committee members during the FOMC policy discussion. We compute the number of statements made during the policy discussion by each voting committee member. This may include pre-prepared statements, but also interjections and questions to other speakers during the session. We regress this number of statements on the $m(data)$ variables, constructed using policy preferences from the previous meeting. Given the count data structure of the number of speeches, we run Poisson regressions (e.g., see Cohn, Liu and Wardlaw (2022) and Chen and Roth (2024) for discussion). We include meeting and member fixed effects in the regressions and standardize the independent variables to have zero mean and unit standard deviation.

Table 9 reports the regression results. With meeting fixed effects (but without member fixed effects), the coefficients range from 0.07 to 0.10, with z-statistics of 2.15 and 3.26. This corresponds to a 7% to 10% increase in the number of statements made by committee members per unit standard deviation their mental models are better matched with the recent data. However, the effect goes to zero when including member fixed effects, indicating that the effect is not driven by a committee member speaking more when their mental model is better aligned with the data; but rather, by committee members whose mental models are better matched on average tending to speak more on average. Nevertheless, one method by which the committee may tilt towards the mental models that match the data better is by such committee members speaking more and having more influence on average.

The chair may also play an important role in aggregating committee members' preferences. Figure 4 plots the eight-meeting rolling average *DirectiveAlignment* of the committee chair and the other voting members of the committee. The chair displays strong alignment with the final decision made by the committee—with an average *DirectiveAlignment* score of 2.90—while other voting members, on average, display lower alignment—with an average *DirectiveAlignment* score of 1.93. Starting with Alan Greenspan's tenure as chair, the committee chair has been especially aligned with the committee's final decision (average *DirectiveAlignment* of 2.98). This result is consistent with the committee chair having substantial influence in shaping the final directive of the committee. The fact that the committee's deci-

	(1)	(2)	(3)	(4)
$m_{i,t-1}(data_t^{quant})$	0.07	-0.02		
	(2.15)	(0.70)		
$m_{i,t-1}(data_t^{text})$			0.10	-0.01
			(3.26)	(0.23)
N	3659	3659	3659	3659
Meeting Fixed Effect	Yes	Yes	Yes	Yes
Member Fixed Effect	No	Yes	No	Yes

TABLE 9: NUMBER OF STATEMENTS AND MENTAL MODELS

Note: The table reports results from Poisson regressions of the number of statements made by a voting committee member i in a given meeting in period t on $m_{i,t-1}(data_t^{quant})$ variables, which capture how well committee member i's mental model matches the macroeconomic data in period t. The $m_{i,t-1}(data_t^{quant})$ variables are constructed as the inner product of a vector committee member i's preferences across macroeconomic and financial variables expressed in the period $t-1$ meeting with a vector of macroeconomic and financial data from period t. Standard errors are clustered by meeting and committee member. z-statistics are reported in parentheses.

sion leans towards the preferences of committee members whose mental models match the recent economic data well is likely to be, in part, a consequence of the chair's guidance.

However, this is not to say that the behavior of other committee members is irrelevant to the decision made. Committee members of the FOMC play a key role in introducing changes to the policy directive that are eventually adopted by the committee.[14] Of changes adopted to the policy directive in our sample, we identify that about 45% explicitly appear in the policy alternatives distributed in the Bluebook or Tealbook prior to the FOMC meeting, while 55% are introduced by committee members during the meeting.[15] Of the 55% that are introduced during meetings, a substantial portion are introduced by committee members other than the chair.

Figure 5 plots the proportion of adopted policy directive changes proposed during FOMC meetings, split by the chair of the committee. Prior to Ben Bernanke, the majority of adopted policy changes were proposed by the chair of the committee. The lowest proportion of directive changes proposed by members other than the chair occurred during the tenure of Alan Greenspan (31%); nonetheless, even then, non-chair members play an important role in the committee's decision. Starting with Ben Bernanke, the vast majority of adopted directive changes are proposed by members other than the chair (upwards of 70%). Indeed, these

[14]The prompts used to assess which committee members introduced changes to the policy directive and whether such changes already appeared in the alternatives are given in Appendix B.2.4.

[15]From our discussions with members of the committee, the alternatives presented in the Bluebooks and Tealbooks reflect input from the Federal Reserve staff, as well as the chair and members of the committee.

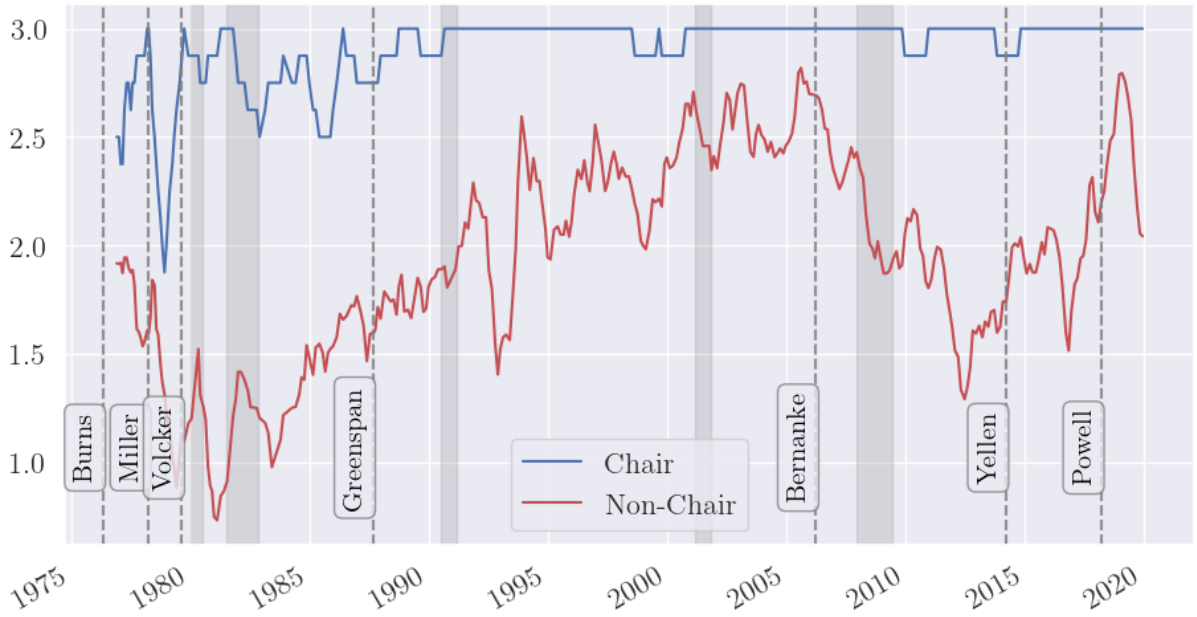


FIGURE 4: COMMITTEE MEMBERS' ALIGNMENT WITH THE POLICY DECISION

Note: The figure plots the eight-meeting rolling average *DirectiveAlignment* of the chair of the committee and of the other voting members of the committee. The *DirectiveAlignment* scores are constructed from the policy discussion section of the FOMC meetings and range from -3, (complete disagreement) to +3 (complete agreement) with the committee's policy decision. The dark grey bands indicate NBER recession periods. The starting date of each chair's tenure is indicated with vertical dashed lines.

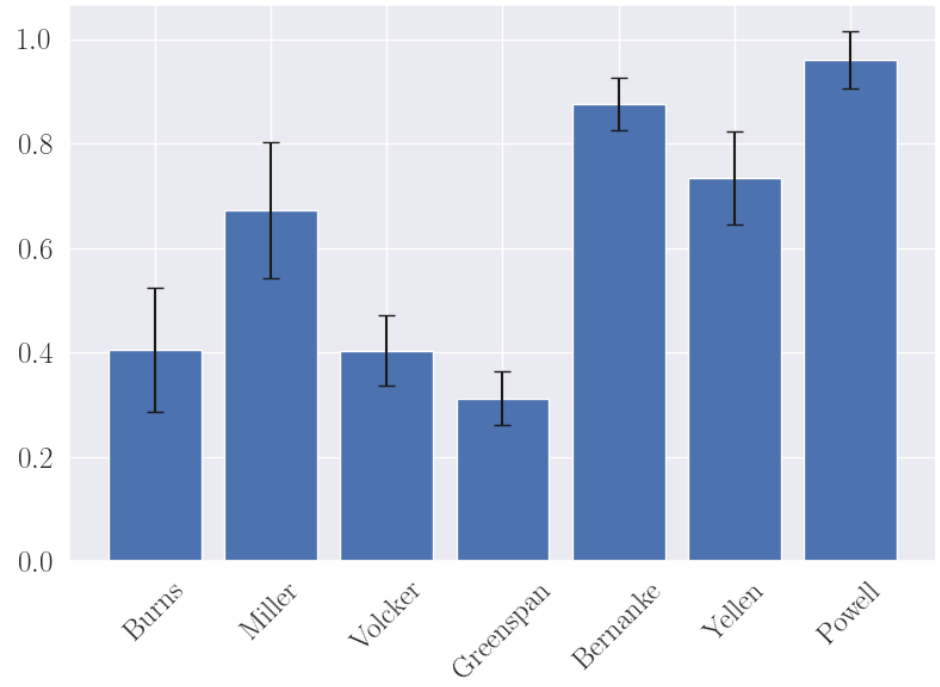


FIGURE 5: CHANGES TO POLICY DIRECTIVE INTRODUCED BY NON-CHAIR MEMBERS

Note: The figure plots the proportion of adopted changes to the policy directive proposed during the meeting by members of the committee other than the chair. We identify 55% of adopted changes are proposed during the meeting, while 45% are reflected in the policy alternatives presented in the Bluebook / Tealbook distributed prior to the meeting. The proportions are plotted separately for each chair's tenure. The bars correspond with plus or minus two standard errors.

patterns across chairs also mirror the proportion of words spoken by the chair of the committee during the policy discussion, presented in Appendix Figure B.1. Chairmen prior to Ben Bernanke tended to take a more dominant role in the policy discussion, while more recent committee chairs have spoken considerably less.

The move in the later period towards members playing a larger role is also supported by anecdotal accounts of the FOMC. Narayana Kocherlakota, the former president of the Federal Reserve Bank Minneapolis, noted that in contrast to Alan Greenspan, who was well known for speaking first and strongly expressing his preferences in the policy discussion of the committee,

> "Bernanke and Yellen both made a point of speaking last. There were exceptions, but in general they made it about wanting to hear everyone's views... [Yellen] wanted more agreement within the committee, and she worked hard to facilitate it. She would move her position to make it happen" (Bowmaker, 2025).

Despite the differing styles across chairs, the results broadly support the idea that committee members may influence the committee's tilt by proposing changes to the policy directive; and

that the chair likely plays an important role in guiding the committee in the changes it adopts.

4.4 Directive Changes and Monetary Policy Surprises

Our evidence of the committee tilting its decision-making towards members whose views better match the data is based on members' alignment with changes to the policy directive. We conclude our empirical analysis by showing that these changes to the policy directive have a strong relationship with monetary policy surprises measured from asset prices.

Following prior work (Gürkaynak, Sack and Swanson, 2005; Nakamura and Steinsson, 2018), we measure monetary policy surprises using the high frequency responses of interest rate futures in a 30-minute window following the FOMC's announcement of its policy decision. Our main measure of monetary policy surprises is mps_t, constructed as the first principal component of changes in the one- to four-quarter ahead Eurodollar futures in the 30-minute window following FOMC announcements, from Swanson and Jayawickrema (2024) and Bauer and Swanson (2023*b*). The sample runs from 1988-2019.[16] We also analyze the changes of n-quarter ahead Eurodollar futures (labeled EDn), as well as the near- and 3-month maturity Fed Funds futures in the 30-minute window following the FOMC announcement (labeled FF and $FF3m$).[17] The sample of Fed Funds futures runs from 1995-2019.

For each meeting, we assign a score from -3 (highly accommodative / dovish) to +3 (highly restrictive / hawkish) to each of the identified changes to the policy directive. We then compute two measures: *DirectiveChange* as the average score across all changes; and *DirectiveChangeExRate* as the average score across all changes excluding target rate changes.[18] We estimate regressions of monetary policy surprises on *DirectiveChange* and *DirectiveChangeExRate*. The *DirectiveChange* variables are standardized to have zero mean and unit standard deviation, and monetary policy surprise variables are in basis points.

The first two columns of Table 10 report results from univariate regressions of mps_t on *DirectiveChange* and *DirectiveChangeExRate*. The coefficients from these regressions are 1.42 and 1.03, respectively, indicating that a one standard deviation change in the *DirectiveChange* measures captures a 1 to 1.4 basis point monetary policy surprise. As reported in the Table, the standard deviation of mps_t is about 5 basis points, indicating an economically substantial relationship between the independent variables and monetary policy surprises.

[16] In the pre-1994 sample, there are multiple mps_t measurements per meeting. This is because, while from 1994 onwards, the committee releases its decision in a press release, in the pre-1994 period, market participants had to infer whether there had been a change in the federal funds rate from the size and type of open market operation conducted by the Fed each morning. Our analysis uses all mps_t measurements for each meeting, but is largely unchanged when only including one, or when only focusing on the post-1994 sample.

[17] Following Bernanke and Kuttner (2005), we scale up the change in the near maturity Fed Funds future by a factor related to the number of days in the month affected by the change. This is due to the contract's settlement price being based on the monthly average Federal funds rate.

[18] There is one fewer observation in *DirectiveChangeExRate* as for one meeting, the only directive decision identified is with respect to the target rate decision.

The mps_t variable captures the component of monetary surprises stemming from both the target rate decision and other elements of the committee's decision, most importantly forward guidance about the future path of interest rates. In columns 3-6 of Table 10, we explicitly separate the surprise of the target rate decision from forward guidance, in the spirit of Gürkaynak, Sack and Swanson (2005). Columns 3 and 4 report results from regressions of the high frequency change in four quarter-ahead Eurodollar futures on the *DirectiveChange* variables, controlling for the contemporaneous change in one quarter-ahead Eurodollar futures. The coefficients on *DirectiveChange* and *DirectiveChangeExRate* are 0.74 and 1.20 in these regressions, compared to the 5.7 bps standard deviation of $ED4$. Columns 5 and 6 report results from regressions of the high frequency change in 3-month ahead Fed Funds futures on the *DirectiveChange* variables, controlling for the contemporaneous change in the near maturity Fed Fund futures contract. The coefficients on *DirectiveChange* and *DirectiveChangeExRate* are 0.66 and 0.63 in these regressions, compared to the 3.7 bps standard deviation of $FF3m$. Taken together, these results indicate that the *DirectiveChange* variables are strongly related to surprise changes to the target rate, as well as to surprise changes in communication about the future path of interest rates.

The relationship between the *DirectiveChange* variables and monetary policy surprises is notable for two reasons. First, our evidence in the previous section is based on comparing members' alignment with the changes to the policy directive with how well their mental models match recent economic data. The relationship with monetary policy surprises indicates that our prior evidence isn't just measuring internal deliberation, but is capturing behavior reflected in financial market movements to the resulting policy.

Second, taken in conjunction with the evidence on the committee's tilt towards different members' models based on the data, the results also suggest that the committee structure may play an important role in generating monetary policy surprises. The committee's tilting behavior is consistent with it not using a fixed policy rule to make a decision, as argued by Hack, Istrefi and Meier (2025), with shifts in the policy stance arising from the shifting influence of committee members coming as a surprise to investors. The link with monetary policy surprises is also complementary to mounting evidence of monetary policy surprises arising from market participants having uncertainty about the Fed's reaction function in response to macroeconomic news (Bauer and Swanson, 2023*a*,*b*; Bauer, Pflueger and Sunderam, 2024). Consistent with this interpretation, columns 7 and 8 of Table 10 report results from regressions of mps_{orth}, measured as the residual of mps after controlling for macroeconomic and financial news arriving prior to the FOMC meeting, on the *DirectiveChange* variables. These results indicate a negligible relationship, suggesting that the component of the *DirectiveChange* variables that is related to monetary policy surprises stems from investors being surprised by the committee's response to macroeconomic and financial news.

	$y = mps$		$y = ED4$		$y = FF3m$		$y = mps_{orth}$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
DirectiveChange	1.42		1.20		0.66		0.36	
	(3.53)		(4.32)		(2.76)		(1.08)	
DirectiveChangeExRate		1.03		0.74		0.63		0.12
		(2.59)		(5.07)		(2.80)		(0.35)
$ED1$			0.89	1.02				
			(8.19)	(17.17)				
FF					0.70	0.73		
					(7.61)	(8.03)		
σ_y	4.95	4.95	5.70	5.70	3.74	3.74	4.63	4.63
N	262	261	262	261	190	189	262	261
R^2	0.08	0.04	0.54	0.81	0.53	0.53	0.01	0.00

TABLE 10: DIRECTIVE CHANGES AND MONETARY POLICY SURPRISES

Note: The table reports results for regressions of high-frequency asset price-based measures of monetary policy surprises on the hawkish/dovish tone of FOMC directive changes. The dependent variables are: mps, the first principal component of changes in one- to four-quarter ahead Eurodollar contracts within a 30-minute post-announcement window; $ED4$, the change in the four-quarter ahead Eurodollar contract; $FF3m$, the change in 3-month ahead Fed Futures yields; and mps_{orth}, computed as residuals from regressing the mps measure on a set of controls that capture incoming macroeconomic and financial news. The independent variables of interest—*DirectiveChange* and *DirectiveChangeExRate* (excluding the policy rate decision)—are the standardized meeting-level averages (scores from –3 for strongly dovish to +3 for strongly hawkish) of changes to the policy directive. The regressions include controls for the high-frequency change in one-quarter ahead Eurodollar futures ($ED1$) and near-maturity Fed Funds futures (FF). The mps and ED measures used in the regressions are from Bauer and Swanson (2023*b*). The table reports the standard deviation of the dependent variables as σ_y; t-statistics computed using robust standard errors are reported in parentheses.

5 Theoretical Framework

Our empirical evidence indicates that committee members have distinct and largely fixed models of the macroeconomy, with the committee tilting its decision towards the models of committee members that are better aligned with recent economic data. The results suggest that the traditional reasoning behind forming committees to conduct policy—to aggregate diverse information across members—is not a sufficient explanation for the behavior of the FOMC.

Guided by our evidence and in order to further explore its implications, we present a theoretical model of group decision-making. We take the perspective that group decision-making provides a way to make decisions in the face of model uncertainty—when the true model of the world is unknown—by bringing together decision-makers with different models.

5.1 Setup

There are two committee members indexed by $i \in \{A, B\}$. Let μ be the true probability measure on a measurable space $(\Omega, \mathcal{F})$, which is unknown to committee members. A *model* is a probability distribution over the state space Ω, where Ω can be taken to represent all possible macroeconomic realizations. The objective of the committee is to select a model, which in turn, maps to its policy decision.

Each committee member i has their own personal, fixed model, m_i, which is the density of a probability measure, M_i, on $(\Omega, \mathcal{F})$. We use lower case to refer to probability densities and upper case to refer to the corresponding measures throughout.

Assumption 1 (Regularity Conditions on Members' Models) *Both members' models satisfy the following regularity conditions:*

(i) **(Partial misspecification).** Members' models are partially misspecified ($\forall i$, $\exists S_i \subseteq \Omega$ such that $\mu(S_i) > 0$ and $M_i(S_i) \neq \mu(S_i)$).

(ii) **(Distinctness).** Members' models are misspecified in distinct ways (there is no constant c such that for all $E \in \mathcal{F}$ with $\mu(E) > 0$, $M_A(E) - \mu(E) = c(M_B(E) - \mu(E))$).

(iii) **(Absolute Continuity).** Members' models assign positive probability to all possible states, and only to those states (models are mutually absolutely continuous with respect to μ).

The committee makes its decision as follows. All members observe data x from a prior period, which is a realization of a random variable—for example, realized CPI and GDP growth—whose distribution depends on the state of the world. For any state ω, the probability of observing x is $\pi(x|\omega)$. Given a model p, the likelihood of observing x, $\mathcal{L}(x|p) = \int_{\omega \in \Omega} \pi(x|\omega) dP(\omega)$, is a weighted average of the conditional probabilities $\pi(x|\omega)$ over the

state space where the weights are given by the model. Given the observed data, likelihood calculations give higher scores to models that assign high probability to states that appear more likely in light of the data.

After observing the data, each member i advocates for a convex combination of committee members' models.

$$v_i = \alpha_i m_i + (1 - \alpha_i) m_j, \text{ where } j \neq i \text{ and } \alpha_i \in [0, 1]. \tag{4}$$

The weight $\alpha_i \in [0, 1]$ reflects how strongly member i advocates for their own model, e.g., how much they speak during the committee meeting, or how actively they propose changes to the policy directive. The committee's model m, which maps to the policy decision of the committee, is given by averaging the members' advocated models:

$$m = \frac{1}{2}(v_A + v_B) = w_A m_A + (1 - w_A) m_B, \tag{5}$$

where $w_A \equiv \frac{\alpha_A + (1 - \alpha_B)}{2}$ is the weight that committee member A's model receives in equilibrium. Committee member i chooses α_i to maximize the objective function

$$U_i(\alpha_i; \alpha_j) = \log \mathcal{L}(x|m) - \lambda KL(m_i \| v_i), \tag{6}$$

where $KL(p\|q) = \int_{\omega \in \Omega} \log\left(\frac{p(\omega)}{q(\omega)}\right) dP(\omega)$ is the Kullback-Leibler divergence, and the coefficient λ is positive.

The first term, $\log \mathcal{L}(x|m)$, captures that committee members want the committee to settle on a model m that explains the data well, all else equal. This term may reflect that members are aware that their models are misspecified, and want to use data as a guide in the face of model uncertainty. Even if committee members are dogmatic that their personal models are correct, such a term may still arise. For example, it may reflect that members have difficulty in advocating for their own model when it is misaligned with the observed data; or it may reflect a need for the committee to justify its decision to constituencies that also observe x, such as financial market participants.

The second term, $-\lambda KL(m_i \| v_i)$, captures a penalty for the extent to which committee member i is unwilling or unable to depart from their own model in their advocacy. The coefficient λ captures how this penalty scales with $KL(m_i \| v_i)$, which captures the deviation of member i's advocacy from their model. A high λ may emerge from a combination of behavioral factors—such as stubbornness or overconfidence—as well as institutional factors—such as the decision-making setting not sufficiently exposing members to the potential merits of others' models.

Remark 1 *The setting here is static. It can be nested into a dynamic setting as a stage game that is played each meeting where the committee makes a decision, with μ representing a conditional distribu-*

tion that members seek to explain.[19] *With dynamics, there also may be additional considerations, for example, learning, reputational effects, preferences for gradualism, and dynamic inconsistencies.*

5.2 Equilibrium

We study pure strategy Nash Equilibria.

Definition 5.1 (Equilibrium) *A Nash equilibrium is a collection of weights,* $\{\alpha_A, \alpha_B\}$*, such that*

$$\alpha_i \in \arg\max_{\alpha^i} U_i(\alpha^i; \alpha_j),$$

where $U_i(\alpha^i; \alpha_j)$ *denotes the utility of committee member* i *given their own chosen weight* α^i *and the other member's best response,* α_j*.*

Proposition 1 (Equilibrium Existence and Uniqueness) *There is a unique equilibrium.*

Proof. All proofs provided in the appendix. □

5.3 Committee Performance

With the equilibrium in hand, we next turn to understanding the conditions under which the committee structure can improve decision-making performance. Throughout our analysis, we evaluate the performance of a model using an ex-ante criterion: the model's (ex-ante) expected log likelihood under the true probability measure, $\mathbb{E}_\mu(\log m)$. We consider this to be a natural measure of performance, as choosing m to maximize $\mathbb{E}_\mu(\log m)$ is equivalent to minimizing the KL divergence between the true model and m, i.e., minimizing the loss of information from using m to approximate the true probability distribution.

Proposition 2 (Potential Committee Outperformance) *There are a set of optimal weights on committee members' models,* $\mathbf{w}^* = (w_A^*, w_B^*)$*, such that the performance of a weighted average of the members' models—its expected log likelihood under the true distribution—exceeds that of even the best committee member. That is, there are a set of optimal weights* $\mathbf{w}^*$ *that satisfy*

$$\mathbb{E}_\mu \log(w_A^* m_A + w_B^* m_B) > \max\{\mathbb{E}_\mu(\log m_A), \mathbb{E}_\mu(\log m_B)\}.$$

Proposition 2 highlights the potential benefits associated with decision-making by committee. Under fairly weak assumptions—that committee members have models that are partially misspecified in different ways—there are benefits to model diversification, since there is a combination of members' models that can outperform even the best committee member's

[19] Such a time-varying μ is consistent with Alan Greenspan's description in a speech to the AEA that "The economic world in which we function is best described by a structure whose parameters are continuously changing. The channels of monetary policy, consequently, are changing in tandem" (Greenspan, 2004).

model. This is true independent of the relative performance of the models. For example, a committee member B may have a model that is substantially more misspecified than committee member A's model, but as long as the models are different, there can be some benefit to combining them.

However, in equilibrium, the committee does not simply choose the optimal weights $\mathbf{w}^*$. First, because the committee settles upon its model using data—which has an element of randomness—the weights chosen by the committee, $\mathbf{w} \equiv (w_A, 1 - w_A)$, will vary. Second, there may be an element of bias, stemming from committee members' stubbornness in advocating for their own models, λ.

To understand the relationship between stubbornness and bias and variance in the committee's choice of model weights, it is useful to consider how the model selected in equilibrium behaves as λ varies. When $\lambda = 0$, the committee chooses the model that best fits the observed data, a choice that leads the weights placed on each member's models to be relatively unbiased relative to the optimal weights w^*, but sensitive to variation in the data. In the limiting case as $\lambda \to \infty$, each member simply advocates for their own model and in equilibrium, the resulting committee model is always $m = \frac{1}{2}m_A + \frac{1}{2}m_B$, regardless of the data. Under these circumstances, there is no variability in the model chosen, but the weights chosen on members' models are biased relative to the optimal choice (except for the degenerate case when $w_A^* = w_B^* = \frac{1}{2}$). For intermediate values of stubbornness, as λ increases, the equilibrium model weights exhibit more bias and less variance relative to the optimal weights.

Whether the potential benefits of committee decision-making are realized or not depends upon the bias and variability of the committee's chosen weights relative to the optimal weights. The committee's model outperforms its best member's model in expectation when the benefits from model diversification are sufficiently large to offset potential bias from the stubbornness of committee members and variability in the committee's decision.

Proposition 3 (Committee Outperformance) *Denote $w_A(\lambda)$ as the equilibrium weight chosen by the committee on model A; Δ as the expected outperformance of the optimal model relative to the best members' model,*

$$\Delta \equiv \mathbb{E}_\mu \log((w_A^* m_A + w_B^* m_B)) - \max\{\mathbb{E}_\mu(\log m_A), \mathbb{E}_\mu(\log m_B)\}; \tag{7}$$

$\beta(\lambda)$ as the expected bias of the committee's weight on model A relative to the optimal weight,

$$\beta(\lambda) = \mathbb{E}_\mu(w_A(\lambda) - w_A^*); \tag{8}$$

$\sigma^2(\lambda)$ as the variance of the committee's weight on model A,

$$\sigma^2(\lambda) = \mathbb{V}_\mu(w_A(\lambda)); \tag{9}$$

and c as the negative Fisher Information evaluated at the optimal model weights,

$$c = \mathbb{E}_\mu \left(\frac{\partial^2}{\partial w_A^2} \log \mathcal{L}(x|w_A m_A + (1 - w_A) m_B)|_{w_A = w_A^*} \right). \tag{10}$$

Under a second-order approximation, if

$$\Delta > \frac{c}{2}(\beta(\lambda)^2 + \sigma^2(\lambda)), \tag{11}$$

then

$$\mathbb{E}_\mu(\log(w_A(\lambda)m_A + (1 - w_A(\lambda))m_B)) > \max\{\mathbb{E}_\mu(\log m_A), \mathbb{E}_\mu(\log m_B)\}. \tag{12}$$

Proposition 3 indicates that the committee may outperform its best member under a few basic conditions. First, there have to be sufficient gains from optimally combining members' models, as measured by Δ. These gains are likely to be higher, for example, if members' models tend to do well in explaining different parts of the state space. Second, committee members must be sufficiently willing to entertain each other's models (λ must not be too high) as to not introduce too much bias into the committee's decision. Third, the data must be sufficiently informative so as to not introduce too much variation into the model selected by the committee when it uses data to guide its choice of model.

Interestingly, while too much stubbornness hinders the committee's ability to improve performance, it is not always optimal for members to be only data-driven ($\lambda = 0$)—in equilibrium, it may be optimal to have committee members exhibit some degree of stubbornness. When the maximum likelihood weights on committee members' models are sufficiently variable, the performance of the committee model is maximized when committee members exhibit some degree of stubbornness.

Corollary 1 (Stubbornness Can Improve Committee Performance) *Denoting $w_A(\lambda)$ as the equilibrium weight chosen by the committee on model A, under a second-order approximation, when*

$$\mathbb{V}_\mu(w_A(0)) > \left(\frac{1}{2} - w_A^* \right)^2, \tag{13}$$

the expected log likelihood of the committee model

$$\mathbb{E}_\mu(\log(w_A(\lambda)m_A + (1 - w_A(\lambda))m_B)) \tag{14}$$

is maximized at $\lambda^* > 0$.

The logic of Corollary 1 is that, up to a second-order, maximizing the expected performance of the committee's model, $\mathbb{E}_\mu(w_A M_A + (1 - w_A)M_B)$, is equivalent to minimizing the

expected mean-squared error of the committee's chosen weights with respect to w^*,

$$\mathbb{E}_\mu\left[(w_A - w_A^*)^2\right] = \underbrace{(\mathbb{E}_\mu(w_A - w_A^*))^2}_{\text{Bias}} + \underbrace{\mathbb{V}_\mu(w_A)}_{\text{Variance}}. \tag{15}$$

Some stubbornness on the part of committee members can be helpful in reducing excess sensitivity to data in selecting a model. This can be understood in Bayesian terms as the committee being willing to put some weight on a prior of equal weighting models and not too much posterior weight on data when the data provide a noisy signal. Stubbornness does come at the cost of pushing the committee towards the biased, equally weighted model, and is particularly valuable when the benefit of reducing the variability w_A is high relative to the bias of weighting both members' models equally.

Remark 2 (Model Variance and Communication) *In our framework, the cost of variation in the weights assigned to each member's model from randomness in the data is statistical. However, there may be other costs, for example, regarding time-inconsistent decisions, and with regard to the central bank's credibility in the eyes of market participants, as emphasized in the debate on rules versus discretion (Kydland and Prescott, 1977; Barro and Gordon, 1983). Another potential cost of this variation comes in the form of communication. Our empirical analysis suggests that the committee's tilting may generate monetary policy surprises, potentially stemming from market participants not understanding the reaction function of the committee in response to macroeconomic news. Intuitively, it may be particularly difficult for the committee to communicate its future policy-making stance when the committee's choice of model exhibits a high degree of sensitivity to incoming data. For example, existing communication devices—such as the published dot plot, which plots anonymized forecasts of the short-term interest rate by each committee member, and the Summary of Economic Projections, which reports members' forecasts of macroeconomic variables—do not communicate how such forecasts are aggregated by the committee in its decision-making.*[20]

5.4 Transparency and Dissent Costs

We consider a simple extension of our model to include costs associated with committee members expressing dissent from the committee's chosen model. This analysis is motivated by prior work suggesting that central bank transparency, in the form of making committee deliberations public via the release of transcripts, increases conformity on the committee (Meade and Stasavage, 2008; Hansen, McMahon and Prat, 2018). We provide additional evidence extending this work in Appendix B.4.

[20]Prior work has recognized that the committee structure may pose challenges for communication due to a 'cacophony of voices' speaking about policy, rather than a single unified voice, e.g., Blinder (2007, 2008). Building on this idea, Vissing-Jorgensen (2021) suggests that the committee structure of the FOMC leads to a 'communication arms race' among members seeking to influence market expectations.

We modify committee members' objective functions in Equation (6) to be

$$U_i(\alpha_i) = \log \mathcal{L}(x|m) - \lambda KL(m_i \| v_i) - \gamma KL(m \| v_i). \tag{16}$$

The newly introduced term, $\gamma KL(m\|v_i)$, captures a dissent cost parametrized by γ that committee members pay to advocate for a model that deviates from the final model settled upon by the committee. When the cost of dissent is sufficiently low, the equilibrium model selected by the committee is unique; otherwise there may be multiple equilibria.

Proposition 4 (Equilibrium Multiplicity) *There is a $\bar{\gamma}$ such that $\gamma < \bar{\gamma}$ implies that the equilibrium is unique, and $\gamma > \bar{\gamma}$ implies that there are multiple equilibria.*

All else equal, a higher cost of dissent can help push the committee toward a model that best fits the economic data, since increased conformity can offset the costs that members face to advocate for a model that is too distinct from their own. However, it is not guaranteed that the committee will select such a model, as increasing the cost of dissent introduces multiple equilibria. When the cost of dissent is sufficiently high, a committee that otherwise would settle on a model that fits the data well may settle for a sub-optimal model, with members settling on such a model simply due to conformity. This is increasingly likely when the data provide little guidance as to what the optimal model might be. In addition to poor performance in explaining the data, multiplicity also imposes an additional cost in a policy making setting; it may make the committee's decision less predictable ex-ante, which, in the context of monetary policy, may impose real costs associated with macroeconomic volatility.

5.5 Implications for Committee Design

Our theoretical framework illustrates tradeoffs associated with committee decision-making, with practical implications for committee design. Our results highlight that while committees may improve decision-making performance in expectation, such improvements require that a committee's members are sufficiently willing to entertain each other's models and that the data are sufficiently informative.

One way to push members in the direction of entertaining each other's models is via structured discussion and deliberation. For example, the FOMC structures its meeting such that committee members first discuss their interpretation of economic and financial conditions, and then they discuss their views on the policy decision, before finally voting. Such deliberation can ensure that members are aware of the models of the other members and the potential merits of their explanations of economic conditions. It may also be important to select members that are willing to entertain one another's models, e.g., that are not too overconfident or stubborn in their viewpoints. At the same time, some degree of stubbornness can be helpful, to stabilize excess sensitivity in the committee's choice of model to the data.

Another institutional design choice is the chair of the committee. In the FOMC, the chair is well-aligned with the final decision made by the committee, indicating strong influence over the committee's decision. The committee chair may play an important role in minimizing the effect of potentially stubborn individual committee members adhering too much to their own models, and in stabilizing the model selected by the committee. Additionally, the chair may also play a role in equilibrium selection—when there is a strong preference for conformity on a committee, the chair may play an important role in which equilibrium model is selected by the committee's members, and ensuring that the committee does not settle on an underperforming model.

Lastly, our framework illustrates novel trade-offs associated with policies that affect the cost of dissent, such as transparency of deliberations. Increasing conformity can help offset the costs associated with members advocating for their own models too much. However, this comes at the potential cost of increasing the possibility of the committee choosing an incorrect model; and more generally of decreasing the predictability of the committee's policy decision-making.

6 Conclusion

In this paper, we study group decision-making at the FOMC. We find that members of the FOMC have distinct models of the economy—each reflecting a different way of interpreting economic data—that lead them to hold persistent differences in policy preferences. In a given period, the committee tilts its decision towards members whose models better match recent economic data. This tilting behavior induces monetary policy surprises.

These findings provide a new perspective on group decision-making, moving beyond the traditional perspective that committees aggregate distinct information, particularly in contexts where most relevant information may be common knowledge. Our evidence suggests that an important element in group decision-making might be that different members supply different models to interpret data, from which the group can draw in making its decision.

These insights have implications for institutional design. The effectiveness of committees depends not just on the diversity of perspectives, but on how they are combined into a final decision. Factors such as the cost of dissent, the role of the chair, and the transparency of deliberations influence whether committees effectively weigh different views or default to consensus. Our theoretical analysis illustrates a novel tradeoff associated with the FOMC's decision-making—tilting towards the models of members that fit the recent data better may improve its expected performance, but also risks excess sensitivity to data. Our findings suggest that, under the right conditions, committees can provide a systematic approach to decision-making under model uncertainty, but there are tradeoffs and design choices to consider.

References

Andre, Peter, Carlo Pizzinelli, Christopher Roth, and Johannes Wohlfart. 2022. "Subjective models of the macroeconomy: Evidence from experts and representative samples." *Review of Economic Studies*, 89(6): 2958–2991.

Andre, Peter, Philipp Schirmer, and Johannes Wohlfart. 2023. "Mental models of the stock market."

Aruoba, S Borağan, and Thomas Drechsel. 2024. "Identifying monetary policy shocks: A natural language approach." National Bureau of Economic Research.

Barahona, Ricardo, S Casella, Kristy AE Jansen, and Vincenzo Pezone. 2024. "Do Teams Alleviate or Exacerbate Overreaction in Beliefs?"

Barro, Robert J, and David B Gordon. 1983. "Rules, discretion and reputation in a model of monetary policy." *Journal of Monetary Economics*, 12(1): 101–121.

Bastianello, Francesca, Paul H Décaire, and Marius Guenzel. 2024. "Mental Models and Financial Forecasts."

Bauer, Michael D, and Eric T Swanson. 2023*a*. "An alternative explanation for the "fed information effect"." *American Economic Review*, 113(3): 664–700.

Bauer, Michael D, and Eric T Swanson. 2023*b*. "A reassessment of monetary policy surprises and high-frequency identification." *NBER Macroeconomics Annual*, 37(1): 87–155.

Bauer, Michael D, Carolin E Pflueger, and Adi Sunderam. 2024. "Perceptions about monetary policy." *Quarterly Journal of Economics*, 139(4): 2227–2278.

Belden, Susan. 1989. "Policy preferences of FOMC members as revealed by dissenting votes." *Journal of Money, Credit and Banking*, 21(4): 432–441.

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

Binetti, Alberto, Francesco Nuzzi, and Stefanie Stantcheva. 2024. "People's Understanding of Inflation." National Bureau of Economic Research.

Blinder, Alan S. 1999. *Central Banking in Theory and Practice.* Mit press.

Blinder, Alan S. 2007. "Monetary policy by committee: Why and how?" *European Journal of Political Economy*, 23(1): 106–123.

Blinder, Alan S. 2008. *The Quiet Revolution: Central Banking Goes Modern.* Yale University Press.

Bobrov, Anton, Rupal Kamdar, and Mauricio Ulate. 2024. "Regional Dissent: Do Local Economic Conditions Influence FOMC Votes?" *American Economic Review: Insights*. Forthcoming.

Bordo, Michael, and Klodiana Istrefi. 2023. "Perceived FOMC: The making of hawks, doves and swingers." *Journal of Monetary Economics*, 136: 125–143.

Bordo, Michael D, Klodiana Istrefi, and Humberto Martínez. 2024. "Rules vs. Discretion: Decoding FOMC Policy Deliberations." National Bureau of Economic Research.

Bowmaker, Simon W. 2025. *Fed Reckoning: Conversations on America's Central Bank.* Yale University Press. Manuscript in preparation.

Brainard, William C. 1967. "Uncertainty and the Effectiveness of Policy." *American Economic Review*, 57(2): 411–425.

Bybee, J Leland. 2023. "The ghost in the machine: Generating beliefs with large language models."

Caballero, Ricardo J, Tomás E Caravello, and Alp Simsek. 2024. "Financial Conditions Targeting."

Chappell Jr, Henry W, Rob Roy McGregor, and Todd Vermilyea. 2004. *Committee Decisions on Monetary Policy: Evidence from Historical Records of the Federal Open Market Committee.* MIT press.

Charness, Gary, and Matthias Sutter. 2012. "Groups make better self-interested decisions." *Journal of Economic Perspectives*, 26(3): 157–176.

Chen, Jiafeng, and Jonathan Roth. 2024. "Logs with zeros? Some problems and solutions." *Quarterly Journal of Economics*, 139(2): 891–936.

Cieslak, Anna, and Annette Vissing-Jorgensen. 2021. "The economics of the Fed put." *Review of Financial Studies*, 34(9): 4045–4089.

Cieslak, Anna, and Michael McMahon. 2023. "Tough talk: The fed and the risk premium."

Cieslak, Anna, Stephen Hansen, Michael McMahon, and Song Xiao. 2023. "Policymakers' uncertainty." National Bureau of Economic Research.

Cogley, Timothy, and Thomas J Sargent. 2005. "The conquest of US inflation: Learning and robustness to model uncertainty." *Review of Economic Dynamics*, 8(2): 528–563.

Cohn, Jonathan B, Zack Liu, and Malcolm I Wardlaw. 2022. "Count (and count-like) data in finance." *Journal of Financial Economics*, 146(2): 529–551.

Condorcet, Marquis de. 1785. "Essay on the Application of Analysis to the Probability of Majority Decisions." *Paris: Imprimerie Royale*, 1785.

Dunn, Wendy, Ellen E Meade, Nitish Ranjan Sinha, and Raakin Kabir. 2024. "Using Generative AI Models to Understand FOMC Monetary Policy Discussions."

Enke, Benjamin, Thomas Graeber, and Ryan Oprea. 2023. "Confidence, self-selection, and bias in the aggregate." *American Economic Review*, 113(7): 1933–1966.

Fos, Vyacheslav, and Nancy R Xu. 2025. "When Do FOMC Voting Rights Affect Monetary Policy?"

Fos, Vyacheslav, Tommaso Tamburelli, and Nancy R Xu. 2024. "Local monetary policy."

Gáti, Laura, and Amy Handlan. 2023. *Monetary communication rules.*

Gerling, Kerstin, Hans Peter Grüner, Alexandra Kiel, and Elisabeth Schulte. 2005. "Information acquisition and decision making in committees: A survey." *European Journal of Political Economy*, 21(3): 563–597.

Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus. 2021. "Five facts about beliefs and portfolios." *American Economic Review*, 111(5): 1481–1522.

Greenspan, Alan. 2004. "Risk and Uncertainty in Monetary Policy." *Speech presented at the Annual Meeting of the American Economic Association, San Diego, CA*.

Gürkaynak, Refet S, Brian Sack, and Eric Swanson. 2005. "The sensitivity of long-term interest rates to economic news: Evidence and implications for macroeconomic models." *American Economic Review*, 95(1): 425–436.

Hack, Lukas, Klodiana Istrefi, and Matthias Meier. 2023. "Identification of systematic monetary policy."

Hack, Lukas, Klodiana Istrefi, and Matthias Meier. 2025. "The Systematic Origins of Monetary Policy Shocks."

Handlan, Amy. 2022. "Text shocks and monetary surprises: Text analysis of fomc statements with machine learning."

Hansen, Lars Peter, and Thomas J Sargent. 2008. *Robustness.* Princeton university press.

Hansen, Stephen, Michael McMahon, and Andrea Prat. 2018. "Transparency and deliberation within the FOMC: A computational linguistics approach." *Quarterly Journal of Economics*, 133(2): 801–870.

Hansen, Stephen, Michael McMahon, and Carlos Velasco Rivera. 2014. "Preferences or private assessments on a monetary policy committee?" *Journal of Monetary Economics*, 67: 16–32.

Herbst, Edward, and Fabian Winkler. 2021. "The factor structure of disagreement."

Istrefi, Klodiana. 2019. "In Fed watchers' eyes: Hawks, doves and monetary policy."

Ke, Shikun. 2024. "Analysts' Belief Formation in Their Own Words."

Kydland, Finn E, and Edward C Prescott. 1977. "Rules rather than discretion: The inconsistency of optimal plans." *Journal of Political Economy*, 85(3): 473–491.

Madeira, Carlos, and João Madeira. 2019. "The effect of FOMC votes on financial markets." *Review of Economics and Statistics*, 101(5): 921–932.

Malmendier, Ulrike, Stefan Nagel, and Zhen Yan. 2021. "The making of hawks and doves." *Journal of Monetary Economics*, 117: 19–42.

Meade, Ellen E, and David Stasavage. 2008. "Publicity of debate and the incentive to dissent: Evidence from the US Federal Reserve." *Economic Journal*, 118(528): 695–717.

Meade, Ellen E, and D Nathan Sheets. 2005. "Regional influences on FOMC voting patterns." *Journal of Money, Credit and Banking*, 661–677.

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

Ottaviani, Marco, and Peter Sørensen. 2001. "Information aggregation in debate: who should speak first?" *Journal of Public Economics*, 81(3): 393–421.

Primiceri, Giorgio E. 2006. "Why inflation rose and fell: policy-makers' beliefs and US postwar stabilization policy." *Quarterly Journal of Economics*, 121(3): 867–901.

Riboni, Alessandro, and Francisco Ruge-Murcia. 2014. "Dissent in monetary policy decisions." *Journal of Monetary Economics*, 66: 137–154.

Romer, Christina D, and David H Romer. 2008. "The FOMC versus the staff: where can monetary policymakers add value?" *American Economic Review*, 98(2): 230–235.

Romer, Christina D, and David H Romer. 2023. "Presidential Address: Does Monetary Policy Matter? The Narrative Approach after 35 Years." *American Economic Review*, 113(6): 1395–1423.

Sargent, Thomas J. 1999. *The conquest of American inflation*. Princeton University Press.

Scanlan, Susannah. 2023. "Attention Allocation and the Factor Structure of Forecasts."

Scharfstein, David S, and Jeremy C Stein. 1990. "Herd behavior and investment." *American Economic Review*, 465–479.

Schonhardt-Bailey, Cheryl. 2013. *Deliberating American monetary policy: A textual analysis*. MIT Press.

Schwartzstein, Joshua, and Adi Sunderam. 2021. "Using models to persuade." *American Economic Review*, 111(1): 276–323.

Schwartzstein, Joshua, and Adi Sunderam. 2024. "Sharing Models to Interpret Data." Mimeo.

Shapiro, Adam Hale, and Daniel J Wilson. 2022. "Taking the fed at its word: A new approach to estimating central bank objectives using text analysis." *Review of Economic Studies*, 89(5): 2768–2805.

Sibert, Anne. 2006. "Central banking by committee." *International Finance*, 9(2): 145–168.

Swanson, Eric, and Vishuddhi Jayawickrema. 2024. "Speeches by the fed chair are more important than fomc announcements: An improved high-frequency measure of us monetary policy shocks."

Taylor, John B. 1993. "Discretion versus policy rules in practice." Vol. 39, 195–214, Elsevier.

Thornton, Daniel L, David C Wheelock, et al. 2014. "Making sense of dissents: a history of FOMC dissents." *Federal Reserve Bank of St. Louis Review*, 96(3): 213–227.

Vissing-Jorgensen, Annette. 2021. "Central banking with many voices: the communications arms race."

Wei, Jason, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. "Chain-of-thought prompting elicits reasoning in large language models." *Advances in neural information processing systems*, 35: 24824–24837.

Wu, Lingxuan. 2023. "Mental Macro-Finance Models: Evidence and Theory."

A Proofs

Proof of Proposition 1

Proof. The game is a potential game with potential function

$$\Phi(\alpha_A, \alpha_B) = \log \mathcal{L}(x|m) - \lambda(KL(m_A \| \alpha_A m_A + (1-\alpha_A)m_B) + KL(m_B \| \alpha_B m_B + (1-\alpha_B)m_A).$$

The potential function is strictly concave given the concavity of $\log \mathcal{L}(x|m)$ and the strict convexity of $KL(m_i\|\cdot)$ in its second argument. By strict concavity, and the fact that the strategy space is convex and compact, the potential function has a unique maximizer, which corresponds with the unique Nash equilibrium. □

Proof of Proposition 2

Proof. The existence of an interior w_A^* (and $w_B^* = 1 - w_A^*$) is guaranteed by the joint assumptions of distinctness and by the concavity of the function $\psi(w_A) = \mathbb{E}_\mu(\log(w_A^* m_A + (1 - w_A^*)m_B))$. □

Proof of Proposition 3

Proof. Taking a second-order Taylor approximation of $\psi(w_A) = \log L(x|w_A m_A + (1-w_A)m_B)$ around w_A^*, observe that

$$\psi(w_A) \approx \psi(w_A^*) - \frac{c}{2}(w_A - w_A^*)^2,$$

where $c \equiv -\psi''(w_A^*) > 0$, and the linear term disappearing as $\psi'(w_A^*) = 0$. Taking expectations, we have that

$$\mathbb{E}_\mu(\psi(w_A)) \approx \mathbb{E}_\mu(\psi(w_A^*)) - \frac{c}{2}\mathbb{E}_\mu\left((w_A - w_A^*)^2\right). \tag{17}$$

Note that

$$\mathbb{E}_\mu\left((w_A - w_A^*)^2\right) = (\underbrace{\mathbb{E}_\mu(w_A - w_A^*)}_{\equiv\beta(\lambda)})^2 + \underbrace{\mathbb{V}_\mu(w_A)}_{\equiv\sigma^2(\lambda)},$$

so,

$$\mathbb{E}_\mu(\psi(w_A)) \approx \mathbb{E}_\mu(\psi(w_A^*)) - \frac{c}{2}(\beta(\lambda)^2 + \sigma(\lambda)^2). \tag{18}$$

By assumption, since $\mathbb{E}_\mu(w_A^* m_A + w_B^* m_B) - \max\{\mathbb{E}_\mu(\log m_A), \mathbb{E}_\mu(\log m_B)\} > \frac{c}{2}(\beta(\lambda)^2 + \sigma(\lambda)^2)$, we have that up to a second order approximation, in expectation, the committee's model's performance exceeds that of the best member's model. □

Proof of Corollary 1

Proof. First, denote

$$w_{A,MLE} = \arg \max_{w_A \in [0,1]} \log \mathcal{L}(x | w_A m_A + (1 - w_A) m_B), \tag{19}$$

i.e., the weight on model A that maximizes the likelihood of the observed data. Then, given the structure of the game, we can write

$$w_A(\lambda) = \alpha(\lambda) w_{A,MLE} + (1 - \alpha(\lambda)) \frac{1}{2}, \tag{20}$$

where $\alpha(\cdot)$ is a continuous and strictly decreasing function that satisfies $\alpha(0) = 1$ and $\lim_{\lambda \to \infty} \alpha(\lambda) = 0$.

Next, observe from Equation (17), maximizing the expected likelihood is equivalent to minimizing $MSE(\lambda) = \mathbb{E}_\mu((w_A(\lambda) - w_A^*)^2)$. From Equation (18), we have that $MSE(\lambda) = \beta(\lambda)^2 + \sigma^2(\lambda)$. We can observe that $MSE(\infty) = (\frac{1}{2} - w_A^*)^2$ and $MSE(0) = (\mathbb{E}_\mu(w_{A,MLE} - w_A^*))^2 + \mathbb{V}_\mu(w_{A,MLE})$. Given the assumption that $\mathbb{V}_\mu > \left(\frac{1}{2} - w_A^*\right)^2$, and from Equation (20) we can see that $w_A(\lambda)$ is continuous in λ, it follows that the $MSE(\cdot)$ is minimized for $\lambda^* > 0$. □

Proof of Proposition 4

Proof. First, we show that there is an equilibrium. The strategy space $S = [0,1] \times [0,1]$ is convex and compact. Each member's utility $U_i(\alpha_i)$ is continuous in α_i and strictly concave due to the concavity of the log-likelihood term and the strict convexity of KL divergences. By the Maximum Theorem, best-response correspondences are non-empty, closed-valued, and upper hemicontinuous. Applying Kakutani's fixed-point theorem to the joint best-response correspondence guarantees existence of a pure strategy Nash equilibrium.

Next, we note that the equilibrium is unique when $\gamma = 0$, as shown in Proposition 1. Then note that as $\gamma \to \infty$, there are multiple equilibria—all strategies $\alpha_A = 1 - \alpha_B$ are equilibria. Lastly, the correspondence that maps γ to the set of equilibria is upper-hemicontinuous, due to the fact that members' best responses are single-valued and continuous. Upper hemicontinuity guarantees the existence of a threshhold $\bar{\gamma}$. □

B Data Appendix

B.1 Transcript Text

FOMC transcripts are provided in PDF form on the FOMC website. We download these PDF files and construct machine readable text files of the meeting transcripts using the Python library PyMuPDF. Speeches are identified in the transcripts by the use of capitalization. For example, speeches by Ben Bernanke are preceded by "MR. BERNANKE" and "CHAIRMAN BERNANKE" during his tenures as a member and as the chair of the board of governors. Reference to Ben Bernanke by other speakers are not capitalized, e.g., 'Chairman Bernanke,' which enables systematic identification of speeches from references to particular people during speeches.

We use the transcripts to create an ordered dataset of speeches during a committee meeting, identified by speaker name. Additionally, prior to 1997, prepared speeches and reports by staff members of the Federal Reserve Board that are presented during the meetings appear in appendices to the transcripts. We identify and append the text of the prepared speeches and reports to our dataset. We also manually inspect and clean names of speakers, to deal with occasional typos and inconsistencies across files. For each meeting, we also identify voters and their votes on the policy decision, obtained from the dataset of Thornton, Wheelock et al. (2014) and append this information to our dataset.

Our analysis particularly focuses on the discussion of the economic and financial situation and the monetary policy discussion within each meeting. Both sections typically begin with staff presentations. We identify these discussions by downloading agendas for each of the meetings from the Federal Reserve website, and passing the agendas and our dataset of meetings to Anthropic Claude, in order to identify the speeches in the dataset that correspond to each part of the meeting. We use the output to construct indicator variables in our dataset for the economic situation and policy discussion sections.

B.2 Prompts

Below we provide the prompts to Claude used to construct our dataset. The prompts structure the output of interest in machine readable XML form, which we then systematically process.

B.2.1 Prompt to Measure Committee Members' Policy Preferences

For each speaker-meeting combination, we run two sets of prompts using the below structure. The first prompt focuses on how macroeconomic factors (concerns about inflation, concerns economic growth and output, and concerns about labor markets) contribute to policy preferences; the second prompt focuses on how financial market factors (concerns about credit market conditions and concerns about the stock market) contribute to the committee member's policy preferences. The template for the prompt is given below.

```
"""
You are an expert economic analyst tasked with interpreting
   speeches from a Federal Open Market Committee (FOMC) meeting.
    Your goal is to understand a specific speaker's policy
   preferences and the key reasons behind those preferences.

Please review the following information:

1. FOMC Meeting Speeches:
<fomc_speeches>
{}
</fomc_speeches>
This is a CSV excerpt where each row represents a statement by a
    speaker. The 'tidyspeaker' column identifies the speaker for
    each row, and the 'within_speech_id' column provides a
   unique identifier for each statement.

2. Policy Alternatives:
<policy_alternatives>
{}
</policy_alternatives>
This is a CSV of the policy alternatives being evaluated during
   the meeting.

3. Speaker of Interest:
```

```
<speaker_of_interest>
{}
</speaker_of_interest>
We are focusing on this speaker's perspective.

4. Potential Preference Factors:
<preference_factors>
{}
</preference_factors>
These are the factors we need to evaluate as potential drivers
   of the speaker's policy preferences.

Your task is to analyze the speaker's policy preferences based
   on their speech. Follow these steps:

1. Determine which policy alternative (indicated by letter)
   aligns most closely with the speaker's preference.

2. Evaluate how each factor contributes to the speaker's policy
   preferences using the following scale:
-3: factor contributes strongly to preference for a smaller rate
    increase or bigger cut
-2: factor contributes moderately to preference for a smaller
   rate increase or bigger cut
-1: factor contributes weakly to preference for a smaller rate
   increase or bigger cut
0: factor not explicitly tied with policy preference
+1: factor contributes weakly to preference for a bigger rate
   increase or smaller cut
+2: factor contributes moderately to preference for a bigger
   rate increase or smaller cut
+3: factor contributes strongly to preference for a bigger rate
   increase or smaller cut

3. Provide explicit quotes that clearly indicate how each factor
    influences their policy preferences.

Before providing your final output, conduct a detailed analysis
```

```
   inside <detailed_analysis> tags. In this analysis:

1. List and count all statements by the speaker of interest,
   including their within_speech_id. Number each statement for
   easy reference.

2. For each policy alternative:
a. List key phrases or words that align with this alternative.
b. Count how many times these phrases or words appear in the
   speaker's statements. List each occurrence with a number.
c. List all relevant quotes from the speech, including their
   within_speech_id.
d. Analyze how these statements align with or diverge from the
   alternative.
e. Consider arguments for and against this being the preferred
   alternative.
f. Make a preliminary judgment on the likelihood of this being
   the preferred alternative.

3. Create a comparison table for policy alternatives with
   columns: Alternative, Key Phrases, Occurrence Count,
   Alignment Strength, Preliminary Judgment.

4. For each preference factor:
Create a table with the following columns:
- Factor
- Relevant Quotes (with within_speech_id)
- Preliminary Judgment (-3 to +3)
- Reasoning
- Confidence Level (Low, Medium, High)

5. Summarize the speaker's overall stance based on the analysis
   so far.

6. After considering all alternatives and factors:
a. Summarize your overall findings.
b. Weigh the evidence for each alternative and each factor.
c. Make and justify your final choice for the preferred
```

```
   alternative.
d. Make and justify your final decisions for each preference
   factor.
e. Double-check that all supporting quotes and reasons directly
   match the given factors and clearly support the speaker's
   policy preference.

After your analysis, provide your final output in the following
   XML format:

<output>
<preferred_alternative>[Letter of most aligned alternative]</
   preferred_alternative>
<preference_drivers>
    <factor1>[-3, -2, -1, 0, +1, +2, or +3]</factor1>
    <factor2>[-3, -2, -1, 0, +1, +2, or +3]</factor2>
    ...
    <factorn>[-3, -2, -1, 0, +1, +2, or +3]</factorn>
</preference_drivers>
<supporting_quotes>
    <factor1_quote>
    <text>[Quote text]</text>
    <within_speech_id>[ID number]</within_speech_id>
    </factor1_quote>
    <factor2_quote>
    <text>[Quote text]</text>
    <within_speech_id>[ID number]</within_speech_id>
    </factor2_quote>
    ...
    <factorn_quote>
    <text>[Quote text]</text>
    <within_speech_id>[ID number]</within_speech_id>
    </factorn_quote>
</supporting_quotes>
</output>

Ensure that your analysis is based solely on explicit statements
    in the speech, not inferences or assumptions. Be thorough in
```

```
    your analysis and provide clear justifications for your
   conclusions. Remember to double-check that all reasons and
   supporting quotes directly match the given factors and
   clearly support the speaker's policy preference before
   finalizing your output."""
```

B.2.2 Prompt to Measure Changes to the Policy Directive

For each meeting, we run a prompt using the below structure, in order to identify changes made to the policy directive; and whether those changes tilted towards being hawkish, neutral, or dovish. To run the prompt, we provide the LLM with all statements made during the policy discussion section of the FOMC meeting, and the policy alternatives being discussed.

```
"""
You are an expert analyst specializing in Federal Open Market
   Committee (FOMC) transcript analysis. Your task is to
   identify and classify successful changes made to the policy
   directive based on FOMC meeting speeches and policy
   alternatives.

First, carefully review the following information:

1. FOMC meeting speeches from the policy discussions section:
<fomc_speeches>
{speeches}
</fomc_speeches>

2. Policy alternatives discussed in the meeting:
<policy_alternatives>
{alternatives}
</policy_alternatives>

Your objective is to analyze these materials and provide a
   comprehensive assessment of successful policy changes. Follow
    these steps:

1. Analyze the policy alternatives and speeches.
2. Identify all changes made to the policy directive or
   statement.
```

```
3. Verify that each identified change was successfully adopted.
4. Classify each successful change on a scale from -3 to 3,
   using these guidelines:
-3: Highly accommodative / strongly dovish
-2: Moderately accommodative / dovish
-1: Slightly accommodative / mildly dovish
    0: Neutral
    1: Slightly restrictive / mildly hawkish
    2: Moderately restrictive / hawkish
    3: Highly restrictive / strongly hawkish
5. Group the changes into these categories:
a. Communication changes
b. Rate changes
c. Other policy changes

Before providing your final output, conduct a thorough analysis
   inside <detailed_analysis> tags. Include the following
   sections:

1. Policy Alternatives Analysis: List and analyze each policy
   alternative, including main points, proposed changes, and key
    economic indicators.

2. Speech Analysis: Analyze each speech, noting main arguments,
   proposed policy changes, and key economic indicators.

3. Economic Indicators and Key Data: Compile a comprehensive
   list of all key economic indicators and data points mentioned
    in both speeches and alternatives.

4. Policy Stance Comparison: Compare the current policy stance
   with previous meetings and analyze arguments for different
   policy stances.

5. Key Quotes: Extract and list key quotes from both speeches
   and policy alternatives that support potential changes.

6. Meeting Timeline: Create a timeline of events and discussions
```

```
    during the meeting.

7. Potential Changes Identification: List all potential changes,
    numbering each and including relevant quotes. For each
   change, list pros and cons for its adoption.

8. Verification of Successful Adoption: For each potential
   change, provide evidence of its successful adoption and
   assign a confidence level (high, medium, low).

9. Uniqueness Check: Review all identified changes to ensure
   each one is distinct. If any changes are similar, combine
   them or differentiate them more clearly.

10. Detailed Change Analysis: For each successfully adopted
   change:
- Draft a clear 1-2 sentence description of the change
- Analyze economic context and potential impact
- List arguments for each classification level (-3 to 3)
- Provide final classification with detailed justification,
   considering both explicit and implicit policy communications

11. Overall Meeting Analysis: Summarize the overall tone and
   combined impact of all successful changes.

After your analysis, present your findings in the following XML
   format:

<fomc_analysis>
<successful_changes>
    <change>
    <description>Brief description of the change (1-2 sentences)
       </description>
    <change_type>Communication/Rate/Other</change_type>
    <classification>Numerical classification (-3 to 3)</
       classification>
    <justification>Brief justification for the classification,
       including:
```

```
    - Evidence of successful adoption
    - Consideration of the macroeconomic context
    - Explicit and implicit policy communications</justification
      >
    </change>
    <!-- Add more successful changes as needed -->
</successful_changes>
</fomc_analysis>

Ensure your analysis is comprehensive, capturing all relevant
   successful changes and providing clear justifications for
   your classifications. Be thorough in identifying all
   successful changes, even those that might seem minor at first
    glance. Consider both explicit policy statements and
   implicit policy communications in your analysis, and always
   take into account the broader macroeconomic context when
   making your classifications.

Remember to provide a clear 1-2 sentence description for each
   change in the <description> tag of your final output.
"""
```

B.2.3 Prompt to Measure Alignment of Committee Members' with the Policy Decision

For each speaker-meeting combination, we run a prompts using the below structure. To run the prompt, we provide the LLM with all statements made during the policy discussion section (by the speaker of interest, as well as by other committee members) of the FOMC meeting, the policy alternatives being discussed, the previously identified policy changes, and the specific speaker of interest.

```
"""
You are an expert analyst tasked with evaluating the alignment
   between a specific speaker's policy preferences and the final
    policy decision made in a Federal Open Market Committee (
   FOMC) meeting. Your analysis will be based on the following
   information:
```

```
1. Speeches made during the FOMC meeting:
<speeches>
{speeches}
</speeches>

2. Policy alternatives discussed in the meeting:
<alternatives>
{alternatives}
</alternatives>

3. The policy changes that were successfully made in the meeting
   :
<policy_changes>
{policy_changes}
</policy_changes>

4. The specific speaker we are focusing on:
<speaker>
{speaker}
</speaker>

Your objective is to thoroughly analyze this speaker's
   statements in the context of the policy alternatives and the
   eventual decision. Then, you will rate how strongly aligned
   the speaker's policy preferences are with the final policy
   decision using this scale:

3: Perfectly aligned - The speaker's preferences fully match the
    final decision
2: Moderately aligned - The speaker's preferences largely align
   with the final decision, with minor differences
1: Weakly aligned - The speaker's preferences somewhat align
   with the final decision, but with significant differences
0: Neutral - The speaker's preferences neither align nor oppose
   the final decision
-1: Weakly opposed - The speaker's preferences somewhat oppose
   the final decision, but with some areas of agreement
-2: Modestly opposed - The speaker's preferences largely oppose
```

```
   the final decision, with minor areas of agreement
-3: Strongly opposed - The speaker's preferences fully oppose
   the final decision

Before providing your final rating and justification, conduct a
   thorough analysis using the following steps. Present your
   analysis inside <detailed_analysis> tags, focusing on
   conciseness and avoiding redundancy:

1. List all policy alternatives discussed in the meeting.
2. Identify and quote key statements from the focus speaker that
    indicate their policy preferences. Number each quote for
   easy reference.
3. Summarize the focus speaker's overall stance based on their
   statements.
4. Create a concise table comparing the focus speaker's stance
   with each policy alternative and the final decision, using a
   numerical scale from -3 to 3 for alignment.
5. Quote and analyze each policy change provided in the <
   policy_changes> section. For each change:
a. Quote the policy change verbatim.
b. Briefly describe the change.
c. Analyze how it aligns with the focus speaker's preferences (
   use a scale from -3 to 3).
6. Calculate an average alignment score for all policy changes.
7. Analyze the focus speaker's overall alignment with the final
   policy decision, considering both the specific policy changes
    and the broader context of the decision.
8. For each point on the rating scale (-3 to 3), list out
   arguments for why this rating might be appropriate,
   considering both specific policy changes and the broader
   context.

After completing your analysis, provide your final output in
   this format:

<rating>[Numerical rating between -3 and 3]</rating>
```

```
<justification>[A paragraph explaining your rating, with
   specific references to the speaker's statements and the final
    policy decision, including all changes agreed upon by the
   committee]</justification>

<proposed_changes>
[List each successful policy change in the following format:]
<change>
<index>[Numerical index of the change, from the first column of
   the policy_changes]</index>
<description>[Brief description of the change]</description>
<alignment>[Numerical value between -3 and 3 indicating the
   focus speaker's alignment with this change]</alignment>
</change>
</proposed_changes>

<speaker_alignment>[A brief analysis of the speaker's overall
   alignment with the final policy decision, considering both
   the specific policy changes and the broader context]</
   speaker_alignment>

Ensure that your output is structured and machine-readable,
   using the XML tags provided. Here's an example of the desired
    output structure (with generic content):

<rating>2</rating>

<justification>The speaker's policy preferences largely align
   with the final decision, with minor differences. [Specific
   examples and reasoning to follow...]</justification>

<proposed_changes>
<change>
<index>1</index>
<description>Generic policy change description</description>
<alignment>2</alignment>
</change>
<change>
```

```
<index>2</index>
<description>Another generic policy change description</
   description>
<alignment>1</alignment>
</change>
</proposed_changes>

<speaker_alignment>Overall, the speaker shows strong alignment
   with the final policy decision. While there are minor
   differences in some specific policy changes, the broader
   context of the speaker's statements indicates general support
    for the direction taken by the committee.</speaker_alignment
   >

Please proceed with your analysis and provide the structured
   output as requested.
"""
```

B.2.4 Prompt to Identify Member that Introduced Directive Change

For each identified change to the policy directive in each meeting, we also identify the member of the committee that introduced the change; whether the change was already part of the alternatives distributed as part of the Tealbook B / Bluebook; and the strongest advocate for the change during the meeting. To do so, we run the below prompts, providing the LLM with all speeches during the policy discussion section by committee members, the policy alternatives, and the previously identified policy changes.

```
"""
You are an expert analyst specializing in Federal Open Market
   Committee (FOMC) transcript analysis. Your task is to
   identify and classify successful changes made to the policy
   directive based on FOMC meeting speeches and policy
   alternatives.

Please review the following information:

1. FOMC meeting speeches from the policy discussions section:
<fomc_speeches>
{speeches}
```

```
</fomc_speeches>

2. Policy alternatives discussed in the meeting:
<policy_alternatives>
{alternatives}
</policy_alternatives>

3. Policy changes:
<policy_changes>
{directive_changes}
</policy_changes>

Your objective is to analyze each policy change and identify the
    following:
a) The speech ID where the change was initially introduced to
   the discussion.
b) The introducer of the change (using the tidyspeaker ID).
c) The name of the member who most strongly advocated for the
   change (using the tidyspeaker ID).
d) Whether the change was already present in one of the policy
   alternatives or newly introduced during the policy discussion
   .

Instructions:
1. Analyze the FOMC speeches, policy alternatives, and policy
   changes.
2. For each policy change, wrap your analysis in <
   change_analysis> tags:
a) Quote relevant parts of speeches that introduce the change.
b) List potential introducers, with arguments for each.
c) Quote relevant parts of speeches that advocate for the change
   .
d) List potential strongest advocates, with arguments for each.
e) Determine if the change was present in policy alternatives by
    quoting relevant parts or noting its absence.
f) Summarize your reasoning for the final choices.

It's okay for this section to be quite long, as we want a
```

```
   thorough analysis for each change.

3. After your analysis, present your findings in the following
   XML structure:

<policy_changes>
<change>
    <index>[Index of the policy change from the policy_changes
       variable]</index>
    <description>[Brief description of the change]</description>
    <introduction>
    <speech_id>[ID of the speech where the change was introduced
       ]</speech_id>
    <introducer>[Tidyspeaker ID of the person who introduced the
        change]</introducer>
    </introduction>
    <strongest_advocate>[Tidyspeaker ID of the strongest
       advocate for the change]</strongest_advocate>
    <in_alternatives>[Yes/No indicator if the change was present
        in policy alternatives]</in_alternatives>
</change>
<!-- Repeat for each policy change -->
</policy_changes>

Example output structure (with generic content):

<policy_changes>
<change>
    <index>1</index>
    <description>Generic policy change description</description>
    <introduction>
    <speech_id>SPEECH_XXX</speech_id>
    <introducer>MEMBER_X</introducer>
    </introduction>
    <strongest_advocate>MEMBER_Y</strongest_advocate>
    <in_alternatives>Yes</in_alternatives>
</change>
</policy_changes>
```

```
Please proceed with your analysis and provide the structured
   output for all policy changes.
"""
```

B.2.5 Prompt to Measure the Economic Data from Tealbook Texts

For each meeting, we measure incoming data pertaining to inflation, growth, labor markets, credit markets, and the stock market, as captured by the text of the Tealbook forecast, using the below prompt template. We run the prompts by grouping together macroeconomic factors in one run (inflation, growth, and labor), and grouping together financial market factors in another run (credit market conditions and the stock market).

```
"""
You are an expert economic analyst tasked with evaluating
   economic and financial conditions based on a Tealbook text.
   Your primary goal is to assess the strength of various
   economic factors and their potential impact on monetary
   policy decisions.

First, carefully read through the following Tealbook text:

<tealbook_text>
{}
</tealbook_text>

Now, consider the following list of factors that need to be
   evaluated:

<factors_to_evaluate>
{}
</factors_to_evaluate>

Your task is to analyze how these factors are discussed in the
   Tealbook text and evaluate their potential impact on monetary
    policy decisions. Follow these steps:

1. For each factor listed above:
   a. Identify relevant statements in the Tealbook text.
```

```
   b. Evaluate the strength of the factor using this scale:
      -3: Strong support for an interest rate cut
      -2: Moderate support for an interest rate cut
      -1: Weak support for an interest rate cut
       0: Neutral or not explicitly discussed
      +1: Weak support for an interest rate hike
      +2: Moderate support for an interest rate hike
      +3: Strong support for an interest rate hike
   c. Provide explicit quotes that support your evaluation.
   d. Explain your reasoning for the evaluation, specifically in
       relation to monetary policy implications.
   e. Assess your confidence level in the evaluation (Low,
      Medium, High).

2. Before providing your final output, wrap your detailed
   analysis in <factor_by_factor_analysis> tags. In this
   analysis:
   a. For each factor:
      - List all relevant quotes from the Tealbook.
      - Analyze each quote's implications for the factor and its
          potential impact on monetary policy.
      - Consider arguments for different strength ratings (-3 to
          +3), explicitly listing pros and cons for each
         potential rating in terms of supporting an interest
         rate cut or hike.
      - Make a preliminary judgment on the factor's strength.
   b. Create a summary table for all factors with columns:
      Factor, Relevant Quotes, Preliminary Strength Rating,
      Arguments For/Against (in relation to monetary policy),
      Confidence Level.
   c. Summarize your overall findings, focusing on the
      collective implications for monetary policy.
   d. Double-check that all supporting quotes directly match the
       given factors and clearly support your strength
      evaluations.
   e. Make and justify your final decisions for each factor's
      strength rating, ensuring they reflect the potential
      impact on monetary policy decisions.
```

```
3. After your analysis, provide your final output in the
   following XML format:

<output>
<factor_evaluations>
    <factor>
        <name>[Factor name]</name>
        <strength>[-3, -2, -1, 0, +1, +2, or +3]</strength>
        <quote>[Quote text]</quote>
        <reasoning>[Your reasoning for the strength rating,
           including monetary policy implications]</reasoning>
        <confidence>[Low, Medium, or High]</confidence>
    </factor>
    [Repeat for all factors]
</factor_evaluations>
<summary>[Brief summary of overall findings and their
   implications for monetary policy]</summary>
</output>

Important Notes:
- Base your analysis solely on explicit statements in the
   Tealbook text. Avoid inferences or assumptions.
- Be thorough in your analysis and provide clear justifications
   for your conclusions, always relating them back to potential
   monetary policy decisions.
- Ensure that all reasons and supporting quotes directly match
   the given factors before finalizing your output.
- The <factor_by_factor_analysis> section should show your
   thought process and consideration of different
   interpretations before reaching your final conclusions, with
   a specific focus on how each factor might influence decisions
    to cut or hike interest rates."""
```

B.3 Additional Tables and Figures

	y =*DirectiveAlignment*				y =*ChangeAlignment*				*y=Dissent*			
$m_{i,t-1}(data_t^{quant})$	0.22	0.12			0.21	0.14			-0.06	-0.03		
	(3.48)	(2.57)			(3.68)	(3.17)			(3.58)	(2.46)		
$m_{i,t-1}(data_t^{text})$			0.15	0.09			0.17	0.13			-0.03	-0.02
			(3.40)	(2.56)			(4.24)	(3.90)			(2.35)	(1.68)
R^2	0.27	0.47	0.26	0.47	0.29	0.46	0.29	0.47	0.16	0.31	0.15	0.30
N	3659	3659	3659	3659	3659	3659	3659	3659	3659	3659	3659	3659
Meeting Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Member Fixed Effect	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Member Type Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

TABLE B.1: MENTAL MODELS AND COMMITTEE ALIGNMENT CONTROLLING FOR MEMBER TYPE

Note: The table replicates the results of Table 7 including Member Typed Fixed Effects, which capture whether a member is a regional Fed president, a governor, or the chairman of the committee.

	y =*DirectiveAlignment*				y =*ChangeAlignment*				*y=Dissent*			
$m_{i,t-1}(data_t^{quant,macro})$	0.19	0.08			0.19	0.11			-0.05	-0.03		
	(2.79)	(1.60)			(3.17)	(2.31)			(2.80)	(1.85)		
$m_{i,t-1}(data_t^{text,macro})$			0.13	0.07			0.15	0.11			-0.03	-0.02
			(2.81)	(1.94)			(3.50)	(2.93)			(2.09)	(1.41)
R^2	0.21	0.46	0.21	0.46	0.23	0.46	0.23	0.46	0.14	0.30	0.13	0.30
N	3659	3659	3659	3659	3659	3659	3659	3659	3659	3659	3659	3659
Meeting Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Member Fixed Effect	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes

TABLE B.2: MENTAL MACROECONOMIC MODELS AND COMMITTEE ALIGNMENT

Note: The table replicates the results of Table 7 constructing $m(data)$ with macroeconomic variables ($inflation_t$, $growth_t$, $labor_t$).

	y =*DirectiveAlignment*				y =*ChangeAlignment*				*y=Dissent*			
$m_{i,t-1}(data_t^{quant,fin})$	0.20	0.09			0.18	0.09			-0.05	-0.02		
	(3.85)	(2.68)			(3.83)	(2.81)			(4.02)	(2.47)		
$m_{i,t-1}(data_t^{text,fin})$			0.13	0.07			0.16	0.10			-0.02	-0.01
			(3.40)	(2.59)			(4.35)	(3.66)			(2.05)	(1.55)
R^2	0.22	0.46	0.21	0.46	0.23	0.46	0.23	0.46	0.14	0.30	0.13	0.30
N	3659	3659	3659	3659	3659	3659	3659	3659	3659	3659	3659	3659
Meeting Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Member Fixed Effect	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes

TABLE B.3: MENTAL MACRO-FINANCE MODELS AND COMMITTEE ALIGNMENT

Note: The table replicates the results of Table 7 constructing $m(data)$ with macroeconomic variables ($stocks_t$, $credit_t$).

	y = *DirectiveAlignment*				y = *ChangeAlignment*			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$m_{i,t-1}(data_t^{quant})$	0.08	0.04			0.09	0.07		
	(2.11)	(1.37)			(2.19)	(2.01)		
$m_{i,t-1}(data_t^{text})$			0.08	0.04			0.11	0.09
			(3.34)	(2.03)			(4.28)	(3.45)
dissent	-2.66	-2.32	-2.68	-2.32	-2.31	-1.97	-2.32	-1.97
	(25.95)	(21.60)	(26.80)	(21.91)	(21.04)	(18.13)	(22.33)	(18.34)
R^2	0.62	0.71	0.62	0.71	0.54	0.64	0.54	0.64
N	3659	3659	3659	3659	3659	3659	3659	3659
Meeting Fixed Effect	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Member Fixed Effect	No	Yes	No	Yes	No	Yes	No	Yes

TABLE B.4: MENTAL MODELS, COMMITTEE ALIGNMENT, AND DISSENT

Note: The table replicates the first eight columns of Table 7, including a 0/1 indicator variable as to whether a committee member dissented in a meeting as an independent variable.

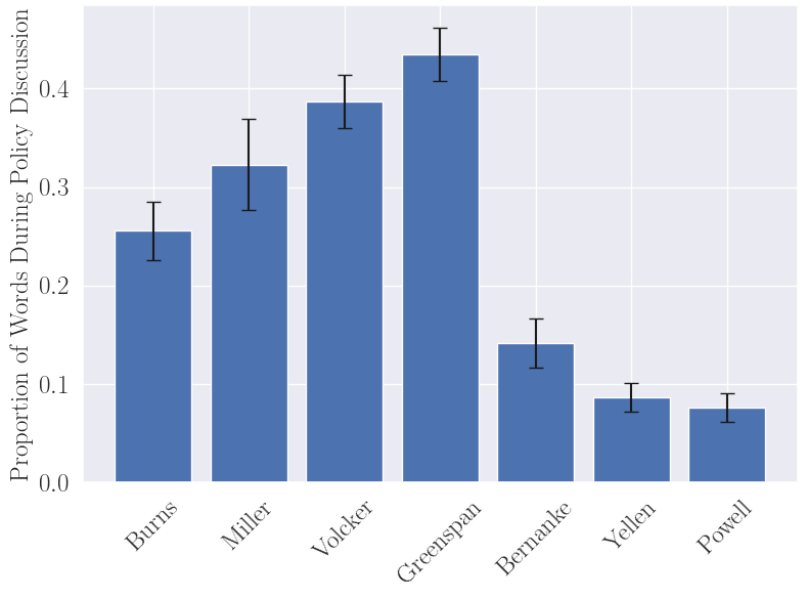


FIGURE B.1: PROPORTION OF WORDS DURING POLICY DISCUSSION

The figure plots share of words during the policy discussion spoken by each of the chairs of the committee.

B.4 FOMC Transparency and the Cost of Dissent

Here we focus on a design choice that may influence conformity on the committee, or, equivalently, the cost of expressing dissent during committee deliberations—the transparency of deliberations.

In 1993, following congressional pressure, the FOMC made the decision to release transcripts of its meetings on a five-year lag, and it has done so ever since. As noted in prior work studying Alan Greenspan's tenure as committee chair, this increased transparency may have the effect of increasing conformity on the committee (Meade and Stasavage, 2008; Hansen, McMahon and Prat, 2018) if members have career concerns, for example, as in the models of Scharfstein and Stein (1990) and Ottaviani and Sørensen (2001). We follow the general approach of Meade and Stasavage (2008), who study dissenting behavior during Alan Greenspan's tenure, and compare the alignment of regional Federal Reserve presidents with the alignment of Federal Reserve Board governors with the policy decision made by the committee before and after the knowledge of the existence of the transcripts and the subsequent decision to release the meeting's transcripts.[21]

The top panel of Figure B.2 plots the average *DirectiveAlignment* of regional Federal Reserve presidents and Federal Reserve governors pre- and post- transparency. The figure reveals a significant increase in the *DirectiveAlignment* of governors that persists throughout the post-transparency era, while there is little change in the average *DirectiveAlignment* for regional Fed presidents. The bottom panel of the figure plots the average $state_match^{text}$ of regional Federal Reserve presidents and Federal Reserve governors, measured using policy preferences expressed in period $t-1$ and the economic state measured in period t. The figure reveals little evidence of a similar pattern regarding alignment with the economic data in the pre- and post-transparency periods, indicating that the change seems to particularly be about the expressed alignment with the decision of the committee. Appendix Figure B.3 displays the plots using the *ChangeAlignment* and $state_match^{quant}$ measures, and suggests similar conclusions.

One interpretation of these results, made by Meade and Stasavage (2008), is that the increased alignment of governors, and the lack of an effect for presidents, is evidence of conformity driven by career concerns. The underlying logic is that as regional Fed presidents are responsible to different regional constituencies (regional boards of directors concerned with regional economic conditions) and do not have guaranteed voting rights in every meeting, they may be less susceptible to career concerns associated with the FOMC decisions than the

[21]The FOMC meetings had been recorded on tapes that, unbeknownst to members of the committee, were transcribed to text to assist with the writing of the meeting minutes. The majority of members of the committee learned of the existence of the transcripts in October 1993, amidst congressional pressure to increase the transparency of FOMC deliberations. The Fed quickly published existing transcripts (on a five-year lag), and was legally required from February 1995 onwards to publish subsequent transcripts. We follow Hansen, McMahon and Prat (2018) in using November 1993 as the first meeting where the transparency effect was known to the committee members.

governors, who work at the Federal Reserve Board under the chair and also tend to serve for shorter terms. We take this interpretation with caution, noting that there may also be other drivers of these results.[22] But we also note that it is interesting that the effect exists using more granular measures of policy preferences and has persisted in subsequent years under different committee chairs, particularly given other evidence regarding conformity induced by transparency during the Greenspan era (Hansen, McMahon and Prat, 2018).

[22]For example, during the Clinton administration, there was a shift towards selecting more academically trained economists as governors.

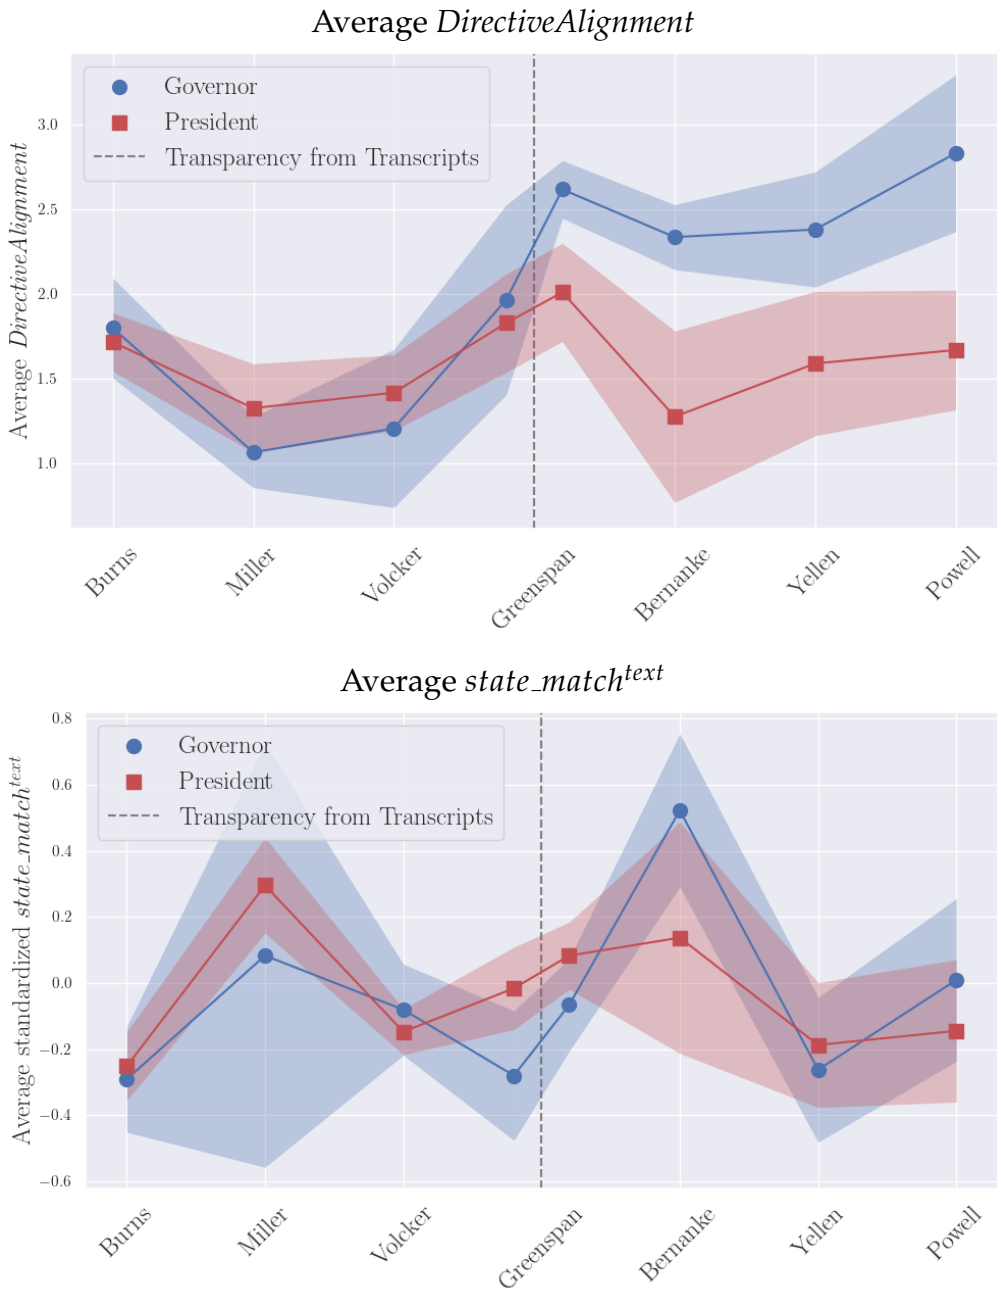


FIGURE B.2: ALIGNMENT WITH THE COMMITTEE DECISION: PRESIDENTS VERSUS GOVERNORS

The top panel of the the figure plots the average *DirectiveAlignment* of Federal Reserve Board governors and regional Federal Reserve presidents during the tenures of each chairperson; for Alan Greenspan's tenure, the figure separately plots the average *DirectiveAlignment* separately for meetings pre- and post-October 1993, which is when the majority of committee members became aware of the existence of transcripts of FOMC meetings that could be released to the public. The bottom panel plots the average $state_match^{text}$ measure for Federal Reserve Board governors and regional Federal Reserve presidents during the tenures of each chairperson, again separating the pre- and post-October 1993 periods. The figure also plots plus or minus two standard errors for the means, estimated by clustering standard errors by speaker and time.

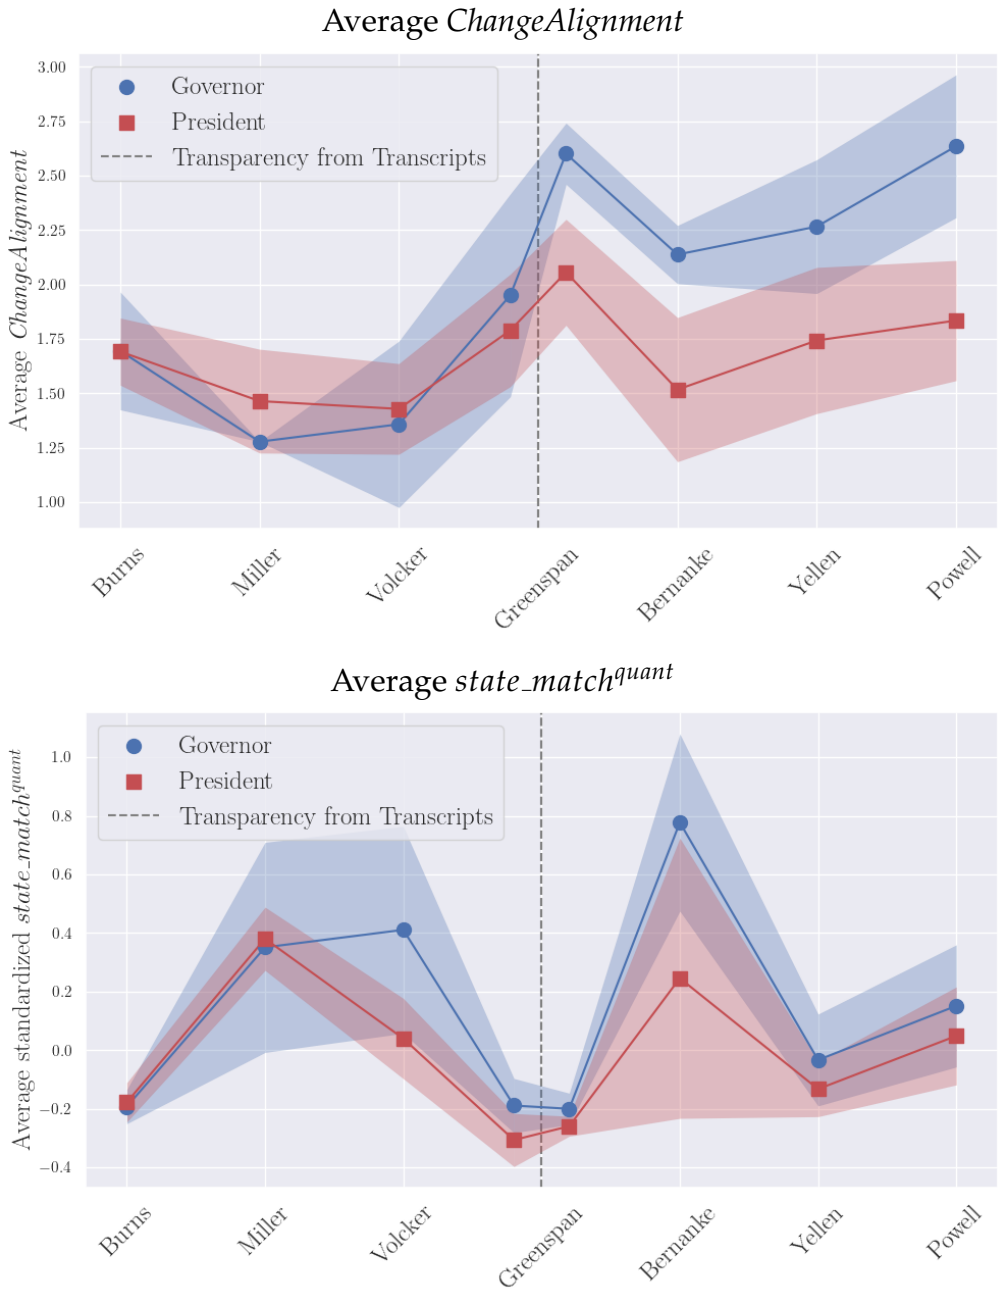


FIGURE B.3: ALIGNMENT WITH THE COMMITTEE DECISION: PRESIDENTS VERSUS GOVERNORS

The figures replicates Figure B.2 using the *ChangeAlignment* and $state_match^{quant}$ measures.

B.5 Framework Classification Methodology

This section provides detailed documentation of our methodology for classifying and analyzing FOMC members' interpretations of economic data. For each theme (inflation, growth, labor markets, and financial conditions), we construct comprehensive taxonomies of potential frameworks by prompting Claude to: Review academic literature and central bank documents to identify canonical theoretical frameworks; codify each framework's core mechanisms and assumptions; develop diagnostic criteria for identification from text; specify key parameters that characterize model implementation.

Tables B.6, B.5, B.7, and B.8 present the full taxonomy for each domain. For example, in the inflation domain we extract four frameworks: the Phillips Curve; Monetarism focusing on money growth and velocity; Cost-Push shocks centered on supply shocks and input costs; Expectations discussions around credibility and anchoring.

For each framework, we specify the core characteristics which define the fundamental theoretical mechanisms and assumptions that characterize each framework (e.g., for Phillips Curve: output gaps drive inflation through resource utilization). These represent the essential features that must be present for a particular framework to be identified. We specify the core phrases which provide linguistic markers that signal the potential use of each framework (e.g., references to "slack", "gap", "utilization"). These are not simple keyword matches but rather contextual indicators that help identify the underlying conceptual framework. We also provide disambiguation rules which specify how to differentiate between similar frameworks and establish the context necessary for proper classification. Lastly, we extract parameters specific to each framework (e.g., Phillips curve slope: flat/steep). These allow us to identify not just which framework is being used, but how it is being applied.

We implement the speaker-meeting framework classification by first parsing meeting transcripts into speaker-specific text blocks.[23] For each speaker-meeting pair, we employ the Claude large language model with the following structured prompt:

```
Input: Speaker's statements, framework taxonomy
Task: Evaluate presence of each model framework
Output format:
{
 "model_name": {
   "base_score": [0-1],       // Basic match
   "context_score": [0-1],    // Contextual appropriateness
   "combined_score": [0-1],   // Weighted average
   "confidence": [0-1],       // Classification confidence
   "parameters": {            // If sufficient evidence exists
```

[23] We impose a filter of (>50 words) for a speaker-meeting.

```
      "param_name": {
        "estimate": "value",
        "confidence": [0-1]
      }
    }
  }
}
```

Claude evaluates the text against each framework's diagnostic criteria and generates probability scores for framework presence. Our classification system assigns each framework a score between 0 and 1, with scores above 0.6 considered significant evidence of usage. This threshold balances sensitivity with specificity in identification. The scoring system considers both the presence of framework indicators and the strength of their manifestation in the text. The LLM also estimates parameter values when evidence exists and assesses the classification confidence.

TABLE B.5: Growth

Model	Core Characteristics	Core Phrases	Disambiguation	Parameters
Potential Output/Gap	• Natural rate exists • Production function foundation • Capacity constraints bind	• Economy operating above/below potential • Sustainable growth rate • Output gap suggests • Capacity utilization indicates • Structural versus cyclical	• Core: Output relative to potential • Context: Sustainable production • Rules: Evaluates capacity constraints	• Measurement: production function/filter • Gap persistence: temporary/persistent • Key constraints: labor/capital
Demand-Side	• Spending determines output • Price rigidities matter • Interest rate channel key	• Aggregate spending patterns • Interest-sensitive sectors • Confidence effects • Stimulus impact • Demand conditions	• Core: Spending drives output • Context: Demand components • Rules: Analyzes spending	• Rigidity: high/low • Interest sensitivity: high/low • Expectation role: key/limited
Supply-Side	• Production capacity primary • Resource constraints bind • Structural factors dominate	• Productive capacity limits • Structural impediments • Efficiency gains • Resource allocation • Productivity trends	• Core: Supply constraints key • Context: Production possibilities • Rules: Capacity focus	• Constraint: labor/capital/tech • Adjustment: fast/slow • Productivity: trend/cycle
Business Cycle	• Cycle-trend decomposition • Shock propagation focus • Pattern recognition key	• Stage of the cycle • Momentum indicators • Recovery pattern • Cyclical behavior • Leading indicators	• Core: Output follows cycles • Context: Cyclical dynamics • Rules: Focuses on position	• Cycle length: standard/extended • Shock persistence: temp/perm • Indicators: leading/coincident

TABLE B.6: Inflation

Model	Core Characteristics	Core Phrases	Disambiguation	Parameters
Phillips Curve	• Output/unemployment gaps • Expectations matter • Supply/demand balance	• Resource slack • Natural rate gap • Output pressure • Market tightness	• Core: Slack-inflation link • Context: Resource use • Rules: Gap analysis	• Slope: flat/steep • Slack measure: various • Expectations: adaptive/rational
Monetarist	• Money growth determines inflation • Direct quantity effects • Velocity relationships	• Excessive money growth • Monetary aggregates • Velocity patterns • Quantity framework	• Core: Money drives inflation • Context: Monetary factors • Rules: Growth analysis	• Aggregate: M2/broad money • Velocity: stable/varying • Transmission: direct/indirect
Cost-Push	• Supply shocks dominate • Input costs drive prices • Sectoral dynamics matter	• Cost pass-through • Supply bottlenecks • Wage pressures • Commodity effects	• Core: Supply drives prices • Context: Cost pressures • Rules: Input tracking	• Key inputs: wages/commodities • Pass-through: high/low • Persistence: temporary/permanent
Expectations-Centric	• Anchoring primary • Credibility crucial • Self-fulfilling possible	• Anchored expectations • Credibility of policy • Survey measures • Market compensation	• Core: Expectations drive • Context: Belief formation • Rules: Anchor analysis	• Formation: backward/forward • Anchoring: strong/weak • Credibility: high/low

TABLE B.7: Labor

Model	Core Characteristics	Core Phrases	Disambiguation	Parameters
Slack-Based	• Cyclical unemployment focus • Gap measures primary • Short-run dynamics	• Labor market slack • Cyclical unemployment • Hiring conditions • Job market tightness	• Core: Market tightness key • Context: Short-run labor • Rules: Utilization focus	• Slack measure: various • Adjustment: fast/slow • Wage response: elastic/rigid
Structural/NAIRU	• Natural rate emphasis • Long-run equilibrium • Structural determinants	• Natural rate level • NAIRU estimates • Structural changes • Long-run trends	• Core: Equilibrium exists • Context: Long-run state • Rules: Natural rate focus	• NAIRU: fixed/varying • Drivers: various factors • Adjustment: quick/slow
Wage Formation	• Wage-price dynamics • Expectation channels • Labor cost focus	• Wage pressure • Price-wage spiral • Labor cost growth • Compensation trends	• Core: Wage dynamics prime • Context: Price formation • Rules: Wage setting focus	• Rigidity: sticky/flexible • Expectations: key/limited • Indexation: high/low

TABLE B.8: Financial Conditions

Model	Core Characteristics	Core Phrases	Disambiguation	Parameters
Interest Rate Channel	• Policy rate transmission • Term structure effects • Asset price valuation	• Yield curve response • Interest-sensitive sectors • Term premium • Borrowing costs	• Core: Rate transmission • Context: Interest effects • Rules: Rate path focus	• Speed: fast/slow • Curve: parallel/twist • Sensitivity: high/low
Credit Channel	• Bank lending focus • Balance sheet effects • Financial frictions	• Credit availability • Lending standards • Balance sheet strength • Credit spreads	• Core: Credit conditions • Context: Lending dynamics • Rules: Credit access	• Frictions: high/low • Channel: strong/weak • Accelerator: active/limited
Risk-Taking Channel	• Portfolio rebalancing • Risk appetite shifts • Search for yield	• Risk appetite • Portfolio shifts • Reach for yield • Asset valuations	• Core: Risk behavior • Context: Portfolio choice • Rules: Risk focus	• Risk: high/low • Rebalancing: fast/slow • Impact: large/small
Financial Stability	• System vulnerabilities • Market functioning • Systemic risk	• Financial stability • Market functioning • Systemic risk • Vulnerabilities	• Core: System stability • Context: Market health • Rules: Stability focus	• Risk: high/moderate/low • Function: normal/stressed • System: stable/vulnerable

B.6 Frameworks and Policy Preferences

To visualize how theoretical frameworks map to policy preferences, we construct belief networks that trace the path from frameworks used to interpret economic data to policy positions. This provides a visual representation of how different frameworks map to policy stances across multiple dimensions: inflation, output, and labor markets.

To construct these network graphs, we start from our member-meeting level measures of mental models and parameter estimates. We identify the most prevalent parameter values for each framework, focusing on parameter estimates that appear with significant frequency and confidence in members' statements. To do this, we identify and keep the two most frequent parameter estimates for each framework-parameter combination. We construct the edge network by calculating the correlation between parameter values and policy preference scores. For visual clarity we apply a minimum threshold of 0.2 to display a connection on the graph. We color hawkish correlations red and dovish correlations blue.

We present these representations in Figures B.4, B.5, B.6 and B.7 below. The networks consist of three layers: frameworks (represented by blue nodes), their key parameters (light blue rectangles), and specific parameter values (white circles). These components connect to hawkish (red) or dovish (blue) policy positions through weighted edges, where edge weights reflect the correlation between a parameter value and policy preferences. We require these correlations to exceed a minimum threshold of 0.2 to exclude incidental associations.

The belief networks reveal theoretically coherent relationships between frameworks and policy preferences. In the inflation domain, Figure B.4, the strongest predictors of hawkish positions align with standard macroeconomic theory. High estimates of Phillips Curve slopes exhibit strong positive correlations (weights above 0.4) with hawkish views, consistent with models where monetary policy has greater employment-inflation tradeoffs. Similarly, weak expectations anchoring pairs with hawkish preferences – consistent with the idea that poorly-anchored expectations increase the risk of inflation persistence and thus warrant more aggressive policy responses. The cross-domain correlations of these parameter values with hawkish views on output and labor markets follows economic logic: members who believe in steeper Phillips Curves and weaker anchoring should rationally favor tighter policy across all dimensions given the higher perceived costs of inflation. The growth network, Figure B.5, displays similarly intuitive patterns. Members who perceive binding supply constraints or persistent output gaps advocate more hawkish positions, consistent with models where operating above potential creates inflationary pressure. There is strong complementarity across frameworks: both high interest rate sensitivity in demand-side models and binding supply constraints in potential output models predict hawkish preferences, reflecting alternative but theoretically valid channels through which growth pressures can necessitate tighter policy. Figures B.6 and B.7 in the Appendix present the structure of labor and financial conditions mental models respectively.

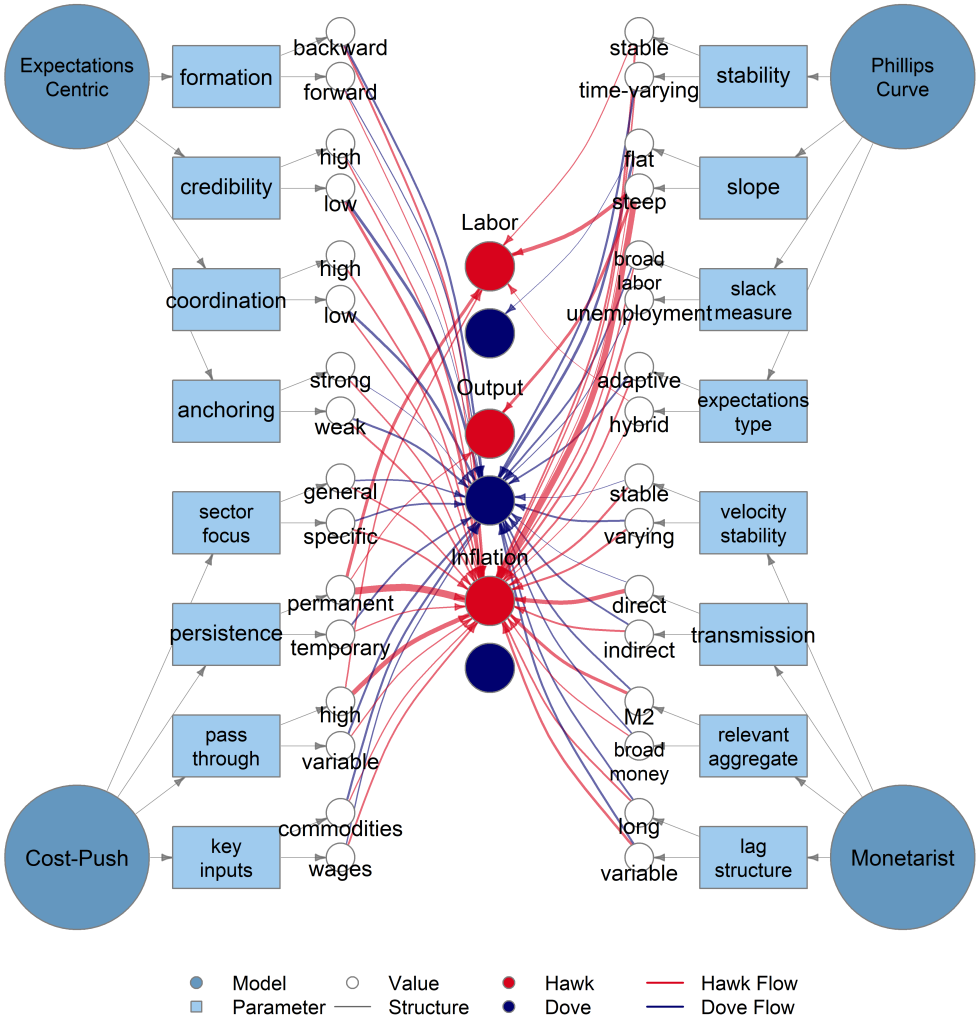


FIGURE B.4: INFLATION FRAMEWORKS AND BELIEF NETWORKS

Note: This figure illustrates the structure of inflation models and policy preferences. Nodes represent economic frameworks (blue circles), their associated parameters (gray rectangles), and parameter estimates (white circles). The central nodes capture policy preferences on inflation, output, and labor market outcomes, with hawkish positions in red and dovish positions in blue. Edges denote the relationships between parameter beliefs and policy preferences, with edge thickness reflecting the strength of association. Hawkish (red) and dovish (blue) connections indicate the directional influence of each parameter on the agent's policy stance.

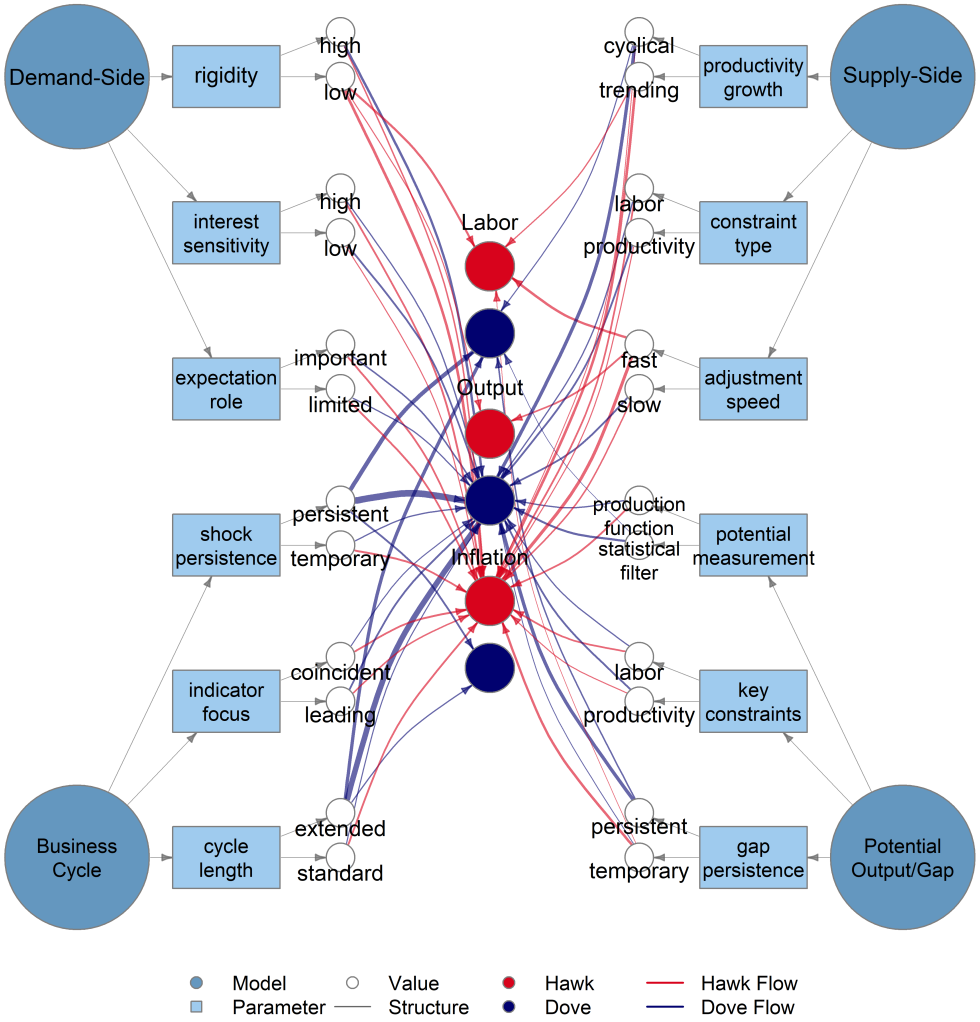


FIGURE B.5: GROWTH FRAMEWORKS AND BELIEF NETWORKS

Note: This figure illustrates the structure of growth models and policy preferences. Nodes represent economic frameworks (blue circles), their associated parameters (gray rectangles), and parameter estimates (white circles). The central nodes capture policy preferences on inflation, output, and labor market outcomes, with hawkish positions in red and dovish positions in blue. Edges denote the relationships between parameter beliefs and policy preferences, with edge thickness reflecting the strength of association. Hawkish (red) and dovish (blue) connections indicate the directional influence of each parameter on the agent's policy stance.

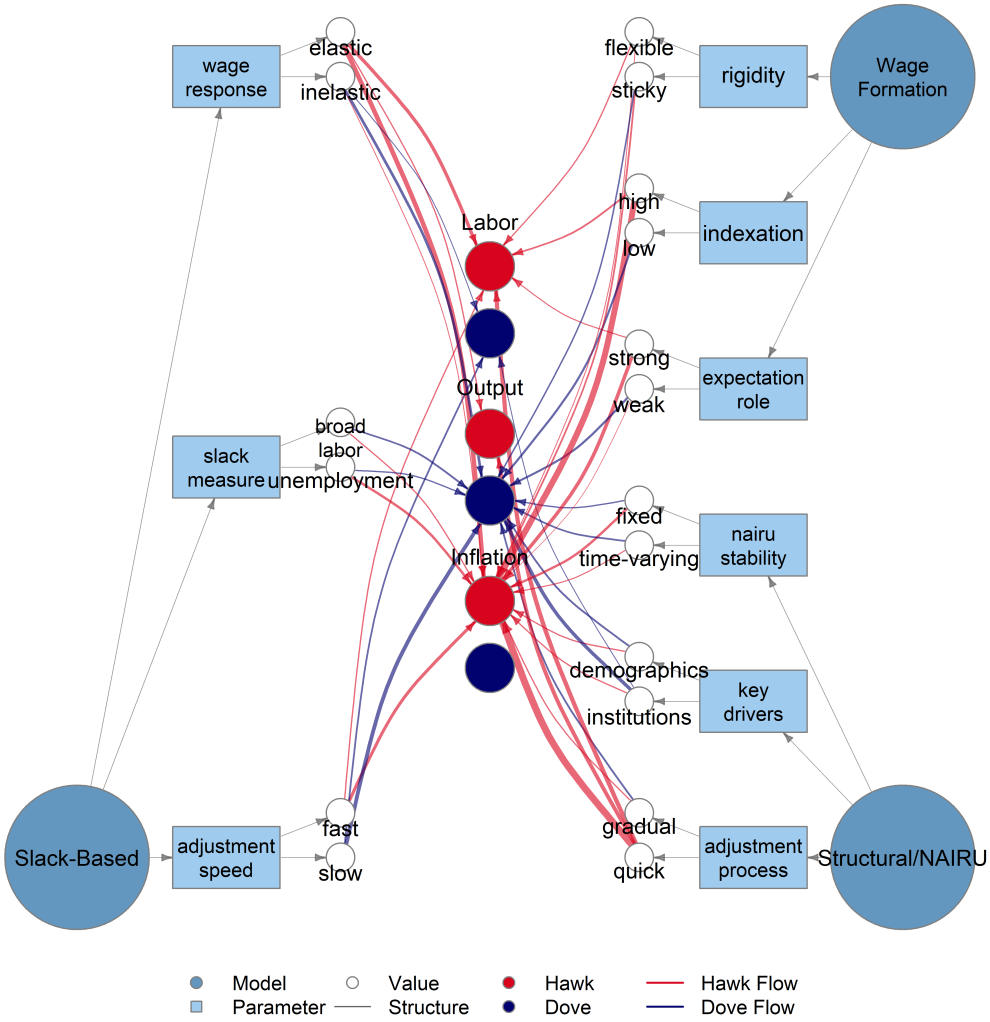


FIGURE B.6: LABOR FRAMEWORKS AND BELIEF NETWORKS

Note: This figure illustrates the structure of labor models and policy preferences. Nodes represent economic frameworks (blue circles), their associated parameters (gray rectangles), and parameter estimates (white circles). The central nodes capture policy preferences on inflation, output, and labor market outcomes, with hawkish positions in red and dovish positions in blue. Edges denote the relationships between parameter beliefs and policy preferences, with edge thickness reflecting the strength of association. Hawkish (red) and dovish (blue) connections indicate the directional influence of each parameter on the agent's policy stance.

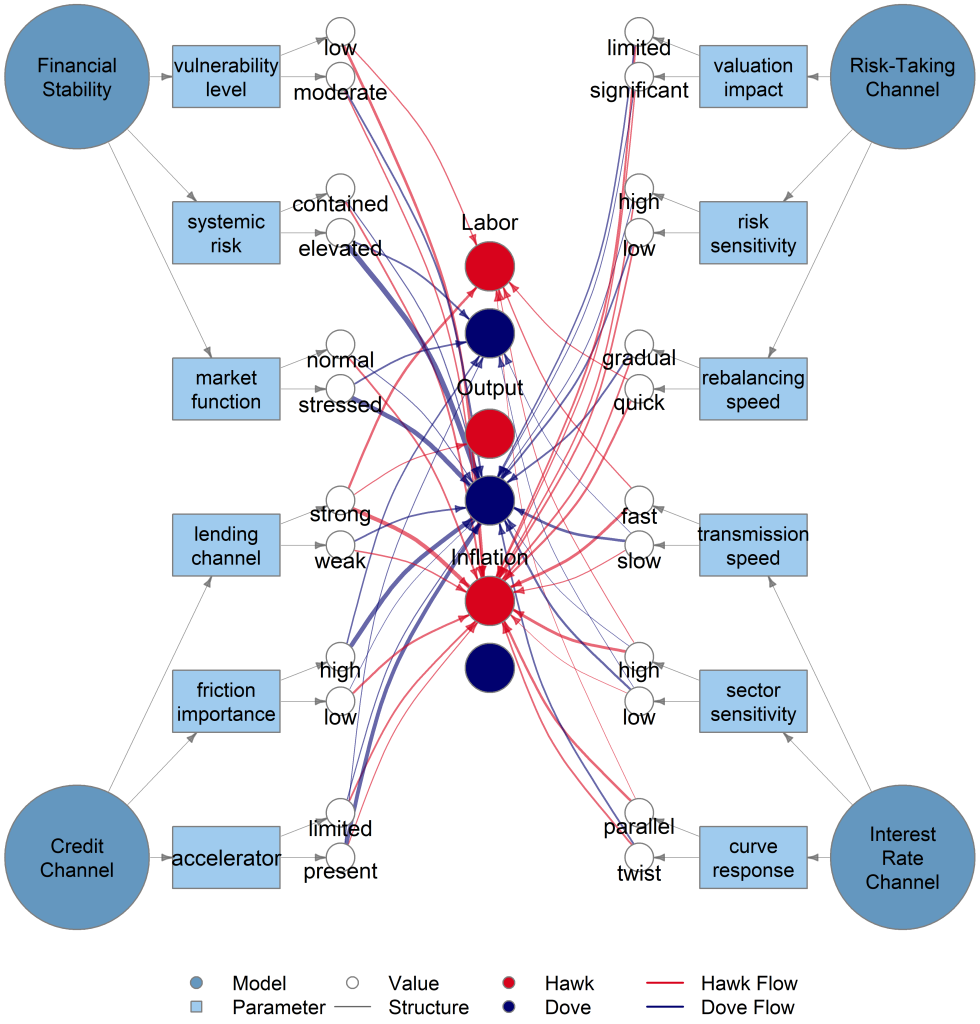


FIGURE B.7: FINANCIAL FRAMEWORKS AND BELIEF NETWORKS

Note: This figure illustrates the structure of financial models and policy preferences. Nodes represent economic frameworks (blue circles), their associated parameters (gray rectangles), and parameter estimates (white circles). The central nodes capture policy preferences on inflation, output, and labor market outcomes, with hawkish positions in red and dovish positions in blue. Edges denote the relationships between parameter beliefs and policy preferences, with edge thickness reflecting the strength of association. Hawkish (red) and dovish (blue) connections indicate the directional influence of each parameter on the agent's policy stance.