Central Banking by Committee

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

February 13, 2025

Preliminary and Incomplete: Not Intended for Distribution

Abstract

We study how groups make decisions using over forty years of meeting transcripts from the Federal Open Market Committee (FOMC). We document that committee members operate with distinct models of the macroeconomy, leading to persistent differences in policy preferences. On average, the committee tilts its decision toward the views of members whose models better align with recent data. Guided by these findings, we construct a theoretical framework of group decision-making under model uncertainty—where the true model is unknown—in which committee members interpret the same data through different models. When members are sufficiently open to each other's models, the committee may outperform all of its members in expectation. We discuss implications for committee design.

*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. We thank Nick Barberis, Bill English, and Alp Simsek for helpful comments and conversations. A substantial portion of this work was completed while Vasudevan was visiting Yale University, whose hospitality he gratefully acknowledges.

1 Introduction

Group decision-making is central to many economic settings, shaping outcomes across a variety of contexts. A longstanding literature, dating back to Condorcet's jury theorem (Condorcet, 1785) has emphasized the notion of the "wisdom-of-the-crowds": that groups can enhance decision quality by pooling diverse information. Yet, decisions are often also made by committee in settings where there may not be clear information-sharing gains. For example, in the context we study, monetary policy committees, Blinder (2008) argues that "you can mostly ignore [differential information] in the monetary policy context because virtually all the data that matter are common knowledge." Indeed, consistent with this argument, prior work has found that, on average, individual policymakers' forecasts add little value over and above those produced by central banking staff (Romer and Romer, 2008). The questions then arise as to why we observe decision-making in groups when the relevant information may already be shared, and how do groups make decisions under such circumstances.

In this paper, we analyze over forty years of verbatim transcripts from meetings of one of the most consequential decision-making groups in the world—the Federal Open Market Committee (FOMC)—to establish how the group decision-making process influences policy actions. We document evidence consistent with individual committee members operating with fixed and distinct "mental models" that they use to interpret economic data, and explore how the committee aggregates its members' diverse models. Guided by the evidence, we construct a theoretical framework taking 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—and we use it to map out novel trade-offs in collective decision-making. Below, we summarize our empirical and theoretical results.

Empirical Results. To conduct our study, we construct a novel dataset capturing the policy preferences of each member of the FOMC for every meeting from March 1976 through December 2019. In particular, we extract the extent to which individual committee members prioritize various macroeconomic and financial factors—inflation, growth, employment, credit market conditions, and the stock market—in shaping their respective policy stances. For each member at each meeting, we use a large language model to translate the member's statements into a set of scores where the magnitude reflects the importance of the factor in their decision-making and the sign indicates whether it pushes them toward a more dovish (negative) or hawkish (positive) stance. Additionally, we construct measures of each committee member's expressed alignment with the committee's final decision.

We validate our dataset by demonstrating that it captures relevant information only available for a short sub-sample of the data—individual committee members' forecasts of inflation, growth, and unemployment from the Survey of Economic Projections—as well as coarse but readily available data on committee members' voting behavior during meetings. Our measures not only track existing measures of committee members' beliefs and preferences well

but also carry the additional advantage of providing a long time-series of measures at a level of granularity not previously available for study.

With the dataset in hand, we examine committee members' preferences across meetings and over time. We document that member fixed effects are able to explain a substantial portion of members' policy preferences as driven by different factors—for example, members that are hawkish due to concerns about inflation are persistently so over the sample, while members that are dovish due to concerns about growth and unemployment remain persistently so. Member fixed effects explain approximately 20% to 30% of variation in individual policy preference measures. 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 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.

The evidence from our newly constructed policy preference measures suggests that each committee member operates with a fixed "mental model" of the macroeconomy, which they use to filter and interpret economic data and to generate forecasts for different variables. Heterogeneity in such mental models leads some members to have persistently higher or lower forecasts across variables, which in turn gives rise to strong fixed effects observed in policy preferences. Indeed, consistent with such an interpretation, we show that committee members' statements during the meetings align with a number of mechanisms emphasized in macroeconomic models—for example, regarding the strength of the trade-off between inflation and unemployment (i.e., the slope of the Phillips curve)—and find that this alignment explains a substantial portion of the cross-member variation in policy preferences. That is, using committee members' statements, we show that their persistent differences in policy preferences can be traced back to the use of specific frameworks for reasoning about macroeconomic variables, reinforcing the idea that they operate with different mental models.

As the key empirical finding, we document how the FOMC aggregates the heterogeneous policy preferences of its members. For each committee member in each meeting, we construct measures of how well their mental model—reflected in their policy preferences scores across variables—aligns with the observed state of the economy. We link individual's alignment with the final decision with measures of the match of their mental models with the state of the economy and find that members whose mental models better match the current state of the economy are also more aligned with the committee's decision. In other words, in making its decision the committee leans towards the views of the members whose mental models better match recent economic data, suggesting a novel source of variation in policy outcomes. This relationship holds both in cross-sectional comparisons across members, as well as in comparisons within members over time.

We find evidence for two distinct channels that could underlie this policy tilt towards

specific committee member's preferences. First, committee members whose models match the state of the economy well on average tend to speak more during the policy discussion section of meetings, suggesting that these members influence the committee's decision via deliberation. Second, we find that the chair of the committee consistently exhibits strong alignment with the committee's decision, underscoring that the chair may play an important role in steering the committee towards models that best fit the data.

We conclude our empirical analysis by examining the effect of a particular design choice of the committee—transparency of its deliberations—has on the committee's aggregation of members' mental models. Building on prior work examining the effect of transparency on dissenting behavior during Alan Greenspan's tenure as chair (Meade and Stasavage, 2008), we find that following increased transparency, Federal Reserve governors' stated alignment with the final policy decision increased substantially and has remained elevated in subsequent decades, compared with little-to-no effect for Federal Reserve presidents. There is no significant difference in how well governors' versus presidents' mental models match economic data pre- and post-transparency. One interpretation, following Meade and Stasavage (2008), is that this differential effect reflects transparency's tendency to increase conformity via career concerns. This is because FOMC governors, who serve under the chair and typically serve shorter tenures, have stronger career concerns associated with the transparency of FOMC deliberations than regional Federal Reserve presidents, who respond to their own respective regional constituencies and do not vote in all FOMC meetings. We present and make note of these results as suggestive evidence that committee design choices can affect conformity—or equivalently, the cost of expressing dissent—a theme that we further explore in our theoretical analysis.[1]

Theoretical Framework. Taken together, our empirical findings suggest that members of the FOMC operate with largely fixed models of the economy, with the committee decision tilting towards the models that best fit the recent data. That members operate with largely fixed models means that the traditional perspective (meaning group decision-making improves decision quality by providing diverse information) is insufficient in accounting for the costs and benefits of decision-making by committee.

Guided by our empirical evidence, we develop a simple and stylized theoretical framework of group decision-making under model uncertainty. Each committee member holds a fixed model that is partially misspecified and distinct from those of others. The committee's goal is to agree upon a model, which then determines its policy decision. Committee members observe data drawn from the true model. Each member advocates for a mixture of the committee members' models where they trade off (1) a desire for a final model that matches the data well; (2) a cost associated with supporting a model that deviates from their own; and

[1] Hansen, McMahon and Prat (2018) also document evidence of transparency increasing conformity; but suggest that it may have also improved discipline and effort by committee members.

(3) a cost of dissent associated with their advocacy deviating from the final model chosen by the committee. The committee's model is determined as the Nash Equilibrium of the members' advocacy game.

When members are sufficiently willing to entertain each other's models, that is, the cost of deviating from their personal model is low, an equilibrium arises in which the committee's model can outperform even the best committee member's model in expectation. This result is a consequence of the feature that each member's model is misspecified in a different way, leading to benefits to model diversification. However, to arrive at this superior model, the committee members must be willing to be consider alternative models. The less willing committee members are to entertain each other's models, the less the committee tilts towards models that best match the data, and the more the committee instead gravitates towards a simple average of the members' individual models.

The cost of dissent also plays an important role in shaping the committee's decision. Raising the cost of dissent fosters conformity. Higher conformity, in turn, may make committee members more willing to consider others' models and can introduce an equilibrium that improves upon all members' models even when committee members strongly prefer their own. However, such conformity can also lead to the emergence of multiple equilibria. Multiplicity can be costly by rendering the committee's decision unpredictable, and by making it possible that the committee settles on an equilibrium model that performs poorly.

The results of the theoretical framework indicate that committee decision-making can be a useful tool, even in the absence of differential information, when the true model of the world is uncertain. The theoretical results also carry practical implications for the design of committees. In particular, they suggest that in the face of model uncertainty, committees may improve overall decision-making performance. To provide such benefits, committees must implement design features that encourage the integration of diverse perspectives and mitigate the potential costs associated with dissent. For example, the structured deliberation process of the FOMC includes an economic round-table discussion during which members present their interpretations of economic data, followed by a policy discussion where potential disagreements are aired and suggested changes to the policy directive are made. This structure likely facilitates the integration of diverse perspectives. Conversely, certain design decisions——such as implementing transparency measures, which can elevate the cost of dissent——have ambiguous effects on committee performance.

Literature Review. Our paper contributes to an empirical and theoretical literature on group decision-making. Previous work has largely taken an informational perspective on group decision-making—that group decision-making may improve the quality of decisions made by aggregating information across a diverse set of sources (Gerling et al. (2005) and Sibert (2006) provide literature reviews). In contrast, we document that while information-sharing may be an aspect of the committee's function, in large part, individual committee

members often operate with different, fixed models, consistent with the ideas voiced in Blinder (1999, 2008). This evidence indicates that information aggregation may not be a sufficient explanation for the function of committees, and we argue that a key benefit of committees is in bringing together members with different, potentially incorrect 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 also relates to a primarily experimental literature in behavioral economics considering how group decision-making may affect decision quality given agents' behavioral biases (Charness and Sutter, 2012; Enke, Graeber and Oprea, 2023).

Our work is also related to a long tradition of work on 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 provides another way that policy makers address model uncertainty, and evaluates potential costs and benefits associated with the committee structure as a tool to deal with model uncertainty.

Finally, 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; Malmendier, Nagel and Yan, 2021; Bordo and Istrefi, 2023), 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 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 across various dimensions over a long sample, and we explicitly tie these preferences in with the macroeconomic models employed by committee members. 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; Ke, 2024; Bastianello, Décaire and Guenzel, 2024). Importantly, in addition to characterizing the mental models used by committee members, we also tie such models into the observed decisions made by the FOMC.

2 Evidence from FOMC Transcripts

To proceed with our empirical analysis, we construct several measures from FOMC transcripts. We summarize relevant institutional details and data sources and then 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 prepares 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 the 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 together 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 Federal Reserve presidents in these sections. The economic and financial situation discussion is led off by a presentation by Federal Reserve Board staff members, 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 of committee members about their policy

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

[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.

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 discussion following this date is with regard to the wording of the statement of the committee.

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 (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 the performance of these models 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.

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.[6] 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 factor for their policy preference, and the sign indicates whether the factor contributes dovishly (-) or hawkishly (+). The measure capture the 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.

[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/.

[6] We also extract speaker preferences with regard to gradualism, forward guidance, consistency with previous policy, and flexibility, but these are not the core focus of our analysis.

To construct the measures, we take each of the member's statements during the economic/financial situation and monetary policy sections of the 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 factor contributes strongly to a member's preferences in the direction of a bigger rate cut (or smaller hike), 0 indicates the factor does not contribute to their policy preference, and +3 indicates that the factor 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 provided score using direct quotes from the speaker. The LLM provides its initial reasoning in text form before making a final determination of its output.

As a concrete example, in the February 1988 meeting of the FOMC, the committee was considering: accommodative policy (Alternative A); maintaining current interest rates (Alternative B); or tightening rates (Alternative C). For Chairman Greenspan, our method provides scores of (+1, -1, -2, -2, -2) along the (inflation, growth, employment, credit, stock market) dimensions, with the LLM indicating that Greenspan appears to favor Alternative B, with a tilt towards accommodation. The LLM supports these scores with a number of quotes, including: [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", "increasing evidence from the initial claims figures that the economy was slowing down very dramatically," "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," and "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."

Figure 1 plots the average of committee members' estimated policy preferences in each meeting for each dimension. 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 factor 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 to 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).[7]

We also validate the ability of the policy preference measures 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 Survey of Economic Projections, which has been conducted prior to four 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 monotonic 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.

2.2.2 Alignment with the Policy Decision

For each meeting, we also construct measures of the alignment of each committee member with the policy decision made by the committee in that meeting. 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 dimensions, for example, targets for monetary aggregates, an important aspect of policy decisions early in our sample; and the inclusion or exclusion of specific words or phrases in the committee's official policy statement, important aspects of the decisions in the later part of our sample.

To construct our measures of alignment for a given committee member in a given meeting, we provide the LLM with the policy discussion section of the meeting and the policy alternatives from the Tealbook B (or the Bluebook). Then, for each speaker, we prompt the LLM to identify quotes and key statements by the speaker identifying their policy preferences. We also prompt the LLM to identify all instances of proposed changes to the policy directive that were adopted in the meeting. Then, 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, 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:

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

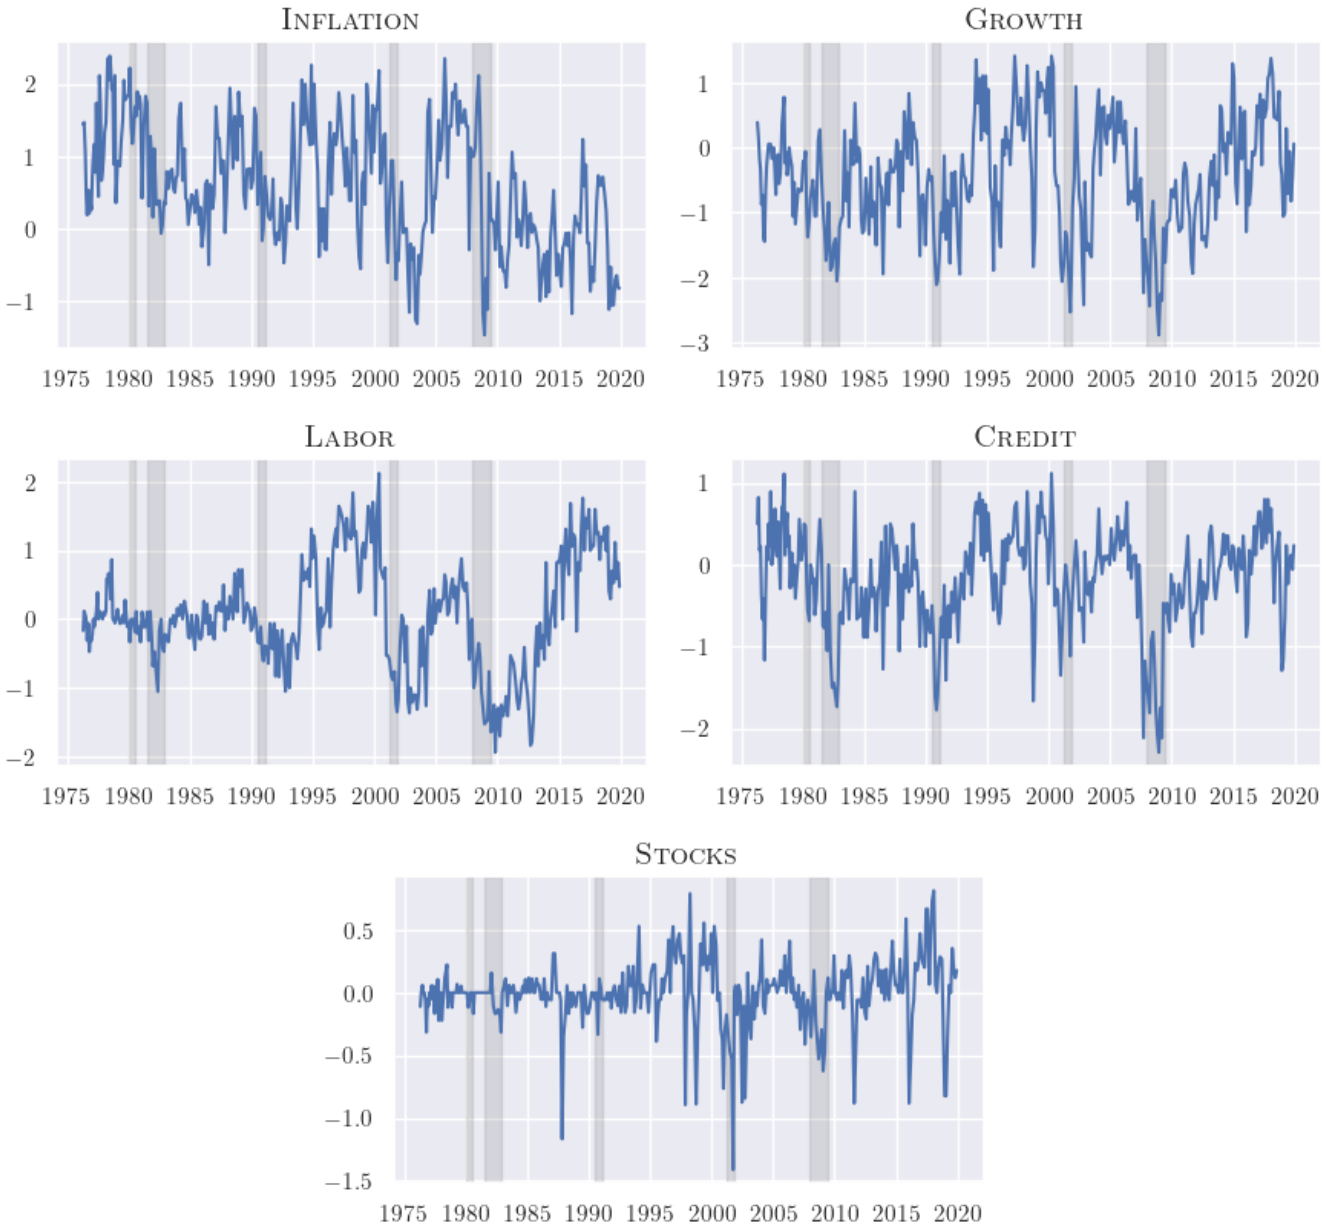


FIGURE 1: COMMITTEE MEMBERS' POLICY PREFERENCES OVER TIME

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 factor, with -3 indicating that a factor contributes strongly to a preference for more dovish policy and +3 indicating that a factor contributes to a preference for more hawkish policy. The dark grey bands indicate NBER recession periods.

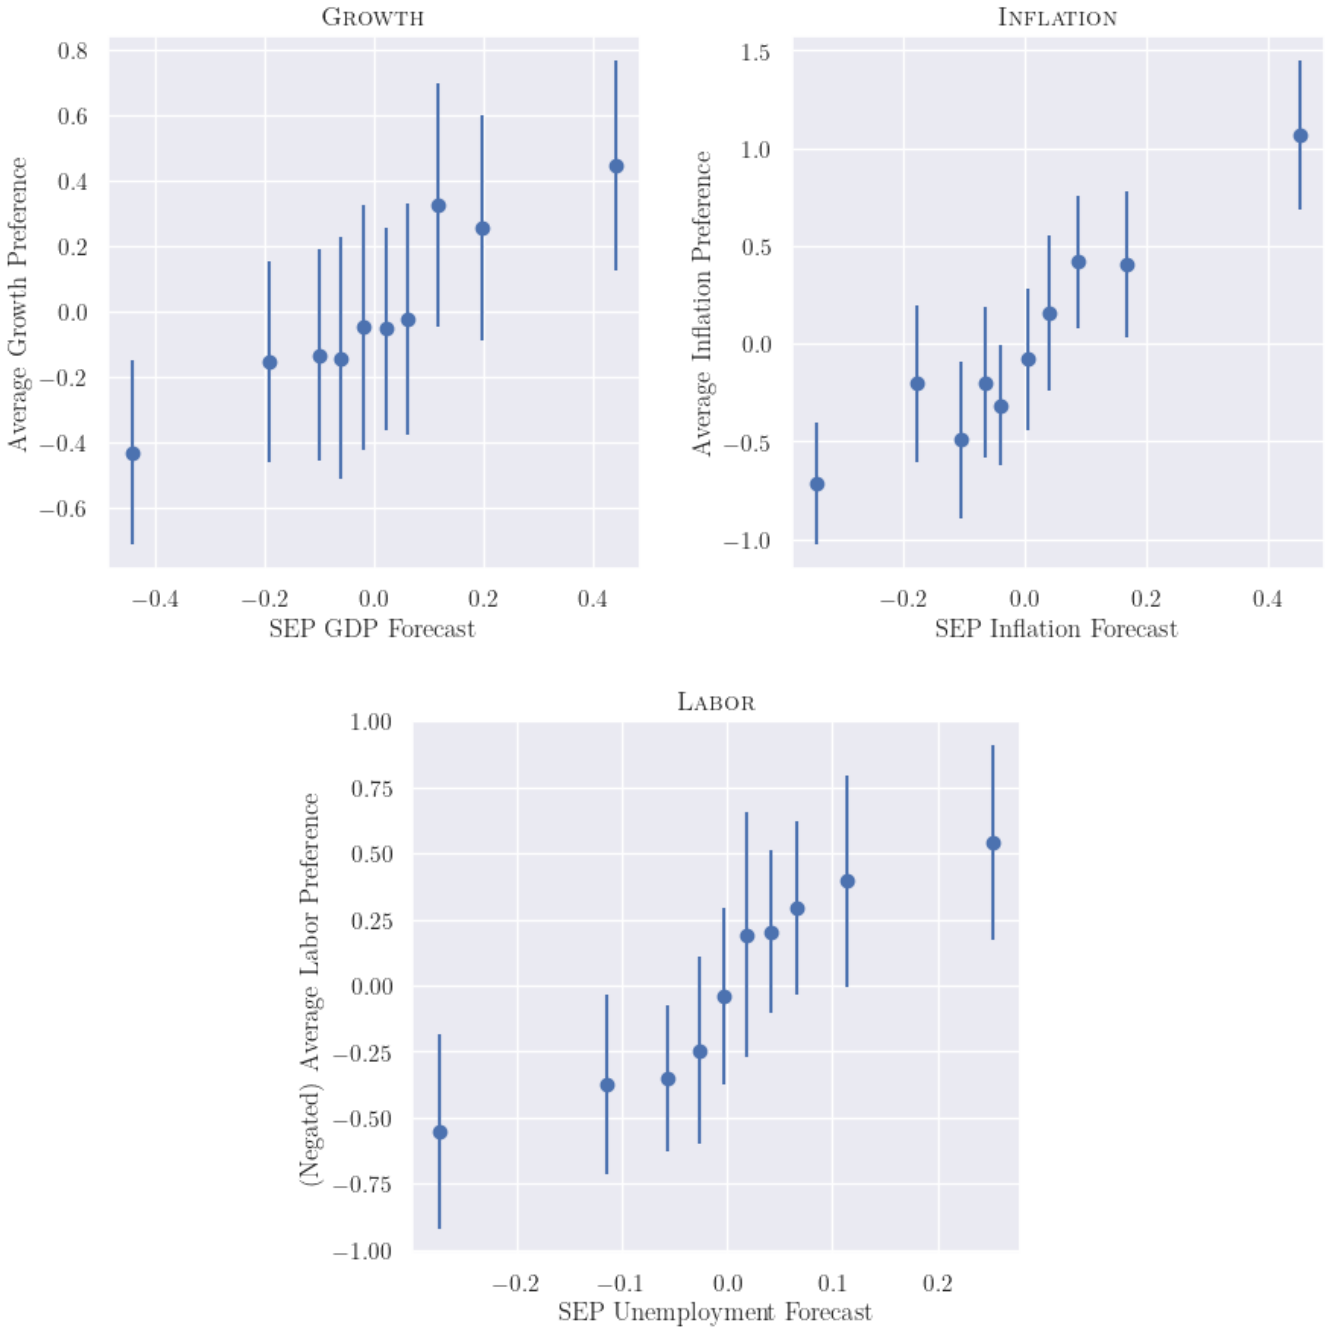


FIGURE 2: COMMITTEE MEMBERS' POLICY PREFERENCES AND SEP FORECASTS

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 Survey 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.

	DirectiveAlignment						
	-3.0	-2.0	-1.0	0.0	1.0	2.0	3.0
Average *ChangeAlignment*	-1.31	-0.41	0.11	0.14	0.93	2.01	2.80
Percentage dissenting votes	98.2%	85.2%	23.1%	0.0	1.9%	0.6%	0.0%
Number of voters	55	182	186	11	418	1361	1596

TABLE 1: DECISION ALIGNMENT AND DISSENT

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.

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 0.83 across observations in our sample.

We compare these measures of policy alignment with a measure commonly used in prior work: dissenting votes for the committee's policy decision.[8] We obtain voting data for FOMC meetings from Thornton, Wheelock et al. (2014). Table 1 groups each voting committee member during 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 dissent scores. Of the voting members of the committee with *DirectiveAlignment* scores of -3.0, -2.0 and -1.0, indicating disagreement with the decision made by the committee, 98.2%, 85.2%, and 23.1% of them dissent in their votes, respectively. Meanwhile, only a negligible portion of voters with positive *DirectiveAlignment* scores dissent in their voting.

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

[8]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). Chappell Jr, McGregor and Vermilyea (2004) also study committee members' target (ranges) for the Federal Funds rate for particular sub-samples of the data.

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 intuitional 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 less coarse measurement than voting dissent. They can also be constructed for non-voting regional Federal Reserve presidents in our sample. We further discuss the time-series variation of these alignment measures later in our analysis.

2.3 Committee Members' Policy Preferences

With these transcript-based measures, we examine committee members' preferences. We document evidence consistent with a substantial portion of committee members' policy preferences across dimensions being driven by member-specific "mental models" of the macroeconomy.

Table 2 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 factors, and ranging from 0.17 to 0.20 for the financial market factors. 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 factors, and from 0.07 to 0.22 for financial market factors. When taking the average across measures, meeting fixed effects explain 30% of preference variation and member fixed effects explain 29%. This finding echoes evidence 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

	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 2: 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 factor induces a hawkish or dovish preference, and the magnitude of the measures captures their importance for the committee member.

to concerns about growth. Jointly, meeting and member fixed effects explain over half of the variation in members' macro factor-driven policy preferences, 24% to 40% of the variation in their financial factor-driven policy preferences, and 56% of the variation in their overall policy preferences.

Next, we explore how members' preferences are related across different factors. Table 3 reports the principal components of the policy preference measures. Panel A reports principal components estimated from member-by-meeting observations. The first principal component loads positively on each of the factors, 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 through fifth principal components capture differential loadings on various factors, such as hawkishness based on inflation versus dovishness based on growth and credit market conditions (the second principal component).

Panel B of the table reports principal components from a matrix of member-level fixed effects for each factor, estimated after controlling for meeting fixed effects.[9] Once again, the first principal component loads positively on all factors, 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

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

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 various factors, 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 preference factors is explained by a hawk-versus-dove dimension. Without controlling for meeting fixed effects, committee members' policy preferences reflect the dominant macroeconomic or financial factors at the time; for example, committee members may be more concerned with inflation more 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.[10]

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. Potential sources of such mental models explored in the context of FOMC committee members might include past experiences (Malmendier, Nagel and Yan, 2021) and past economics education (Bordo and Istrefi, 2023).

2.4 Policy Preferences and Mental Models

We further explore the notion that committee members' persistent policy preference represent mental models through which they intepret economic data by systematically linking these preferences with different economics frameworks used by the committee's members. To do so, we first construct a taxonomy of economic frameworks spanning inflation, growth, labor markets, and financial conditions. We then analyze members' statements during meetings to identify their use of these frameworks. Our analysis reveals that different theoretical frameworks lead to systematically different policy preferences through economically interpretable channels. These mental models explain most of the variation captured by member fixed effects, suggesting persistent policy differences reflect genuine differences in economic thinking.

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

Panel A: Principal Components of Full Data					
Principal Component Loadings	PC1	PC2	PC3	PC4	PC5
Growth	0.55	-0.40	0.13	0.72	-0.01
Inflation	0.53	0.82	-0.20	0.08	0.01
Labor	0.42	-0.03	0.77	-0.48	-0.04
Credit	0.47	-0.40	-0.59	-0.48	-0.21
Stocks	0.12	-0.10	-0.09	-0.12	0.98
Proportion of Variance Explained	0.63	0.14	0.12	0.07	0.04

Panel B: Principal Components of Member Fixed Effects					
Principal Component Loadings	PC1	PC2	PC3	PC4	PC5
Growth	0.52	-0.07	-0.56	0.61	0.2
Inflation	0.55	0.25	0.76	0.20	0.16
Labor	0.37	0.69	-0.32	-0.43	-0.32
Credit	0.52	-0.68	0.02	-0.40	-0.33
Stocks	0.12	-0.04	-0.12	-0.50	0.85
Proportion of Variance Explained	0.84	0.07	0.04	0.03	0.01

TABLE 3: PRINCIPAL COMPONENTS OF COMMITTEE MEMBERS' POLICY PREFERENCES

Note: The table reports the results from principal components analysis of the inflation, growth, labor, credit, and stock market policy preference measures. Panel A reports results where observations are FOMC committee member by meeting preferences. Panel B reports results where observations are committee member fixed effects for each factor, estimated from regressions of the factors on member and meeting fixed effects.

We begin by constructing taxonomies of theoretical models within different themes – inflation, growth, labor, and financial conditions – identifying the mechanisms, diagnostic features, and key parameters that distinguish different models. 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.[11]

To map members' statements to these frameworks, we prompt the LLM to evaluate the statements of each member during each meeting against the diagnostic criteria of the relevant taxonomy. The LLM assigns probability scores indicating the presence of each model in a given member's statements and estimates key model-specific parameter values when sufficient evidence exists. This procedure generates a vector of parameter values for each meeting, capturing members' expressed views on economic mechanisms. For example, for inflation, we infer whether a member believes in a steep or flat Phillips Curve, whether inflation expectations are well-anchored or adaptive, and whether supply shocks play a dominant role. For growth models, we assess the strength of monetary transmission and the extent to which output is constrained by supply-side factors. To ensure these estimates reflect meaningful differences in economic reasoning rather than incidental references, we apply a confidence threshold of 0.6, restricting the analysis to cases where the classification is well-supported, as assessed by the LLM, by the text. This approach allows us to test whether persistent variation in policy preferences stems from stable differences in economic thinking. If mental models are a source of systematic heterogeneity in how members interpret macroeconomic conditions, then these classifications should explain a substantial portion of the variation in policy preferences attributed to member fixed effects. With these measures, we examine how mental models map to members' policy preferences and quantify their explanatory power relative to the fixed effects documented in the previous section.

First, to visualize how theoretical frameworks map to policy preferences, we construct belief networks that trace the path from economic models to policy positions. This provides a visual representation of how different theoretical views map to policy stances across multiple dimensions: inflation, output, and labor markets. We construct these networks by first identifying the most prevalent parameter values for each theoretical framework, focusing on parameter estimates that appear with significant frequency and confidence in members' statements. The networks consist of three layers: theoretical 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

[11] Section C in the Appendix provides the complete taxonomies and details on the measure construction.

exclude incidental associations.[12]

The belief networks reveal theoretically coherent relationships between frameworks and policy preferences. In the inflation domain, Figure 3, 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 4, displays similarly intuitive patterns. Members who perceive binding supply constraints or persistent output gaps advocate more hawkish positions (correlation weights exceeding 0.35), 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.[13]

[12]Section C.3 provides details on the figure construction.

[13]Figures C.1 and C.2 in the Appendix present the structure of labor and financial conditions mental models respectively.

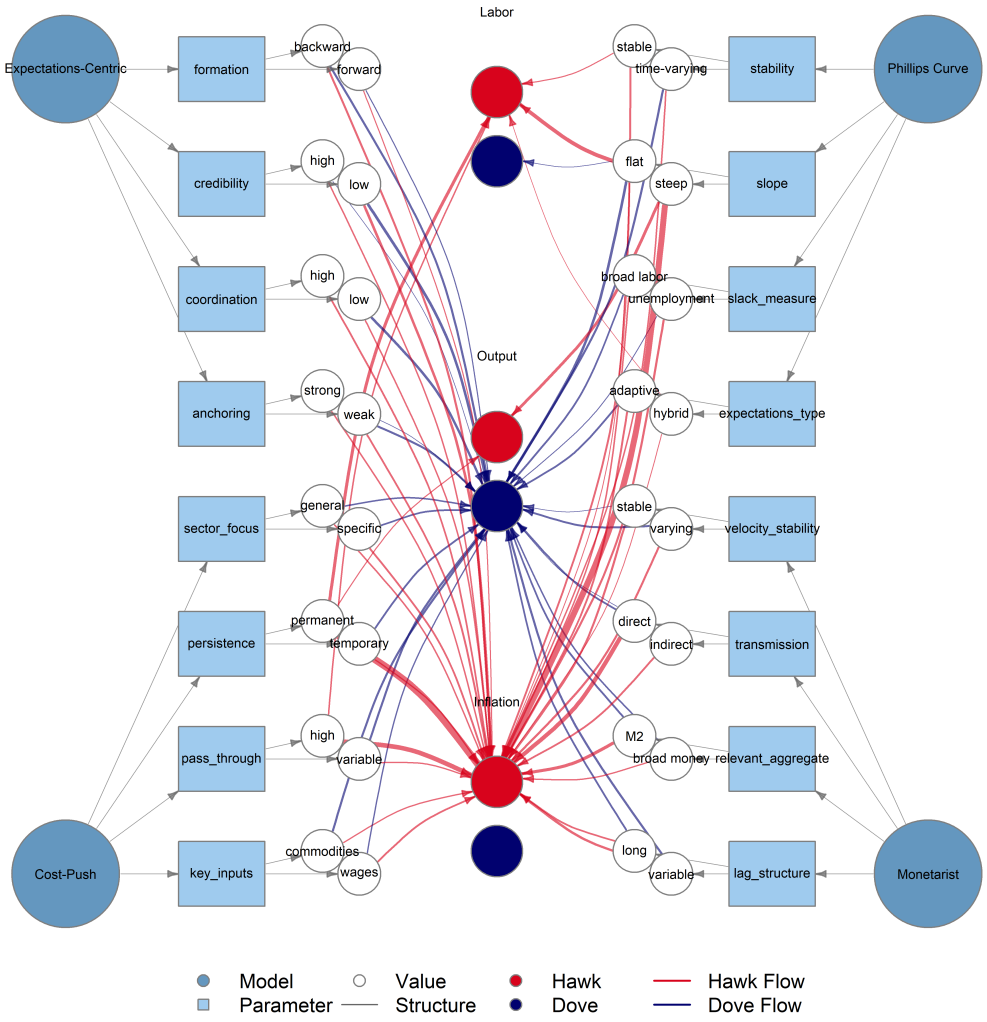


FIGURE 3: INFLATION MODELS AND BELIEF NETWORKS

Note: This figure illustrates the structure of inflation models and policy preferences. Nodes represent economic models (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.

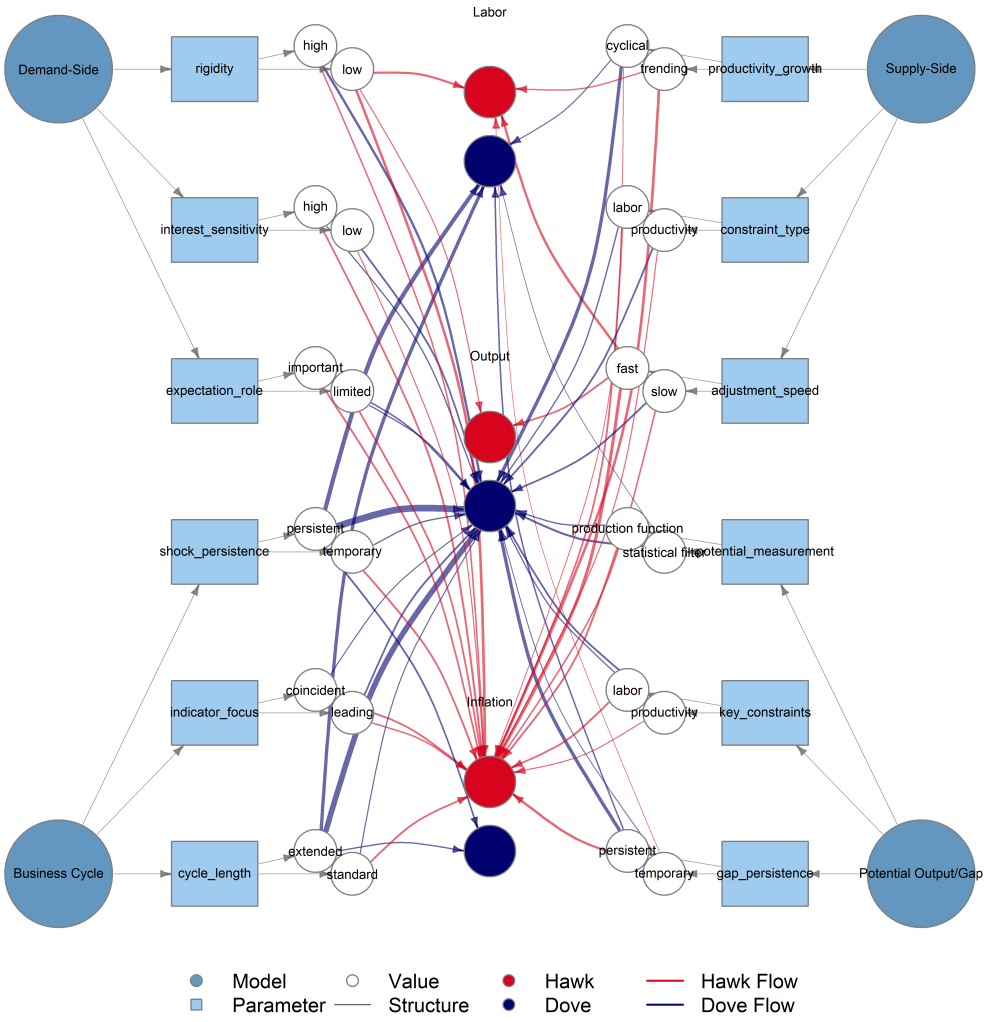


FIGURE 4: GROWTH MODELS AND BELIEF NETWORKS

Note: This figure illustrates the structure of growth models and policy preferences. Nodes represent economic models (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.

Fixed Effect Decomposition. To assess whether persistent differences in policy preferences reflect stable differences in members' economic frameworks, we extend the fixed effects analysis from Table 2. Specifically, we evaluate whether member fixed effects can be explained by systematic variation in members' theoretical frameworks. If mental models drive policy preferences, then replacing member fixed effects with direct measures of members' economic beliefs should capture much of the observed heterogeneity.

We estimate regressions in which the dependent variable is the policy preference measure for inflation, growth, and labor markets, as in Table 2. Rather than relying solely on member fixed effects, we introduce categorical variables representing members' views on key model parameters, such as the steepness of the Phillips Curve and the degree of expectations anchoring. This allows us to quantify how much of the persistent variation in preferences is attributable to differences in economic reasoning rather than other individual-specific characteristics. The parameter-level specifications follow:

$$\gamma_{it}^{k} = \alpha + \sum_{m} \beta_m \cdot \mathbb{1}(\theta_{it}^{m}) + \varepsilon_{it}^{k} \tag{1}$$

where γ_{it}^{k} is the policy score along dimension $k \in \{Inflation, Output, Labor\}$ of member i in meeting t, $\mathbb{1}(\theta_{it}^{m})$ is an indicator (categorical variable) for the parameter values from different models m, specific to each member i and meeting t.

	Inflation	Growth	Labor
Panel A: Baseline Fixed Effects			
Time FE Only	0.31	0.35	0.36
Member FE Only	0.31	0.20	0.23
Panel B: Parameter-Level Analysis			
Parameters Only	0.31	0.35	0.34
Parameters + Member FE	0.47	0.43	0.44
Parameters + Time FE	0.47	0.48	0.50

TABLE 4: Mental Models and Policy Preferences

Notes: The table reports R^2 values from regressions of inflation, growth, and labor market policy preference measures on meeting fixed effects and FOMC committee member fixed effects (Panel A), and on parameters from members' expressed theoretical frameworks (Panel B). The policy preference measures range from -3 to 3, with positive values indicating hawkish preferences and negative values indicating dovish preferences. The magnitude of these measures captures their importance for the committee member. Parameters capture specific theoretical views expressed by members during FOMC meetings, including beliefs about key economic mechanisms and relationships.

Table 4 reports the results.[14] Panel A replicates the baseline fixed effects estimates. Meet-

[14]Table B.5 in the appendix presents the results showing adjusted R^2. The results are similar.

ing fixed effects explain between 31% and 36% of preference variation, while member fixed effects explain between 20% and 31%, consistent with the substantial role of individual differences documented in the previous section. Panel B reports the results from specifications following Equation 1 of policy preferences on members' mental models. Parameter specifications account for 31% to 34% of variation, comparable to the explanatory power of member fixed effects alone. When combined with member fixed effects, the explanatory power increases substantially to 47% for inflation, 43% for growth, and 44% for labor markets, suggesting that a significant share of the observed individual variation in policy preferences is rooted in stable differences in economic frameworks. The additional explanatory power of member fixed effects may reflect additional sources of heterogeneity, for example differences in risk tolerance, loss functions, or personal weighting of policy trade-offs, that are not fully captured by model parameters; or alternatively, measurement error.

The similar explanatory power achieved when combining parameters with time fixed effects (47% to 50%) further supports the view that mental models are a distinct and meaningful source of variation in policy preferences, complementing rather than substituting for the shared response to macroeconomic conditions. Overall, these findings provide evidence that the persistence of member fixed effects is not merely an artifact of individual idiosyncrasies but rather reflects fundamental differences in how members conceptualize economic mechanisms.

2.5 The Committee Decision and Aggregating Models

We next analyze how the FOMC aggregates the preferences of its members. We document evidence that in any given meeting, the committee's decision tilts towards the committee members whose models are best aligned with the recent macroeconomic data.

We proceed using the policy preference measures across the economic and financial dimensions. Specifically, we represent an individual member's mental model as a vector of their policy preferences along each dimension, i.e. as

$$m_{i,t} = \begin{pmatrix} \gamma_{i,t}^{\text{inflation}} \\ \gamma_{i,t}^{\text{growth}} \\ \gamma_{i,t}^{\text{labor}} \\ \gamma_{i,t}^{\text{credit}} \\ \gamma_{i,t}^{\text{stocks}} \end{pmatrix}, \tag{2}$$

where $\gamma_{i,t}^{k}$ is a measure of committee member i's policy preference in meeting t for factor k. We construct two proxies for $\gamma_{i,t}^{k}$ that we focus on in our main analysis. The first is simply committee member i's policy preference as measured in the period $t-1$ meeting for factor k. The second is the fixed effect for member i on factor k, estimated over the full sample

after controlling for meeting fixed effects. Both proxies capture the persistent component of committee members' policy preferences across dimensions.[15]

For committee member i in the meeting in period t, we then estimate how well i's model matches the state of the macroeconomy based on the inner product of their mental model with a vector capturing the state of the economy in that period. We define the state of the economy in period t as the vector

$$s_t = \begin{pmatrix} inflation_t \\ growth_t \\ labor_t \\ credit_t \\ stocks_t \end{pmatrix}, \tag{3}$$

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 the variable of interest as

$$state_match_{i,t} = m_{i,t}^T s_t. \tag{4}$$

More positive values of $state_match_{i,t}$ indicate that committee member i's mental model is better aligned with the state of the economy.

We construct two measures of the state of the economy: one using quantitative data and one using textual data. For our first measure, s_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).[16] 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 across dimensions.

[15]While the evidence in the previous section suggests that members may use potentially different frameworks to arrive at the same policy preferences, from a decision-making perspective, the relevant component of members' mental models are the policy preferences they induce across dimensions as captured by this representation.

[16]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 in the history due to excess volatility of the expanding window z-score of headline CPI forecasts at the beginning of the sample. Our results are similar simply using GDP deflator forecasts throughout the full sample.

For our second measure, s_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.2.3. This text-based measure has the additional advantage of incorporating the various metrics discussed in the Tealbook A in formulating its score for each factor, for example focusing on business and consumer credit across segments when considering credit market conditions.

With measures of the state of the economy in hand, we construct $state_match_{i,t}^{text}$ and $state_match_{i,t}^{quant}$ for each committee member i in period t as the inner products of the member's mental model for t, $m_{i,t}$, with s_t^{quant} and s_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 *state_match* 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 state are also more aligned with the committee's decision on average. 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 the state of the economy—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 state. 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 state.

Table 5 reports the results of the regressions in which we measure the mental model of a committee member using the vector of their lagged policy preferences. 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 0.14 to 0.25 and t-statistics ranging from 3.46 to 4.00, 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 state of the economy. When including member fixed effects, the regression coefficients decline, ranging from 0.09 to 0.12, with t-statistics ranging from 2.40 to 3.22. The decline in coefficients suggests that part of the relationship between the *state_match* 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, the coefficients range from -0.018 to -0.061, indicating that a one standard deviation decrease in *state_match* corresponds to a 1.8% to 6.1% increase in dissent probability, depending on the specification. 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 state of the macroeconomy.

Table 6 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 5. 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 regression results reported in Table 5. 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.[17] These fixed effects do not substantially influence the results. Tables B.2 and B.3 report regression results in which *state_match* variables are constructed using only the traditional macroeconomic variables and only the financial variables, respectively. The results reveal that both the macroeconomic and financial state of the economy 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 dissenting behavior being highly related to alignment with the committee's final decision. However, the regressions also indicate that the alignment variables are associated with additional information about committee members' policy preferences over and above dissent, particularly in the cross-sectional dimension.

[17] For examples, see Belden (1989) and Chappell Jr, McGregor and Vermilyea (2004).

	y =*DirectiveAlignment*				y =*ChangeAlignment*				*y=Dissent*			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
$state_match^{quant}$	0.25	0.12			0.18	0.09			-0.06	-0.03		
	(3.97)	(2.53)			(3.60)	(2.44)			(3.68)	(2.39)		
$state_match^{text}$			0.16	0.09			0.14	0.09			-0.03	-0.02
			(3.46)	(2.40)			(4.00)	(3.22)			(2.35)	(1.62)
R^2	0.22	0.45	0.21	0.45	0.23	0.42	0.22	0.42	0.15	0.30	0.13	0.30
N	3660	3660	3660	3660	3660	3660	3660	3660	3660	3660	3660	3660
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 5: MENTAL MODELS AND POLICY ALIGNMENT

Note: The table reports results from regressions of how much the committee's policy decision is aligned with committee member i's policy preference in period t on $state_match_{i,t}$ variables, which capture how well committee member i's mental model matches the macroeconomic data in period t. The $state_match_{i,t}$ variables are constructed as the inner product of a vector committee member i's preferences across macroeconomic and financial factors expressed in the period $t-1$ meeting with a vector of macroeconomic and financial data from period t. 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)
$state_match^{quant}$	0.23	0.14			0.16	0.11			-0.05	-0.03		
	(3.20)	(2.46)			(2.87)	(2.60)			(3.05)	(1.69)		
$state_match^{text}$			0.16	0.11			0.14	0.12			-0.03	-0.02
			(2.82)	(2.59)			(3.20)	(3.41)			(2.00)	(1.95)
R^2	0.22	0.45	0.21	0.45	0.23	0.42	0.22	0.42	0.15	0.30	0.13	0.30
N	3660	3660	3660	3660	3660	3660	3660	3660	3660	3660	3660	3660
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 6: FIXED EFFECTS, MENTAL MODELS, AND POLICY ALIGNMENT

Note: The table reports results from regressions of how much the committee's policy decision is aligned with committee member i's preferences in period t on $state_match_{i,t}$ variables, which capture how well committee member i's mental model matches the macroeconomic data in period t. The $state_match_{i,t}$ 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 factors on meeting and member fixed effects, with a vector of recent macroeconomic data. 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 match the state of the economy better are better aligned with the policy decision made by the committee indicates that the committee tilts towards the policy preferences of members whose mental models better match the state of the economy.

2.6 Preference Aggregation Mechanisms

We next explore the mechanisms through which the committee aggregates its individual members' preferences such that the final decision tilts towards the preferences of committee members whose mental models are better aligned with recent macroeconomic data. We present evidence that the tilt may be driven in part by specific committee members with mental models that tend to match the data well speaking more. 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 *state_match* 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 7 reports the results from the regressions. 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 state of the economy. 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 state; 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 on average and having more influence on the committee.

The chair may also play an important role in aggregating committee members' preferences, by playing the role of "consensus-builder." The first panel of Figure 5 plots the average *Alignment* of the chair of the committee (2.81) versus the average *Alignment* of non-chair voters (1.81), indicating that the chair of the committee tends to be substantially more aligned with the policy decision than other committee members. The second panel of figure plots the average *Alignment* for different chairs, and, for those that served on the committee before

	(1)	(2)	(3)	(4)
$state_match^{quant}$	0.07	-0.02		
	(2.15)	(0.73)		
$state_match^{text}$			0.10	-0.01
			(3.26)	(0.32)
N	3660	3660	3660	3660
Meeting Fixed Effect	Yes	Yes	Yes	Yes
Member Fixed Effect	No	Yes	No	Yes

TABLE 7: 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 $state_match_{i,t}$ variables, which capture how well committee member i's mental model matches the macroeconomic data in period t. The $state_match_{i,t}$ variables are constructed as the inner product of a vector committee member i's preferences across macroeconomic and financial factors 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.

becoming chair, plots their average *Alignment* as non-chairs. The panel illustrates that, with the exception of Ben Bernanke, chairs of the committee tend to become considerably more aligned with the committee after becoming chair, and that the chair has become increasingly aligned with the committee's decision over time.

One interpretation of this result is that the committee chair has substantial influence in shaping the final directive of the committee, as reflected by their strong alignment with its decision. The fact that the committee's decision 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 in that direction. The form that such guidance takes may vary by chair. For example, Figure 6 plots the proportion of words spoken by each chair during policy discussions during their tenure. While it was common for chairmen prior to Ben Bernanke to heavily dominate the discussion with the dominance of the chair reaching its peak with Alan Greenspan, starting with Bernanke, the committee chair has spoken considerably less. As it pertains to influencing the direction of the committee, Don Kohn, the long-term secretary of the FOMC and a former vice chairman of the committee, noted about Alan Greenspan:

> "Greenspan became more and more dominant over time in part because he was right on [a number of things]...So when you have this guy who has proved himself to have better analysis intuition over the economy over a long period of time, I think there is a tendency to defer to him,"

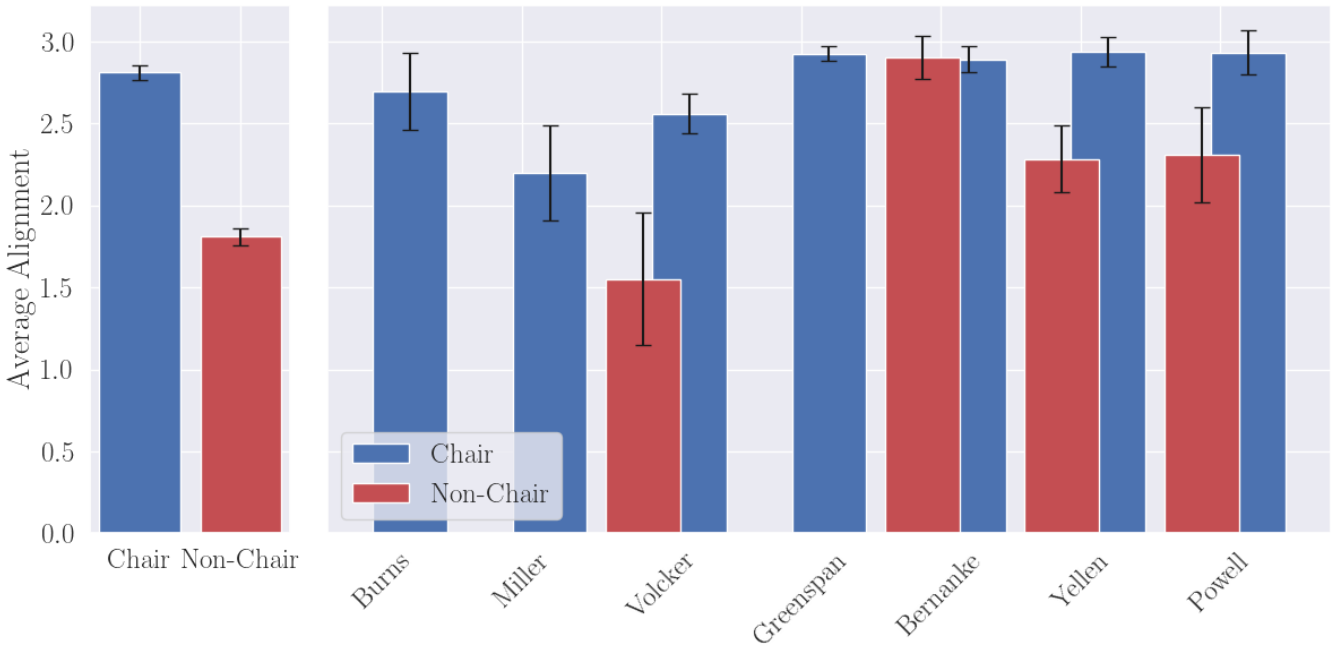


FIGURE 5: CHAIR'S ALIGNMENT WITH THE POLICY DECISION

Note: The first panel of the figure plots the average *Alignment*—which captures alignment with the committee's final policy decision, with values of -3 indicating complete opposition and +3 capturing perfect alignment—for the chair of the committee and for non-chair voters across meetings. The second panel of the figure plots the average *Alignment* score of chair of the committee in blue; and in red, plots the average alignment score of those chairmen as governors or as regional Fed presidents. The bars correspond with plus or minus two standard errors.

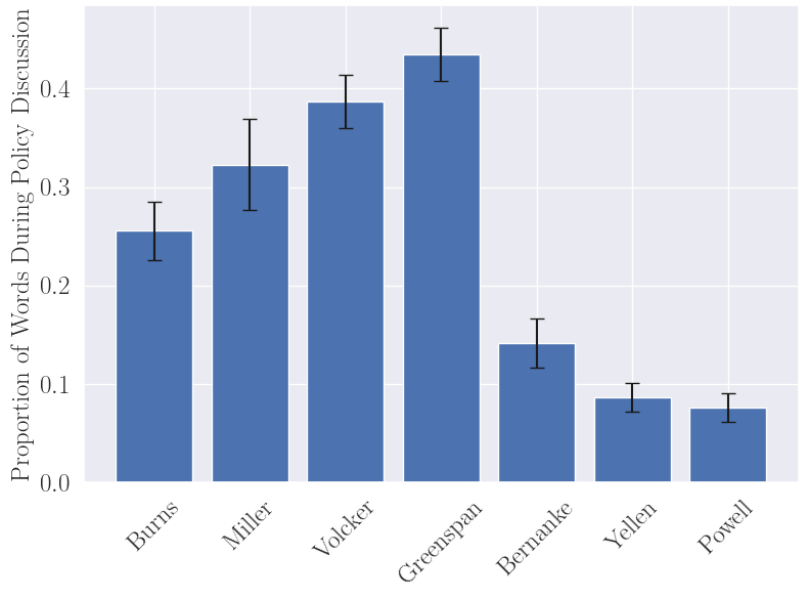


FIGURE 6: PROPORTION OF WORDS DURING POLICY DISCUSSION

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

while he noted about Ben Bernanke:

> "He would say I would like to hear your views. So he doesn't lead off... he does a very good job of basically having the committee move when he thinks it should move, although he listens to people and [decides] what he thinks can be modified by what people say" Schonhardt-Bailey (2013).

Despite the differing styles, speaking versus allowing others to speak, the results from Figure 5 nevertheless indicate the influence that the chair has on the final decision, and suggests they may play an important role in the committee tilting its decision towards models more consistent with the data.

2.7 FOMC Transparency and the Cost of Dissent

We conclude our empirical analysis by briefly discussing how committee design decisions may influence the aggregation of individual committee members' preferences, a point that we further explore in our theoretical analysis. In particular, 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.[18]

The top panel of Figure 7 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.1 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 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.[19] 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).

[18]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.

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

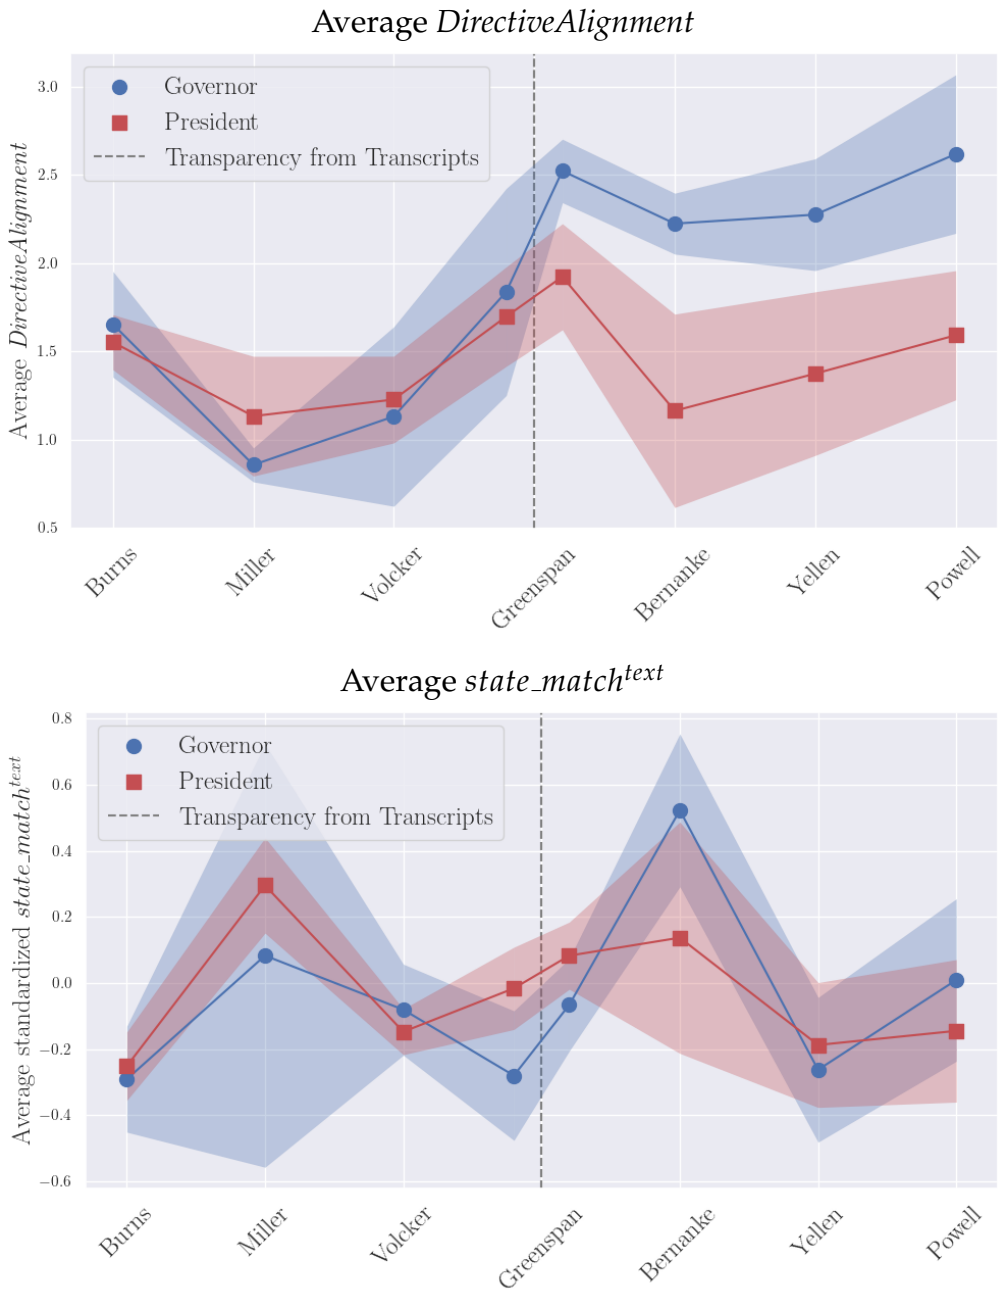


FIGURE 7: ALIGNMENT WITH THE COMMITTEE DECISION: PRESIDENTS VERSUS GOVERNORS

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.

3 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 the state of the economy. This evidence suggests 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 empirical 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.

3.1 Model Setup

There are two committee members, $i \in \{A, B\}$. Let μ be the true probability measure on a measurable space $(\mathcal{D}, \mathcal{F})$, where μ is unknown to the committee members. Each member i has a *model* of the world, m_i, which is the density of a probability measure, M_i, with respect to μ.

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

(i) **(Partial misspecification).** Members' models are partially misspecified (for $i \in \{A, B\}$, $\exists S_i \subseteq \mathcal{D}$ 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 $E \subseteq \mathcal{D}$ 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 events, and only to those events (models are mutually absolutely continuous with respect to μ).

The committee makes its decision as follows: all members observe economic data $D \subseteq \mathcal{D}$, drawn according to μ. After observing the data, each member i advocates for a model by choosing a weight $w_i^A \in [0, 1]$ on m_A, and weight $(1 - w_i^A)$ on m_B. For member A, the weight w_i^A can be thought of as reflecting how strongly they advocate for their personal model during the committee meeting relative to model B, e.g., in the case of the FOMC, how many statements they make during the policy discussion. The model advocated by member i is the convex combination:

$$v_i = w_i^A m_A + (1 - w_i^A) m_B.$$

The committee's model m, which maps to the policy decision of the committee, is given by averaging the members' advocated models:

$$m = \frac{1}{2}v_A + \frac{1}{2}v_B = wm_A + (1-w)m_B,$$

where

$$w = \frac{1}{2}\left(w_A^A + w_B^A\right).$$

In choosing their advocacy, committee member i maximizes the utility function:

$$U_i(w_A^A, w_B^A) = \log \mathcal{L}(D|m) - \lambda_1 KL(m_i \| v_i) - \lambda_2 KL(m \| v_i),$$

where $\mathcal{L}(x|p) = p(x)$ is the likelihood of the data x evaluated using density p, $KL(p\|q) = \int_D p(x) \log\left(\frac{p(x)}{q(x)}\right) d\mu(x)$ is the Kullback-Leibler divergence, and the coefficients λ_1 and λ_2 are positive.

The first term, $\log \mathcal{L}(x|p)$, captures that committee members want a model 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 due to a need for the committee to justify its decision to constituencies that also observe D, such as financial market participants.

The second term, $-\lambda_1 KL(m_i\|v_i)$, captures a penalty associated with the extent to which committee members are unwilling or unable to depart from their own models. The coefficient λ_1 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_1 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.

The third term, $-\lambda_2 KL(m\|v_i)$, reflects a penalty associated with expressing dissent. The coefficient λ_2 captures how this penalty scales with $KL(m\|v_i)$, the divergence between a member's advocated density v_i and the committee's final model m. As λ_2 increases, members increasingly conform in advocating for the same model. Again, in practice, λ_2 may reflect a combination of behavioral and institutional factors that induce conformity.

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 distribution that members seek to explain given a history of data D.*[20] *With dynamics, there also may be additional considerations, for example, members may have a preference for gradualism or may care*

[20]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).

about reputational effects.

3.2 Equilibrium

We study pure strategy Nash Equilibria.

Definition 3.1 (Equilibrium) *A Nash equilibrium is a pair of weights on model A,* $(w_A^A, w_B^A) \in [0,1] \times [0,1]$*, such that:*

$$w_A^A \in \arg\max_{w^A \in [0,1]} U_A(w^A, w_B^A) \quad \textit{and} \quad w_B^A \in \arg\max_{w^A \in [0,1]} U_B(w^A, w_A^A),$$

where $U_i(w^A, w_j^A)$ *denotes the utility of committee member* $i \in \{A, B\}$ *given their own chosen weight* w^A *and the other member's weight* w_j^A*.*

Proposition 1 (Equilibrium Existence) *There is an equilibrium.*

Proof. All proofs provided in the appendix. □

3.3 Committee Decision-Making versus Individual Decision-Making

With the equilibrium in hand, we next turn to understanding the conditions under which the committee structure can improve decision-making performance.

Proposition 2 (Committees Outperformance) *If committee members are sufficiently open to entertaining others' models, then there is an equilibrium in which the committee's model outperforms the best member's model in expectation. That is,* $\exists \bar{\lambda}_1 > 0$ *such that for all* $\lambda_1 < \bar{\lambda}_1$*,*

$$\mathbb{E}_\mu\Big[\log m\Big] > \max\Big\{\mathbb{E}_\mu\big[\log m_A\big],\, \mathbb{E}_\mu\big[\log m_B\big]\Big\}.$$

Proposition 2 can be understood by the fact that there is always a combination of members' models that can outperform either one of their models; this is due to the fact that members' models are partially misspecified in different ways, so there are benefits to model diversification. In the case when members are perfectly able and willing to fully consider other committee members' models ($\lambda_1 = 0$), then the committee members combine together their models in a way that maximizes the log likelihood of the observed data. In expectation, this leads to a superior model to either of their individual models. As it becomes increasingly difficult for members to entertain others' models (higher λ_1), members increasingly advocate for their own model, and with sufficiently high λ_1, they may advocate for their personal model too much. In the limit as $\lambda_1 \to \infty$, each member simply advocates for their own model and ignores the data; the resulting equilibrium model selected is an average of the members' models.

Proposition 2 illustrates that there are potential benefits to group decision-making—improved model performance—but these are only obtainable when members are sufficiently willing and able to entertain each other's models. When they are not able to do so, moving to the committee structure can *decrease* the expected decision-making performance. For example, with high λ_1, introducing a committee member B whose models performs strictly worse than member A's model will decrease the expected log likelihood of the model chosen, because the committee will have to balance between the models of A and B. But with a low λ_1, given that member B's model might diversify the mistakes of member A's model, there is still a benefit, even when member B's model performs worse than member A's in expectation.

The existence of an equilibrium that leads the committee's model to outperform its members in explaining the data is independent of the cost of dissenting, λ_2. However, this cost of dissent does shape whether the committee reaches such an equilibrium or not.

Proposition 3 (Equilibrium Multiplicity) *When the cost of dissent is sufficiently low, the equilibrium model selected by the committee is unique; otherwise there may be multiple equilibria. That is there is a $\bar{\lambda}_2$ such that $\lambda_2 < \bar{\lambda}_2$ implies that the equilibrium is unique, and $\lambda_2 > \bar{\lambda}_2$ 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.

3.4 Implications for Committee Design

Our theoretical framework illustrates potential costs and benefits 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. One way to push members in the direction of entertaining each others' models is via structured discussion and deliberation. For example, since the Greenspan era, the FOMC has consistently structured its discussion such that committee members 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 their committee 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 overconfident or stubborn in their viewpoints.

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. 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.

4 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 the models of members that better match the state of the world.

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 a potential benefit in making decisions in groups is that different members supply different models from which the group can draw on 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 findings suggest that, under the right conditions, committees can provide a systematic approach to decision-making under model uncertainty, adapting to new information without being overly constrained by any single viewpoint.

References

Andre, Peter, Carlo Pizzinelli, Christopher Roth, and Johannes Wohlfart. 2022. "Subjective models of the macroeconomy: Evidence from experts and representative samples." *The Review of Economic Studies*, 89(6): 2958–2991. 5

Andre, Peter, Philipp Schirmer, and Johannes Wohlfart. 2023. "Mental models of the stock market." 5

Bastianello, Francesca, Paul H Décaire, and Marius Guenzel. 2024. "Mental Models and Financial Forecasts." 5

Belden, Susan. 1989. "Policy preferences of FOMC members as revealed by dissenting votes." *Journal of Money, Credit and Banking*, 21(4): 432–441. 5, 12, 25

Blinder, Alan S. 1999. *Central banking in theory and practice.* Mit press. 5

Blinder, Alan S. 2008. *The quiet revolution: Central banking goes modern.* Yale University Press. 1, 5

Bobrov, Anton, Rupal Kamdar, and Mauricio Ulate. 2024. "Regional Dissent: Do Local Economic Conditions Influence FOMC Votes?" *American Economic Review: Insights*. Forthcoming. 12

Bordo, Michael, and Klodiana Istrefi. 2023. "Perceived FOMC: The making of hawks, doves and swingers." *Journal of Monetary Economics*, 136: 125–143. 5, 15

Bordo, Michael D, Klodiana Istrefi, and Humberto Martínez. 2024. "Rules vs. Discretion: Decoding FOMC Policy Deliberations." National Bureau of Economic Research. 5

Brainard, William C. 1967. "Uncertainty and the Effectiveness of Policy." *American Economic Review*, 57(2): 411–425. 5

Caballero, Ricardo J, Tomás E Caravello, and Alp Simsek. 2024. "Financial Conditions Targeting." 9

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. 5, 12, 25

Charness, Gary, and Matthias Sutter. 2012. "Groups make better self-interested decisions." *Journal of Economic Perspectives*, 26(3): 157–176. 5

Chen, Jiafeng, and Jonathan Roth. 2024. "Logs with zeros? Some problems and solutions." *Quarterly Journal of Economics*, 139(2): 891–936. 28

Cieslak, Anna, and Annette Vissing-Jorgensen. 2021. "The economics of the Fed put." *Review of Financial Studies*, 34(9): 4045–4089. 5, 9

Cieslak, Anna, Stephen Hansen, Michael McMahon, and Song Xiao. 2023. "Policymakers' uncertainty." National Bureau of Economic Research. 5

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. 5

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. 28

Condorcet, Marquis de. 1785. "Essay on the Application of Analysis to the Probability of Majority Decisions." *Paris: Imprimerie Royale*, 1785. 1

Dunn, Wendy, Ellen E Meade, Nitish Ranjan Sinha, and Raakin Kabir. 2024. "Using Generative AI Models to Understand FOMC Monetary Policy Discussions." 7

Enke, Benjamin, Thomas Graeber, and Ryan Oprea. 2023. "Confidence, self-selection, and bias in the aggregate." *American Economic Review*, 113(7): 1933–1966. 5

Fos, Vyacheslav, and Nancy R Xu. 2025. "When Do FOMC Voting Rights Affect Monetary Policy?" 12

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. 4

Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus. 2021. "Five facts about beliefs and portfolios." *American Economic Review*, 111(5): 1481–1522. 13

Greenspan, Alan. 2004. "Risk and Uncertainty in Monetary Policy." *Speech presented at the Annual Meeting of the American Economic Association, San Diego, CA*. 35

Hansen, Lars Peter, and Thomas J Sargent. 2008. *Robustness*. Princeton university press. 5

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. 3, 5, 6, 32

Herbst, Edward, and Fabian Winkler. 2021. "The factor structure of disagreement." 15

Ke, Shikun. 2024. "Analysts' Belief Formation in Their Own Words." *Available at SSRN 5025830*. 5

Madeira, Carlos, and João Madeira. 2019. "The effect of FOMC votes on financial markets." *Review of Economics and Statistics*, 101(5): 921–932. 13

Malmendier, Ulrike, Stefan Nagel, and Zhen Yan. 2021. "The making of hawks and doves." *Journal of Monetary Economics*, 117: 19–42. 5, 15

Meade, Ellen E, and David Stasavage. 2008. "Publicity of debate and the incentive to dissent: Evidence from the US Federal Reserve." *The Economic Journal*, 118(528): 695–717. 3, 13, 32

Meade, Ellen E, and D Nathan Sheets. 2005. "Regional influences on FOMC voting patterns." *Journal of Money, Credit and Banking*, 661–677. 12

Ottaviani, Marco, and Peter Sørensen. 2001. "Information aggregation in debate: who should speak first?" *Journal of Public Economics*, 81(3): 393–421. 32

Primiceri, Giorgio E. 2006. "Why inflation rose and fell: policy-makers' beliefs and US post-war stabilization policy." *The Quarterly Journal of Economics*, 121(3): 867–901. 5

Riboni, Alessandro, and Francisco Ruge-Murcia. 2014. "Dissent in monetary policy decisions." *Journal of Monetary Economics*, 66: 137–154. 12

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. 1

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. 5

Sargent, Thomas J. 1999. *The conquest of American inflation*. Princeton University Press. 5

Scanlan, Susannah. 2023. "Attention Allocation and the Factor Structure of Forecasts." 15

Scharfstein, David S, and Jeremy C Stein. 1990. "Herd behavior and investment." *American Economic Review*, 465–479. 32

Schonhardt-Bailey, Cheryl. 2013. *Deliberating American monetary policy: A textual analysis*. MIT Press. 5, 31

Schwartzstein, Joshua, and Adi Sunderam. 2021. "Using models to persuade." *American Economic Review*, 111(1): 276–323. 5

Schwartzstein, Joshua, and Adi Sunderam. 2024. "Sharing Models to Interpret Data." Mimeo. 5

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. 5

Sibert, Anne. 2006. "Central banking by committee." *International Finance*, 9(2): 145–168. 4

Taylor, John B. 1993. "Discretion versus policy rules in practice." Vol. 39, 195–214, Elsevier. 7

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. 12, 44

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. 7

Wu, Lingxuan. 2023. "Mental Macro-Finance Models: Evidence and Theory." 15

A Proofs

Proof of Proposition 1

Proof. The strategy space $S = [0,1] \times [0,1]$ is convex and compact. Each member's utility $U_i(w_A^A, w_B^A)$ is continuous in w_i^A and quasi-concave due to the concavity of the log-likelihood term and the 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. □

Proof of Proposition 2

Proof. We proceed in two steps. First, we show that there is a $w^* \in (0,1)$ such that

$$\mathbb{E}_\mu \left[\log \left(w^* m_A + (1-w^*) m_B\right)\right] > \max\Big\{\mathbb{E}_\mu\big[\log m_A\big],\, \mathbb{E}_\mu\big[\log m_B\big]\Big\}.$$

The existence of an interior w^* is guaranteed jointly by the assumption of distinctness and by the concavity of the function $\psi(w) = \mathbb{E}_\mu \left[\log \left(w m_A + (1-w) m_B\right)\right]$. Next consider the case of $\lambda_1 = 0$, and focus on symmetric equilibria ($w_A = w_B$), which eliminates the penalty associated with $\lambda_2 KL(m \| v_i)$; any symmetric equilibrium remains an equilibrium regardless of the choice of λ_2. Then the equilibrium model is equivalent to the solution of $\max_w \log \mathcal{L}(D|wm_A + (1-w)m_B)$. Taking the expectations, and applying the result from step 1, we obtain that in expectation, the expected log likelihood of the committee's model exceeds the expected log likelihood of either member's model. The existence of a threshold $\bar{\lambda_1} > 0$ is guaranteed by the continuity of commitee members' best responses in λ_1. □

Proof of Proposition 3

Proof. First, we show that the equilibrium is unique when $\lambda_2 = 0$. With $\lambda_2 = 0$, the game can be written as a potential game, with the potential function $\Phi(w_A^A, w_B^A) = \log \mathcal{L}(D\|m) - \lambda_1(KL(m_A\|v_A) + KL(m_B\|v_B))$. Φ is strictly concave in w_i^A given the strict concavity of the log likelihood and strict convexity of the KL divergence in its second term. By strict concavity, and since the strategy space $[0,1] \times [0,1]$is compact and convex, there is a unique maximizer, which corresponds with the unique Nash equilibrium. Next we note that as $\lambda_2 \to \infty$, there are multiple equilibria—all strategies $w_A^A = w_B^A$ are equilibria. Lastly, the correspondence that maps λ_2 to the set of equilibria is upper-hemicontinuous, due to the fact that members' best responses are single-valued and continuous. Upper hemi-continuity guarantees the existence of a threshhold $\bar{\lambda}_2$. □

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 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, 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 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, please
    conduct a thorough analysis using the following steps.
   Conduct your analysis inside <detailed_analysis> tags:

1. List all policy alternatives discussed in the meeting.
2. Identify and quote key statements from the speaker that
   indicate their policy preferences. Consider the speaker's
    tone and emphasis in these statements.
3. Summarize the speaker's overall stance based on their
   statements.
4. State the final policy decision, focusing on all changes
   made to the statements that were agreed upon by the
   committee. Quote key policy statements from the final
```

```
   decision.
5. Compare the speaker's stance with each policy alternative
   , highlighting areas of alignment and disagreement.
6. Compare the speaker's stance with the final policy
   decision, highlighting areas of alignment and
   disagreement.
7. List ALL proposed changes that were adopted in the policy
    discussion, regardless of who proposed them. For each
   change, include:
a) A brief description of the change
b) Who proposed it
c) How it aligns with the focus speaker's preferences (use a
    scale from -3 to 3)
8. Analyze the focus speaker's overall alignment with all
   successful proposed changes, including those made by
   other speakers.
9. Summarize the overall trend of the speaker's alignment
   with successful proposed changes.
10. Consider each point on the rating scale (-3 to 3) and
   briefly note why it might or might not be appropriate.

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 proposed change in the following
   format:]
<change>
<description>[Brief description of the change]</description>
<proposer>[Speaker name or "Other" if not proposed by the
   focus speaker]</proposer>
```

```
    <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 all successful proposed changes,
       including those made by other speakers]</
       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 only minor
       differences. [Specific examples and reasoning to follow
       ...]</justification>

    <proposed_changes>
    <change>
    <description>Example change description</description>
    <proposer>Example proposer</proposer>
    <alignment>1</alignment>
    </change>
    </proposed_changes>

    <speaker_alignment>Brief analysis of overall alignment with
       successful proposed changes...</speaker_alignment>

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

B.2.3 Prompt to Measure the State of the Economy from Tealbook Texts

For each meeting, we measure the state of the economy as it pertains 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*			
$state_match^{quant}$	0.22	0.13			0.15	0.10			-0.06	-0.03		
	(3.82)	(2.68)			(3.29)	(2.62)			(3.58)	(2.46)		
$state_match^{text}$			0.15	0.09			0.12	0.09			-0.03	-0.02
			(3.42)	(2.58)			(3.94)	(3.48)			(2.35)	(1.68)
R^2	0.27	0.46	0.26	0.46	0.27	0.43	0.27	0.43	0.16	0.31	0.15	0.30
N	3660	3660	3660	3660	3660	3660	3660	3660	3660	3660	3660	3660
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 5 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*			
$state_match^{quant}$	0.20	0.10			0.15	0.08			-0.05	-0.03		
	(3.04)	(1.97)			(2.94)	(2.04)			(2.80)	(1.85)		
$state_match^{text}$			0.13	0.08			0.11	0.08			-0.03	-0.02
			(2.86)	(2.01)			(3.25)	(2.69)			(2.10)	(1.41)
R^2	0.21	0.45	0.21	0.45	0.22	0.42	0.22	0.42	0.14	0.30	0.13	0.30
N	3660	3660	3660	3660	3660	3660	3660	3660	3660	3660	3660	3660
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 5 constructing *state_match* with macroeconomic variables ($inflation_t$, $growth_t$, $labor_t$).

	y =*DirectiveAlignment*				y =*ChangeAlignment*				*y=Dissent*			
$state_match^{quant}$	0.20	0.09			0.14	0.06			-0.05	-0.02		
	(4.26)	(2.52)			(3.58)	(2.04)			(4.02)	(2.47)		
$state_match^{text}$			0.12	0.06			0.11	0.06			-0.02	-0.01
			(3.39)	(2.42)			(3.77)	(2.69)			(2.05)	(1.55)
R^2	0.21	0.45	0.20	0.45	0.22	0.42	0.22	0.42	0.14	0.30	0.13	0.30
N	3660	3660	3660	3660	3660	3660	3660	3660	3660	3660	3660	3660
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 5 constructing *state_match* with macroeconomic variables ($stocks_t$, $credit_t$).

	y = *DirectiveAlignment*				y = *ChangeAlignment*			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$state_match^{quant}$	0.09	0.05			0.05	0.04		
	(2.75)	(1.78)			(1.61)	(1.35)		
$state_match^{text}$			0.08	0.05			0.07	0.05
			(3.45)	(2.16)			(3.01)	(2.36)
Dissent	-2.63	-2.28	-2.65	-2.29	-2.07	-1.73	-2.07	-1.73
	(26.47)	(23.01)	(27.19)	(23.17)	(20.17)	(17.87)	(21.56)	(17.94)
R^2	0.61	0.69	0.61	0.69	0.46	0.56	0.47	0.56
N	3660	3660	3660	3660	3660	3660	3660	3660
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 5, including a 0/1 indicator variable as to whether a committee member dissented in a meeting as an independent variable.



FIGURE B.1: ALIGNMENT WITH THE COMMITTEE DECISION: PRESIDENTS VERSUS GOVERNORS

The figures replicates Figure 7 from the main text using the *ChangeAlignment* and $state_match^{quant}$ measures.

	Inflation	Growth	Labor
Panel A: Baseline Fixed Effects			
Time FE Only	0.264	0.306	0.323
Member FE Only	0.302	0.184	0.216
Panel B: Parameter-Level Analysis			
Parameters Only	0.303	0.327	0.326
Parameters + Member FE	0.450	0.407	0.417
Parameters + Time FE	0.432	0.436	0.460

TABLE B.5: Mental Models and Policy Preferences, Adjusted R^2

Notes: The table reports adjusted R^2 values from regressions of inflation, growth, and labor market policy preference measures on meeting fixed effects and FOMC committee member fixed effects (Panel A), and on parameters from members' expressed theoretical frameworks (Panel B). The policy preference measures range from -3 to 3, with positive values indicating hawkish preferences and negative values indicating dovish preferences. The magnitude of these measures captures their importance for the committee member. Parameters capture specific theoretical views expressed by members during FOMC meetings, including beliefs about key economic mechanisms and relationships.

C Mental Model Classification Methodology

This appendix provides detailed documentation of our methodology for classifying and analyzing FOMC members' mental models.

C.1 Construction of Mental Model Taxonomies

For each theme (inflation, growth, labor markets, and financial conditions), we construct comprehensive taxonomies of potential mental models by prompting Claude to:

(i) Review academic literature and central bank documents to identify canonical theoretical frameworks

(ii) Codify each framework's core mechanisms and assumptions

(iii) Develop diagnostic criteria for identification from text

(iv) Specify key parameters that characterize model implementation

Tables C.2, C.1, C.3, and C.4 present the full taxonomy for each domain. For example, in the inflation domain we distinguish between:

- Phillips Curve models emphasizing resource utilization gaps

- Monetarist models focused on money growth and velocity
- Cost-Push models centered on supply shocks and input costs
- Expectations-Centric models prioritizing credibility and anchoring

For each model, we specify:

(i) **Core Characteristics** define the fundamental theoretical mechanisms and assumptions that distinguish each model (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.

(ii) **Core Phrases** provide linguistic markers that signal the potential use of each model (e.g., references to ”slack”, ”gap”, ”utilization”). These are not simple keyword matches but rather contextual indicators that help identify the underlying conceptual framework.

(iii) **Disambiguation** rules specify how to differentiate between similar models and establish the context necessary for proper classification. These rules ensure consistent model identification across different speakers and meetings.

(iv) **Parameters** capture variations in how models are implemented or interpreted (e.g., Phillips curve slope: flat/steep). These allow us to identify not just which model is being used, but how it is being applied.

C.2 Classification Implementation

We implement the speaker-meeting mental model classification through the following procedure:

Data Preparation:

(i) Parse meeting transcripts into speaker-specific text blocks

(ii) Filter for substantive contributions (¿50 words)

(iii) Clean and standardize speaker identifiers

(iv) Extract relevant meeting sections (economic situation and policy discussions)

Model Classification: For each speaker-meeting pair, we employ the Claude large language model with the following structured prompt:

```
Input: Speaker's statements, mental model taxonomy
Task: Evaluate presence of each model framework
Output format:
{
 "model_name": {
   "base_score": [0-1],      // Basic model match
   "context_score": [0-1],   // Contextual appropriateness
```

```
    "combined_score": [0-1],  // Weighted average
    "confidence": [0-1],      // Classification confidence
    "parameters": {           // If sufficient evidence exists
      "param_name": {
        "estimate": "value",
        "confidence": [0-1]
      }
    }
  }
}
```

The classification system:

(i) Evaluates text against each model's diagnostic criteria

(ii) Generates probability scores for model presence. Our classification system assigns each model a score between 0 and 1, with scores above 0.6 considered significant evidence of model usage. This threshold balances sensitivity with specificity in model identification. The scoring system considers both the presence of model indicators and the strength of their manifestation in the text.

(iii) Estimates parameter values when evidence exists

(iv) Assesses classification confidence

TABLE C.1: 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	• Stage of the cycle • Momentum indicators • Recovery pattern • Cyclical behavior	• Core: Output follows cycles • Context: Cyclical dynamics • Rules: Focuses on	• Cycle length: standard/extended • Shock persistence: temp/perm • Indicators:

TABLE C.2: 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 C.3: 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 C.4: 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

C.3 Network Construction

This section details how we construct the belief networks underlying Figures 3, 4, C.1 and C.2.

We start from our member-meeting level measures of mental models and parameter estimates. Then we:

(i) Filter for substantive parameter estimates (combined score ¿0.6)

(ii) Remove cases with missing values

(iii) For each model-parameter combination, identify the two most frequent parameter estimates

Edges weights are constructed in two tiers:

(i) Structural edges (model-parameter-value):

- Uniform weight of 1
- Visualized width of 1.2 units
- Gray color to indicate structure

(ii) Belief edges (value-preference):

- Calculate correlation between parameter value and hawk-dove scores
- Apply minimum threshold of 0.2
- Normalize to [1.2, 6.2] range using:

$$width = (score - min_score)/(max_score - min_score) * 5 + 1.2$$

- Color red for hawkish correlations, blue for dovish
- Apply 0.6 alpha transparency

Layout Algorithm:

(i) Position nodes in hierarchical left-to-right structure:

- Mental models (left): Blue circles
- Parameters: Gray rectangles
- Parameter values: White circles
- Hawk-dove preferences (right): Red-blue circles

(ii) Use controlled spacing to ensure consistent presentation

(iii) Implement parallel processing for multi-theme generation

We present the network figures for Labor and Financial conditions below.

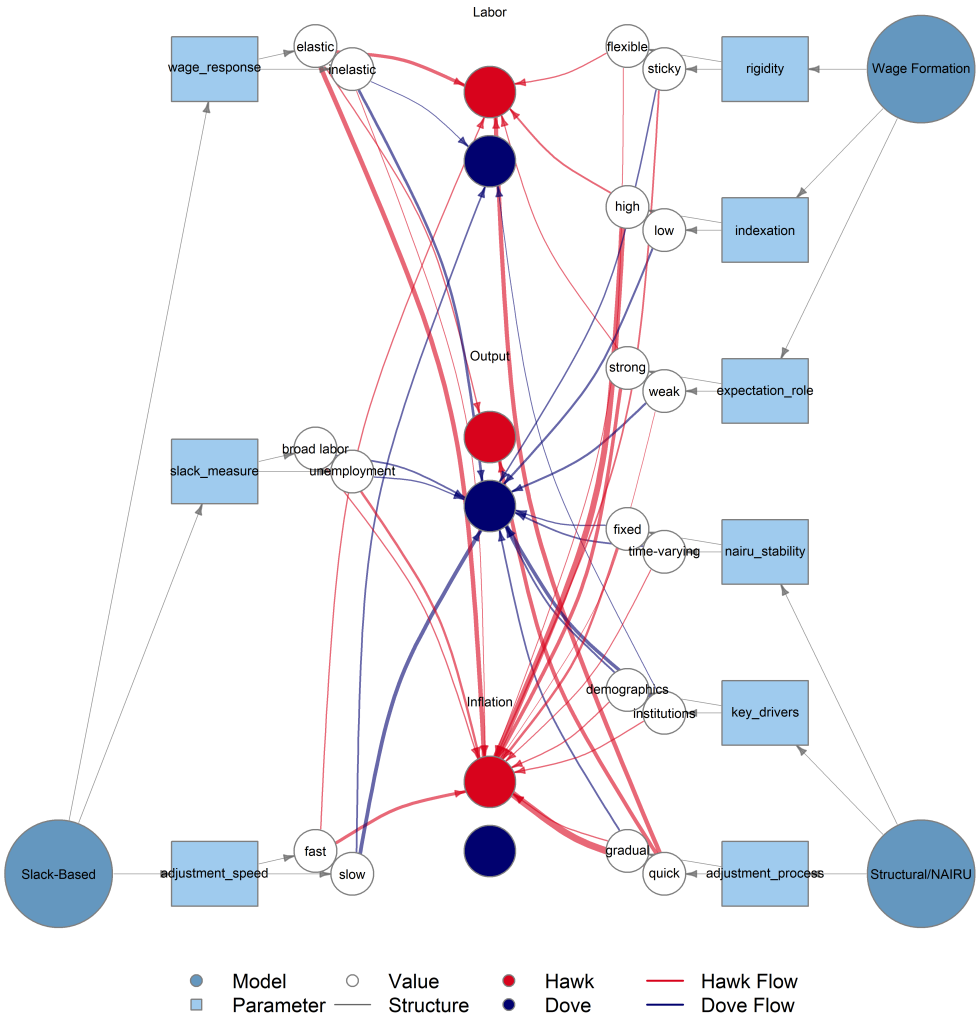


FIGURE C.1: LABOR MODELS AND BELIEF NETWORKS

Note: This figure illustrates the structure of labor models and policy preferences. Nodes represent economic models (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.

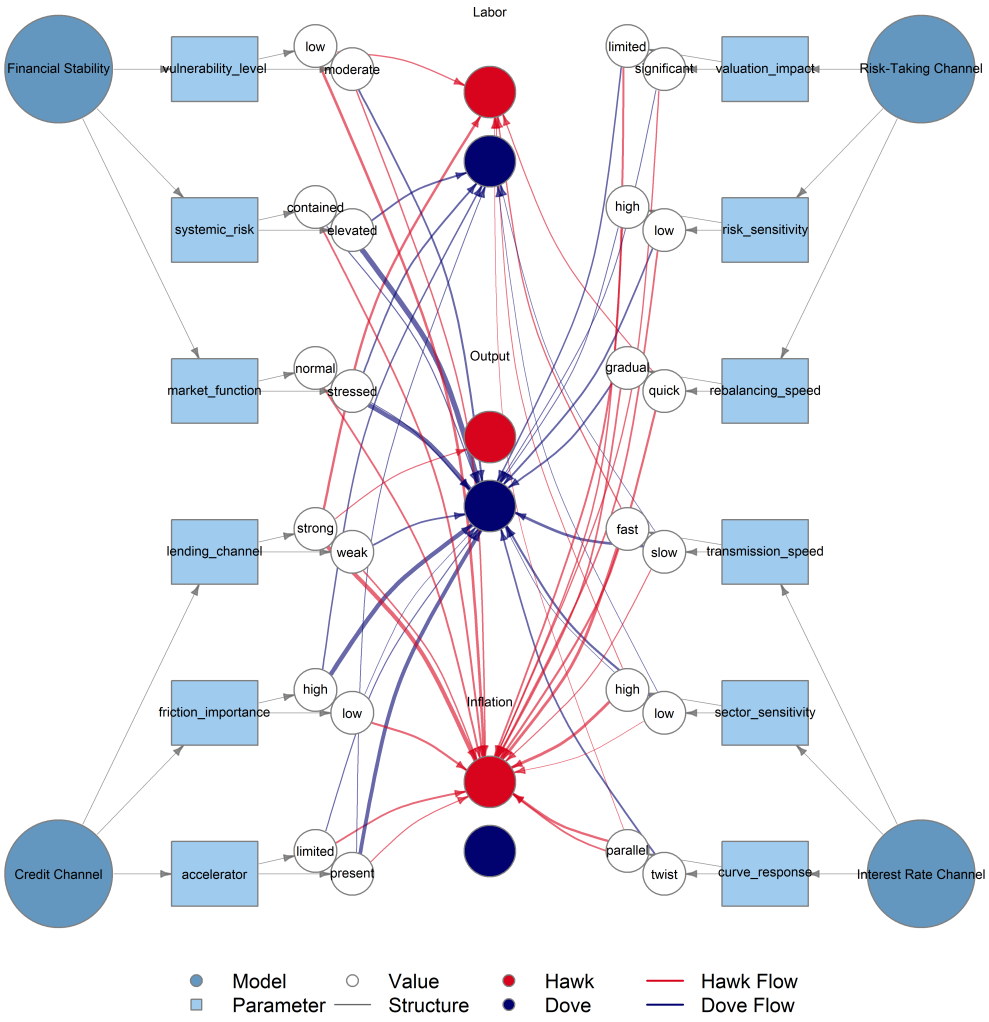


FIGURE C.2: FINANCIAL MODELS AND BELIEF NETWORKS

Note: This figure illustrates the structure of financial models and policy preferences. Nodes represent economic models (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.