Sequential Deliberation in Collective Decision-Making: The Case of the FOMC*

Gabriel López-Moctezuma†

May 24, 2018

Abstract

A process of deliberation precedes almost every collective decision. Yet, beyond scarce evidence coming from field and laboratory experiments, few studies have analyzed the role played by sequential deliberation in policy-relevant committees. To fill this gap, I estimate an empirical model of policy-making that incorporates social learning via deliberation. In the model, committee members speak in sequence, allowing them to weight their own information and biases against recommendations made by others. I quantify the extent of social learning using historical transcripts from the Federal Open Market Committee (FOMC). I find the process of deliberation significantly changes individual behavior and aggregate monetary policy: a typical FOMC member would modify her policy recommendation in 36% of the meetings after listening to previous speakers, compared to the scenario where members exclusively follow their private information. I found modest gains of changes in the order of speech on the quality of monetary policy, suggesting that the observed deliberation order within the FOMC is effective at reducing the probability of mistakes. Incorporating sequential learning explains the observed pattern of individual recommendations and observed monetary policy extremely well and improves the fit over behavioral models that ignore deliberation.

*Special thanks to Isaac Baley, Graeme Blair, Ted Enamorado, Matías Iaryczower, Gabriel Katz, John Londregan, Paula Mateo, Adam Meirowitz, Tom Palfrey, Tom Romer, Carlos Velasco-Rivera, and members of the Imai Research Group, the Political Economy Colloquium, and the Political Methodology Colloquium at Princeton University for their feedback. Finally, I am indebted to Henry Chappell Jr. for generously sharing his data with me.

†Assistant Professor, California Institute of Technology. E-mail: glmoctezuma@caltech.edu.

1 Introduction

In almost all relevant policy-making bodies such as courts, legislative committees, bureacratic agencies, international organizations, among others, decisions are commonly preceded by some form of communication among individual members. In all these cases, deliberation provides a unique opportunity for participants to arrive at more reasoned judgments (Habermas [1996]; Macedo [2010]), enhance the legitimacy of the collective decision (Gutmann and Thompson [1996]), encourage the cooperation among participants (Goeree and Yariv [2011]), and affect collective decision-making by influencing others (Landa and Meirowitz [2009]). Thus, along with voting, deliberation has been the most important feature of committee decision-making to assess the influence of individual members on aggregate policies.

Empirically, quantifying the impact of deliberation on policy-making bodies has faced important limitations which prevent us from understanding to what extent individual members learn from each other, whether they act upon this information, and how much this learning process affects policy outcomes. One practical limitation is that the communication protocols of policy-making bodies are rarely available to the outside public. Second, even if obtainable, communication among policy-makers is usually unstructured, which makes it harder to disentangle the influence of individual participants throughout the deliberation process. These reasons explain why an overwhelming portion of the existing empirical literature needs to rely on field and laboratory experiments to assess whether the presence of deliberation has an effect on the efficiency of policy choices and its interaction with the voting mechanism (Dickson, Hafer and Landa [2008]; Dickson, Hafer and Landa [2015]; Goeree and Yariv [2011]; Karpowitz and Mendelberg [2011]; Humphreys, Masters and Sandbu [2006]).[1]

In this paper, I overcome these limitations by introducing the effect of social learning into an empirical model of committee policy-making that accounts for heterogeneity in members' biases and expertise (Iaryczower and Shum [2012]). In the model, members are privately informed about the true state of the world and speak openly in front of the rest of the committee about their desired policy. The deliberation protocol is sequential, a feature that captures the nature of debate associated with most deliberative committees (Van Weelden et al. [2008]). In this way, by the time their turn to speak arises, members have already learned the content of the statements made by previous speakers and incorporate this information using Bayes' rule. (Banerjee [1992]; Bikhchandani, Hirshleifer and Welch [1992]; Smith and Sørensen [2000]). Therefore, this process of sequential learning captures how members' private information interacts with information obtained via deliberation to form a belief about the true state of the world.

I structurally estimate the model with a novel Bayesian approach that directly recovers members' preference and expertise parameters, while incorporating the informational value of deliberation contained in the statements of early speakers. This approach allows me to quantify

[1] An exception is provided by Iaryczower, Shi and Shum [2014] who, under a structural approach, quantify the effects of deliberation on decision-making at appellate courts in the U.S. Their focus is on the effect of deliberation on aggregate choices and not on the analysis of any particular communication protocol, as they lack data on the deliberation process.

the effects of learning from sequential deliberation on the behavior of committee members, which would not be possible with reduced-form methods, given the non-experimental nature of the data.

I estimate the model using data from deliberation records of the Federal Open Market Committee (FOMC), the body in charge of implementing monetary policy in the United States. The FOMC is an ideal case to analyze the role of communication in collective decision-making for several reasons. First, the decisions that the FOMC implements have relevant policy implications, as they influence inflation expectations, affecting households' and firms' strategies. Second, policy deliberation at any given FOMC meeting follows a sequential deliberation process, in which members voice their opinions in a fixed order of speech. I exploit the availability of historical deliberation transcripts and the variation in their speaking order to disentangle the contribution of individual members throughout the policy debate. Finally, studying the FOMC allows me to account for the common information committee members possessed in real time about the underlying state of the economy in the form of staff forecasts and real-time economic indicators.

The results from the estimation using the history of policy recommendations for the period 1970-2008 suggest substantial effects of deliberation as an information-sharing mechanism. With the model estimates at hand, I assess the value of deliberation in the FOMC measured by the probability that any given FOMC member would switch her behavior after listening to previous recommendations compared to the counterfactual case of no deliberation. On average, a FOMC member changes her policy recommendation after listening to previous speakers in 36% of the meetings along their tenure.

In terms of the performance of the FOMC, I find that the observed speaking order has been effective at reducing the probability of mistakes in monetary policy-making with respect to counterfactual deliberation orders. In particular, the best counterfactual speaking order, with more experienced FOMC members speaking first, decreases the probability of mistakes only by 3% with respect to the observed order.

Introducing sequential learning from deliberation turns out to fit the data remarkably well, explaining 92% of observed policy recommendations. This is a significant improvement over alternative behavioral models that fail to account for sequential learning in committee decision-making: an ideological model that characterizes members' behavior according to their preference divergence is able to explain 86% of observed recommendations, while a simultaneous model that incorporates heterogeneity in the quality of information across members but rules out the possibility of learning is able to explain 76% of observed recommendations.

The better performance of the *sequential deliberation model* comes from the fact that it allows ideology to interact with the value of information contained in member's private signals and in the previous recommendations made by other FOMC members. In addition, it is able to disentangle the effect of private information from that of the history of previous recommendations, providing expertise estimates that discount learning.

The rest of the paper is organized as follows. Section 2 places the contribution of this work with respect to the available literature. Section 3 introduces and develops the *sequential deliber-*

ation model. Section 4 describes the data and relevant institutional characteristics of the FOMC for the empirical analysis. Section 5 describe the estimation procedure and the identification of the empirical model. Section 6 presents the results from the estimation, discusses the counterfactual simulations and compares the relative performance of the *sequential deliberation model*. Finally, section 7 presents concluding remarks.

2 Related Literature

The empirical model estimated in this paper follows the theoretical literature that has emphasized the importance of heterogeneity in biases and in the quality of information as key factors behind the differential impact of deliberation in commitee-decision making under a common value framework (e.g., Austen-Smith and Feddersen [2005]; Austen-Smith and Feddersen [2006]; Coughlan [2000]; Doraszelski, Gerardi and Squintani [2003]; Gerardi and Yariv [2007]; Van Weelden et al. [2008]).

By developing an empirical model to explain the heterogeneity in individual behavior of committee members and assessing the extent of social learning within the FOMC, this paper builds on and extends the framework developed in Iaryczower and Shum [2012], which incorporates differences in the quality of private information into a purely spatial ideological model to explain decision-making in the U.S. Supreme Court. In the context of monetary policy, Hansen, McMahon and Velasco-Rivera [2014] estimated this model to the voting patterns of Bank of England's monetary policy committee to explain differences in ideological biases and expertise between internal and external committee members.

The presence of both preferences and private information in the model captures relevant features of monetary policy making that have been emphasized in the empirical literature on monetary policy decision-making (Blinder [2007]; Gerlach-Kristen [2006]). The preference biases of committee members can be interpreted as the relative costs of over- or under-predicting the true state of the economy, which is consistent with the different views of committee members regarding the tradeoff between inflation and unemployment. The quality of private information captures the expertise of committee members to gauge inflationary pressures. This expertise can be a function of the privileged data that members oftentimes use to discuss monetary policy. This private information can be acquired through business contacts in members' regions or through early access to certain economic indicators. Moreover, the heterogeneity in private information can capture differences in the amount of resources that members possess regarding their technical staff and the quality of the forecasts they produce.

Conditional on members' ideology and expertise, I incorporate the process of deliberation as a key feature of collective decision-making. In the model, the structure of debate can have important consequences, as it shapes members' inferences about the uncertain state of the economy. This feature arises because members, after listening to early speakers, weight the information and the potential for bias contained in previous recommendations against their own according to Bayes rule. This behavioral model incorporates Bayes-rational individuals as first introduced by Banerjee [1992] and Bikhchandani, Hirshleifer and Welch [1992] in the social learning literature, and later extended by Smith and Sørensen [2000] to allow for a continuum of signals and for heterogeneity in preferences.

There is a sizable empirical literature applying the social learning framework in economics.[2] In a political economy application, Knight and Schiff [2010] include social learning in an empirical model of sequential voting in primary elections. In the particular case of FOMC deliberations, Chappell, McGregor and Vermilyea [2012] use the policy recommendations for the period under Arthur Burns as Chairman to investigate the presence of Bayesian-updating in a "reduced-form" framework. The main limitations of their study, which prevents them to find any evidence of learning, is the assumptions that members have the same quality of information and preferences do not interact either with the value of private information or with the history of previous recommendations.

3 The Model

In the model there are T monetary policy meetings, $t = 1, \dots, T$. At any given meeting t, each committee member $i = 1 \dots, N$ offers a policy recommendation $r_{it} \in \{0, 1\}$ to the committee Chairman C, who then proposes a policy directive $d_t \in \{0, 1\}$ that is implemented as the committee's decision. In this setting 0 represents the lowest of two possible rate changes and 1 the highest.[3]

Member i's preferences over the policy directive (d_t) depends on the state of the economy, $\omega_t \in \{0, 1\}$, that encompasses unknown inflationary pressures, where $\omega_t = 1$ represents the high inflation state (consistent with a high interest rate) and $\omega_t = 0$ is the low inflation state (consistent with a low interest rate).

With full information, members want the directive to match the state, $d_t = \omega_t$. With this specification, we are assuming that members behave as if pivotal in the deliberation stage, that is, they act as if their recommendations actually change the policy directive. Members if pivotal, condition on the information that other members' behavior gives them about the underlying state of the economy (Austen-Smith and Banks [1996]). I will define the pivotal event for member i in meeting t as PIV_i^t.

The payoffs of $d_t = \omega_t = 0$ and $d_t = \omega_t = 1$ are normalized to zero. However, members disagree on the costs of implementing the incorrect directive (i.e., mismatching the state). Member i suffers a cost $\pi_i \in (0, 1)$ when the proposed directive is the low policy rate in a high inflation state ($d_t = 0$ when $\omega_t = 1$) and of $1 - \pi_i$ when the policy directive is a high rate in a low inflation state ($d_t = 1$ when $\omega_t = 0$). Accordingly, $1 - \pi_i$ can be thought of as member i's threshold of evidence above which she is willing to recommend the higher rate. Thus, $\pi_i > \frac{1}{2}$

[2] For a literature review see Bikhchandani, Hirshleifer and Welch [1998].

[3] In the context of the FOMC, d_t can be thought of as the policy proposal that the Chairman puts to a formal vote, which historically, has always coincided with the implemented policy at every FOMC meeting. Therefore, by abstracting us from modeling the final voting stage in this committee, we do not lose much in terms of explaining the actual influence of individual members in the policy-making process. A model that takes into account the presence of dissents in the voting stage would be relevant to explain monetary policy in a dynamic setting (across meetings), where dissents may have an effect on future actions of fellow members (Riboni and Ruge-Murcia [2014]).

reflects her bias towards the higher policy rate (i.e., member i is an inflation "hawk"), while $\pi_i < \frac{1}{2}$ reflects her bias towards the lower policy rate (i.e., member i is an inflation "dove").

I model the sequence of deliberation from the policy go-around, as follows. The inflation state ω_t is released but unobserved to committee members. In addition, the sequential order of speech is exogenously given to FOMC members. Members of the committee are ordered according to that sequence: member i offers her preferred policy option in rank $n(i)_t$, according to $p_t : N \to N$.

Prior to giving a policy recommendation, member i forms beliefs on ω_t by relying on four sources of information. First, there is public available information captured in members' common prior beliefs about the state of the economy, $\rho_t \equiv Pr[\omega_t = 1]$. Second, member i observes an informative private signal $s_{it}|\omega_t \sim \mathcal{N}(\omega_t, \sigma_i^2)$. Conditional on the state ω_t, these signals are statistically independent, with σ_i as a measure of the informativeness or precision of member i's information, which I label member i's expertise or ability interchangeably (i.e., lower σ_i denotes higher expertise or ability). Third, member i observes the history of recommendations when it is her turn to speak. We denote the relevant history for member i at meeting t, $\boldsymbol{x}_{n(i)_t} = (r_{1,t}, \ldots, r_{(n(i)_t-1),t}) \in \{0,1\}^{(n(i)_t-1)}$. The history for the member who speaks first is empty, $\boldsymbol{x}_{1,t} = \emptyset$. In this way, member i can potentially weight previous recommendations against her private information to update her prior belief on the state of the world ω_t. Fourth, given that members care about the committee's decision, member i's also conditions her recommendation on the information contained in the event that she is pivotal for the directive d_t, PIV_i^t. With this information at hand, the strategy for member i is defined by a map $\gamma_{it} : \mathbb{R} \to (0,1)$, where $\gamma(s_{it}) \equiv Pr(r_{it} = 1|s_{it})$.

I assume that members' types (π_i, σ_i) are public information and recommendations are heard by all committee members.[4] Note that by the normality assumption on s_{it}, the likelihood ratio

$$L(s_{it}) \equiv \frac{Pr[s_{it}|\omega_t = 1]}{Pr[s_{it}|\omega_t = 0]} = \frac{\phi(\frac{s_{it}-1}{\sigma_i})}{\phi(\frac{s_{it}}{\sigma_i})} = e^{\frac{2s_{it}-1}{2\sigma_i^2}}, \tag{1}$$

is increasing in s_{it}. This Monotone Likelihood Ratio Property implies that the equilibrium strategies are in cutoff points, where $\gamma(s_{it}) = 1$ if $s_{it} > s_{it}^*$ and $\gamma(s_{it}) = 0$, otherwise (Duggan and Martinelli [2001]). In particular, given the information contained in s_{it}, member i recommends the higher rate change, $r_{it} = 1$, whenever the posterior distribution $Pr[\omega_t = 1|s_{it}, \boldsymbol{x}_{n(i)_t}, PIV_i^t] \geq 1 - \pi_i$ and $r_{it} = 0$, otherwise. By basic manipulation of Bayes' rule, this condition can be written as

$$Pr[\omega_t = 1|s_{it}, \boldsymbol{x}_{n(i),t}, PIV_t^i]$$

[4] The only difference between the Chairman (C) and the rest of the committee, is that the former observes both her private signal s_{Ct}, and the full vector of reports of the N committee members $\boldsymbol{x}_{Ct} = (r_{1t}, \ldots, r_{Nt})$ when choosing the policy directive d_t at the end of the policy go-around.

$$= \frac{Pr[\omega_t = 1]Pr[s_{it}|\omega_t = 1]\prod_{j=1}^{n(i)_t-1} Pr[r_{jt}|\boldsymbol{x}_{n(j)_t}, \omega_t = 1]Pr[PIV_t^i|\omega_t = 1]}{\sum_\omega Pr[\omega_t]Pr[s_{it}|\omega_t]\prod_{j=1}^{n(i)_t-1} Pr[r_{jt}|\boldsymbol{x}_{n(j)_t}, \omega_t]Pr[PIV_t^i|\omega_t]}$$

$$= \frac{1}{1 + \left(\frac{1-\rho_t}{\rho_t}\right)\left(\frac{Pr[PIV_t^i|\omega_t=0]}{Pr[PIV_t^i|\omega_t=1]}\right) L(s_{it})^{-1}\prod_{j=1}^{n(i)_t-1}\Psi(s_{jt}^d)} \geq 1 - \pi_i;$$

Manipulating the normal density and solving for s_{it}, $r_{it} = 1$ whenever

$$s_{it} \geq \frac{1}{2} + \sigma_i^2\left[\log\left(\frac{1-\pi_i}{\pi_i}\right) + \log\left(\frac{1-\rho_t}{\rho_t}\right) + \sum_{j=1}^{n(i)_t-1}\log\left(\Psi(\boldsymbol{x}_{jt})\right) + \log\left(\frac{Pr[PIV_t^i|\omega_t = 0]}{Pr[PIV_t^i|\omega_t = 1]}\right)\right], \tag{2}$$

where the value of member j's recommendation on member i's equilibrium behavior, $\Psi(\boldsymbol{x}_{jt})$, is defined as

$$\Psi(\boldsymbol{x}_{jt}) \equiv \left[\frac{\gamma_{jt,0}(s_{jt}^*)}{\gamma_{jt,1}(s_{jt}^*)}\right]^{r_{jt}}\left[\frac{1-\gamma_{jt,0}(s_{jt}^*)}{1-\gamma_{jt,1}(s_{jt}^*)}\right]^{1-r_{jt}}. \tag{3}$$

Let s_{it}^* denote the value of s_{it} such that $s_{it} = s^*(\pi_i, \sigma_i, \boldsymbol{x}_{n(i)_t}, PIV_i^t, \rho_t)$. The effect of pivotality on s_{it}^* given in equation (2) depends on both the strategy profile of subsequent speakers and its effect on the Chairman equilibrium cutoff. Therefore, it is a function of the order of speech at each meeting given by p_t. The analytical expression for the pivotal event becomes convoluted as the order of speech decreases, as it needs to account not only for how member i's recommendation affects the Chairman' cutoff directly, but also indirectly through her effect on subsequent speakers. Both of these pieces of information are incorporated into the Chairman's posterior update on the state ω_t.

Consider as an example, the pivotality event for members i and j when they speak at the last ($n(i)_t = N$) and next-to-the-last positions ($n(j)_t = N-1$), respectively. These pivotality events are depicted in Figure 1. In the lower panel of Figure 1, member j directly affects the cutoff of member i in the interval B where $s_{it} \in [s_{it}^*(r_{jt} = 1), s_{it}^*(r_{jt} = 0)]$. The higher panel of Figure 1 shows the influence of both members i and j on the Chairman's cutoff. First, member j's recommendation affects s_{ct}^* directly in the green and blue intervals, for any given recommendation of member i. The red interval shows the change in s_{Ct}^* that results from both members switching from a high recommendation, $s_{Ct}(r_{jt} = 1, r_{it} = 1)$, to a low one, $s_{Ct}(r_{jt} = 0, r_{it} = 0)$.

In general, with the equilibrium cutoff pinned down, the probability of $r_{it} = 1$ in state ω_t can be written as

$$\gamma_{it,\omega_t}(s_{it}^*(\boldsymbol{x}_{it})) \equiv 1 - \Phi\left(\frac{s_{it}^*(\boldsymbol{x}_{it}) - \omega_t}{\sigma_i}\right). \tag{4}$$

Notice how the signal cutoff, s_{it}^*, varies across both members and meetings. First, differences in cutoffs across FOMC members can be explained by members' heterogeneity in both preferences, $\{\pi_i\}_{i=1}^N$, and expertise, $\{\sigma_i\}_{i=1}^N$. Second, movements across meetings in the cutoff

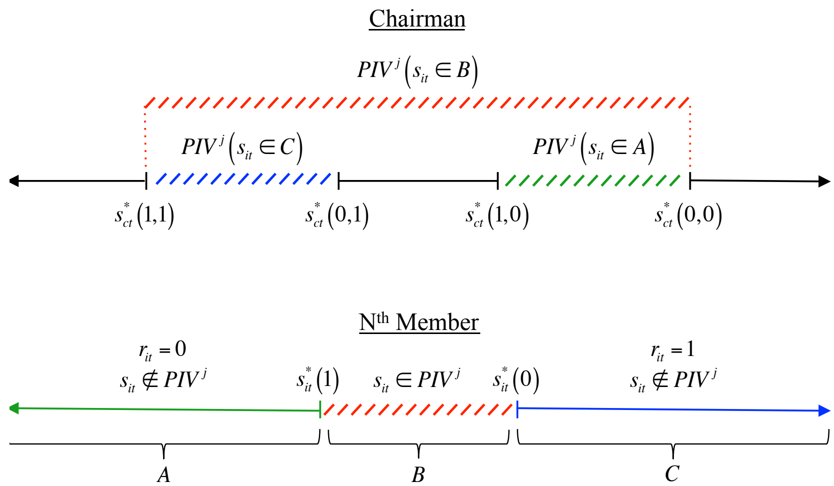

Figure 1: Pivotal Event for the Next-to-the-Last Member in the Speaking Order.

are captured by changes in the common prior (ρ_t). Finally, variation over both members and meetings is also explained by the differences in the order of speech that modifies the history of policy recommendations and the characteristics of subsequent speakers for any given member at each meeting.

Since behavior in this model is completely characterized by the signal cutoff, s^*_{it}, I can write the likelihood of observing the vector of recommendations and the Chairman decision at meeting t, $\boldsymbol{r}_t = (r_{1t}, \ldots, r_{Nt}, d_t)$, as

$$Pr[\boldsymbol{r}_t] = \sum_{\omega} \rho_t^{\omega_t}(1-\rho_t)^{1-\omega_t} \prod_{i=1}^{N+1} \gamma_{it,\omega_t}(s^*_{it})^{r_{it}}[1-\gamma_{it,\omega_t}(s^*_{it})]^{1-r_{it}}. \tag{5}$$

The likelihood in equation (5) as a function of equilibrium cutoffs, implicitly accounts for the history of previous recommendations in the sequential deliberation process given in equation (3). To better understand the role of this relevant parameter, consider a two-member committee where member i takes a policy decision ($n(i)_t = 2$), right after member j provides her policy recommendation ($n(j)_t = 1$). Under this scenario, the influence of member j on the equilibrium cutoff s^d_{it} can be written as

$$\log(\Psi(s^d_{jt})) = \begin{cases} \log(\gamma_{jt,0}) - \log(\gamma_{jt,1}) & \text{if } r_{jt} = 1 \\ \log(1-\gamma_{jt,0}) - \log(1-\gamma_{jt,1}) & \text{if } r_{jt} = 0. \end{cases}$$

Suppose, for instance, that member j recommends a high policy rate (i.e., $r_{jt} = 1$). The value of information for member i given by this action will depend on the relative likelihood that member j's recommendation matches the high state inflation (i.e., $\log(\gamma_{jt,0}) - \log(\gamma_{jt,1})$). If member j's probability of matching the state is equal to the probability of mismatching it, then deliberation would provide no informational value (i.e., $\log(\gamma_{jt,0}) - \log(\gamma_{jt,1}) = 0$).

Suppose instead that after listening to member j recommending the high rate ($r_{jt} = 1$), her probability of correctly matching the high state is larger than the probability of incorrectly recommending $r_{jt} = 1$ when $\omega_t = 0$ (i.e., $\log(\gamma_{jt,0}) - \log(\gamma_{jt,1}) < 0$). The additional information embedded in this recommendation will reduce member i's equilibrium cutoff in equation (2), making her more prone to follow member j's recommendation (i.e., $r_{it} = 1$).[5]

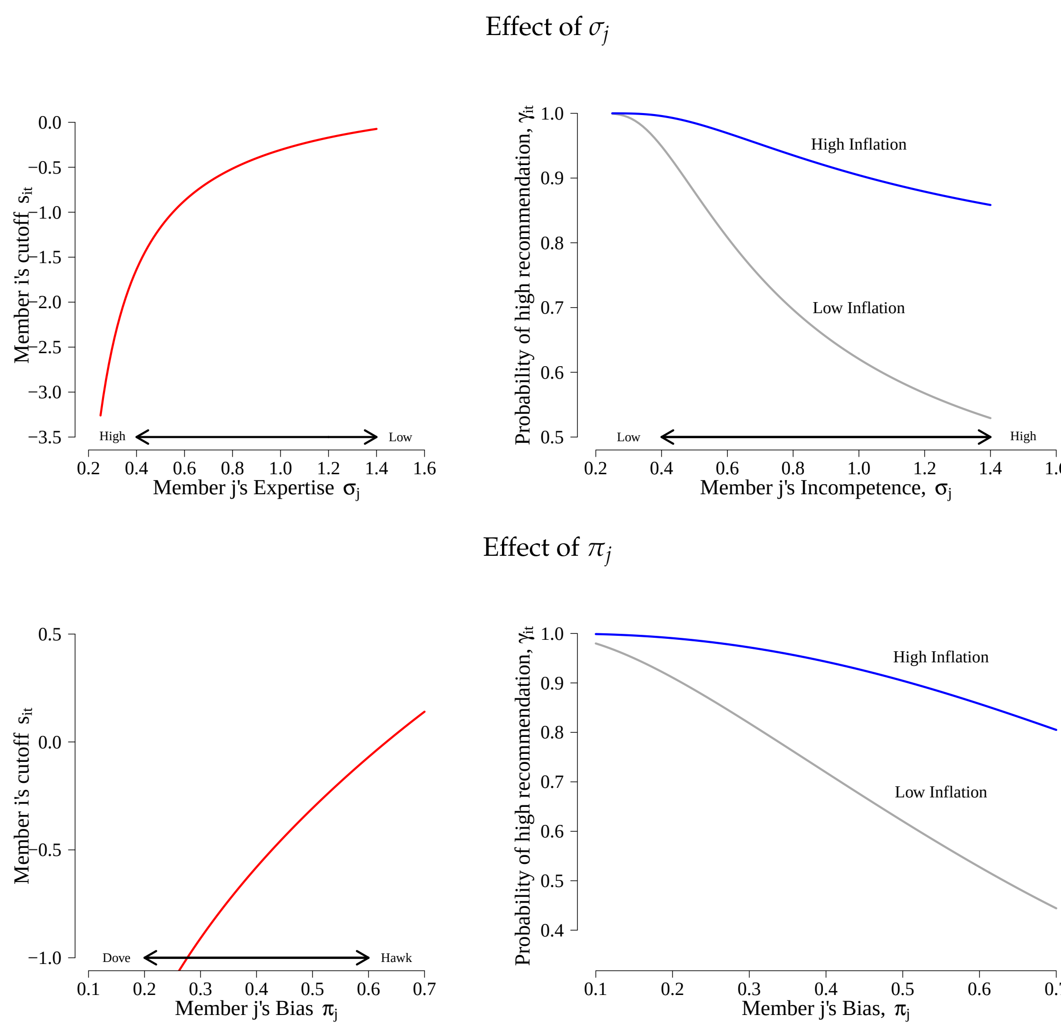

Figure 2: Potential Effect of Policy Recommendations on Subsequent Speakers. This figure presents the effect of varying a committee member j's expertise, σ_j (top panel), and ideological bias, π_j (bottom panel), on subsequent speaker i's equilibrium cutoff, s_{it}, and probability of following j's recommendation, γ_{it}. The changes in both expertise (from 0.25 to 1.4) and bias (from 0.1 to 0.7) come from the estimated parameters' distribution of the empirical model below.

It is important to emphasize that the magnitude of the shift in s^*_{it} after listening to member j's recommendation hinges on member j's expertise (σ_j) and bias (π_j). In particular, s^*_{it} is monotonic in both σ_j and π_j, but with different behavioral implications given their effect on member

[5] Notice also that the value of information of member j's recommendation can also work in the other direction: if $\log(\gamma_{jt,0}) - \log(\gamma_{jt,1}) > 0$, this would also provide member i with more information about the true state of the economy, ω_t, increasing the probability that member i goes against member j by recommending the low policy rate $r_{it} = 0$.

i's recommendation probabilities. Consider the upper panels of Figure 2, which display the effect of varying the expertise of member j on member i's optimal cutoff, s^*_{it}, and on its respective probabilities, $\gamma_{it,0}$ and $\gamma_{it,1}$. Notice that, as s_{jt} becomes more informative, the probability that member i mismatches both states of the economy diminishes, which makes her recommendation more influential on member i, reducing her cutoff, s^*_{it}, and increasing her probability of recommending the high rate, irrespective of the actual inflation state, ω_t.

Regarding the case of the effect of member j's bias (π_j), the lower panel of Figure 2 shows that, as member j becomes more "hawkish", member i increasingly discounts member j's recommendation ($r_{jt} = 1$) and increases the probability of recommending the opposite policy $r_{it} = 0$. This is because, as the bias of member j becomes more "hawkish", she will be more likely to match the high state while mismatching the low state.

4 Data and FOMC Institutional Background

Examples of the relevance of incorporating the sequential deliberation process to explain monetary policy-making can be extracted from the deliberation transcripts of the FOMC. Take for example the monetary policy meeting of March 1994 under Alan Greenspan as Chairman. During the policy deliberation portion of the discussion, Philadelphia district president Ed Boehne was the first member to speak and stated a recommendation in favor of tightening the policy rate 50 basis points, which was 25 basis point higher than the median policy the staff previously proposed and the one Chairman Greenspan stated as his preferred one. After him in the speaking order came district presidents Parry and Broaddus from San Francisco and Richmond district banks, respectively. Both members followed Boehne in his recommendations. More importantly, in making the case for his proposal <president Broaddus stated:

> *Let me just say that I agree 100 percent with Ed Boehne. He said it very well; he really reflected my position completely[...]. But my own feeling is the same as Ed Bohne's–that the risks are at least as great in not taking this action; I think there is a good chance that we would be seen as too cautious and too tentative.*

By accounting for the information contained in previous recommendations, the proposed empirical model is able to assess whether Broaddus' recommendation would have been different in the counterfactual scenario where he did not learn about Boehne's statement. More importantly, in the case that his recommendation contains additional information about the state of the world, the *sequential deliberation model* is able to attribute this effect to learning and not to the quality of Broaddus' private information, giving a more precise assessment of his ability as policymaker.

In this section, I explain the data collected from FOMC transcripts used for the estimation of the empirical model and provide the institutional context under which monetary policy is implemented in the FOMC.

4.1 FOMC Institutional Background

By the Banking Act of 1935, monetary policy decisions in the U.S. are the sole responsibility of the FOMC, which usually meets around eight times a year to set the short-term rate (i.e., Federal Funds rate) for open market operations - sales and purchases of government securities.[6] The current structure of the FOMC is depicted in Figure 3 and consists of seven members of the Board of Governors, including the committee's Chairman, as well as the twelve presidents of district Reserve Banks located throughout the country. All board members along with five of the twelve district presidents have voting rights at any given meeting.[7] Nevertheless, the remaining seven non-voting district presidents attend committee meetings, participate in the discussions, and contribute to the committee's assessment of the economy and policy options.[8]

The institutional appointment process of FOMC members differs between board governors and district presidents. The former are appointed by the President of the United States and ratified by the Senate to serve staggered fourteen-year terms.[9] The latter are chosen to serve five-years renewable terms by their own boards of directors with the consent of the Board of Governors. The board of directors of each district's Bank consists of nine members representing three different sectors: banking, agriculture and commerce, and a mix of academia and other members of the general public.

FOMC meetings throughout the period under study follow a standard protocol with four main stages. First, the staff offers an outline of economic conditions and forecasts regarding the current state of the economy nationwide. The presentation on the current state of the economy prepared by the staff is contained in a report that members receive before each meeting labeled the *Greenbook*, which includes data on the national economy, as well as the staff projections for the U.S. economy in the short and medium term. After the staff's presentations, individual members discuss their own impressions of the state of the economy, emphasizing first regional economic conditions in the case of district presidents, and then, the national and international economic situation.[10] The discussion of economic conditions is usually followed by the policy go-around. At this stage, the staff presents possible policy alternatives and their consequences to inform the committee as it proceeds to select a policy directive. Then, individual members verbally express their preferred policy position sequentially, with an order that varies across meetings. Finally, the Chairman crafts a directive that is brought to a formal vote by majority rule. In this stage, members can only agree or disagree with respect to the directive. In the case

[6] The Federal Funds rate is the rate at which commercial banks lend funds overnight with one another and is a crucial determinant of other rates with longer maturity.

[7] From the latter group, the district president of the Federal Reserve Bank of New York has a right to vote at every meeting, and four of the remaining district presidents serve one-year terms as voting members on a rotating basis.

[8] For the purposes of this paper, the term "member" is used for board governors, as well as both voting and non-voting presidents. The rotating voting seats are filled from the following four groups of Banks, one district president from each: Boston, Philadelphia, and Richmond; Cleveland and Chicago; Atlanta, St. Louis, and Dallas; and Minneapolis, Kansas City, and San Francisco.

[9] One of the seven governors is appointed Chairman by the U.S. President for a four-year term subject to a Senate confirmation.

[10] The *Beigebook* contains a summary of the economic conditions pertaining each of the twelve districts as organized by district presidents.

of disagreement, FOMC includes a brief statement in the minutes indicating the direction of the disagreement, from which it can be inferred whether members dissent because they want "easier" (i.e., higher policy rate) or "tighter" (i.e., lower policy rate) monetary policy.

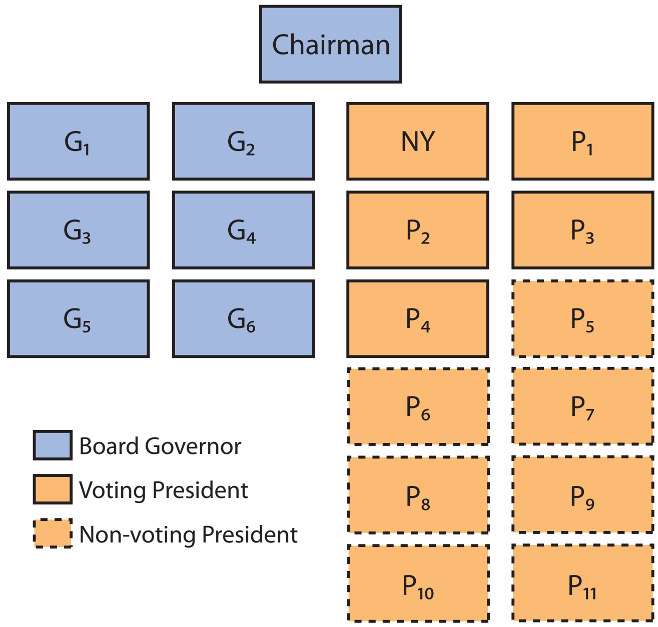

Figure 3: Structure of the FOMC.

4.2 Data

In principle, given the structure of FOMC meetings, I can analyze the information contained in both policy recommendations and voting records. In practice, FOMC voting records are not very informative to explain members' behavior. This is because dissenting votes are extremely rare in the policy-making history of the FOMC, as can be observed in Figure 4. The light blue bars in this figure show the yearly evolution of the number of dissenting votes with respect to the Chairman's policy proposal for the period 1966-2008, which covers five different chairmen. For the period under study, dissents represent, on average, only 5.8% of the total number of votes cast. The rare instances of dissent within the FOMC are also comparatively low with respect to those in other central banks. For example, Riboni and Ruge-Murcia [2014] find that dissents are significantly more frequent in the monetary policy committees of the Bank of England and the Sveriges Riksbank than at the Federal Reserve. Moreover, there has not been a single instance in FOMC's history where the Chairman' policy directive is on the losing side of the vote.[11] Therefore, the Chairman's policy directive invariably coincides with the implemented policy rate at any given meeting. This feature has been noted by Swank and Visser

[11] In addition, dissenting voting records do not provide information about the behavior of non-voting committee members, who nevertheless, attend FOMC meetings, discuss monetary policy, and ultimately express their desired policy in front of the rest of the committee at the deliberation stage.

[2007], among others, who argue that the FOMC as a whole is known to appreciate showing a united front to the market observers regarding the voting decision that is immediately released to the public after each meeting. This consensus-seeking desire constraints chairmen to offer policy proposals that can obtain at least a majority of votes. Nevertheless, as has been shown by Chappell, McGregor and Vermilyea [2005], the alignment between the Chairman's proposal and the policy recommendations of either the FOMC median or mean voters is consistent with a policy directive that is influenced by FOMC members.

The limitation of voting records to characterize the FOMC has been noted since the 1960's, despite the fact that all the work that followed on the topic well into the 2000's, focused precisely on these records, as this quote from Yohe [1966] summarizes:

> *The reasons are not at all clear for the almost uncanny record of the Chairman in never having been on the losing side of a vote on the policy directive. While there is no evidence to support the view that the directive always voted upon and passed on the first ballot merely reflects the Chairman's own preference, there is also no evidence to refute the view that the Chairman adroitly detects the consensus of the committee, with which he persistently, in the interest of System harmony aligns himself.*

(William Yohe, *"A Study of the Federal Open Market Committee Voting"*, cited in Chappell, McGregor and Vermilyea [2005].)

Fortunately, records of FOMC deliberations contained in FOMC transcripts provide us with the discussion that leads to a policy adoption, in which FOMC members share their views about the future state of the economy and voice their preference for a particular policy rate. All of this, before votes are cast and officially recorded.

The amount of information one can extract from the deliberation process can also be seen in Figure 4, where the dark blue bars show the yearly evolution of the amount of voiced dissent, measured as differences in the voiced policy recommendation of each member with respect to the Chairman's directive during the policy go-around. Just by looking at the discrepancies in dissent between deliberation and voting stages, one can draw a different picture of members' behavior than the one that can be extracted solely from voting patterns. For instance, the proportion of voiced dissent with respect to the Chairman's proposal reaches an average of 33% over the period under study. This increase represents almost a fivefold jump in disagreement with respect to what can be found from looking at voting records.

The voiced policy recommendations shared by FOMC members in the policy go-around, as well as the record of their order of speech at every meeting under study, are obtained from the verbatim transcripts of FOMC meetings. To systematically code the recommendations and speaking order of each committee member from textual records, I followed the efforts of Chappell, McGregor and Vermilyea [2012] who collected these voiced interest rate recommendations and a record of the speaking order for the period under Arthur Burns as a Chairman between 1970 and 1978. I complemented and extended these data myself by collecting, whenever possible, the desired policy rate and speaking order of every FOMC committee member during the Chairmanship of G. William Miller (1978-1979), the Greenspan years (1987-2006), and the

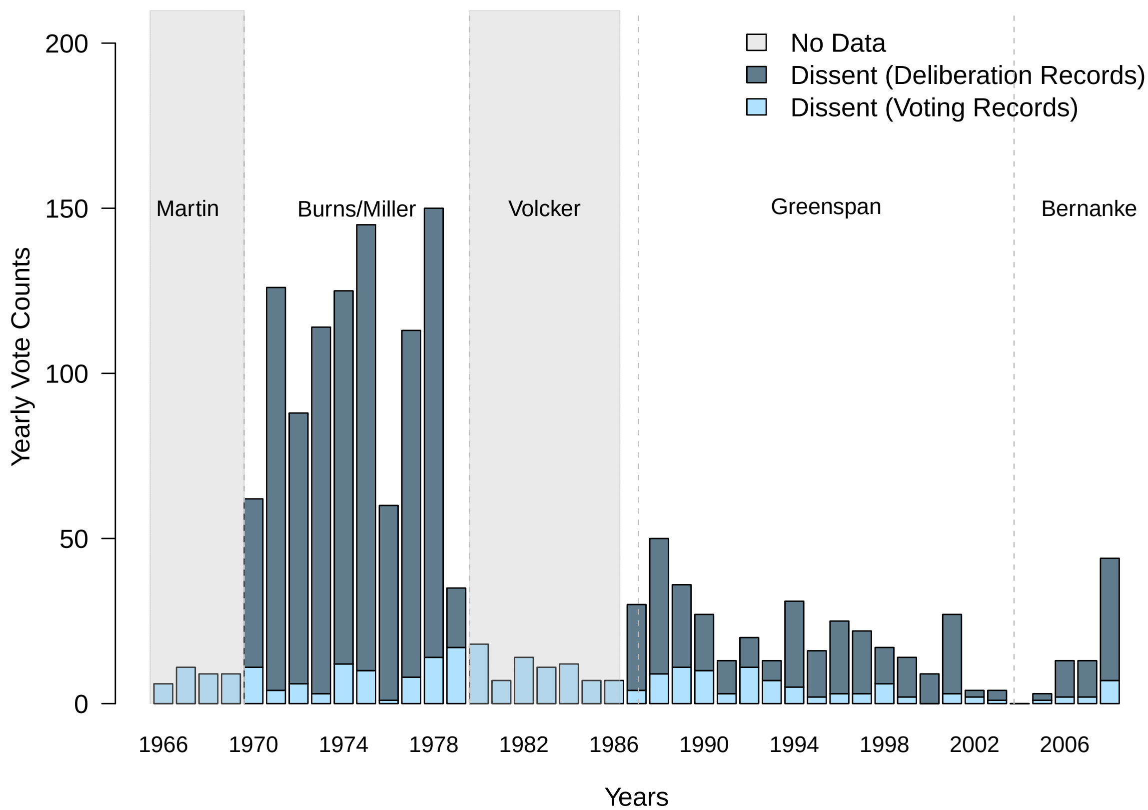

Figure 4: History of Dissents at the FOMC. This figure presents the yearly counts of dissents across committee members and meetings for both voting records (light blue) and policy recommendations expressed during the pre-vote deliberation stage (dark blue) for the period 1966-2008 encompassing five FOMC Chairmanships. Policy recommendations are unavailable for the Chairmanships of Martin (1966-1970) and Volcker (1980-1986) highlighted in gray.

Bernanke period (2006-2008).

From the available transcripts, I excluded the period under Volcker (1979-1986) because, during his tenure as Chairman, the FOMC changed its main policy instrument from a Fed Funds rate to a borrowed reserves instrument that directly targeted the money supply, making the coding and comparison across periods infeasible. I also excluded the available meetings held during 2009 under Bernanke given that, as a consequence of the economic crisis of 2008, the Fed Funds rate reached the zero lower bound in December 2008 and remained at this level throughout the following year.[12]

I classify members' desired policy rates into binary (low *vs* high rate) recommendations, by first establishing a benchmark policy with which members' preferred rates could be compared. For this purpose, I rely on the policy scenarios suggested and distributed by the staff to FOMC

[12] In addition, since the financial crisis monetary policy has taken a turn towards unconventional instruments that target the balance sheet of the central bank through the purchase of mortgage-backed securities and other securitized assets.

members in advance of each meeting and then summarized just before the policy go-around takes place.[13]

I quantify a composite benchmark from these different alternatives by computing the median proposed policy offered by the staff at each meeting. Then, based on the textual records of deliberations, I code members as recommending a high policy rate whenever their desired Fed funds rate target is equal or higher than the staff median proposal and a low policy rate, otherwise. In those instances in which desired rates are not observable, I imputed a binary recommendation if members expressed a leaning direction or assenting preference with respect to the staff proposal, or to the recommendation of other members who explicitly expressed a desired rate.

I examine the policy recommendations of all members who sat on the FOMC for the period under study, excluding from the analysis those who participated in less than 10% of all meetings under consideration. In total, the sample comprises 265 monetary policy decisions made by 57 voting and nonvoting FOMC members for a total of 3,490 policy recommendations. Table 1 presents the distribution of policy recommendations, along with the average macroeconomic conditions during each of the Chairmanships under consideration. As can be seen from this table, the sample of policy recommendations analyzed here were made under very diverse economic conditions, which coincide with changes in the identity of the FOMC Chairman. On the one hand, the Burns and Miller regimes were characterized by high and increasing levels of inflation, paired with a strong slowdown in economic growth; whereas the Greenspan years coincide with a period of sustained output growth with low and stable inflation, a prosperous period that ended abruptly during the Bernanke regime, with the largest economic crisis since the Great Depression, albeit under a period where inflation remained anchored at low levels.

Period	Meet	Rec	Size	Unan %	$r_{it} = 0$	$r_{it} = 1$	Fed Funds	Inf	GDP	Unem	M1
Burns ('70-'78)	99	1203	12	44.44	27.18	72.82	6.44	5.55	3.62	6.32	5.79
Miller ('78-'79)	11	138	13	18.18	34.78	65.22	7.97	7.35	4.58	5.93	5.87
Greenspan ('87-'06)	132	1917	15	62.12	28.12	71.88	4.93	4.06	2.48	5.63	3.95
Bernanke ('06-'08)	23	232	10	39.13	50.43	49.57	4.05	3.33	1.59	5.07	0.85
All data ('70-'08)	265	3490	13	51.70	29.98	70.02	5.54	4.69	2.92	5.85	4.45

Note: Author's calculations. Meet denotes the total number of meetings per period. Rec denotes the number of recommendations per period. Size refers to the median size of the committee for each period. Unan % is the percentage of unanimous recommendations per period. $r_{it} = 0(1)$ refers to the percentage of low (high) rate recommendation per period. Fed Funds, Inf, GDP, and Unem refer to period averages for the Fed Funds rate, quarterly forecasts for inflation, real GDP growth, and civilian unemployment, as presented in the *Greenbook* by the staff of the Board of Governors. M1 denotes the three-month moving average money growth around the date of FOMC meetings, also provided in the *Greenbook*.

Table 1: Policy Recommendations by Chairmanship, 1970-2008

[13]This data is contained in the *Blue Book* provided to members around a week in advance of FOMC meetings.

5 Estimation and Identification

I describe the procedure to estimate the *sequential deliberation model* outlined in section 3 and then discuss identification issues.

5.1 Estimation

I directly recover both preference and expertise parameters from the likelihood function in equation (5). This contrasts with the two-step approach developed by Iaryczower and Shum [2012] that first estimates a flexible "reduced-form" version of individual choice probabilities while controlling for individual and time-varying covariates. Then, it recovers the structural parameters by solving for the equilibrium conditions of the voting game they analyze.

The benefit from the "direct" approach in this context is that it does not rely on estimates from reduced-form voting probabilities, which makes it insensitive to the robustness of "first-stage" parameters. I implement a Bayesian estimation of the structural parameters that easily incorporates a hierarchical structure that exploits variation across members and meetings. Finally, it allows me to directly estimate parameter uncertainty, as it approximates the full posterior distribution, instead of relying on modal approximations, such as the Delta method, or quasi-Bayesian simulations.

The "direct" estimation approach comes at a cost, as it calculates the recommendation probabilities across committee members over different meetings for every trial value of the parameters, which can be computationally intensive. For this reason, I implement the approximation of the posterior distribution with an efficient Markov chain Monte Carlo (MCMC) algorithm, *via* the Hamiltonian Monte Carlo method (Homan and Gelman [2014]).[14] This technique includes ancillary parameters that allow the algorithm to move further in the parameter space at each iteration, providing faster mixing, even in high dimensions.

The estimation algorithm of the empirical model requires two main related steps: first, the computation of the equilibrium condition, and the subsequent construction of the likelihood ("the inner loop"), and second, the estimation of the parameter vector ("the outer loop"). The estimation of the model is done sequentially at every meeting t using the observed speaking order of committee members. In this manner, I am able to incorporate the effect of both sequential learning and pivotality, $\sum_{j=1}^{n(i)_t-1} \log\left(\Psi(\boldsymbol{x}_{jt})\right)$ and $\log\left(\frac{Pr[PIV_t^i|\omega_t=0]}{Pr[PIV_t^i|\omega_t=1]}\right)$, respectively, when updating the optimal cutoff.

Equilibrium Condition (Inner Loop): Fix a parameter vector $\boldsymbol{\theta} \equiv \{\{\pi_i, \sigma_i\}_{i=1}^{N+1}, \rho_t\}$. For member in order $n(i)_t = 1, \ldots, N$:

1. Solve for the equilibrium condition in equation (2).
2. Given s_{it}^*, compute $\gamma_{it,0}(s_{it}^*)$ and $\gamma_{it,1}(s_{it}^*)$ using equation (4).
3. Compute $\sum_{j=1}^{n(i)_t-1} \log\left(\Psi(\boldsymbol{x}_{jt})\right)$ using equation (3).
4. Compute the increment of the likelihood at every meeting t from equation (5).

[14] The estimation of the joint posterior distribution is implemented in the software STAN (Team [2015]).

Approximation of the Joint Posterior Distribution (Outer Loop): Given the likelihood function in equation (5), I write the posterior distribution of the vector of parameters ($\boldsymbol{\theta}$) as a proportion of the product of the likelihood and its prior distribution

$$
\begin{aligned}
&Pr\left[(\boldsymbol{\theta},\boldsymbol{\lambda})|\boldsymbol{r}_t\right]\\
\propto &Pr(\boldsymbol{\theta},\boldsymbol{\lambda})Pr[\boldsymbol{r}_t|\boldsymbol{\theta}]\\
= &Pr(\boldsymbol{\lambda})Pr(\boldsymbol{\theta}|\boldsymbol{\lambda})\prod_{t=1}^{T}\sum_{\omega}\rho_t^{\omega_t}(1-\rho_t)^{1-\omega_t}\prod_{1}^{N+1}\gamma_{it,\omega_t}(s_{it}^*)^{r_{it}}[1-\gamma_{it,\omega_t}(s_{it}^*)]^{1-r_{it}},
\end{aligned}
$$

where I aggregate the increments to the likelihood over FOMC meetings and $\boldsymbol{\lambda}$ denotes the vector of hyperparameters in the model.

1. I allow for heterogeneity in the common prior beliefs by estimating ρ_t to vary as a function of meeting characteristics $\boldsymbol{X_t}$ that are available to committee members before the sequential deliberation process of the policy go-around takes place:

$$
\rho(X_t) = \frac{exp(X_t'\delta)}{1+exp(X_t'\delta)}; \qquad \delta \sim N\left(0,(9/4)I\right), \tag{6}
$$

   where δ is a fixed coefficient that is normally distributed. The value imposed on the variance is consistent with an uninformative prior for $\rho_t \approx \frac{1}{2}$. $\boldsymbol{X}_t$ is a vector of meeting-level predictors obtained from the *Greenbook* and distributed to FOMC members prior to each meeting. This vector includes the level of the Federal Funds Rate the week prior to each meeting (*previous policy*), recent money growth (M1) calculated as the mean of the last three available monthly figures prior to each meeting. Finally, I include two-quarter ahead staff forecasts of the inflation rate ($\mathbb{E}$(Inflation)), unemployment ($\mathbb{E}$(Unemployment)), and GDP growth ($\mathbb{E}$(RGDP Growth)).

   I account for a switch in the transparency of FOMC deliberations, since prior to November 1993 FOMC members were not aware that meeting deliberations were being recorded and eventually published. After November 1993, meeting discussions took place under the assumption that every individual statement and comment would be publicly available within five years after each meeting. Thus, to measure this transparency change, I include an indicator variable (*transparency*) that takes the value of one after November 1993 and zero, otherwise.

   To fully control for changes in the composition of the FOMC over time and for different agenda-setting power across chairmen, I include an indicator variable for the identity of the FOMC Chairman at the time of each meeting (Burns, Miller, Greenspan, or Bernanke). These Chairman effects are relevant in this context because they capture differences in the deliberation protocol across FOMC regimes, specifically regarding the intervention of the Chairman in the policy go-around. For instance, Burns sometimes spoke early, stating a preference for a particular policy rate while Greenspan routinely spoke right after the staff, suggesting a specific proposal. Bernanke, did not intervened during the policy go-around, waiting after all members spoke to craft a policy directive. This informal influence from the Chairman to the rest of the FOMC is an important component of agenda setting power that shapes not only the voting stage of the decision-making process within the FOMC, but also the flow of the debate, which is accounted for in the empirical model.[15]

2. For the estimation of pivotality effects on the equilibrium threshold, I let $Pr[PIV_t^i|\omega_t]$ be a function of covariates of member i at meeting t and of the subset of members who speak after her (i.e., $\{n(i)_t + 1, \ldots, N+1\}$). I choose a functional form to constrain $Pr[PIV_t^i|\omega_t] \in [0,1]$ as follows:

$$\log\left(\frac{Pr[PIV_t^i|\omega_t = 0]}{Pr[PIV_t^i|\omega_t = 1]}\right) = \alpha_1 Early_{it} + \alpha_2 Late_{it} + \boldsymbol{X}_{it}\boldsymbol{\beta}'. \tag{7}$$

The vector $\boldsymbol{X}_{it}$ of member-meeting level covariates includes member i's experience at meeting t (*Rookie*), which is given in the form of a binary indicator that takes the value of one if member i served in less than 34 meetings, which represents the 25th percentile of a member's tenure in the sample, and zero, otherwise.

For the covariates of remaining members speaking after member i, I include the fraction of them who are rookies ($Experience_{remain}$), who are district presidents ($Pres_{remain}$), and who are Democrat-appointed governors (Dem_{remain}). I also include the average fraction of their past career (i.e., before the FOMC) spent in private financial institutions (Fin_{remain}), as economists (Eco_{remain}), and within the ranks of the Federal Reserve (CB_{remain}). To measure the past career experience of FOMC members, I employ and expand the measure of career backgrounds created by Adolph [2013] which partitions central bankers' past jobs into seven mutually exclusive categories, namely: financial (i.e., private banking jobs), government (i.e, bureaucrats outside the Federal Reserve and the Treasury Department), finance ministry (i.e., bureaucrats in the Treasury Department), central bank (i.e., staffers within the Federal Reserve System), economics (i.e, academic economists), business (i.e, private sector excluding banks), and other (e.g., international organization officials). The career experience for each category is computed as the fraction of the FOMC member's career spent in that job category up to the date of her most recent appointment as a FOMC member.[16] Finally, this specification includes random effects for the order of speech of member i ($\alpha_0 1$ and α_2) where *Early* (*Late*) takes the value of one (zero) if member i speaks in the first (second) half of the policy go-around and zero, otherwise.

3. For the remaining structural parameters and their respective hyperparameters, I choose

[15] In the empirical estimation I also allow chairmen Burns and Greenspan to speak more than once at any given meeting whenever they voiced a policy recommendation during the policy go-around in addition to their policy directive at the end of the sequential deliberation process.

[16] Figure 26 in the Appendix show the distribution of career experience across job categories for the sample under study.

the following distributional assumptions based on their constrained scale in the model:

$$\begin{aligned}
\pi_i &\sim Beta(\alpha_\pi, \beta_\pi), \text{ for } i = 1, \ldots, 57. \\
\sigma_i &\sim Cauchy(0, \tau_\sigma) \text{ for } i = 1, \ldots, 57. \\
\alpha_\pi\ \beta_\pi &\sim U(0, 10), \\
\tau_\sigma &\sim Cauchy(0, 2).
\end{aligned}$$

4. I obtain posterior samples of the vector of parameters from its posterior marginal density at each iteration $m = 1, \ldots, M$. I run three parallel chains with dispersed initial values for $10,000$ iterations each with an initial warm-up period of 5,000 iterations and thinning of 100. I assess convergence for each parameter based on the potential scale reduction factor, $\hat{R}$ (Gelman and Rubin [1992]) and through a visual inspection of the trace plots. Appendix G contains the traceplots for the model hyperparameters and the main parameters of interest, as well as a set of sampling statistics relevant for the diagnosis of the Hamiltonian Monte Carlo sampler.

5.2 Identification

Having laid out the estimation procedure, the formal identification of the parameters ρ_t and of the equilibrium probabilities γ_{it,ω_t} in the likelihood function of equation (5) is given by the fact that, conditional on the unobserved state ω_t, the observed vector of recommendations, $\boldsymbol{r}_t$, is drawn from a finite mixture distribution with mixing parameter equal to the common prior (ρ_t). Under the Bayesian framework, the estimation of mixture models transforms its complex structure by simpler conditional ones using latent variables or unobserved indicators, as given in this case by the state of the economy, ω_t, that specifies the mixture component from which policy recommendations are drawn. The identification is solved by imposing distributional assumptions on prior parameters and sampling ω_t from its full conditional distribution.[17]

The identification of the structural parameters contained in equilibrium cutoffs s^*_{it} is as follows. In the case of the common prior, ρ_t, the identification comes from the presence of a common value ω_t in the empirical model. In particular, the prior is identified from the frequency with which the majority of FOMC members recommend the high rate. This is due to the fact that high values of the common prior induce higher signals for all FOMC members at any single meeting. Thus, as the instances where the majority of members choosing the high rate increase, the estimated value of ρ_t also increases.

For the preference parameter π_i, the identification comes from the assumption on preference differences. Changes in the common prior, ρ_t, induce increases in the probability of voting for the high rate, but to a larger degree for members with a high value of π_i. Therefore, low variability in the pattern of recommendations over meetings for a particular FOMC member will be estimated as more extreme bias.

[17] As can be seen from the visual inspection of the traceplots for each parameter of interest in appendix G, there does not seem to be evidence of label-switching, which is a common problem of other Bayesian mixture models.

The identification of members' expertise (σ_i) comes from the common value feature of the model. This is because increases in the common prior, keeping preferences (π_i) fixed, will induce a higher signal correlation in which members with higher ability will be better able to predict the true state of the economy. In the data, a member with an observed pattern of recommendations that follows the majority over time, will be estimated as having a high expertise (i.e., low σ_i). Analogously, a member whose pattern of recommendations tends to disagree with this majority will be estimated as having a low quality of information (i.e., high σ_i).

As can be seen in equation (2), the value of deliberation is directly identified from the non-linear function Ψ that maps $(\rho_t, \{pi_i, \sigma_i\}_{i=1}^N)$ into a social learning parameter that varies both across members and over meetings.

To separately identify the pivotality effect from the value of deliberation, I exploit the variation in members' order of speech across meetings that allows me to treat $\log\left(\frac{Pr[PIV_t^i|\omega_t=0]}{Pr[PIV_t^i|\omega_t=1]}\right)$ as a primitive to be identified and estimated directly from covariates of remaining members.[18] The availablity of a different speaking order across meetings potentially changes the composition of subsequent speakers for any given member, which is a source of variation in the pivotal event beyond the one induced by $(\rho_t, \{pi_i, \sigma_i\}_{i=1}^N)$.

To clarify this strategy, consider the equilibrium cutoff in the extreme case where the speaking order is the same across meetings. Under this scenario, it would not be possible to estimate changes in the pivotality effect, as these would be fixed across meetings and not separable from the effect of members' biases and expertise. Fortunately, FOMC speaking order varied substantially across both members and meetings. I observe FOMC members sharing their policy recommendation in different speaking positions along their tenure. According to anecdotal evidence, there was no prescribed order to speak before each meeting's policy go-around took place. Laurence Meyer, former board governor, labeled the order assignment as "the wink system", in which each FOMC member would wink at the FOMC deputy secretary her ideal position on the policy go-around at any given meeting. However, the FOMC secretary, who is a member of the staff of the Board of Governors, would decide the final speaking order at his discretion. Then, the Chairman would call upon the FOMC in the order of that list without members knowing in advance which exact speaking position they were going to be called upon(Meyer [1998]).

Figure 5 plots the mean speaking order of each FOMC member in the sample, along with the distribution of speaking order along members' tenure. Overall, it can be seen that, with the exception of members McDonough and Hayes who spoke first at 80% of the meetings they were part of, there was a substantial variation in the speaking order across meetings. In fact, members spoke, on average, just 8% of the meetings in any particular speaking order (with a 6% standard deviation).[19]

[18] One could in principle solve for the equilibrium cutoff points for every trial value of the parameters. However, given the potential for multiple equilibria, one could end up selecting a solution of cutpoints not played in the data. By directly estimating the pivotality effect from the data, we avoid this problem.

[19] The reason behind the small variation for New York presidents McDonough and Hayes comes from the fact that, during the Burns and early Greenspan years, the New York district president, who also serves as the FOMC

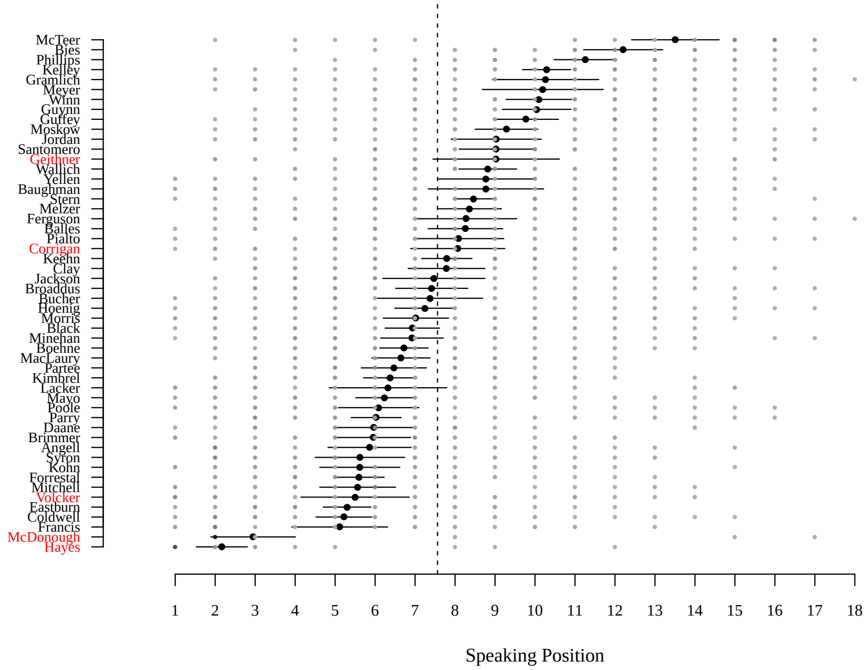

Figure 5: Speaking Order Across FOMC Members. This figure shows the mean speaking order during the policy go-around of each FOMC member along their tenure. Solid lines depict 95% confidence intervals. Dots depict the distribution of speaking positions for each FOMC member across meetings. Darker colors denote a higher relative frequency for a particular speaking position. For instance, a black point depicts a speaking position where a member spoke 100% of their time along their tenure. Members highlighted in red are FOMC vice-chairmen. Vice-chairmen Hayes, McDonough and Volcker usually were granted the right to provide their policy recommendations early during the policy go-around.

6 Results

I begin by describing the results for the effect of meeting-level covariates, $\boldsymbol{X_t}$, to predict the common prior (ρ_t) which tracks the evolution of the unobserved state of the economy, ω_t. Figure 6 displays the expected prior estimated from equation (6), under hypothetical values of the explanatory variables. In particular, these counterfactual scenarios are constructed by changing each covariate from its 10th percentile value to its 90th percentile value in the sample, while keeping the rest of explanators at their median. The main takeaway point from these results is that the economic indicators included in the specification have a large and significant influence in predicting the common prior, ρ_t. These effects go in the expected direction in terms of the tradeoff between inflation on the one hand, and output growth and unemployment, on the other. These effects highlight the importance of some of these indicators as proxies of inflationary pressures and economic growth. For instance, increments in expected

Vicechairman was usually granted the informal right to speak first in the sequence. Still, even for these members, we can find meetings where they provided their recommendations late in the policy go-around.

output growth, $\mathbb{E}$(RGDP Growth) and expected inflation $\mathbb{E}$(RGDP Growth), as captured by staff forecasts, are associated with larger inflationary pressures and therefore, with predicted increments in the common prior ρ_t. On the other hand, increments in expected unemployment, $\mathbb{E}$(Unemployment), are perceived by FOMC members as diminishing inflationary pressures, while increasing the negative risks for economic growth. Consistent with a reversion to the mean effect, high levels of the prevailing policy rate (*Previous Policy (FFR)*) are negatively associated to larger inflationary states.

The effect of chairman individual effects are calculated with respect to the Burns chairmanship (1970-1978), which is the omitted category. The Greenspan years up to the transparency reform of 1993 were more prone to higher inflationary states, while the post-transparency years (1994-2008) have reversed this trend, as they have been a period where inflation has been well anchored at low levels.

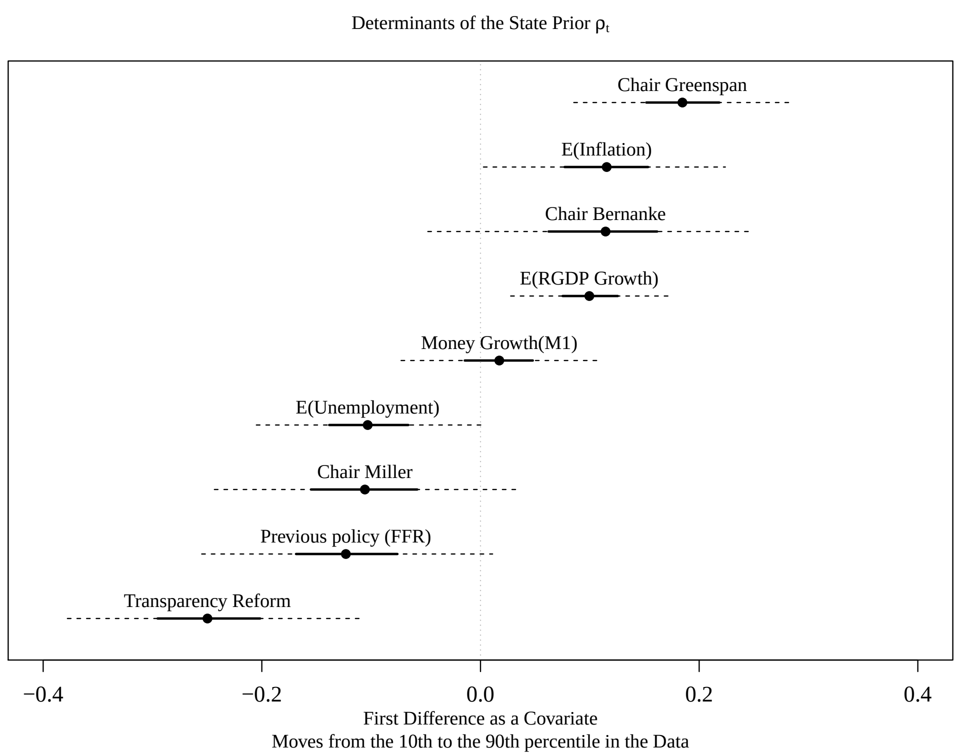

Figure 6: Determinants of the Prior (ρ_t) on the Unobserved State of Inflation $\omega_t = 1$. The figure provides the effect of increasing each of the displayed covariates on the common prior $\rho_t \equiv Pr(\omega_t = 1)$. Solid circles give the posterior median, with vertical solid lines corresponding to the interquartile range of the posterior distribution and dashed lines corresponding to 95% posterior credible intervals. The counterfactual increase in the covariate of interest is a change from the 10th to the 90th percentile in the sample. For each estimate, all other covariates are set at their median sample values.

The effect of both meeting-level covariates and chairman individual effects ultimately map into a predicted common prior about the state of the economy that captures both the effect of objective economic indicators, as well as the interpretation of FOMC members about these effects. Figure 7 shows the evolution of the predicted common prior (ρ_t) over the period under

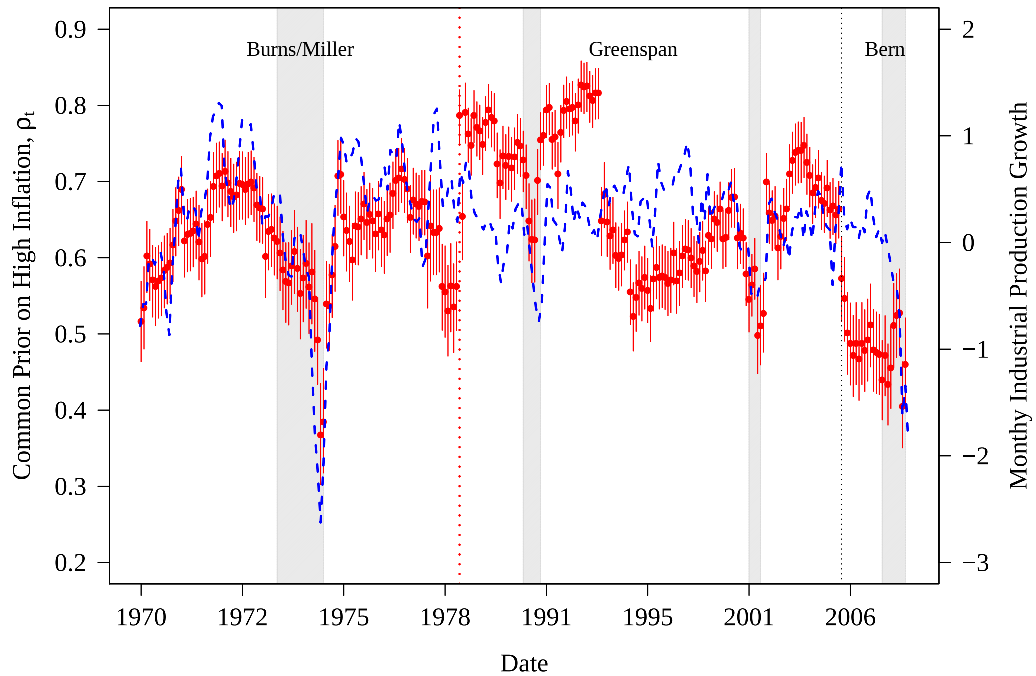

Figure 7: Evolution of Prior ρ_t Over Time. The figure provides the estimated value of the common prior $\rho_t \equiv Pr(\omega_t = 1)$ across meetings. Solid circles give the posterior median, with horizontal solid lines corresponding to the interquartile range of the posterior distribution. Gray areas correspond to economic recessions in the sample measured as the period from peak to trough of a business cycle according to the National Bureau of Economic Research. The blue dashed line on the secondary right axis depicts the monthly industrial production growth during the same period obtained from real-time data series collected at the Philadelphia Federal Reserve (`www.philadelphiafed.org`).

study. The first thing to notice is that the estimated common prior follows the actual trade-off between inflation and output growth for the period under study remarkably close. For instance, the estimated ρ_t sharply decreases in periods with deteriorating output and unemployment, while increasing following economic expansions and higher inflation risks. In fact, the gray shades in this Figure are evidence that sustained declines in the estimated common prior (ρ_t) are closely associated with the presence and duration of all four economic recessions that hit the U.S. economy in the sample under study, as measured by the National Bureau of Economic Research (NBER). In addition, the common prior closely follows fluctuations in output growth, as captured by monthly changes in industrial production, which is an indicator that is not part of the covariates employed in the estimation specification of ρ_t.[20]

Figure 8 summarizes the findings related to the posterior estimates of preference biases ($\{\pi_i\}_1^N$). The top panel provides the ranking of members according to the magnitude of these biases. Solid circles give the posterior median, with horizontal lines corresponding to the in-

[20] In fact, the data on monthly industrial production is not part of the information set of FOMC members at any given meeting, as it published after monetary policy meetings take place.

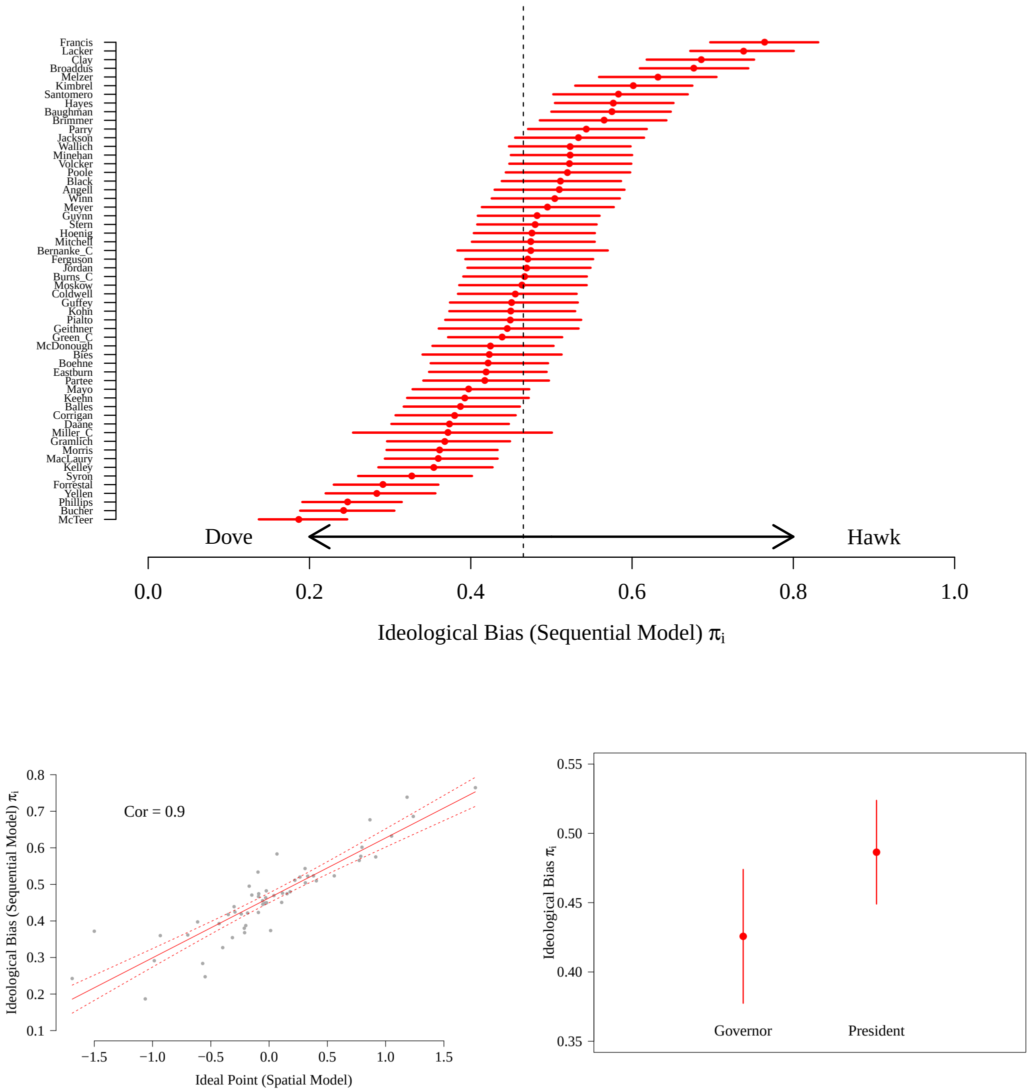

Figure 8: Preference Bias Estimates, π_i. The top panel provides posterior summaries of the ideological bias, π_i for each FOMC member during the periods 1970-1979 and 1987-2008. Solid circles give the posterior median, with horizontal lines corresponding to the interquartile range of the posterior distribution. Committee members who participated in more than 30 meetings are included. The bottom left panel of the Figure provides the result of a linear fit along with 95% confidence intervals between ideological biases π_i, as recovered by the *sequential deliberation model* and ideal points, z_i, as recovered by the *spatial ideological model*. The bottom right panel of the figure provides posterior summaries of the ideological bias of FOMC committee members, π_i, aggregated by appointment, where a member is either a district president board governor.

terquartile range of the posterior distribution. The posterior median of ideological biases across members ranges between 0.19 for Dallas district president Robert McTeer to 0.76 for Darryl Francis, district president of the Saint Louis Fed during the Burns period. The distribution of estimated ideological biases shows a high degree of polarization, with 46% of them being statistically different from 0.5 (with a 90% confidence level), which captures ideologically neutral members.

Members' preference biases, $\{\pi_i\}_1^N$, are defined in the empirical model as members' thresholds of evidence above which they are willing to recommend a higher policy rate. In the data, however, these biases seem to be mostly associated to an ideological dimension in terms of a "hawk-dove" spectrum, which has been popularly used to classify central bankers in general, and FOMC members in particular. To check this issue, the lower left panel of Figure 8 plots the correlation between members' estimated biases and the ideal points of FOMC members that are obtained from fitting a *spatial ideological model* to the pattern of policy recommendations. The estimation of the *spatial ideological model* is a Bayesian version of a multilevel Item Response Theory (IRT) model fitted to the policy recommendations of FOMC members. Under this model, committee members recommend the policy alternative that is closer to their preferred policy or ideal point, $z_i \in \mathbb{R}$. The details of the estimation can be found in appendix A.[21]. The fact that members' biases, π_i, are extremely similar to the rank order of ideal points (with a Pearson (Spearman) correlation of 0.90 (0.91)) indicates tha this ideological interpretation fits well the profile of FOMC members.

Consider district president Tom Melzer, who is placed on the highest tail of the distribution of members' biases. He has been recognized by the press and other central bankers as one of the most "hawkish" members in the history of the FOMC. In fact, at almost every opportunity he had, he stated his views on monetary policy that can be summarized in the following quote from one of his speeches: *"In my opinion, the main contribution the Fed can make to the economy in the long run is to keep inflation low and inflation uncertainty to a minimum. This means maintaining a consistent policy over a long period of time with a credible commitment to low inflation."* (Melzer [1994]). In contrast, on the "dovish" extreme, the spatial model places former governor and current chairwoman of the FOMC Janet Yellen, who has been characterized as a "dovish" member by the media, given her policy views that can be summarized in the following extract from one of her interventions during a FOMC meeting, *"...I would agree that the Fed probably cannot achieve permanent gains in the level of unemployment by living with higher inflation. But the Federal Reserve can, I think, make a contribution on the employment side by mitigating economic fluctuations-by stabilizing real activity."*(Yellen [1995]).

The difference between a board governor, such as Janet Yellen and a district president like Francis or Melzer in terms of their bias differences goes beyond the anecdotal. In fact, the relative ordering of members' preferences is systematically correlated with their appointment process, as can be confirmed from the evidence depicted in the lower right panel of Figure 8. This figure provides posterior summaries of the ideal points of FOMC members aggregated by

[21] This technique has been widely used in the political science literature to empirically analyze spatial models of voting (e.g., Clinton, Jackman and Rivers [2004]; Jackman [2000]; Poole and Rosenthal [2000])

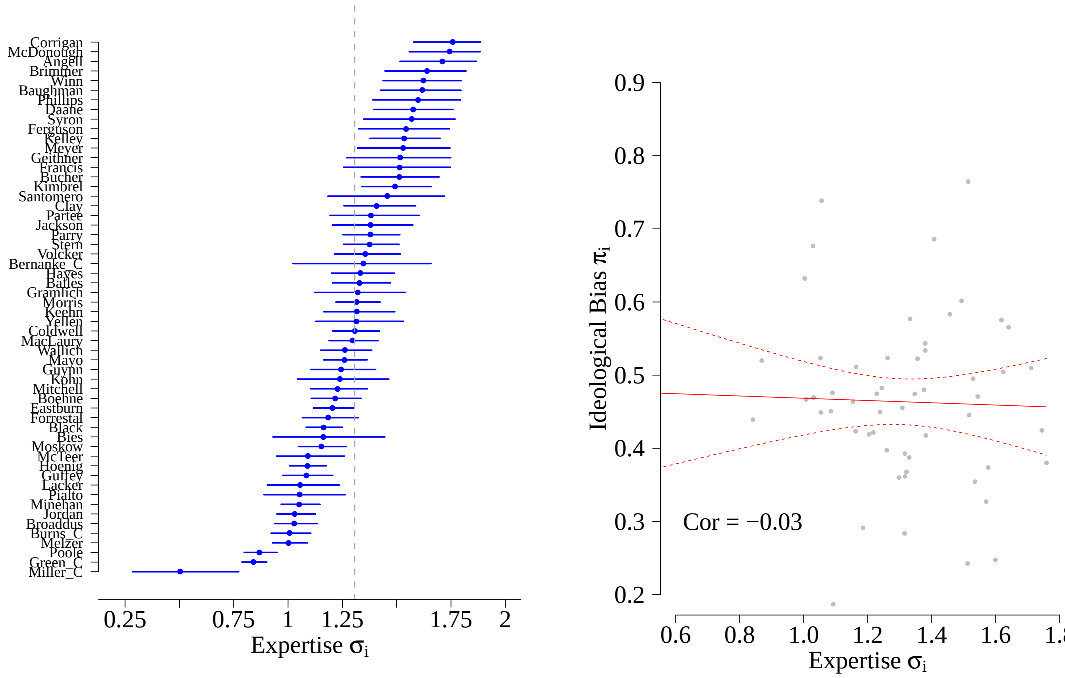

Figure 9: Expertise Estimates, σ_i. The left panel of the figure provides posterior summaries of the measure of expertise of FOMC committee members, σ_i. Solid circles give the posterior median, with horizontal lines corresponding to the interquartile range of the posterior distribution. Committee members who participated in more than 30 meetings are included. The right panel of the figure provides the results of a linear fit, along with 95% **confidence intervals between the estimated posterior median ideological biases, π_i, and individual expertise, σ_i, as recovered by the *sequential deliberation model*.**

appointment. I find that board governors, who are appointed by the President, are 19% more "dovish" than Federal Reserve presidents, who are appointed by district boards of directors comprised of regional banking and industry interests. As shown in in column (3) of Table 3 in the Appendix, this finding is statistically robust to controlling for both the party of the President who appointed board governors and the career experience of FOMC members.[22]

Overall, these results partially confirm previous studies on the FOMC that have explained the differences between board governors and district presidents in terms of the political pressure through appointment that the Executive exerts on board governors (Chang [2003]). According to this argument, U.S. presidents, who are presumed as having biased preferences towards the real side of the economy, appoint central bankers with similar preferences to implement "dovish" policies. However, the observed behavior of FOMC members while deliberating monetary policy cannot be characterized exclusively as a mere reflection of members' ideology. At best, this characterization is incomplete, unless two important features of monetary policy deliberation are considered: first, there is the notion that monetary policy entails

[22]This table also replicates this exercise using the ideal points, z_i, from the *spatial ideological model* (column (1)), as well as those biases obtained from a model I label the *simultaneous model* (column (2)), which incorporates members' private information, but assumes that recommendations are made simultaneously, ruling out the possibility of information transmission through sequential deliberation (i.e., ignores learning from previous speakers as well as pivotality effects). The details and posterior estimates of the model without deliberation can be found in appendix B.

implementing a policy that seeks to match the true state of the economy, an issue that hinges on efficiently interpreting current economic conditions in an environment of pervasive uncertainty; second, there is the important feature of deliberative committees, such as the FOMC, in which the structure of debate itself can have important consequences in the decision-making process, mainly in shaping members' inferences about the uncertain state of the world.

The estimates of FOMC members' expertise in gauging the state of the economy are depicted in in the left panel of Figure 9. From this plot, it can be observed a sizable amount of heterogeneity across FOMC members, with an interquartile range that goes from a signal quality (σ_i) of 0.5 for Chairman Miller to 1.76 for New York district president Corrigan. The dispersion in members' expertise is a fundamental component to expand our understanding of committee decision-making in general, and of the FOMC in particular, which was missing in previous empirical work on the topic. Mainly, it places the ideological divisions in perspective, showing that preference differences cannot account for the total variation in the behavior of committee members. In fact, members' ideological biases are not systematically related to their expertise, as can be seen in the right panel of Figure 9. The null relationship between biases and ability comes from the observed behavior of FOMC members and not from any modeling assumptions, as the theoretical framework does not impose any covariance structure between members' preferences and their information structure.[23]

A policy relevant result obtained from the heterogeneity of expertise estimates across FOMC members is that three out of four FOMC chairmen in the data, with the exception of Ben Bernanke, are ranked in the top five of members' expertise distribution. The importance of this result comes from the fact that the policy directive that is selected at FOMC meetings is ultimately crafted by the Chairman and as such, his expertise to track the evolution of relevant economic indicators and to efficiently incorporate the information and discount the biases that other members bring to the discussion is crucial in determining the quality of the implemented policy.

Given the interaction of private information and sequential learning, assessing the expertise of committee members is drastically changed once members are allowed to learn from each other. Compared to a *simultaneous model* where parameters are recovered without accounting for the deliberation process, there is a substantial difference in members' expertise estimates, from an average of 0.48 for the *simultaneous model* to 1.3 for the *sequential deliberation model*.[24]. In addition, the rank order of members' expertise estimates changes drastically when members are allowed to learn from previous recommendations compared to the case without learning. This is shown in Figure 10, which compares the expertise estimates for both behavioral models. The reason behind these discrepancies arises from the relevance of deliberation as an information-sharing mechanism. The *simultaneous model* interprets any potential value of information contained in the history of recommendations of early speakers as a higher

[23]When I split the sample of FOMC members by appointment, and in contrast to the preference biases, there does not seem to be any statistically significant difference between board governors and district presidents as shown in Table 4 in the Appendix.

[24]The estimation results for the *simultaneous model* can be found in Appendix B.

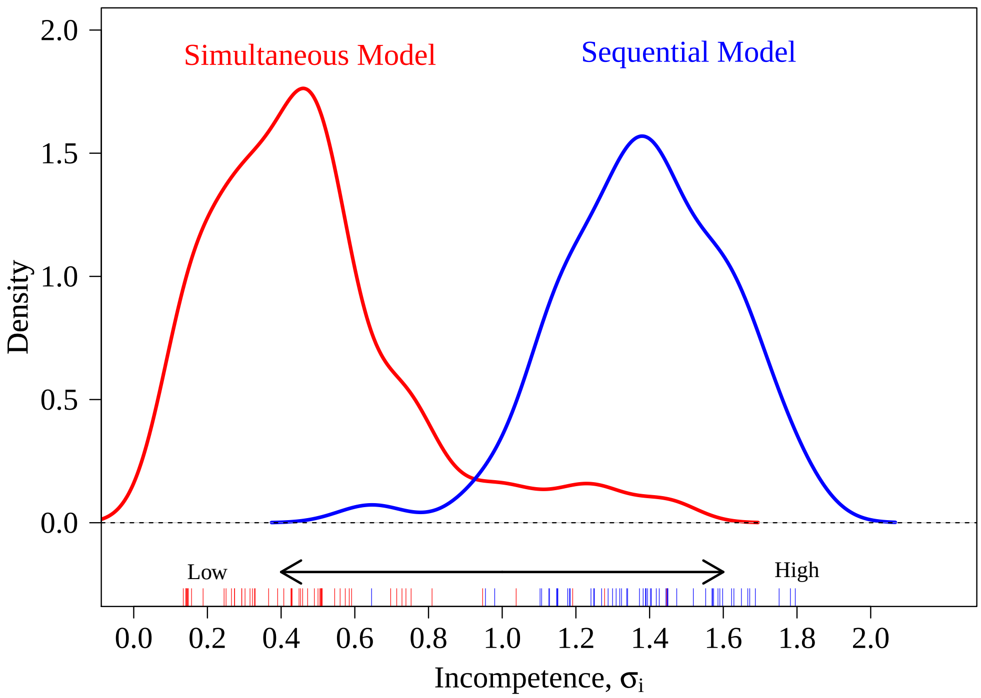

Figure 10: Expertise under the Sequential Deliberation and Simultaneous Models. The figure provides the distribution of expertise (σ_i) estimates under the two different behavioral models. Under the *simultaneous model*, members provide their recommendations taking into account their bias and private information. Under the *sequential deliberation model*, members also take into consideration the recommendations of previous speakers and the probability of being pivotal on the Chairman's proposal.

quality of their private information. In contrast, by explicitly accounting for the presence of information complementarities through sequential deliberation, the *sequential deliberation model* attenuates the relevance of private information in the quality of members' recommendations and distinguishes it from the social learning embedded in the deliberation process.

6.1 The Value of Deliberation

With the structural estimates at hand, I present a measure to quantify the value of deliberation within the FOMC. This measure is defined as the frequency with which a member would change her policy recommendation after incorporating the policy proposals of previous speakers, compared to the scenario where she exclusively follows her private information.

To compute the value of deliberation, I generate two policy recommendations for each of the N members across the T meetings under consideration, r_{it}^{D} under the actual *sequential deliberation model* and r_{it}^{S} for the counterfactual committee without deliberation. The procedure is as follows:

For each posterior draw $m = 1, \ldots, M$ from the parameter distribution:

1. Draw the state of the economy ω_t conditional on ρ_t.
2. Draw a signal for each member from $\mathcal{N} \sim (\omega_t, \sigma_i^2)$.
3. Draw a policy recommendation for each member:
   - $r_{it}^{D} = 1$ if $s_{it} \geq s_{it}^{*} \equiv (\pi_i, \sigma_i, \boldsymbol{x}_{it}, PIV_i^t, \rho_t)$.
   - $r_{it}^{S} = 1$ if $s_{it} \geq s_{it}^{**} \equiv (\pi_i, \sigma_i, \rho_t)$.
4. The value of deliberation by member is given by $\frac{1}{T}\sum_{t=1}^{T} \mathbf{1}_{r_{it}^{S} \neq r_{it}^{D}}$.
5. The value of deliberation by meeting is given by $\frac{1}{N}\sum_{i=1}^{N} \mathbf{1}_{r_{it}^{S} \neq r_{it}^{D}}$.

Given the Bayesian estimation framework, I am able to analyze the posterior distribution of the value of deliberation and estimate its uncertainty by computing p_{th} percentiles from the M posterior random draws.

The estimated value of deliberation takes a value of zero whenever members' policy recommendations are identical with or without deliberation. The value of deliberation takes a value of one whenever members' policy recommendations differ from the one without deliberation.

The left panel of Figure 11 presents the average value of deliberation in the FOMC for each member across meetings. The average value of deliberation across FOMC members is 36% on average, with a substantial variation across members. On one extreme, an inflation hawk such as president Francis would switch his policy recommendation after listening to other members only 6% of the time. On the other extreme, board governors Phillips and Chairman Bernanke would change their recommendations after incorporating previous recommendations 48% of the time.

Notice that from the group of chairmen, Bernanke and Miller are in the top five of members' ranking according to the value of deliberation, whereas Greenspan and Burns rank in the middle of the distribution. In fact, Greenspan is the chairman with the lowest value of deliberation (39%), which is consistent with previous accounts of Greenspan's dominance, in which he played a predominant role during FOMC deliberations by steering the policy directive closer to his initial leanings and expertise rather than to the opinion of other FOMC members. A clear example of this behavior comes from the February 1994 meeting, as former FOMC member Alan Blinder notes "[...] *when the Fed began a cycle of interest rate increases by moving the Federal funds rate up* 25 *basis points. The transcript of that meeting (which is now public) shows that a clear majority of the committee favored moving up by* 50 *basis points. Greenspan, however, insisted not just on* 25 *basis points, but on a unanimous vote for that decision. He got both.*" (Blinder [2008])

The right panel of Figure 11 partially confirms the dominance of Greenspan over the FOMC as a whole during his early years as Chairman. This plot shows the time trend of the value of deliberation averaging over FOMC members. From Greenspan's appointment up to the transparency reform of FOMC transcripts in 1993, the average value of deliberation decreased to its lowest point historically, from 35% during the Burns/Miller period to 22%. However, from 1990 onwards there has been an upward trend in the value of deliberation that peaked in the late Greenspan's years (i.e. around 2004).

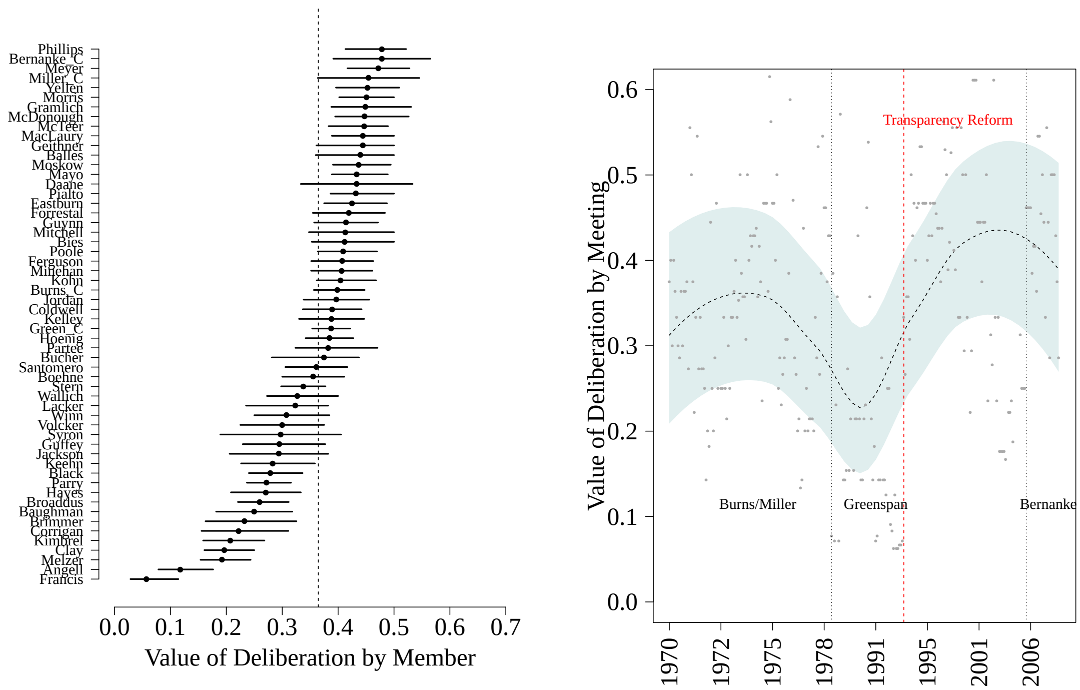

Value of Deliberation by Member   Value of Deliberation by Meeting

Figure 11: The Value of Deliberation in the FOMC. The left panel shows the posterior distribution of the value of deliberation aggregated by member, with points corresponding to the median and solid lines to the interquartile range of the posterior distribution. The right panel of the figure plots the smoothed (by 2nd degree local polynomial) time trend of the value of deliberation by meeting. The shaded area corresponds to the posterior interquartile range. The value of deliberation is defined as the probability that a member i in meeting t change her policy recommendation after incorporating the recommendations made by previous speakers, as well as her pivotality effect compared to the scenario where she follows her private information.

Figures 12 and 13 present the correlates of the value of deliberation by member and meeting, respectively. Figures 12 shows that the variation of the value of deliberation across members can be systematically explained by differences in members' biases and expertise. The left panel of this figure indicates that inflation "hawks" tend to rely more on their private information than "neutral" and "dove" members. Moving from the most "dovish" member (i.e., McTeer) to the most "hawkish" one (i.e., Francis), reduces the value of deliberation from 43% to 13%. The right panel of this figure is evidence that FOMC members with a high degree of expertise tend to incorporate others' recommendations more often than those with less expertise. Moving from the member with highest expertise (i.e., Miller) to the member with the lowest expertise (i.e, Corrigan) decreases the value of deliberation from 45% to 29%.

In terms of the variation of the value of deliberation over time, Figure 13 shows that beyond differences across chairmen, the value of deliberation is higher during low inflationary states as can be seen from the fact that an increased in expected inflation reduced the value of deliberation around 11%, whereas an increase in expected unemployment increase the value of deliberation around 12%.

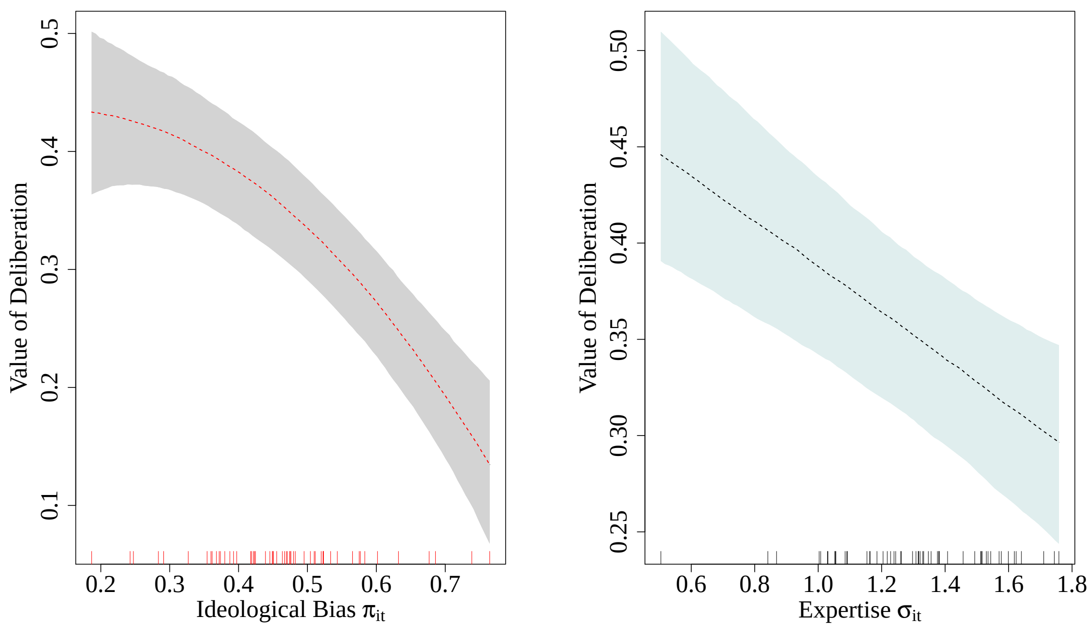

Figure 12: Covariates of Value of Deliberation by Member. The left panel shows the relationship between members' value of deliberation and their ideological biases, π_i. The right panel shows the relationship between members' value of deliberation and their expertise, σ_i.

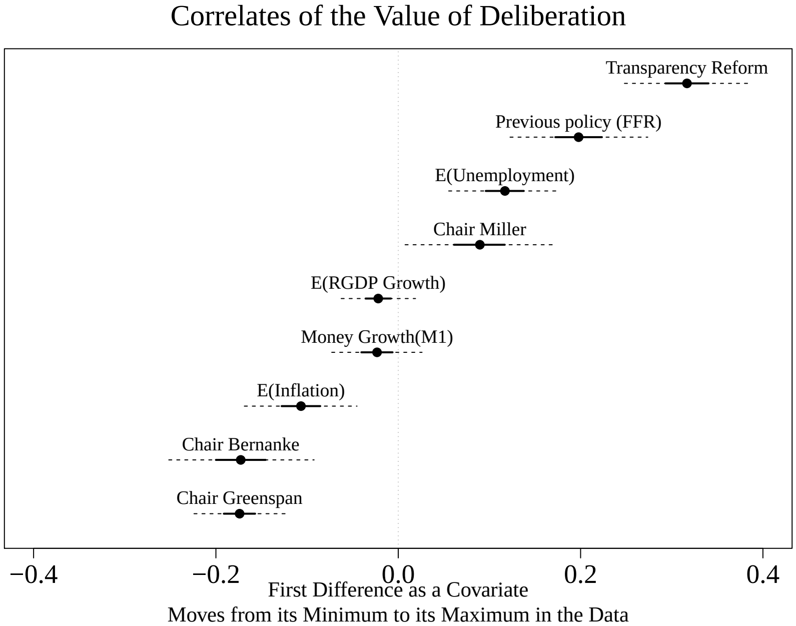

Figure 13: Covariates of Value of Deliberation by Meeting. The figure shows the relationship between the estimated value of deliberation by meeting and time-varying covariates. Solid circles give the posterior median, with horizontal solid lines corresponding to the interquartile range of the posterior distribution and dashed lines corresponding to 95% posterior credible intervals. The counterfactual increase in the covariate of interest is a change from its minimum to its maximum in the sample. For each estimate, all other covariates are set at their median sample values.

6.2 Social Learning *vs* Pivotality

The value of deliberation shown above is calculated under a model where policymakers weight their own information against the recommendation of previous speakers, as in a pure social learning framework and in addition, account for the potential effect that their decision might have on the chairman's proposal, as included in their pivotality effect. Thus, up to this point it is unknown whether the information they obtain from the deliberation process comes mainly from social learning or from strategic considerations. To shed light on this issue, I compute the relative weights that both considerations have on members' equilibrium cutoffs, s_{it}.

For each member's cutoff expressed in equation (2), I calculate the average magnitude of social learning and pivotality that is given by the absolute value of $\sum_{j=1}^{n(i)_t-1} \log\left(\Psi(\boldsymbol{x}_{jt})\right)$ and $\log\left(\frac{Pr[PIV_t^i|\omega_t=0]}{Pr[PIV_t^i|\omega_t=1]}\right)$, respectively, as a proportion of the optimal cutoff's magnitude

$$\left|\log\left(\frac{1-\pi_i}{\pi_i}\right)\right| + \left|\log\left(\frac{1-\rho_t}{\rho_t}\right)\right| + \left|\sum_{j=1}^{n(i)_t-1} \log\left(\Psi(\boldsymbol{x}_{jt})\right)\right| + \left|\log\left(\frac{Pr[PIV_t^i|\omega_t=0]}{Pr[PIV_t^i|\omega_t=1]}\right)\right|. \quad (8)$$

Figure 14 presents the results of this exercise. The top panel compares the densities of the relative weights of social learning and pivotality in members' equilibrium cutoffs. The lower panel disaggregates these relative weights by member. Overall, it can be seen that, on average, 51% of FOMC members' behavior is driven by the effect that deliberation exerts on their recommendations. From this, committee members assign a very small weight to strategic considerations and instead, place a larger emphasis on the information provided by previous speakers when providing their own recommendations.

The average weight that FOMC members put on pivotality considerations is 10% on average. Not only the average magnitude is small, but also its variation across FOMC members, which goes from a minimum weight of 4% for chairman Burns to a maximum of only 14% for governor Mitchell. On the other hand, the relative weight of social learning on FOMC members' behavior is four times larger than that of pivotality, as it accounts for 41% of the total magnitude of member's cutoffs. Moreover, there is a significant dispersion of the weight members put on social learning. Consistent with the value of deliberation presented above, the weight that inflation "hawks", such as presidents Francis and Hayes assign to social learning is less than 20%, which contrasts to the weights of 58% and 69% that Governor Meyer and Chairman Bernanke assign to previous recommendations, respectively.[25]

[25]Figure 24 in Appendix D shows that, at least for the last and next-to-the-last speakers, estimating pivotality effects using the flexible specification with covariates shown in equation (7) provides similar results to those that would be obtained by exactly computing the analytical posterior updates of the Chairman as given by equations (**??**) and (**??**), for the last and next-to-the-last speakers, respectively.

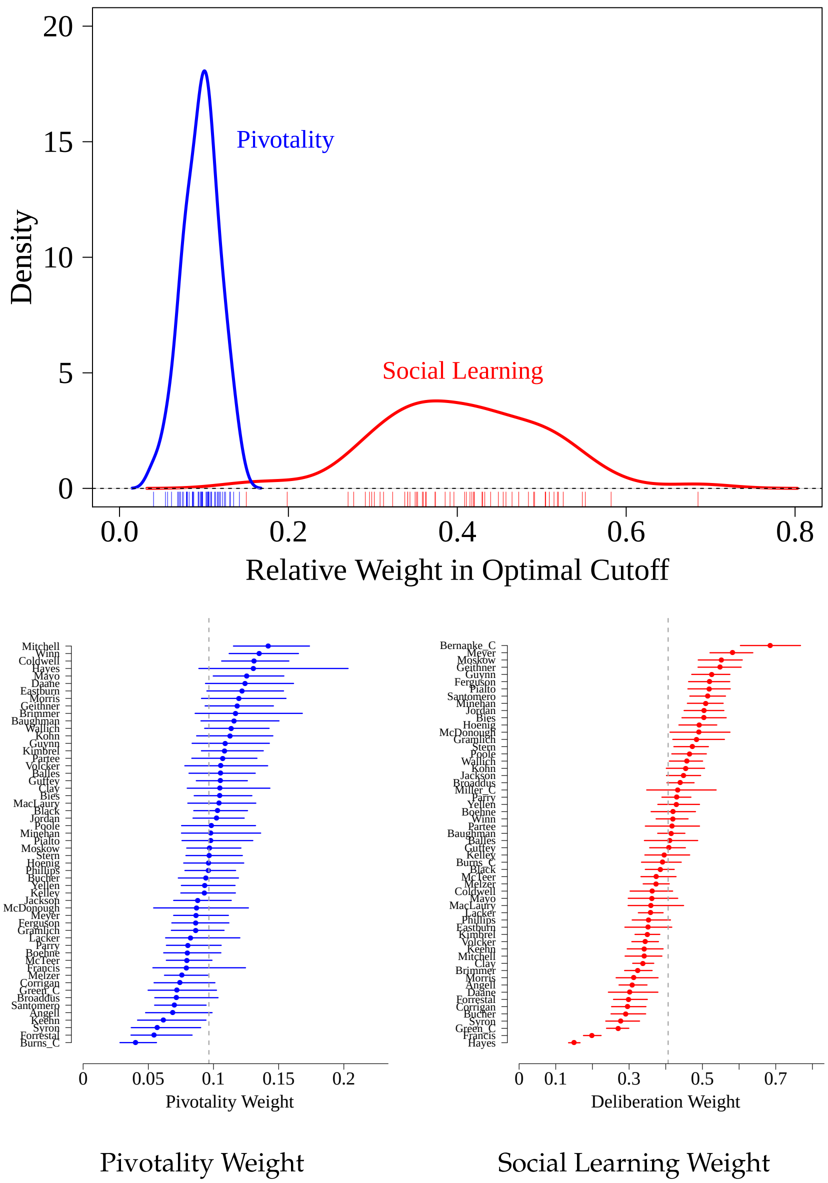

Figure 14: Relative Weights of Social Learning and Pivotality in Committee Members' Optimal Cutoffs. The top panel of the figure shows the distribution across committee members of the median weight of both pivotality and sequential learning in member's optimal cutoffs. The bottom panel shows these relative weights by committee member. For each member, solid points denote the median of the posterior distribution with lines depicting the interquartile range of the posterior distribution. The relative weight of social learning and pivotality are computed as the absolute value of $\sum_{j=1}^{n(i)_t-1} \log\left(\Psi(x_{jt})\right)$ and $\log\left(\frac{Pr[PIV_t^i|\omega_t=0]}{Pr[PIV_t^i|\omega_t=1]}\right)$, respectively as a proportion of $|\log\left(\frac{1-\pi_i}{\pi_i}\right)| + |\log\left(\frac{1-\rho_t}{\rho_t}\right)| + |\sum_{j=1}^{n(i)_t-1} \log\left(\Psi(x_{jt})\right)| + |\log\left(\frac{Pr[PIV_t^i|\omega_t=0]}{Pr[PIV_t^i|\omega_t=1]}\right)|$.

6.3 Order of Speech and Correct Decisions

In any given policy go-around, the FOMC chairman, after listening to the sequence of policy recommendations of individual members, arrives at a policy directive that is officially voted by majority rule. As this directive has obtained at least a majority of votes at every meeting in the history of the FOMC, I compute a measure of the quality of decision-making at the FOMC that is given by the probability that the chairman proposes a policy directive that is consistent with the true state of the economy, meaning proposing a high policy rate when the state is high (i.e., $\gamma_{it,1}$ when $\omega_t = 1$) and proposing a low policy rate when the state is low (i.e., $1 - \gamma_{it,0}$ when $\omega_t = 0$). A natural question to ask given this measure of performance is how this quality is affected by the deliberation process in place at the FOMC? Then, I ask whether the performance of the FOMC decision-making process would have been different if deliberation was modified according to members' characteristics. In particular, I modify the speaking order according to members' biases, expertise and career experience. The counterfactual speaking orders are given as follows: from the least to the most biased member and viceversa, from the least to the most expert member and viceversa, and by their experience as central bankers, where members with a longer career within the ranks of the Federal Reserve speak first.[26]

The *ex-ante* probability that the chairman proposes a correct policy directive can be expressed as $\rho_t \gamma_{Ct,1}(s^*_{Ct}) + (1 - \rho_t)(1 - \gamma_{Ct,0}(s^*_{Ct}))$. This measure can be computed under any committee composition and history of recommendations observed by the chairman, $\boldsymbol{x}_{Ct} = (r_{1t}, \ldots, r_{Nt})$.

Notice that for each of the counterfactual scenarios under consideration, the computation of the probability of the correct decision needs to account for changes, not only in the value of social learning from previous recommendations, but also in the pivotality considerations of FOMC members, as this is a function of the order of speech and of the characteristics of subsequent speakers. Therefore, given the estimated coefficients recovered from the covariate specification in (7), I recalculate the pivotality effects at the counterfactual value of each covariate in the pivotality specification.

Figure 15 summarizes the results from these counterfactual simulations. The evidence from this figure indicates that, compared to the case where the Chairman takes the decision in isolation (No Deliberation), the observed order of speech at FOMC meetings increases the probability of a correct policy decision by 6%. Moreover, from the counterfactual rankings considered, ordering members by either their ideological biases (from most neutral to most biased), their expertise (from most expert to least expert), or their experience as central bankers (from most experienced to least experienced), improves the decision-making quality of the FOMC by 9% on average, with respect to the case of no deliberation.

The small gains in precision from the best counterfactual rankings (i.e, by Fed experience and least biased member first) with respect to the observed speaking order, can be explained by the fact that the actual order of speech that is observed in the data partially incorporates the information contained in some of these counterfactual scenarios. Table 2 estimates the corre-

[26] Members with no experience at the Fed are ranked alphabetically in the counterfactual simulation.

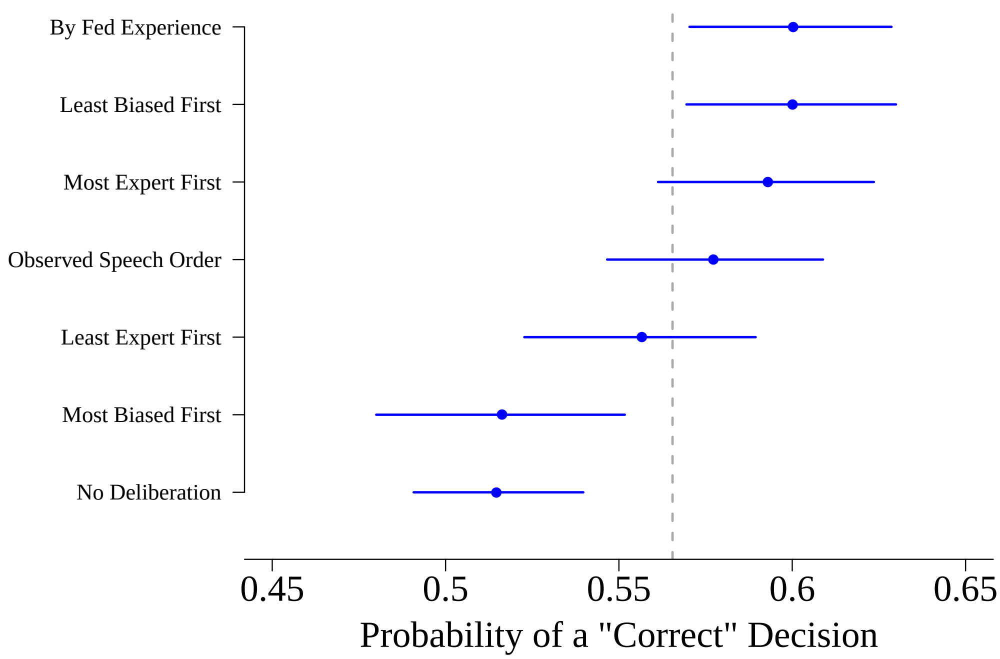

Figure 15: Observed *vs* Counterfactual Committees. This igure shows the estimated posterior probability that the chairman proposes a policy directive consistent with the true state of the economy for different speaking orderings of committee members. Solid points denote the median of the posterior distribution and lines depict the interquartile range of the posterior distribution. The scenario of *No Deliberation* corresponds to the case where the chairman only takes into account his private information. The scenarios *Most (Least) Biased First* ranks committee members in descending (ascending) order given by $|\pi_i - 0.5|$. The scenario *Most (Least) Expert First* ranks committee members by expertise in ascending (descending) order given by σ_i. The scenario *By Fed Experience* rank members according to the proportion of their career spent within the ranks of the Federal Reserve in descending order.

	Dependent variable:		
	Order of Speech		
	(1)	(2)	(3)
Ideological Bias	1.661^{***}	1.511^{***}	1.568^{***}
	(0.451)	(0.452)	(0.454)
Expertise	-0.031	-0.077	-0.071
	(0.138)	(0.139)	(0.139)
Experience at the Fed		-0.301^{***}	-0.296^{***}
		(0.071)	(0.071)
Rookie FOMC Member			0.390^{***}
			(0.071)
Observations	3,621	3,621	3,621

*p<0.1; **p<0.05; ***p<0.01
Note: Standard errors in parentheses

Table 2: Order of Speech Correlates. The dependent variable is a multinomial indicator that takes the value of 1 if the speaking position is in the first third of the policy go-around, takes the value of 2 if the speaking position is in the second third of the policy go-around and takes the value of 3 if the speaking position is in the last third of the policy go-around. Ideological Bias is given by $|\pi_i - 0.5|$. Expertise is given by the value of σ_i. Experience at the Fed is the fraction of a member's career before joining the FOMC spent within the ranks of the Federal Reserve.

lation between members' characteristics-such as bias magnitude, expertise, career experience within the Fed, and experience as FOMC members- and the observed order of speech. This order is categorized into *Early*, *Middle*, and *Late* speaking positions whenever members speak during the first, second, and third portion of the policy go-around, respectively. As can be seen from this Table, it is the case that more neutral members systematically speak early compared to more biased FOMC members. In particular, reducing the bias from 0.19 to 0.02 (i.e., sample interquartile range) increases the probability of speaking early in around 5%. In terms of members' experience as central bankers, it is the case that going from a member with no past central bank experience to one with a entire career within the Fed is associated with an increase of 6% in the probability of speaking early. In the same sense, FOMC members with longer tenure tend to speak 8% more often in early positions than rookie members.

6.4 Model Fit Comparison

The relevance of incorporating the deliberation process in explaining monetary policy-making can be assessed not only by quantifying the value of information transmitted *via* deliberation, but also on whether it improves our understanding of the actual pattern of behavior of committee members better than competing behavioral models available in the literature where deliberation is ignored. For this purpose, I evaluate the explanatory power of the *sequential deliberation model* based on goodness-of-fit metrics that quantify the extent of improvement that

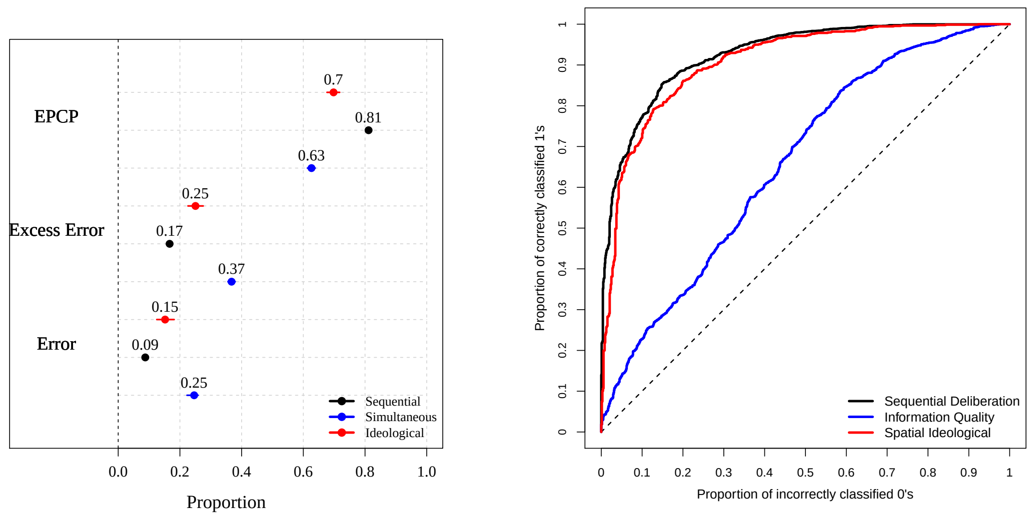

Goodness-of-Fit Statistics

ROC Curves

Figure 16: Fit Measures Across Models. The left panel presents goodness-of-fit statistics for the *sequential deliberation*, *spatial ideological* and *simultaneous* models. I estimate Bayesian versions of the percentage of Error (Error), the Excess Error Rate (Excess Error) and the expected proportion of correctly predicted recommendations (EPC), averaged by both committee member and meeting. The right panel shows a model comparison based on ROC curves, where the 45-degree line corresponds to a random prediction model.

incorporating deliberation provides in explaining observed behavior heterogeneity. For comparison purposes, I use both the *spatial ideological model* and the *simultaneous model* introduced above, which have been previously employed for explaining voting behavior in general and in the FOMC, in particular.[27] The first indicator I use is the percent of (in)correctly classified recommendations (*Error*). The second indicator is the excess error rate (*Excess Error*) proposed by Bafumi et al. [2005], defined as the proportion of error beyond what would be expected, given the model's predicted values. The third indicator is the expected percent of correctly predicted recommendations (*EPCP*), which was proposed by Herron [1999] to alleviate the coarse classification rule in fitted probabilities of binary outcomes, which can over-estimate the true fit of the model.[28]

The left panel of Figure 16 presents a summary of the posterior distribution of each of the three goodness-of-fit measures. The fit of the *sequential deliberation model* is significantly better than any of the two other models in explaining the observed patterns of recommendations, irrespective of the performance metric used. The differences in explanatory power are substantial. The *sequential deliberation model* is able to correctly predict 92% of the individual recommendations in the sample, compared to 76% for the *simultaneous model* and 86% for the *spatial ideological model*. Also, using the expected proportion of correctly predicted recommendations (*EPCP*) as metric, we have that the *sequential deliberation model* correctly predicts 82% of

[27]Details and additional results of the estimation of these models can be found in Appendices A and B.

[28]Under a Bayesian framework, the goodness-of-fit indicators are a function of the model parameters and as such, inherit the uncertainty coming from random sampling, which allows me to provide credibility intervals to these performance measures.

recommendations, whereas the *spatial ideological* and the *simultaneous* models correctly predict 73% and 63% of recommendations, respectively.

The absolute excess error rates are also considerably lower under the *sequential deliberation model* than under the other two behavioral models. In particular, they are around 7% and 20% lower than those for the *spatial ideological* and the *simultaneous* models, respectively.

Another way to compare the performance across models is by plotting their receiver operating characteristic (*ROC*) curves, which are a graphical summary of the correctly classified recommendation rates against the incorrectly classified ones, for different cutoffs c for which $r_{it} = 1$ if $\hat{Pr}(r_{it} = 1) > c$. As can be seen in the right panel of Figure 16, the curve for the *sequential deliberation model* dominates the other two curves for any given cutoff c. The area under the ROC curve can also be used to assess the accuracy of each model. In this respect, the *sequential deliberation model* dominates as well the other two models with an area of 93%, *versus* 90% and 66% for the *ideological* and *simultaneous* models, respectively.

The comparison across models presented above focuses on in-sample fit, for which I contrast observed *versus* classified policy recommendations using the entire data to estimate the structural parameters. However, to assess the predictive accuracy across competing models it is necessary to estimate out-of-sample predictions using within-sample fits. This exercise of predictive accuracy is infeasible for the *spatial ideological model* because an estimate of the location of policy alternatives is needed in order to fit it. Nevertheless, I can provide out-of-sample prediction accuracy for both the *simultaneous* and *sequential deliberation* models. In this respect, Figure 25 in Appendix E shows the results from this exercise *via* leave-one-out cross-validation (LOO), which compares observed recommendations across meetings and members with respect to predicted recommendations based on a training sample that excludes the data from member i at meeting t at every iteration. [29] The results from this exercise show that, similar to the in-sample results, out-of-sample predictive rates are 92% and 76% for the *sequential deliberation* and *simultaneous* models, respectively.

In conclusion, the *sequential deliberation model* fits the observed patterns of FOMC recommendations really well both in- and out-of-sample. The accuracy of this model is substantially better than alternative frameworks that ignore learning associated with the structure of debate.

7 Conclusion

Deliberation is a fundamental component of collective decision-making. Policymakers invest a huge portion of their time and effort expressing their own views and listening to others' arguments regarding the appropriate policy that should be implemented. The relevance and potential consequences of deliberation on collective choices have been explored in previous

[29] To avoid computing an exact LOO, which would require re-fitting the model a total of 3490 times (i.e., the total number of observations in the sample), I estimate an approximate LOO using Pareto smoothed importance sampling (PSIS) (Vehtari, Gelman and Gabry [2015]). This approach provides a computationally feasible and reliable estimate of LOO by resampling the joint posterior density with importance weights that are smoothed with a Pareto distribution to minimize their instability.

theoretical and empirical work. However, less is known regarding the particular mechanisms that affect the behavior of policymakers throughout the deliberation process. I quantify the role of social learning as a fundamental mechanism of deliberation in policy-relevant institutions. In particular, I measure the influence that individual participants exert on others throughout the deliberation process. To do this, I estimate an empirical model of policy-making that incorporates the role of learning by exploiting the sequential nature of deliberation. This approach allows me to estimate changes in the behavior of committee members as they listen other members advocating for policies under different speaking orders. Moreover, this approach provides, for any given committee composition, the optimal order of speech that maximizes the quality of information transmission.

Explaining the patterns of recommendations from deliberation is particularly important in policy-making institutions where voting records and implemented policies are not informative of the underlying heterogeneity in members' behavior. This is the case of the FOMC, where the policy proposal that is put to a vote reflects the policy recommendations that members provide at the deliberation stage. The results of the empirical model using deliberation records of FOMC meetings change the common characterization of this committee in terms of ideological differences and, instead, emphasize the role of information acquisition as a key determinant of members' heterogeneity. Second, it quantifies the value of deliberation in terms of the information it provides to committee members *vis-à-vis* their own private information. Third, it accounts for the observed pattern of behavior better than alternative explanations.

The empirical results presented in this article, by quantifying social learning effects from sequential deliberation, should inform future research on the relevance of learning as an information-transmission mechanism behind real-world deliberation. This empirical model can be used to explain the behavior of members in other deliberative policy-making bodies such as legislative committees, courts, and international organizations, where members are asked to speak in order to the issue in turn.

This analysis can be extended in several directions. One avenue of further research would be to incorporate reputational concerns into the current framework of sequential deliberation, where individuals would care not only about matching their actions to the state of the world, but also about being considered well informed. With this additional dimension I would be able to incorporate a dynamic component to the deliberation process as well as additional counterfactual exercises related to changes in the publicity of debate and transparency of information that have drawn attention in both theoretical and empirical literature (Meade and Stasavage [2008]; Ottaviani and Sørensen [2001], Visser and Swank [2007]), but that have not focused on the learning mechanism embedded in the deliberation process.

References

Adolph, Christopher. 2013. *Bankers, bureaucrats, and central bank politics: The myth of neutrality*. Cambridge University Press.

Austen-Smith, David and Jeffrey S Banks. 1996. "Information aggregation, rationality, and the Condorcet jury theorem." *American political science review* 90(01):34–45.

Austen-Smith, David and Timothy Feddersen. 2005. *Deliberation and voting rules*. Springer.

Austen-Smith, David and Timothy J Feddersen. 2006. "Deliberation, preference uncertainty, and voting rules." *American Political Science Review* 100(02):209–217.

Bafumi, Joseph, Andrew Gelman, David K Park and Noah Kaplan. 2005. "Practical issues in implementing and understanding Bayesian ideal point estimation." *Political Analysis* 13(2):171–187.

Banerjee, Abhijit V. 1992. "A simple model of herd behavior." *The Quarterly Journal of Economics* pp. 797–817.

Bikhchandani, Sushil, David Hirshleifer and Ivo Welch. 1992. "A theory of fads, fashion, custom, and cultural change as informational cascades." *Journal of political Economy* pp. 992–1026.

Bikhchandani, Sushil, David Hirshleifer and Ivo Welch. 1998. "Learning from the behavior of others: Conformity, fads, and informational cascades." *The Journal of Economic Perspectives* pp. 151–170.

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

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

Chang, Kelly H. 2003. *Appointing central bankers: The politics of monetary policy in the United States and the European monetary union*. Cambridge University Press.

Chappell, Henry W, Rob Roy McGregor and Todd A Vermilyea. 2012. "Deliberation and learning in monetary policy committees." *Economic Inquiry* 50(3):839–847.

Chappell, H.W., R.R. McGregor and T. Vermilyea. 2005. *Committee decisions on monetary policy: Evidence from historical records of the Federal Open Market Committee*. The MIT Press.

Clinton, Joshua, Simon Jackman and Douglas Rivers. 2004. "The statistical analysis of roll call data." *American Political Science Review* 98(02):355–370.

Coughlan, P.J. 2000. "In defense of unanimous jury verdicts: Mistrials, communication, and strategic voting." *American Political Science Review* pp. 375–393.

Dickson, Eric S, Catherine Hafer and Dimitri Landa. 2008. "Cognition and strategy: a deliberation experiment." *The Journal of Politics* 70(04):974–989.

Dickson, Eric S, Catherine Hafer and Dimitri Landa. 2015. "Learning from Debate: Institutions and Information." *Political Science Research and Methods* pp. 1–24.

Doraszelski, Ulrich, Dino Gerardi and Francesco Squintani. 2003. "Communication and voting with double-sided information." *Contributions in Theoretical Economics* 3(1).

Duggan, John and César Martinelli. 2001. "A Bayesian model of voting in juries." *Games and Economic Behavior* 37(2):259–294.

Gelman, Andrew and Donald B Rubin. 1992. "Inference from iterative simulation using multiple sequences." *Statistical science* pp. 457–472.

Gerardi, Dino and Leeat Yariv. 2007. "Deliberative voting." *Journal of Economic Theory* 134(1):317–338.

Gerlach-Kristen, Petra. 2006. "Monetary policy committees and interest rate setting." *European Economic Review* 50(2):487–507.

Goeree, Jacob K and Leeat Yariv. 2011. "An experimental study of collective deliberation." *Econometrica* 79(3):893–921.

Gutmann, Amy and Dennis Frank Thompson. 1996. *Democracy and disagreement*. Harvard University Press.

Habermas, Jurgen. 1996. *Between Facts and Norms: Contributions to a Discourse Theory of Law and Democracy*. Cambridge, MA: MIT Press.

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

Herron, Michael C. 1999. "Postestimation uncertainty in limited dependent variable models." *Political Analysis* 8(1):83–98.

Homan, Matthew D and Andrew Gelman. 2014. "The no-U-turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo." *The Journal of Machine Learning Research* 15(1):1593–1623.

Humphreys, Macartan, William A Masters and Martin E Sandbu. 2006. "The role of leaders in democratic deliberations: results from a field experiment in São Tomé and Príncipe." *World Politics* 58(04):583–622.

Iaryczower, Matias and Matthew Shum. 2012. "The value of information in the court: Get it right, keep it tight." *The American Economic Review* 102(1):202–237.

Iaryczower, Matias, Xiaoxia Shi and Matthew Shum. 2014. Can Words Get in the Way? The Effect of Deliberation in Collective Decision-Making. Technical report mimeo, Princeton University.

Jackman, Simon. 2000. "Estimation and inference are missing data problems: Unifying social science statistics via Bayesian simulation." *Political Analysis* 8(4):307–332.

Karpowitz, Christopher F and Tali Mendelberg. 2011. "An experimental approach to citizen deliberation." *Cambridge handbook of experimental political science* pp. 258–272.

Knight, Brian and Nathan Schiff. 2010. "Momentum and Social Learning in Presidential Primaries." *Journal of Political Economy* 118(6):1110–1150.

Landa, Dimitri and Adam Meirowitz. 2009. "Game theory, information, and deliberative democracy." *American Journal of Political Science* 53(2):427–444.

Macedo, Stephen. 2010. "Why Public Reason? Citizens Reasons and the Constitution of the Public Sphere." *Unpublished Manuscript* .

Meade, E.E. and D. Stasavage. 2008. "Publicity of Debate and the Incentive to Dissent: Evidence from the US Federal Reserve*." *The Economic Journal* 118(528):695–717.

Melzer, T.C. 1994. "A Low Inflation Policy is a Pro-Growth Policy." *The Regional Economist* 8(July).

Meyer, Laurence. 1998. "Come with Me to the FOMC." *Federal Reserve Bank of Minneapolis The Region* 12:6–15.

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

Poole, Keith T and Howard Rosenthal. 2000. *Congress: A political-economic history of roll call voting*. Oxford University Press.

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

Smith, Lones and Peter Sørensen. 2000. "Pathological outcomes of observational learning." *Econometrica* pp. 371–398.

Swank, O. and B. Visser. 2007. "On committees of experts." *Quarterly Journal of Economics* 122(1):337–372.

Team, Stan Development. 2015. "Stan: A C++ Library for Probability and Sampling, Version 2.8.0.".
URL: *http://mc-stan.org/*

Van Weelden, Richard et al. 2008. "Deliberation rules and voting." *Quarterly Journal of Political Science* 3(1):83–88.

Vehtari, Aki, Andrew Gelman and Jonah Gabry. 2015. "Efficient implementation of leave-one-out cross-validation and WAIC for evaluating fitted Bayesian models." *arXiv preprint arXiv:1507.04544* .

Visser, Bauke and Otto H Swank. 2007. "On Committees of Experts." *The Quarterly journal of Economics* 122(1):337–372.

Yellen, Janet. 1995. “Interview with Janet Yellen.”.

Yohe, William P. 1966. “A Study of Federal Open Market Committee Voting, 1955-64.” *Southern Economic Journal* pp. 396–405.

A Spatial Ideological Model

Under the spatial model, committee members are perfectly informed about the characteristics of the alternatives under consideration and have Euclidean preferences that can be represented in a one-dimensional space by points on the real line. Each committee member has an ideal point or prefered outcome $z_i \in \mathbb{R}$ and, for any two available policy rates in meeting t, $d_t^0 = 0$ and $d_t^1 = 1$, she prefers $d_t^1 = 1$ if and only if $d_t^1 = 1$ is closer to z_i than $d_t^0 = 0$. Conditional on this behavioral assumption, one only needs the ideal policy of committee members and the ideological location of the policy choice under consideration to confidently predict the observed pattern of policy recommendations across members and meetings.

I estimate the spatial model following a standard operationalization in the literature that assumes commitee members have quadratic utility functions over the policy space with an additive idyosincratic shock, $U(d) = -(z_i - r)^2 + \eta_{ir}$ (Clinton, Jackman and Rivers [2004]). Given this functional form, a committee member recommends the high policy rate, $r_{it} = 1$ whenever $U(1) > U(0)$ and recommends the lower rate, $r_{it} = 0$, otherwise. Assuming that the errors η_{i0} and η_{i1} are jointly normal, with $\eta_{i1} - \eta_{i0} \sim N(0, \tau_t^2)$, we can write $Pr(r_{it} = 1) = Pr(U(1) > U(0)) = \Phi(\lambda_t [z_i - \kappa_t])$, where $\lambda_t \equiv \frac{r_t^1 - r_t^0}{\tau_t}$ and $\kappa_t \equiv \frac{r_t^1 + r_t^0}{2}$.

I estimate the structural parameters of interest, z_i, κ_t, and λ_t, for $i = 1 \ldots, 57$ and $t = 1, \ldots 265$, by fitting a Bayesian version of a multilevel ideal point model (Bafumi et al. [2005]). In particular I assume $z_i \sim N(0,1)$, $\kappa_t \sim N(\boldsymbol{X_t \beta_{ideol}}', \sigma_\kappa^2)$, $\lambda_t \sim LN(0, \sigma_\lambda^2)$, $\sigma_\kappa^2, \sigma_\lambda^2 \sim Cauchy(0,1)$ for $i = 1, \ldots, 57$, and $t = 1, \ldots, 265$.

Notice that all the model parameters are globally identified, as we are constraining the ideal points (z_i) to have mean zero and standard deviation one. In addition, I solve for the reflection invariance problem that plagues ideal point models by constraining the average gap parameter λ_t to be positive, which is a reasonable assumption for the FOMC decision-making process, because it is clear that a positive recommendation corresponds to higher interest rates. Finally notice that the midpoint of the policy location κ_t id a function of covariates, as in the case of the estimation of the prior ρ_t in the *sequential deliberation model*. Therefore, the number of parameter estimates of the *spatial ideological model*

I approximated the posterior distribution of the parameters of interest,

$$\{z_i\}_{i=1}^{57}, \{\kappa_t, \lambda_t\}_{t=1}^{265}, \sigma_\kappa^2, \; \sigma_\lambda^2,$$

with an application of Markov Chain Monte Carlo (MCMC) *via* the Hamiltonian Monte Carlo method as in Homan and Gelman [2014]. I obtained posterior samples of the parameters from their posterior marginal density at each iteration $m = 1, \ldots, M$. I ran three parallel chains with dispersed initial values for 10,000 iterations with an initial warm-up period of 5,000 iterations. I assessed convergence for each parameter based on the potential scale reduction factor, $\hat{R}$ (Gelman and Rubin [1992]).

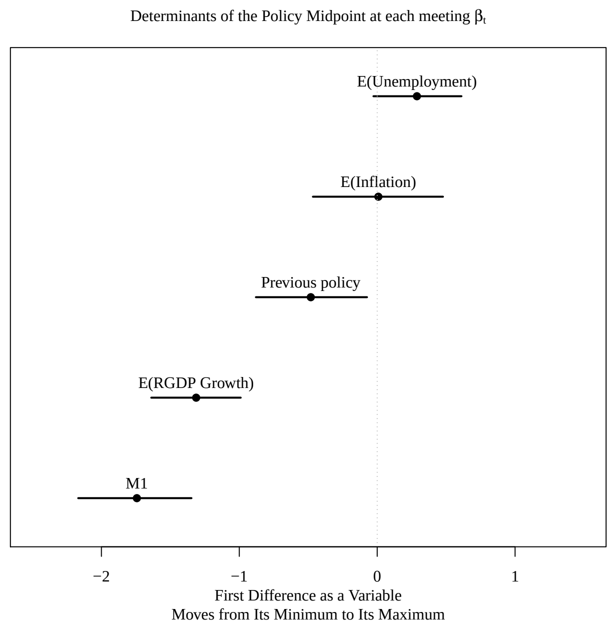

Figure 17: Determinants of the Midpoint parameter (β_t) (*Ideological Model*). This figure provides the effect of increasing each of the displayed covariates on the midpoint β_t. Solid circles give the posterior median, with horizontal solid lines corresponding to the interquartile range of the posterior distribution and dashed lines corresponding to 95% posterior credible intervals. The counterfactual increase in the covariate of interest is a change from the 10th to the 90th percentile in the sample. For each estimate, all other covariates are set at their median sample values.

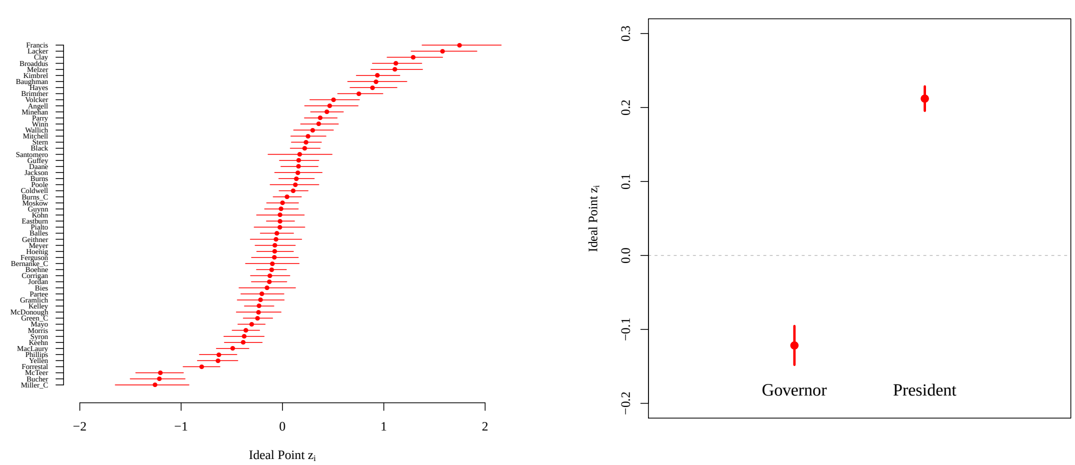

Figure 18: Ideal Point Estimates. The left panel of the figure provides posterior summaries of the ideal points recovered from the *spatial ideological model* for each FOMC committee member, z_i, during the periods 1970-1979 and 1987-2008. Solid circles give the posterior median, with horizontal lines corresponding to the interquartile range of the posterior distribution. Committee members who participated in more than 30 meetings are included.The right panel figure provides posterior summaries of the ideal points of FOMC committee members, z_i, aggregated by appointment, where a member is either a president of a regional Federal Reserve Bank or a member of the Board of Governors.

B Simultaneous Model

The main difference of the *simultaneous model* with respect to the *sequential deliberation model* comes in the optimal cutoff s^*_{it}, which in the latter case does not consider the value of sequential deliberation:

$$s^*(\pi_i, \sigma_i, \rho_t) \equiv \frac{1}{2} + \sigma_i^2 \left[\log\left(\frac{1-\pi_i}{\pi_i}\right) + \log\left(\frac{1-\rho_t}{\rho_t}\right) \right]. \tag{9}$$

I estimate the model following the same algorithm as in the case of the *sequential deliberation model*.

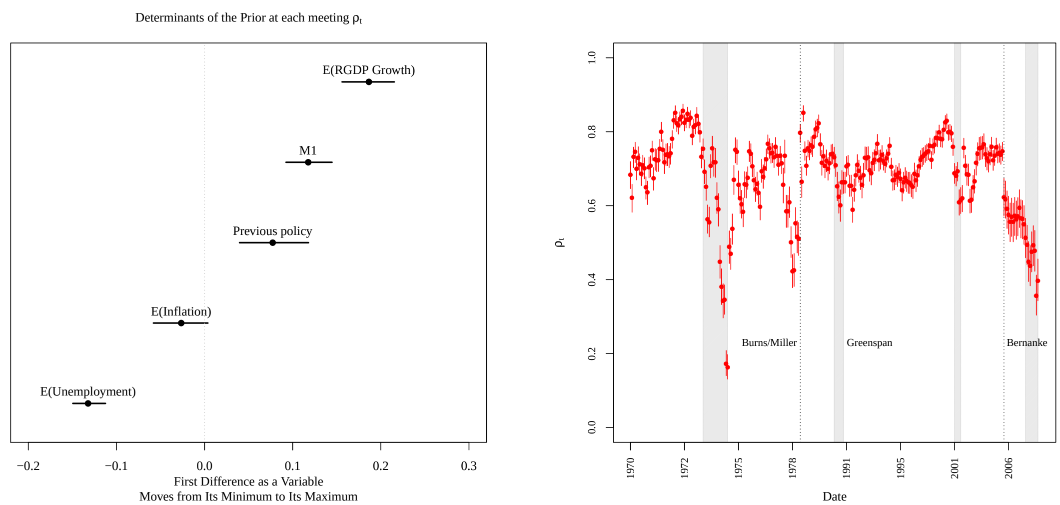

Figure 19: Determinants and Evolution of the Prior (ρ_t) on the unobserved state of inflation $\omega_t = 1$ at each meeting t (*Simultaneous Model*). The left panel of the figure provides the effect of increasing each of the displayed covariates on the common prior $\rho_t \equiv Pr(\omega_t = 1)$. Solid circles give the posterior median, with horizontal solid lines corresponding to the interquartile range of the posterior distribution and dashed lines corresponding to 95% posterior credible intervals. The counterfactual increase in the covariate of interest is a change from the 10th to the 90th percentile in the sample. For each estimate, all other covariates are set at their median sample values. The right panel of the figure provides the estimated value of the common prior $\rho_t \equiv Pr(\omega_t = 1)$ across meetings. Solid circles give the posterior median, with horizontal solid lines corresponding to the interquartile range of the posterior distribution. Gray areas correspond to economic recessions in the sample measured as the period from peak to through of a business cycle according to the National Bureau of Economic Research.

As the equilibrium behavior of committee members in the *simultaneous model* is driven by common information, ideological biases, and private signals, assessing the value that the latter have in members' pattern of recommendations, would imply isolating its contribution from that of the rest of the parameters. For this purpose, Iaryczower and Shum [2012] quantified a measure of the value of private information by computing the probability that member i gives a different policy recommendation from the one she would have given in a conterfactual scenario, had she only weighted the common prior against her signal. This "FLEX" score for member i at meeting t can be written as

$$FLEX_{it} = \begin{cases} \rho_t(1-\gamma_{it,1}) + (1-\rho_t)(1-\gamma_{it,0}) & \text{if } \rho_t > 1-\pi_i \\ \rho_t\gamma_{it,1} + (1-\rho_t)\gamma_{it,0} & \text{if } \rho_t \leq 1-\pi_i \end{cases}$$

I compute the posterior median distribution of $FLEX$ scores for each member and meeting of the FOMC, and present a summary of the results in Figure 22.

In terms of the variation across members, the left panel of Figure 22 shows that, on average, FOMC members have tended to follow their initial leanings when giving a policy recommendation, motivated solely by their preference biases and the common prior they observe. This is the result of estimating an average "FLEX" across FOMC members around 0.3, which implies that an FOMC member would have reverted his recommendation 30% of the time due to the information contained in their private information. Nonetheless, the dispersion on the value of information across members is sizable. On the one hand, we can see district president Francis, an extreme "hawk" ($\pi_i = 0.8$) with a very low ability ($\sigma_i = 1.2$), who obtains almost no value out of his private signal ($FLEX_i = 0.02$). On the other hand, district president McTeer, who is estimated as the second most "dovish" member in the committee ($\pi_i = 0.15$), with a medium level of expertise ($\sigma_i = 0.40$), has a median "FLEX" score of 0.58, which implies that the probability of giving a different recommendation than the one he would have given in the abscence of private information is about 58%.

In the right panel of Figure 22, I track the evolution of the median "FLEX" score over time for the period under study. From this plot, we can see one of the main substantive findings that come out of the quality of information model regarding behavior within the FOMC, namely, that at least since the Volcker Revolution from 1979, the FOMC has became increasingly more responsive to their information and at the same time, has

placed less emphasis on their ideological leanings.

The evolution of the decision-making towards a more informative process has been substantial, as it can be assessed from a comparison of "FLEX" scores during the Burns and Miller years with respect to last available information under Bernanke as Chairman. On average, the value of information more than doubled in almost 30 years of monetary policy making from a "FLEX" score from around 0.2 to around 0.45.

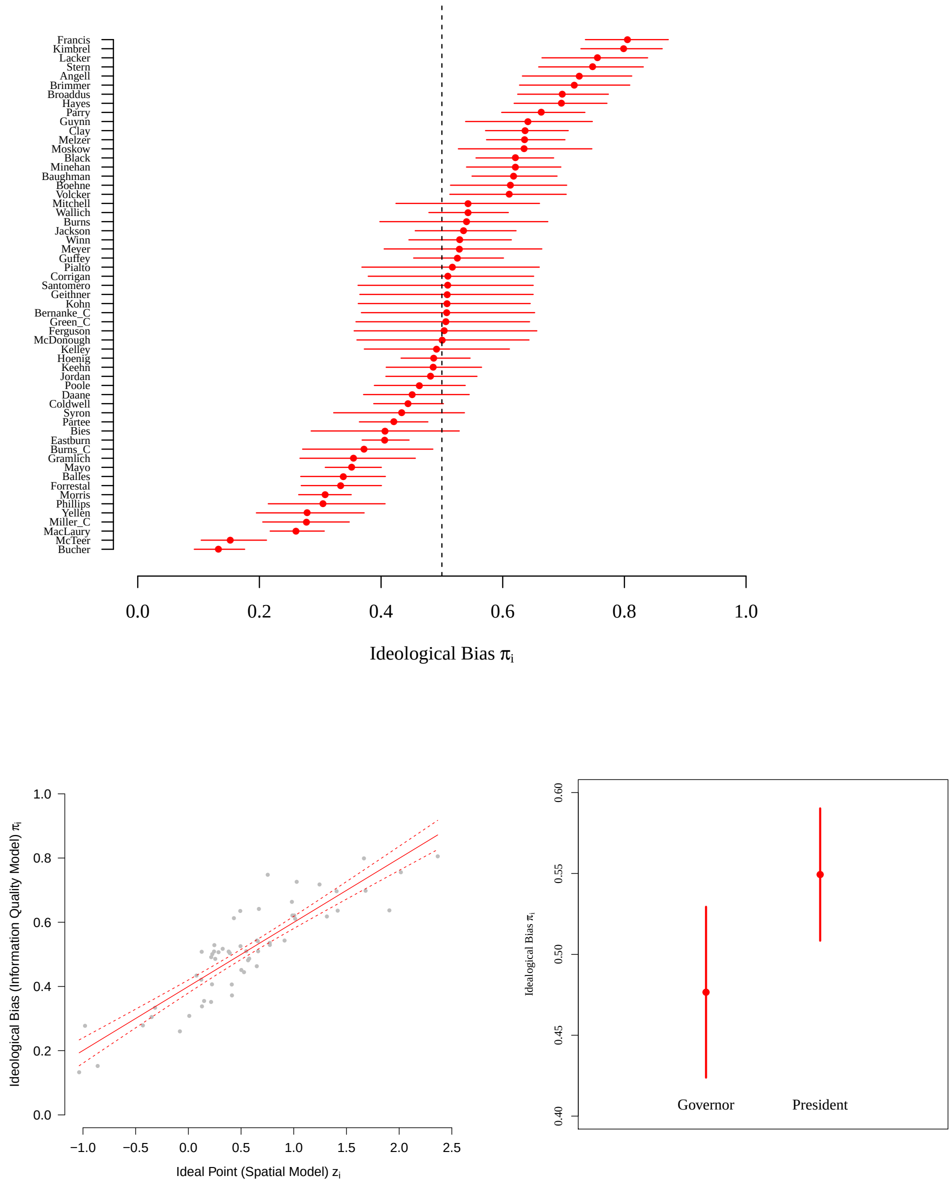

Figure 20: Ideological Estimates, π_i for the *Simultaneous Model*. The left panel of the figure provides posterior summaries of theideological bias, π_i for each FOMC committee during the periods 1970-1979 and 1987-2008. Solid circles give the posterior median, with horizontal lines corresponding to the interquartile range of the posterior distribution. Committee members who participated in more than 30 meetings are included. The middle panel of the figure povides the result of a linear fit along with 90% confidence intervals between ideological biases π_i, as recovered by the *Simultaneous Model* and ideal points, z_i, as recovered by the Spatial model. The right panel of the figure provides posterior summaries of the ideal points of FOMC committee members, z_i, aggregated by appointment, where a member is either a president of a regional Federal Reserve Bank or a member of the Board of Governors.

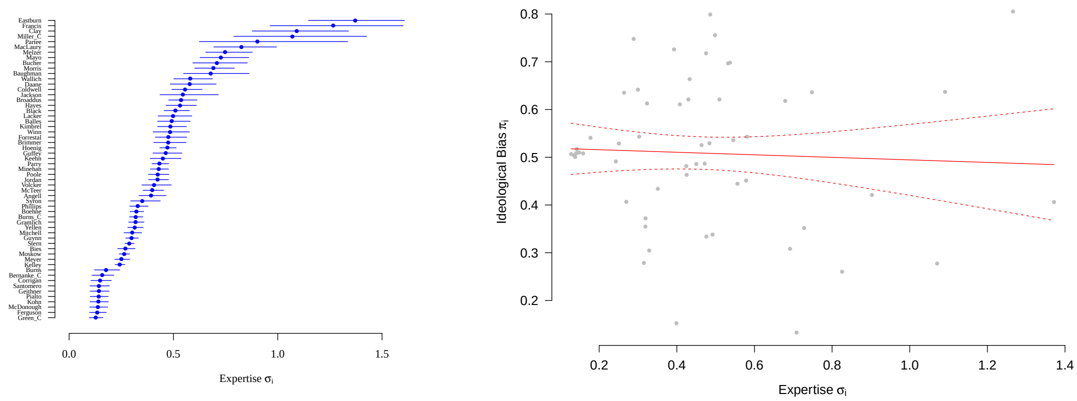

Figure 21: Expertise Estimates (*Simultaneous Model*). The left panel of the figure provides posterior summaries of the measure of expertise of FOMC committee members, σ_i. Solid circles give the posterior median, with horizontal lines corresponding to the interquartile range of the posterior distribution. Committee members who participated in more than 30 meetings are included. The right panel of the figure provides the results of a linear fit, along with 90% confidence intervals between ideological biases π_i and individual expertise, σ_i, as recovered by the *simultaneous model*.

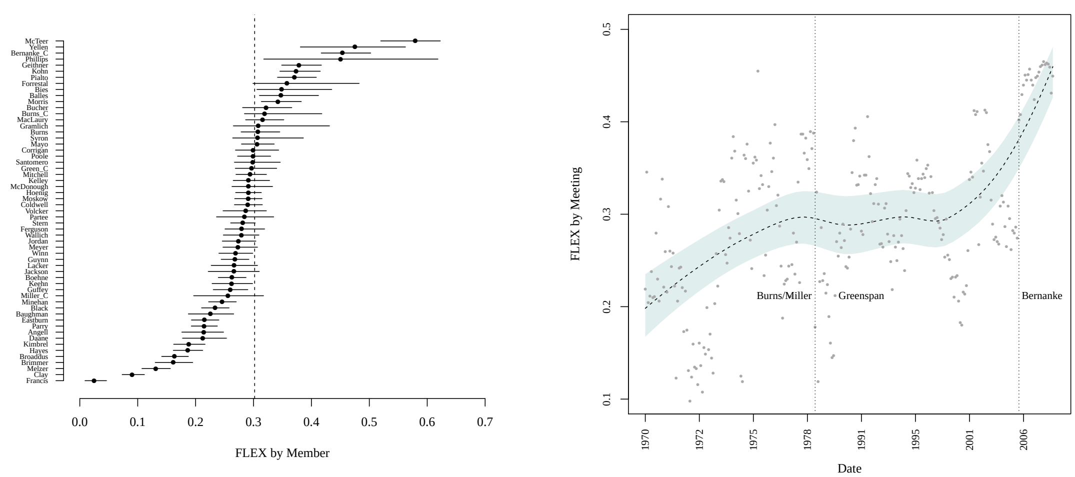

FLEX by Member

FLEX over Time

Figure 22: FOMC's FLEX Scores for the *Simultaneous Model*. The left panel of the figure provides posterior summaries of the median FLEX Score by FOMC Member. The dashed line represents the mean value across FOMC members. The right panel of the figure plots the smoothed (by 2nd degree local polynomial) time trend of the median FOMC by year. The shaded area corresponds to the posterior interquartile range.

C Differences in Bias and Expertise by Members' Characteristics

	Dependent variable:		
	Spatial Ideological	*Simultaneous*	*Sequential Deliberation*
	(1)	(2)	(3)
Fed President	0.412*	0.086*	0.077**
	(0.205)	(0.048)	(0.036)
Democrat Appointment	0.158	0.059	0.048
	(0.199)	(0.046)	(0.035)
Financial Experience	0.032	0.030	−0.002
	(0.425)	(0.099)	(0.075)
Government Experience	0.144	0.114	−0.136
	(1.032)	(0.242)	(0.183)
Treasury Experience	−0.612	−0.146	−0.122
	(0.772)	(0.181)	(0.137)
Central Bank Experience	0.025	0.008	−0.039
	(0.401)	(0.094)	(0.071)
Economics Experience	0.183	−0.0004	0.016
	(0.420)	(0.098)	(0.075)
Constant	−0.341	0.439***	0.424***
	(0.366)	(0.086)	(0.065)
Observations	56	56	56
R^2	0.116	0.126	0.154
Adjusted R^2	−0.013	−0.001	0.030
F Statistic (df = 7; 48)	0.899	0.990	1.244

*p<0.1; **p<0.05; ***p<0.01
Note: Standard errors in parentheses

Table 3: Bias Correlates for Different Behavioral Models

	Dependent variable:	
	Simultaneous	*Sequential Deliberation*
	(1)	(2)
Fed President	0.059	−0.040
	(0.090)	(0.083)
Democrat Appointment	0.047	0.011
	(0.088)	(0.080)
Financial Experience	−0.159	0.038
	(0.187)	(0.171)
Government Experience	−0.798*	−0.073
	(0.455)	(0.416)
Treasury Experience	−0.276	0.241
	(0.340)	(0.311)
Central Bank Experience	−0.163	−0.019
	(0.177)	(0.161)
Economics Experience	−0.370*	0.054
	(0.185)	(0.169)
Constant	0.639***	1.311***
	(0.161)	(0.147)
Observations	56	56
R^2	0.153	0.034
Adjusted R^2	0.029	−0.107
F Statistic (df = 7; 48)	1.234	0.238

*p<0.1; **p<0.05; ***p<0.01
Note: Standard errors in parentheses

Table 4: Expertise Correlates for Different Behavioral Models

D Implied *vs* Exact Pivotality Effects

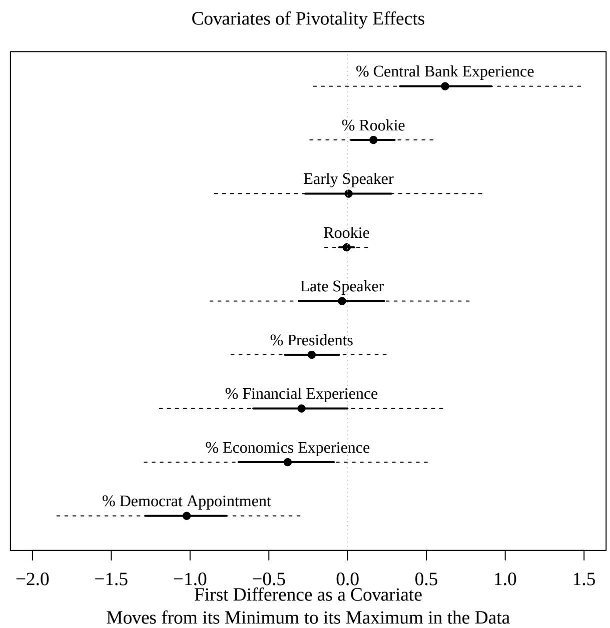

Figure 23: Determinants of Pivotality Effects. This figure provides the effect of increasing each of the displayed covariates on the Pivotality Effect. Solid circles give the posterior median, with horizontal solid lines corresponding to the interquartile range of the posterior distribution and dashed lines corresponding to 95% posterior credible intervals. The counterfactual increase in the covariate of interest is a change from the 10th to the 90th percentile in the sample. For each estimate, all other covariates are set at their median sample values. The pivotality effect is estimated in the model as $\log\left(\frac{Pr[PIV_t^i|\omega_t=0]}{Pr[PIV_t^i|\omega_t=1]}\right) = \alpha_1 Early_{it} + \alpha_2 Late_{it} + \boldsymbol{X}_{it}\boldsymbol{\beta}'$. The matrix $\boldsymbol{X}_{it}$ of member-meeting level predictors includes member i's experience in the form of an indicator variable (*Experience*), that takes the value of one if member i has served in less than 34 meetings and zero otherwise. Here 34 meeting represents the 25th percentile of term length in the sample and allows me to classify members as experienced ($Rookie = 0$) and inexperienced ($Rookie = 1$). I include the fraction of remaining speakers after member i who are inexperienced (% Rookie), the fraction of remaining speakers who are Federal Reserve presidents ($\%Presidents$), the fraction of remaining speakers who are Democrat-appointed governors (% Democrat Appointment). For the remaining speakers after member i at meeting t I also include the average fraction of their past career (i.e., before the FOMC) in private financial institutions (% Financial Experience), as economists (% Economics Experience), within the ranks of the Federal Reserve (% Central Bank Experience). Finally, this specification include random effects for the order of speech of member i (α_0 and α_1) where *Early* (*Late*) takes the value of one (zero) if member i speaks in the first (second) half of the policy go-around and zero otherwise.

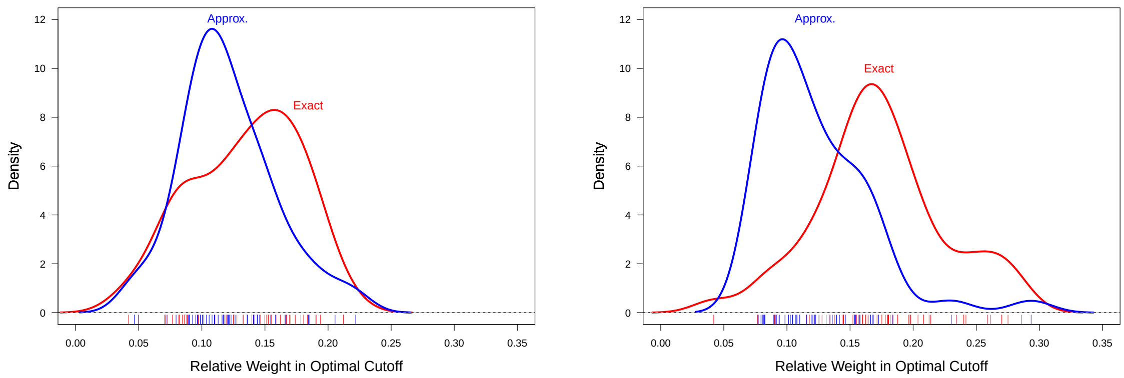

Next-to-Last Speaker

Last Speaker

Figure 24: Implied *vs* Exact Pivotality Effects for Last and Next-to-Last Speakers. The left panel of the figure provides the distributions across committee members of the relative weight of pivotality for the next-to-last spaeker calculated both as a function of covariates (blue line) and computed exactly from the posterior update of the Chairman (red line). The right panel shows the results of the same exercise for the last speaker.

E Goodness-of-Fit with Leave-One-Out Cross Validation

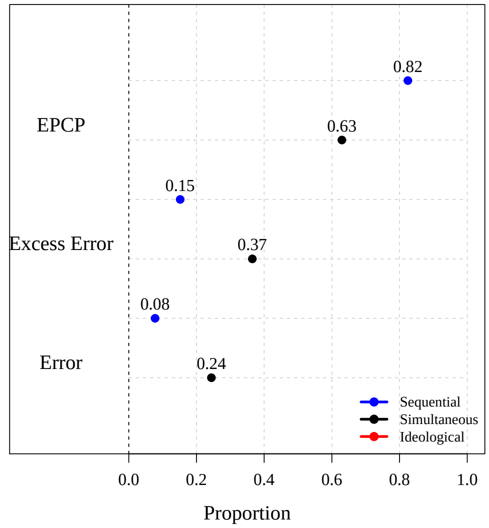

Figure 25: Fit Measures Across Models with LOO Cross Validation. The Figure presents goodness-of-fit statistics for the *sequential deliberation* and *simultaneous* models. We estimate Bayesian versions of the percentage of Error (Error), the Excess Error Rate (Excess Error) and the expected proportion of correctly predicted recommendations (EPC), averaged by both committee member and meeting..

F Career Experience within the FOMC

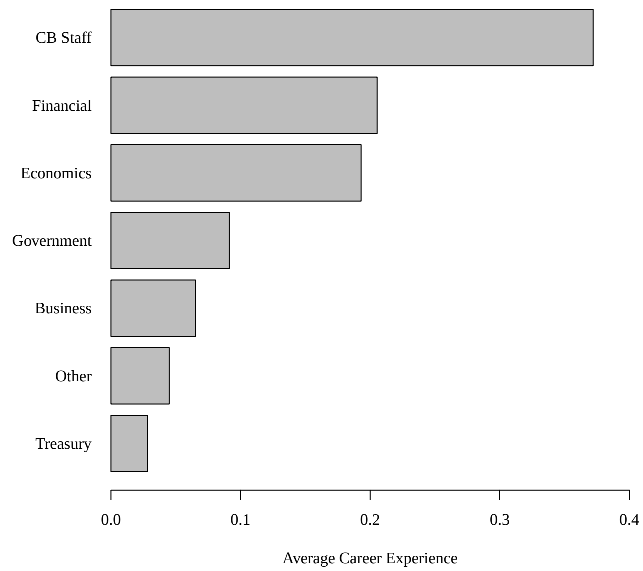

Figure 26: Average Career Experience by Job Category. The Figure shows the mean career experience score for each job category. Career experience is the fraction of a member's career prior to FOMC membership.

G MCMC Statistics

Acceptance Rate	Stepsize	Treedepth	Frogsteps	Divergent Transitions
Min. :0.00	Min. :4.9e-04	Min. : 1	Min. : 1	Min. :0.00
1st Qu.:0.79	1st Qu.:5.0e-02	1st Qu.: 4	1st Qu.: 15	1st Qu.:0.00
Median :0.90	Median :5.7e-02	Median : 5	Median : 31	Median :1.00
Mean :0.83	Mean :6.3e-02	Mean : 5	Mean : 37	Mean :0.74
3rd Qu.:0.96	3rd Qu.:7.1e-02	3rd Qu.: 6	3rd Qu.: 56	3rd Qu.:1.00
Max. :1.00	Max. :1.4e+01	Max. :11	Max. :2047	Max. :1.00

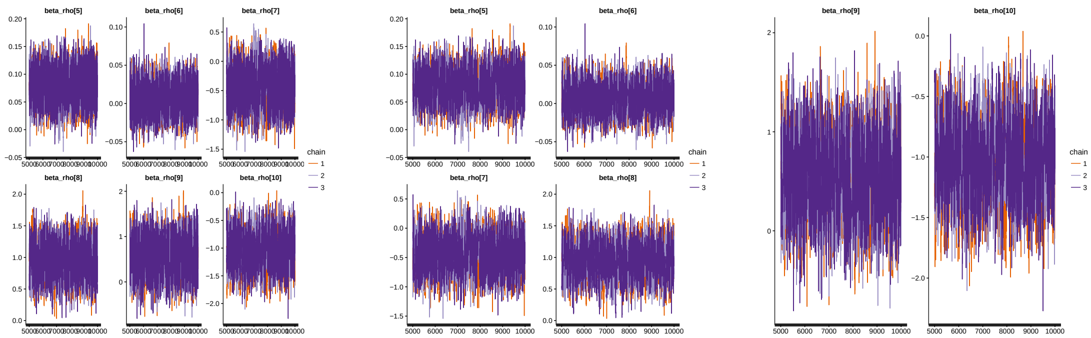

Figure 27: Trace Plots of Determinants of ρ_t. The figure shows the post-burn traceplots for three chains of $10,000$ iterations each with a burn-in of $5,000$ iterations and a thining of 100 iterations.

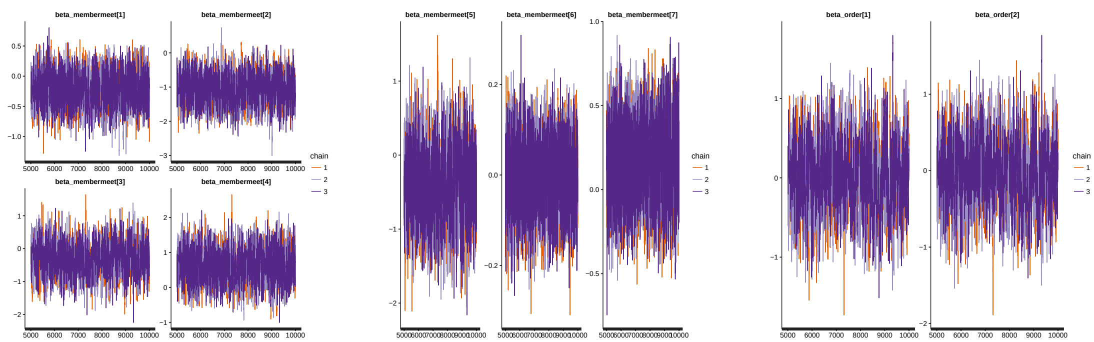

Figure 28: Trace Plots of Determinants of $\log\left(\frac{Pr[PIV_t^i|\omega_t=0]}{Pr[PIV_t^i|\omega_t=1]}\right)$. The figure shows the post-burn traceplots for three chains of $10,000$ iterations each with a burn-in of $5,000$ iterations and a thining of 100 iterations.

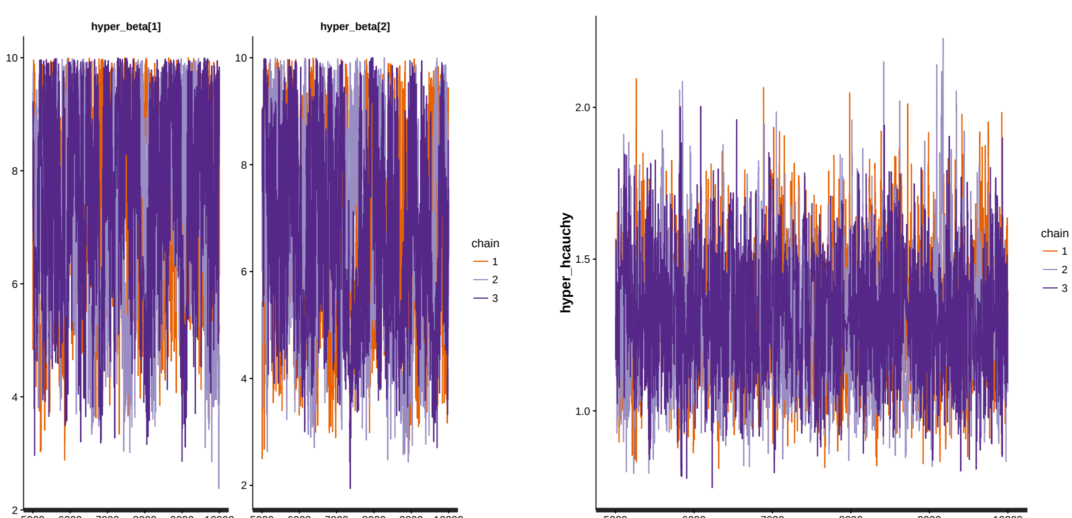

Figure 29: Trace Plots of Hyperparameters. The figure shows the post-burn traceplots for three chains of $10,000$ iterations each with a burn-in of $5,000$ iterations and a thining of 100 iterations.