A Textual Taylor Rule: Estimating Central Bank Preferences Combining Topic and Scaling Methods[*]

NICOLE BAERG & WILL LOWE

Abstract

Scholars often use voting data to estimate central bankers' policy preferences but consensus voting is commonplace. To get around this, we combine topic-based text analysis and scaling methods to generate theoretically motivated comparative measures of central bank preferences on the U.S. Federal Open Market Committee leading up to the financial crisis in a way that does not depend on voting behavior. We apply these measures to a number of applications in the literature. For example, we find that FOMC members that are Federal Reserve Bank Presidents from districts experiencing higher unemployment are also more likely to emphasize unemployment in their speech. We also confirm that committee members on schedule to vote are more likely to express consensus opinion than their off schedule voting counterparts.

Political economists have long studied the role of central banks and central bankers in the governance structure of modern countries. Despite their specialty, central banks share similarities with many other political institutions. For example, in the U.S., a small committee of individuals set interest rate policy according to majority rule. Because central bank committees often take "roll call" votes during committee meetings like in legislatures, researchers commonly apply spatial voting models to help understand central bank decision making. With informative voting, researchers can locate central bankers on a single common dimension, based on the correlations in their voting behavior (see for e.g. Chappell et al., 2000; Chang, 2003; Gerlach-Kristen and Meade, 2010; Hix et al., 2010; Adolph, 2013; Eijffinger et al., 2017; Ainsley, 2017).

Previous research has presented some possible problems with applying spatial voting models to central bank decision making. First, committee members often vote according to consensus and consensus voting seems to describe most central bank policy outputs irrespective of how the committee is organized; this makes computation of members' preferences difficult because of invariant or sparse data (Riboni and Ruge-Murcia, 2010). Second, the status of committee members, such as duration of tenure on the committee, whether the member is appointed by the president or elected by his or her district, and whether a member has dissented before, is shown to predict whether a member will dissent or not (Meade and Sheets, 2005a; Meade, 2005); this makes it hard to distinguish members' preferences from member hierarchy. Worse, any willingness to cast a dissenting vote is made harder by a dominant committee Chair (Crawford and Haller, 1990; Shapiro and Varian, 2013) as well as tendencies towards "group-think" due to members' socialization, such as education and vocation, that may lead members to draw similar conclusions (Johnson, 2016; Kaplan, 2017; Hallerberg and Wehner, 2018). In short, sparse data and differences across individuals in their ability and willingness to dissent makes spatial methods designed for estimating preferences from roll-call votes particularly challenging for studying central banking.

Instead of using voting data to estimate committee members' relative preferences, we use a method that combines topic modeling and positional analysis based on a textual interpretation of a theoretical model called the Taylor rule. We use Federal Open Market Committee (FOMC) meeting transcripts between 2005 and 2008 to estimate committee members' preferences using the relative emphasis on policy topics that U.S. central bankers are mandated to consider: *inflation* and *unemployment/output*.

Our main contribution is to provide a strategy for estimating central bankers' preferences using what policymakers say in meetings rather than how they vote in roll calls. Instead of using real time economic data such as committee members' economic forecasts to infer relative emphasis (El-Shagi and Jung, 2015), we compute relative positions by using how much a member emphasizes inflation relative to output and unemployment in a policy meeting. Methodologically, our approach highlights the value in explicitly including speech data and human judgment to assist in the interpretation of unsupervised learning approaches increasingly common in political science (see also Roberts et al., 2014; Lucas et al., 2015). An unusual aspect of our application is to leverage common understanding of the micro-foundations of FOMC deliberation, enshrined in the Fed's dual mandate, to focus an unsupervised and often hard to interpret machine learning model (Chang et al., 2009).

Our results have several important implications for the politics of committee decision-

making and political economy. First, we show that despite their consensus voting, central bankers express much larger variation in policy positions in speech than they do in votes. This is important as a shift in FOMC preferences has important implications not only for the U.S. economy and its citizens, but, due to the prominence of the U.S. in the global financial system, that can dramatically shape international markets as well. Second, the fact that we find greater variation in speech data than in voting data speaks to the similarity that central bank committees have with other political institutions. Our findings add to mounting evidence in other contexts, such as legislatures and political parties, where political actors similarly express greater diversity in speech opinions than they do in formal voting behavior (Proksch and Slapin, 2012; Schwarz et al., 2017). Third, market actors increasingly use speech data to forecast future FOMC behavior. CNN, for example, recently reported that *Minutes from the Federal Open Market Committee meeting in September show that a December interest rate hike is likely*, providing qualitative analysis of committee members' inflation preferences derived from committee meeting minutes. As markets increasingly turn to what members say in addition to how they vote, the need to understand similarities and differences in members' revealed preferences during policy deliberations is all the more pressing.[1]

1   Central Bank Preferences

The individuals who sit on the FOMC are mandated by the U.S. Congress to fulfill three goals: maximize employment, stabilize prices, and keep long-term interest rates steady. Although all members on the FOMC share the same mandate, committee members have different preferences. A useful ranking of central bank preferences is based on their inflation preferences – or how sensitive policymakers are to increases in inflation risks. Borrowing terms used to describe individuals' appetites for war, those that worry relatively more about inflation risks are called *Inflation Hawks* whereas those that worry relatively less about inflation risks are called *Inflation Doves*. Ordering individual central bank committee members according to their preferences along a single common dimension, "how much do I care about inflation risks relative to output and unemployment risks," is important as it allows researchers to compare preferences across actors and over time and to test theories about committee decision-making. Indeed members themselves talk about inflation this way: in January 31 2007, Bank President Geithner (New York) reported his inflation concern strictly using such terms, "We see somewhat less downside risk to growth and somewhat less upside risk to inflation than we did, but the overall balance of risks in our view is still weighted toward inflation." The fact that committee members have such a mandate and talk in these terms provides an unusual theoretical consensus on a low dimensional structure to understand members' preferences over a wide variety of apparently disparate policy topics.[2]

---

[1] Jeff Cox, "Fed minutes: December rate hike all but certain despite low inflation," *CNBC*, October 11, 2017

[2] Unlike most legislative scaling applications there are both substantive as well as statistical reasons for selecting a single dimension, and agreement about what the dimension represents. For discussion on the statistical benefits of low-dimensionality, see Poole and Rosenthal (1997); Martin and Quinn (2002); and Clinton et al. (2004). For a paper that treats estimation when dimensionality is higher, see Lauderdale and Clark (2014).

If members of the FOMC are all tasked with the same objective, why would central bankers have different preferences? First, a member's characteristics may affect their appetite for inflation, such as who appoints them (Chang, 2003; Meade and Sheets, 2005a; Chappell Jr et al., 1993; Chappell et al., 1995; Havrilesky and Gildea, 1991, 1995; Tootell, 1996) and their educational training (Hallerberg and Wehner, 2018; Kaplan, 2017; Johnson, 2016). Second, central bankers have their own career histories, goals, and objectives (Havrilesky and Gildea, 1991; Eichler and Lähner, 2014; Adolph, 2013) and we expect their career incentives to influence their policymaking. Inter-committee politics may also affect their behavior. Schonhardt-Bailey (2013) shows that hallway conversations as well as explicit deliberation allows central bankers to influence and persuade their colleagues. Finally, gender may also play a role with female members holding more Hawkish preferences than their male counterparts (Diouf and Pépin, 2017).

Previous measures of central bank preferences predominately depended on analyzing voting behavior (see for e.g., Chappell, Havrilesky, and Roy McGregor, 2000; Chang, 2003; Gerlach-Kristen and Meade, 2010; Hix, Høyland, and Vivyan, 2010; Adolph, 2013; Eijffinger, Mahieu, and Raes, 2017; Ainsley, 2017; Bennani, Farvaque, and Stanek, 2018). Following the roll call voting literature (Poole and Rosenthal, 1997; Carroll et al., 2009; Clinton et al., 2004), scholars use committee votes to help infer committee members' policy preferences. For example, Chappell et al. estimate monetary policy reaction functions for individual FOMC members using expressed preferences in the transcripts as their data source. These authors code member preferences into three categories: favoring interest rate tightening, favoring interest rate easing, or consenting to the proposed policy rate (Chappell et al., 2000, 72).

However, while it seems natural to characterize central bankers as inflation "Hawks" who vote in favor of interest rate tightening and "Doves" who vote for loosening, this behavior may be a better measure of whether and to what extent policymakers coordinate, rather than their preference similarity or dissimilarity (Schelling, 1978). This is especially likely because in central bank decision-making, public consensus outcomes are partly determined by actors' desire to minimize apparent divergences of opinion to avoid economic uncertainty and to stabilize expectations (Woodford, 2005; Rosa and Verga, 2007). Similarly, members willingness to cooperate often depends on others in the group and especially on the willingness of the committee Chair to allow for public dissent in the first place (Crawford and Haller, 1990; Riboni and Ruge-Murcia, 2010; Shapiro and Varian, 2013).

Consequently, we observe very little variation in central bank voting behavior, which makes identifying individual committee member preferences statistically challenging. This is especially so in the United States. For example, between 1948 and 2014, FOMC members cast on average only one dissenting vote per meeting, which is less than 10 percent of all votes cast (Thornton and Wheelock, 2014). Between 2005 to 2008, the period that we consider in this paper, there were just seven official dissenting votes out of 255 official roll call votes, about two percent of all votes cast. Finally, dissenting votes are also only cast by a select subset of the committee members. Indeed one member, Bank President Lacker (Richmond), cast four out of the seven dissenting votes between 2005 and 2008. Figure 1 shows the small number of dissents over the period we examine in this paper, demonstrating little variation to actually estimate preferences. For those cases where there is enough variation to estimate members' preferences, or for those

cases where members' expressed preferences in deliberations prior to the roll call are used to supplement the data, the recovered estimates arguably remain sensitive to information provided by those rare events where the committee member dissents. Thus, while each committee member has the right to vote according to his or her own preferences (or the preferences of his/her underlying constituency), official central bank decision making, at least as it is reported by official voting, most often confirms the status quo policy rate and likely reflects policy coordination as well as committee members' preferences.[3]
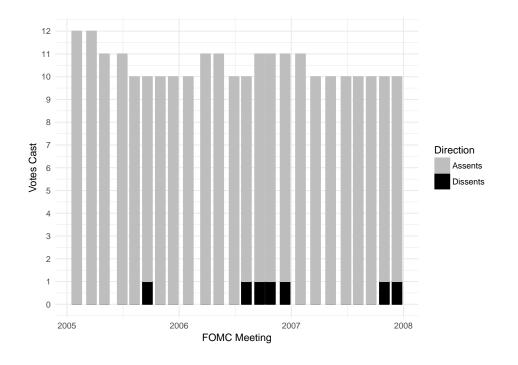

Figure 1: Voting on the FOMC between 2005 and 2008

One solution to this problem is to compute a measure of implied dissents in speech and add this data to observed dissents in votes. Using supplemental information from meeting transcripts, it is possible to identify those cases where members express dissent in the meeting deliberations but refrain from doing so officially when the actual roll call vote is cast (Chappell et al., 2000; Eijffinger et al., 2017; El-Shagi and Jung, 2015). With the inclusion of more cases, this measurement strategy leads to a greater numbers of dissents and relaxes some of the problems associated with sparse data. Unfortunately, however, this approach brings up new theoretical and measurement questions. First, why do some members dissent in speech but not in votes? Second, are dissents in speech and dissents in votes the same? For example, if a member is willing to dissent both in speech and in votes and another is only willing to express dissent in speech, is the first member relatively more extreme in position location or in position intensity or in both? Do dissents in speech predict future dissenting voting behavior or are members' speech and voting behaviors independent?

---

[3]Committee members most often vote to maintain rather than increase or decrease interest rates, and conventional policy moves are often restricted to 25 basis points byway of norms.

To avoid such problems with vote data, we construct estimates for FOMC members by using a relative-position, textual interpretation of interest rate policy-setting based on a monetary policy rule called the "Taylor rule." If central bankers' policy decisions are based on an expressed trade-off between higher values of inflation and lower values of unemployment/output, then we can use the relative amount of time central bankers spend talking about each topic to develop measures of FOMC members' preferences.

As a first step, we measure topics using a version of Latent Dirichlet Allocation (LDA) augmented to capture topical phrases as well as single words. We then apply human judgment to identify only those topics that policymakers are mandated to care about (inflation, unemployment/output). Finally, we use the expressed topic mentions of inflation and output and unemployment to scale actors' relative positions on a single dimension measuring inflation aversion. While other scholars have applied topic models, or similar approaches, using FOMC transcript data, to our knowledge, this paper is the first to start from a theoretical model of topic emphasis and then to combine the results from expressed topic emphasis with scaling models common in the political science literature.[4]

2   Estimating Central Bank Preferences

The Taylor rule is a simple approximation or guideline for what central bankers ought to do given changes in the economy developed by Stanford economist John Taylor. The original rule provides an aggregate policy recommendation for changing the federal funds rate as output and inflation change over time (Taylor, 1993). Importantly, the Taylor rule represents an explicit reaction function for the monetary policy committee; it shows how the federal funds rate should change in a response to changes in inflation (deviations from target) and changes in output (the output gap).[5]

$$r = \pi + \frac{1}{2}(y) + \frac{1}{2}(\pi - 2) + 2, \tag{1}$$

The outcome variable r is the desired federal funds rate, π is the rate of inflation over the previous four quarters, and y is the percent deviation of real GDP to target GDP. We assume that members are stating their beliefs about forward inflation and output/unemployment in their speeches, and that the topics members emphasize are consistent with their preferences, because they are all using something resembling the Taylor rule to guide them.

Importantly, the Taylor rule, or some derivation of the Taylor rule, is used as a guideline in most central banks around the world. The Board of Governors and the research staff at regional Federal Reserve Banks consult with FOMC members about projections from Taylor rule estimates and policymakers are equipped with economic projections from different Taylor rule specifications. For example, Appendix B in the 2008 "Bluebook" presents six scenarios of the future economy, three of which are based on some

---

[4]Previous work by Schonhardt-Bailey (2013); Acosta and Meade (2015) models committee member rhetoric using related exploratory scaling models. Our approach is distinguished by a theoretically motivated confirmatory rather than an unsupervised exploratory approach.

[5]For the rest of the paper we refer to output/unemployment so as to capture either the Taylor rule or the Phillips Curve, which models this trade-off according to unemployment.

derivative of the Taylor rule.[6] Even when there is uncertainty about the true model of the economy, the Taylor rule is expected to deliver good outcomes in a variety of models. A version of the Taylor rule has been used, for example, to estimate interest rate decisions by the ECB's Governing Council (Hayo and Méon, 2013) and in the United States (Plosser, 2014; Blinder et al., 2005).

In addition to the theoretical and empirical evidence, and as aforementioned, policymakers themselves talk about unemployment/output and inflation in Taylor rule terms whilst deliberating over monetary policy. For example, in one meeting on December 11, 2007, Board Governor Kroszner explained that "each incremental step we take with respect to policy easing potentially has higher and higher costs with respect to inflation. There are no free lunches here." Similarly, even for those cases when the Taylor rule is not used explicitly, members still use these two topics to anchor their thinking. Board Governor Gramlich on June 30, 2005 notes, "We don't necessarily follow the explicit Taylor rule outcomes of the Bluebook but we look at inflation and unemployment and try to make them hit our target values, at least over some time horizon."

The Taylor rule expresses a linear trade-off between inflation and output/unemployment, of the form $r = \beta_0 + \beta_1 \pi + \beta_2 y$ where, according to the weights of equation (1), $\beta_0 = 1$, $\beta_1 = 1.5$ and $\beta_2 = 0.5$. The original Taylor rule is an aggregate policy rule and does not envisage that each individual member will have a different inflation sensitivity. But if we allow for such individual variation we can then estimate rather than stipulate β_1 and β_2 for each member. Expressed as a trade-off between inflation and output concerns, we can represent the balance of an individual member's considerations built into the rule by $\log \frac{\beta_1}{\beta_2}$. Again, for illustration, consider the original specification, $\log \frac{\beta_1}{\beta_2} = \log \frac{1.5}{0.5} \approx 1$. The interpretation is that an individual has a uniquely preferred proportional emphasis of inflation versus output sensitivity.[7] In order to be consistent with the ideal point literature, where point predictions are scaled on a single left-right dimension, we use inflation sensitivity β_1 as the numerator and β_2 as the denominator. This naturally scales members that are more sensitive to inflation risks to the right of zero position as "Inflation Hawks" and those to the left of the zero position as "Inflation Doves." This allows us to systematically locate "Inflation Doves" on the left of the scale, "Inflation Hawks" on the right, and to locate other actors between.

The different sensitivity of FOMC committee members to inflation output risks are, of course, unobserved. With informative voting, we might be able to infer positions (Clinton et al., 2004; Lauderdale and Clark, 2014), but because votes are relatively uninformative in the context of FOMC committee meetings we need to look elsewhere for information. We show how to get relative emphasis from the amounts of words members deploy on inflation topics versus output topics. We model the counts of words and phrases in the inflation and output topics (c_1 and c_2) respectively as:

$$[c_1, c_2] \sim \text{Binomial}(p_i, N) \tag{2}$$
$$p_i = P(c_1 \mid N) \tag{3}$$
$$N = c_1 + c_2 \tag{4}$$

---

[6]FOMC Bluebook

[7]The original Taylor rule implies that deviation from the inflation target is $\beta_1/\beta_2 = 3$ times more important than the output gap.

Topics not related to inflation or output can be ignored because they give no information about β_1/β_2.

Clearly there are other systematic determinants of the proportion (p_i) of inflation-related speech, so while more elaborate multidimensional models are certainly possible, we limit our attention to the mandated topics as those are the specified parameters in our Taylor rule.[8] In order to account for other (time varying) concerns, we include meeting random effects m_t. We also assume that speakers' positions are exchangeable and model them as draws from a population of committee members. These effects provide our position estimates.

$$log \frac{p_{i,t}}{1 - p_{i,t}} = a + s_i + m_t \tag{5}$$

Interpreting unobserved terms in the logistic regression model can be motivated either directly, as a form of 'logit scaling' from policy position scaling (Lowe et al., 2011), or indirectly, by noting that that these quantities are rescaled versions of the positions that would be estimated from automated text scaling (Slapin and Proksch, 2008) using an association model (Goodman, 1985).[9]

2.1 Measurement

In theory, we could estimate ideal points using historical transcript data starting in 1993, when the speakers became aware that their meetings were being recorded. However, text-based ideal point estimation implicitly depends on a stable mapping between words and phrases, and the topics they indicate. Because economic and political changes over such a long period make this unlikely, we take a conservative approach. For example, Acosta and Meade (2015) argue that the FOMC's December 2008 meeting represents a semantic break from previous statements because of changes in the federal funds rate and the unique discussion of the asset purchases program. We restrict the sample time period to the three most recent years for which transcripts are available, leading up to but not including the U.S. financial crisis.

Our unit of analysis is the speaker-meeting. We construct a speaker-meeting "document" by concatenating each individual's speech contributions within each FOMC meeting for the meetings in our sample (2005-2008). Meetings are held approximately every six weeks or eight times a year, leading to a total of twenty-four meetings, which span two days. Each speaker speaks an average of 1,020 words per meeting across all meetings. In this paper we focus on actors that have a formal policymaking role and remove speeches made by staff members and other meeting participants that are either (1) not members of the Board of Governors or (2) not Bank Presidents. We do include in the sample members' speech contributions even when they are not "on schedule to vote" in the roll-call.[10]

---

[8]One possible concern is that as the housing crisis approaches, members might be more troubled by housing rather than concerned with output or inflation. As shown in table A.1, however, members talk about inflation eight times more than they do housing which means that adding housing as a topic will do little to change the results.

[9]See Lowe (2017) for a discussion of the connections between these approaches.

[10]We return to the issue of rotation voting in the section *When do Members Dissent?*.

To estimate topic prevalence we use the topical n-grams model (Wang et al., 2007) to identify those topics that central bankers talk about in FOMC meetings and to count how often language associated with a given topic is used.[11] The n-gram topic model is an extension of traditional LDA (Blei et al., 2003) topic model that allows for both words and phrases to be associated with each other within topics. Specifically, the model posits an unobserved variable 'between' tokens indicating whether any two tokens should be treated as a phrase in one of the topics. This allows us to account for economic phrases such as *price elasticity*, which would otherwise be returned as *price*, and *elasticity*. Because many of the topics that FOMC members discuss reference underlying economic theories and concepts with associated jargon, the use of topic phrases is particularly helpful for organizing and classifying topics. It also allows us to infer phrases of arbitrary length (i.e. inflation-expectations and consumer price index are of length two and three) and is flexible enough such that word pairs and phrases in one topic are not necessarily present in other topics (i.e. price-elasticity can belong to topic 1 and price to topic 2).

The Gibbs sampler used to fit the model draws from the posterior distribution of topic assignments within each document.[12] We use the posterior samples from the Gibbs sampler directly to provide our counts. The use of actual topic counts rather than normalized counts is relevant because it allows us to be more certain (or less certain) of the topic balance – and therefore position – of a speaker-meeting document depending on its length; while the relative emphasis on inflation in a speaker-meeting "document" with seventy-five inflation-related and twenty-five output/unemployment-related terms is the same as one with fifteen inflation-related and five output-related terms, we are more certain in the first case than the second case.

Once we have an estimate of our topic words and topic phrases, we use the topic words and phrases generated by the n-gram topic model to identify by hand our topics of interest. The remaining topics are discarded. In this data the most prevalent topic is inflation, with the word *inflation* itself used 2,924 times whereas the words *unemployment* and *output* are used 245 and 260 times. The word inflation alone as a word (without its related phrases such as *core inflation*, *headline inflation*, and *inflation expectations*) accounts for 22% of the top topic words. Perhaps unsurprisingly, even Dovish central bankers on the FOMC spend most of their time talking about inflation. Table A.1 in the Supplementary Appendix shows the top topic words and their counts for all topics.

In order to make sure that the results from the n-gram model that we use to extract topics are robust to different model choices, we run the n-gram model for different numbers of topics. Running the topic model for $\{5, 10, 25, 50\}$ topics, we find that $K = 25$ topics are sufficient to identify the topics at a suitable granularity. From our estimated twenty-five topics, we only keep three covering different aspects of output and unemployment and one topic on inflation.

We expect this topic aggregation strategy to work well when K is relatively large, so that we have a larger number of relatively small but precisely estimated topics to

---

[11]We use a slightly adjusted version of the implementation in MALLET

[12]We set the iterations at 750 times. We also set the optimize-interval to ten as suggested, which turns on hyperparameter optimization, which allows the model to better fit the data by allowing some topics to be more prominent than others, and use the default burn-in which is twice the optimize interval. For further details see MALLET Topic Modeling.

aggregate. Many authors find that statistical measures of overall model fit are negatively correlated with substantive interpretability ([Chang et al.](#), 2009), but theoretically driven aggregation allows us to choose less biased models with good representation of our two core theoretical concepts

To give a sense of what words and phrases are generated by the topics we label output and inflation, we show the most likely five words and most likely five phrases associated with each topic:

- **Output/Unemployment:** productivity, compensation, energy, measured, hour, market psychology, large trucks, filter estimate, price elasticity, weekend strains

- **Inflation:** inflation, percent, year, time, don, basis points, core inflation, monetary policy, inflation expectations, energy prices

Allowing n-grams to be generated by topics provides considerable help for experts attempting to identify which concept is captured by a topic. This is also helped by the fact that FOMC members work with a shared technical vocabulary. The benefit is a model that distinguishes between *federal reserve, federal funds rate,* and *federal government,* which would alternatively be returned as *federal, reserve, government, funds, rate* with uni-grams, which are more ambiguous and therefore, more threatening to topic identification.[13]

3   Results and Discussion

In this section, we present results from combining the topic model and the scaling model and discuss our findings.

Figure 2 shows the computed average estimates for each FOMC members' ideal point across the sample period, including meeting random effects.[14] Each individual's estimate is a point prediction of their relative inflation preferences from a mixed model treating speakers as a draw from a wider population of central bankers, each with their own intercept. Thus we estimate member specific reaction functions based on inflation versus employment and output emphasis. This is different from previous studies that estimate each member's implied desired federal funds rate using a fixed-effect model from implied voting decisions ([Chappell et al.](#), 2000; [Bennani et al.](#), 2018) or a spatial model based on voting decisions ([Eijffinger et al.](#), 2017). The position of each speaker's intercept on the $x - axis$ measures the degree of inflation aversion: that is, how willing a central banker is to forgo an increase in economic activity (output and jobs) for a reduction in inflation, with those scaled on the left hand side being relatively more *Inflation Dovish* than those *Inflation Hawkish* types located on the right hand side.

The results are interesting in that they confirm previous studies while offering some important new insights. The first, and perhaps most obviously, dissenting Bank President Lacker (Richmond) is estimated as having more moderate inflation preferences

---

[13]Our validation checks that more counts of inflation words indicates greater concern for rising (rather than falling) inflation is shown in Figure 1 in the Supplementary Appendix.

[14]We also estimated a model with no meeting random effects and find similar results.

compared to expectations based on his interest-rate tightening, dissenting, voting behavior. Lacker expressed official dissent to proposed changes to the federal funds rate a number of times during his tenure at the Fed. In 2006, for example, during our sample period, Lacker dissented in the August, September, October, and December meetings. In these meetings, the FOMC agreed to keep interest rates at the status quo rate whereas Lacker voted for additional interest rate tightening. According to vote based measures, his persistent dissents (four times) and the call for rising rates would push him towards to the extreme on Hawkishness. Estimates for Lacker based on implied voting, for example, have him the farthest right on the scale for 2007 (Eijffinger et al., 2017). Using our method, while we find Lacker to be Hawkish when we restrict the sample to only 2007, his estimated preferences for the three years place him much closer to the median.

What might explain this discrepancy? In reading the texts of the transcripts, we see that Lacker appeals to his reputation for inflation Hawkishness to discuss output and unemployment more freely in his speech. In the December 12, 2006 meeting, for example, which is the period right before he dissents, he says: "In past meetings I have expressed misgivings about whether our strategy is going to bring inflation down fast enough. So I won't belabor those again today." His next topic is one that expresses concern about output and unemployment in his home district. By examining voting decisions alone, countervailing concerns, such as concerns about the economy in a member's home district,are subsumed by the rarity (and therefore power) of dissenting votes.
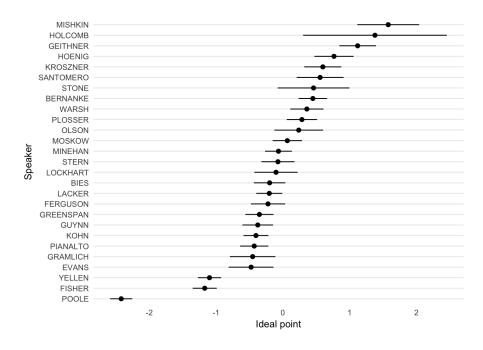

Figure 2: Estimated Fixed Ideal Points from Full Transcripts with meeting random effects

Another interesting finding is that Bank President Geithner (New York) is estimated by our method as very Hawkish, significantly more Hawkish than both President Lacker and Chairman Greenspan, despite never casting a dissenting vote. This is important for two reasons. First, as President of the New York Fed, Geithner is perhaps most likely to represent the interests of the financial sector than other Bank Presidents; his Hawkishness may suggest that constituency effects matter. Second, President Geithner had a

career in government before and after being on the FOMC board. His vocational expectations may also shape his preferences on inflation (Eichler and Lähner, 2014; Adolph, 2013). As the representative of the largest district, the New York Federal Reserve Bank President always votes on the FOMC. If Geithner is a hawk, and more hawkish than the average Bank President on the FOMC, this means that during his tenure at the FOMC, the bias of Bank Presidents voting on the FOMC committee was likely to be skewed rightwards even though inflation was close to or at the target rate of 2%.

Lastly, Chairman Greenspan straddles the middle position when compared against his two successors, Chairman Bernanke and Chair Yellen. Surprisingly, Bernanke is estimated as relatively more inflation adverse than either Greenspan or Yellen, whereas unsurprisingly, Yellen is estimated as more Dovish. Table A.2 in the Supplementary Appendix reports the full set of static ideal point estimates for FOMC members.
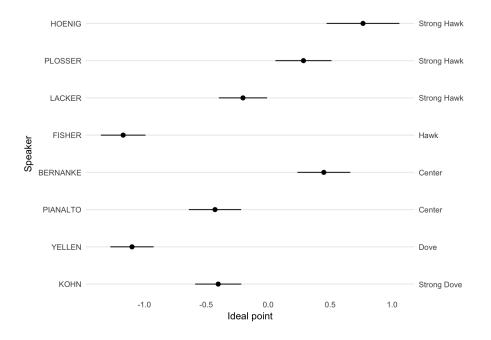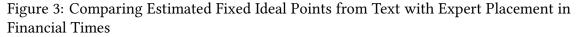
3.1   Convergent Validity

Do these relative positions really uncover noticeable differences in members' inflation preferences as expressed in FOMC meetings? How do our measures compare with other measures in the literature? In order to assess the validity of our measures, we first compare our measures with a subset of "expert rankings" published in the newspaper, the *Financial Times* (FT) in 2009.[15] Because the ranking is from 2009, the FT rankings are out of sample (recall that our data ends in 2008). We then select and rank FOMC members on an ordered categorical scale from 1 to 5 according to their FT labels: "Strong Dove," "Dove," "Center," "Hawk," or "Strong Hawk." The categories are mutually exclusive. Figure 3 shows our ordering for the eight committee members that match across the two samples. We find a positive relationship between our measure and the expert ranking (correlation 0.50). Both President Lacker (Richmond) and President Fisher (Dallas) are located more to the left and mid-point in our sample compared to the FT. Another difference is that we estimate Chair Bernanke as being more Hawkish than the expert ranking, who locate him closer to the midpoint.

The main point of departure between our estimate and the expert placement is with President Fisher, and so we discuss him further. President Fisher is one of the few members of the FOMC that dissents during his tenure on the FOMC committee, although he does not dissent during our sample but the following year, five times from January to May 2008, which is behavior likely incorporated in the FT ranking. Fisher again dissents twice in 2011 (Thornton and Wheelock, 2014). In terms of understanding the validity of our measures for those individuals that dissent, there are two possible scenarios: (1) our model ignores dissenters' positions (i.e. puts no weight on members' voting behavior) and therefore suffers from underfitting or (2) alternatively, the expert ranking over-estimates members' positions (i.e. puts too much weight on voting behavior) and therefore suffers from overfitting. Because our method can estimate members' position irrespective of whether or not a member is on schedule to vote, our model is more likely to accurately capture the regularities in the data and generalize well to unseen data. This coupled with the fact very few members dissent and do so rarely, (recall that in our data the percentage of dissents is 2% and 10% historically) suggests that our method is valid for uncovering most members positions on the FOMC most of the time.

---

[15]Krishna Guha, "US Federal Reserve: An eclectic aviary," *Financial Times*, November 12, 2009.

That being said, our estimate for President Fisher does little to suggest that he will cast a series of interest-rate tightening votes that go against the consensus of the committee the following year. As mentioned in the introduction, further theoretical and empirical research explaining why we observe stark variations in how members speak and how they vote is an interesting avenue worth exploring.

Figure 3: Comparing Estimated Fixed Ideal Points from Text with Expert Placement in Financial Times

How do our measure compare with estimates from statistical models in the literature rather than the expert rankings by the FT? Figure 4 presents our estimates (column 1) alongside estimates of ideal points based on implied voting behavior by Eijffinger et al. (2017) (column 2) and estimated interest-rates by Bennani et al. (2018) (column 3), only this time ordered according to our measure; the remaining people are those members not included in the FT ranking but estimated in the other papers.

The most striking feature of our data compared to the other measures is that using textual information gives us significantly more variation across members than what is uncovered using voting information. This is perhaps unsurprising as, the alternative measures use voting information and central bank voting is relatively invariant. The impact of a limited number of official dissents on estimated ideal points (in this case, the preferred interest rate) is most stark in column 3, where members are, more or less, indistinguishable from one another. To compute these measures, the authors derive members' policy differentials, "or the difference between the desired interest rate and the Fed's actual policy rate"[16] The estimates are based on a sample of members' voting decisions, between 1994 and 2008 (Bennani et al., 2018). Unfortunately, their dataset does not include estimates for Fisher and Lacker, making it difficult to compare those individuals across all measures. President Yellen (San Francisco) is estimated as relatively more hawkish than expected.

---

[16]The authors make these estimates available in their appendix Table A.2, p.8.
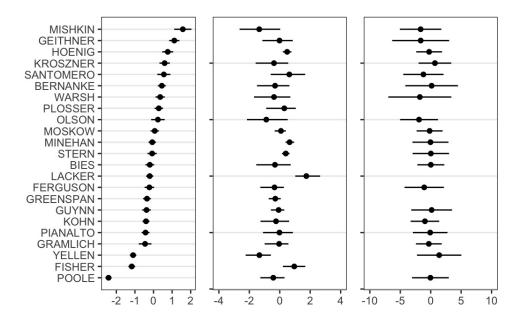
Figure 4: Comparing Ideal points from Text, Implied Voting and Preferred Interest Rate measures

Column 2 presents estimates by Eijffinger et al. (2017), who use actual and implied votes as coded from the transcripts and a Bayesian static spatial voting model over the period 1989-2007. Because it is Bayesian, the estimated uncertainty bounds are 95% credible intervals rather than confidence intervals and the midpoint estimate are not necessarily located in the middle.[17] Comparing their estimates with column 3, one sees the utility in including additional information in implied votes. The uncertainty on some (but not all) of the members shrink as a consequence of more dissents. The estimate for Yellen in column 2 compared to column 3, for example, is also much more similar to our estimate and to the expert placement. None of the measures are well correlated, however. Our measure (column 1) is uncorrelated with the vote based (column 2) measure (-0.06) and negatively correlated with the interest-rate measure (column 3) (-0.48). The two vote-based measures (column 2 and column 3) are also uncorrelated (0.04). What about if we examine the correlation with the ranking by the FT as well? Here we see greater consistency, at least across the measures in column 1 and column 2. The correlation with the FT and our text-based measure (column 1) is (0.50) and with the implied vote measure (column 2) (0.72). Meanwhile, the interest-rate measure (column 3) is also uncorrelated with the FT expert ranking (0.00).

3.2 Predictive Validity

This section applies our new estimates to existing findings in the literature as a way to demonstrate predictive validity.

---

[17]The data was generously provided by the authors.

3.2.1 Evaluating Regional Economic Performance

One possible explanation for why central bankers have different preferences is they hold preferences related to their home-districts and that their home districts have different economic conditions (see for e.g., Chappell et al., 2000; Bennani et al., 2018; Meade and Sheets, 2005b). A related literature also examines whether members' personal incentives, such as career motivations, are an important predictor of members' policy decisions. Adolph (2013), for example, examines how FOMC members might make decisions today to position themselves for future employment in the financial sector. Consequently, members that anticipate future links with finance have incentives to express more Hawkish tones while policymakers on the FOMC (see also Eichler and Lähner, 2014).

To the best of our knowledge, while evidence for regional and professional biases have been linked to members' voting decisions, to date, scholars have not tested such theories using speech data. As a first application, we replicate these findings by (1) examining the correlation between the topics that FOMC members talk about and regional level economic data and (2) examining whether or not members express views consistent with lenders or borrowers in their region.

As before, the unit of analysis is speaker-meeting contributions. One way that the sample is different from that above, however, is that we have to restrict the sample to include only regional district banks and the Bank Presidents who preside over these districts.[18] To assess whether or not regional level unemployment is correlated with how much a member talks about unemployment and output, we collect information for the regional unemployment rate in the member's district at the current period. As a second measure of unemployment, we also examine one year ahead forecast unemployment. The projections are unemployment rates as reported in the Federal Reserve Greenbook, which are published two weeks before the FOMC meeting. Importantly, the inclusion of forecasts from the FOMC meeting means that we use real-time (as opposed to revised) data, which is the same as what the members themselves had at the time of the meeting.

Next, we also collect real-time current estimates of national inflation, measured according to consumer prices (CPI) and personal consumption expenditures (PCE). Second, and like the unemployment rate, we collect CPE and PCE inflation forecasts for the current period and for one year ahead. The examination of future conditions is especially important as monetary policy has an eight to twelve month lag and so committee members may base their topics of speech today on expectations of future movements in the economy (Gordon and Leeper, 1994).

Further, to reexamine the link between the financial sector and the predicted "Hawkishness" of a committee member, we use the logged dollar amount of non-performing loans in a member's district as a proxy for the size of the financial sector in the members' district. We think that this particular indicator is a good test as there are countervailing expectations about the size of debt and predicted Hawkishness of the Bank President. On the one hand, high levels of non-performing loans may make Bank Presidents more "Dovish." This is because an increase in inflation lowers the real (as opposed to nominal) value of outstanding loans. If the amount of non-performing loans is nega-

---

[18]This means that members that are Governors are left out of the analysis as their constituency is expected to be the national economy rather than their regional district.

tively related to inflation talk, we might expect that Bank Presidents lean towards representing borrowers rather than lenders in their districts. Alternatively, if the amount of non-performing loans is negatively related (or unrelated) to Bank Presidents' inflation talk, we might expect that Presidents with larger financial sectors are more likely to be Hawkish, which is more in line with the interests of lenders. Another benefit of using this particular measure is that the measure is behavioral, or dependent on the actions of financial actors located in the members' district, rather than a policy measure, such as interest rates, that the Presidents themselves have control over.

Table 1 presents our regression results. The dependent variable is the share of inflation talk over output and unemployment talk. The independent variables, as discussed above, are (a) current-period national inflation; (b) current-period national unemployment; (c) the difference between district-level unemployment and national level unemployment in the current-period; (d) the log of the current level of non-performing loans; (e) one year ahead projected national inflation; and (f) one year ahead projected national unemployment. The projections are made by staff members at the Board of Governors two weeks before the FOMC meeting and are given to the meeting participants. The statistical model used is a mixed-effects logistic regression model where the dependent variable is the proportion of the raw counts expressed by each speaker at every meeting he or she attends.

Model (1) confirms that national-level inflation is positively (but not statistically significantly) related to a greater share of inflation talk. Meanwhile, national unemployment is negatively (and statistically significantly) related to more inflation talk, also as expected. Similarly, a rise in the gap between unemployment at the district level and at the national level is negatively related to inflation; this confirms that district representatives are sensitive to both the national economy and also their districts. In sum, model (1) shows that when unemployment conditions move higher, either nationally or regionally, inflation talk declines.

What about when we examine the role of non-performing loans on inflation and output/unemployment speech? Here we find that district Bank Presidents talk more about inflation when the level of non-performing loans in their district is higher. In other words, there is a positive correlation between district-level debt and inflation talk. If the level of non-performing loans is a good proxy for the size of the financial sector, such a result might imply that Bank Presidents are more likely to represent creditor rather than debtor interests in their districts. Like in model (1), model (2) continues to show that unemployment is negatively and significantly related to inflation talk. Taken together, this model provides further evidence that FOMC Bank Presidents act as if governed by a *Textual Taylor Rule*.

Finally, model (3) examines forward looking projections of the economy rather than conditions on the ground at the time of the FOMC meeting. In this model, we find that while not only is the regional unemployment gap and projected unemployment negatively related with inflation talk, so too is projected future inflation. This is particularly puzzling as why would projected inflation be negatively related to talk about inflation today if monetary policy works only with a lag? Interestingly, the mean level of current period inflation is significantly higher than the four quarter ahead forecast. This is true irrespective of whether or not we examine CPI or PCE inflation values. On average, the current CPI inflation average over the series (3.01) is similar to the actual average (3.4;

3.2; 2.8) reported by the Bureau of Labor Statistics,[19] yet, the projected one year ahead forecast is substantively lower than the current inflation values almost all of the time. The fact that future inflation is systematically predicted as being lower than current inflation is likely what is generating the negative relationship between future inflation and inflation talk in model (3).

In summary, previous studies have examined the role of district-economic performance and financial-sector connectivity as possible predictors of central bank preferences. As with most of the previous literature, such studies omit members' speeches. In this application, we show how inflation and unemployment are related to members speech in accordance with a *Textual Taylor Rule*; members speak more about inflation when unemployment goes down as well as the reverse. We also find that district Bank Presidents are more likely to favor inflation speech when debts in their region are increasing. Such a finding suggests that Bank Presidents may better represent the interests of creditors rather than debtors in their districts.

3.2.2 When do Members' Dissent?

At any given time, the total number of FOMC committee members officially included in the roll call vote is twelve. These twelve people include seven members of the Board of Governors of the Federal Reserve System; the President of the Federal Reserve Bank of New York; and four of the remaining eleven Reserve Bank Presidents. The four Bank Presidents serve one-year terms and voting rotates according to an agreed calendar (i.e. members do not select themselves into any given meeting). Despite lacking voting privileges at some (but not all) meetings, the remaining Reserve Bank Presidents attend meetings, make economic assessment of the economy, and express their individual policy positions for both the policy rate and the official policy statement in the transcripts prior to the roll call.

Using a natural experiment to analyze the effect of transparency on members' incentives to dissent, Meade and Stasavage (2008) find evidence that voting members are much less likely to dissent than non-voting members. We retest this finding using our data, taking into account a member's estimated ideal point from text as well as whether the member is on schedule to vote or not. If we find that members behave differently with their speech and with their votes when on schedule to vote, we can assume that voting and non-voting members have different audiences. Recall that while voting decisions are made readily available, publication of transcripts are with a 5 year lag.

Following the procedure by Meade and Stasavage (2008), using the transcripts, we code dissent as a binary measure based on whether a policymaker voiced disagreement with the Chair's policy proposal during the policy debate before the formal roll-call vote. Further, to take a more generous measure of "dissent," we code a member as dissenting even if the member expresses dissent only in speech and not in the roll call; this gives us a larger number of dissents to work with (32 vs. 7). Furthermore, instead of just looking at dissenting votes that protest changes in the policy interest rate, we also examine whether members propose changes to the official policy statement which is publicized after the decision; this includes an even larger number of dissents (118 out of 425 total). Following Meade and Stasavage (2008), we expect a negative relationship between those

---

[19]Bureau of Labor Statistics Inflation

Table 1: Regression Results for FOMC Bank Presidents and District Level Economic Information

	Dependent variable:		
	(Inflation/Output&Unemployment)		
	(1)	(2)	(3)
National Inflation Rate	0.007		
	(−0.010, 0.024)		
log(Non-Performing Loans)		0.415***	
		(0.314, 0.517)	
National Unemployment Rate	−0.466***	−0.318***	
	(−0.585, −0.347)	(−0.440, −0.195)	
National Inflation 1 year ahead			−0.456***
			(−0.567, −0.344)
National Unemployment 1 year ahead			−0.741***
			(−0.992, −0.489)
Difference in Regional to National Unemployment	−0.253***	−0.245***	−0.371***
	(−0.366, −0.141)	(−0.358, −0.133)	(−0.489, −0.252)
Constant	6.035***	−0.866	8.453***
	(5.338, 6.732)	(−2.680, 0.947)	(7.041, 9.865)
Observations	175	175	175
Log Likelihood	−1,098.516	−1,074.790	−1,092.171
Akaike Inf. Crit.	2,207.032	2,159.580	2,194.342
Bayesian Inf. Crit.	2,222.856	2,175.404	2,210.166

Note: *p<0.1; **p<0.05; ***p<0.01

members on rotation to vote and expressing a dissenting opinion, regardless of their policy position. The rationale is that consent is more likely expressed by voting members due to strategic incentives to minimize economic and policy uncertainty that results from dissenting opinions.

We test this hypothesis by regressing dissent on member ideal points and voting status, with meeting random effects to control for over-time changes in the economy. The results are shown in Table 2. Model (1) predicts dissent over the policy interest rate, and model (2) predicts dissent over the text of the official policy statement.

Table 2: Predicting Dissents in Policy Rates and Policy Statements with Ideal Points and Voting Status

	Dependent variable:	
	Dissent Policy Rate	Dissent Statement
	(1)	(2)
est	−0.063	−0.145
	(0.153)	(0.089)
vote	−0.738*	−0.753***
	(0.404)	(0.243)
Constant	−2.610***	−0.650**
	(0.424)	(0.256)
Observations	405	403
Log Likelihood	-100.618	-222.952
Akaike Inf. Crit.	209.236	453.904
Bayesian Inf. Crit.	225.252	469.900

Note:	*p<0.1; **p<0.05; ***p<0.01

We find evidence in support of the hypothesis that the rotation system engenders consensus opinion among *those members on schedule to vote.* This confirms theories arguing that official voting behavior likely reflects strategic behavior towards coordination in addition to the expression of policy positions.[20] Our findings suggest that being on schedule to vote makes you less likely to dissent across both categories (i.e. strategic voting). We also find that being more inflation Hawkish makes you less likely to dissent to the policy statement, however, this is not statistically significant. So, according to these results, members' voting position (on rotation or off rotation) helps predict whether or not a member will dissent, thus suggesting strategic rather than

---

[20]In a separate analysis, we also compare members' estimated positions from the full meeting and positions from the economic-go around when members report on their district and where members can request when to speak (at the start, middle, or end) – thus allowing for strategic behavior. We find no evidence that members take on different positions within a meeting. Results are available upon request from the authors.

sincere voting.

4   Conclusion

This article shows the usefulness in combining topic and scaling methods to estimate central bank policy positions. We develop a new way to estimate central bank preferences on the FOMC based on topics that committee members are mandated to talk about. We extract topic-counts from committee meeting transcripts and compute relative topic emphasis to help identify differences in actors' policy positions. Our main contribution is that we offer an estimate of FOMC members' preferences based on private policy deliberations that occur before members agree (or disagree) to coordinate their policy actions by roll-call voting and we do so for all members, irrespective if they are on cycle to vote or not. By using textual information as data to help uncover each members' relative inflation emphasis, we offer ideal-point estimates based on a members' professed inflation aversion.

We show greater variation in FOMC members' preferences than what we see in voting behaviors. This supports recent evidence in other domains where researchers also uncover greater variation in speech than in votes. We also find that by only analyzing voting data, researchers likely overemphasize a dissenter's dissimilarity with his or her colleagues. This latter point is especially important because both researchers and market participants use voting behavior to predict changes in key interest rates. One fruitful item for further research, therefore, is to examine whether there exists an similar relationship between dissents in speech and votes and changes to key interest rates. It is not clear whether the pathway of expressed dissent goes from speech to votes or if these two policy behaviors are independent. Do committee members start off making verbal protests in private only to build up to official dissent or do policymakers use these two behaviors so as to satisfy different goals?

Finally, our results also provide important policy implications. The Bank of England recently changed its procedures and will make permanent recordings and transcripts of Monetary Policy Committee meetings (Warsh, 2014). While some have warned that by doing so, committee members may stifle their debate, our findings suggest that publication of the meetings need not stifle diversity in deliberations amongst members.

Moving away from the FOMC and central banking, our approach also provides an important tool for researchers that are interested in estimating actors' preferences in political institutions where expressed dissents are rare. As we show in this paper, matching human judgment with unsupervised learning techniques allows us to estimate members' preferences despite the fact that institutional norms encourage consensus decisions. Our approach can be directly applied to any type of political institution with a dual mandate. For example, those interested in bureaucratic politics might look to specific agencies, such as the U.S. Housing and Urban Development, and determine how the agency prioritizes low-income housing versus mortgage insurance. At the international level, another application might be the Stability and Growth Pact, where Member States commit themselves to staying within certain debt and deficit spending limits. Here one might examine whether different EU Member States prioritize one fiscal limit over the other.

References

Acosta, Miguel and Ellen E. Meade. 2015. Hanging on Every Word: Semantic Analysis of the FOMC's Postmeeting Statement. Feds note, Board of Governors of the Federal Reserve System, Washington, D.C.

Adolph, Christopher. 2013. *Bankers, Bureaucrats, and Central Bank Politics: The Myth of Neutrality.* Cambridge University Press.

Ainsley, Caitlin. 2017. The politics of central bank appointments. *The Journal of Politics 79*(4), 1205–1219.

Bennani, Hamza, Etienne Farvaque, and Piotr Stanek. 2018. Influence of regional cycles and personal background on fomc members' preferences and disagreement. *Economic Modelling 68*, 416 – 424.

Blei, David M., Andrew Y. Ng, and Michael I. Jordan. 2003. Latent Dirichlet Allocation. *Journal of Machine Learning Research 3*, 993–1022.

Blinder, Alan S, Ricardo Reis, et al. 2005. Understanding the Greenspan standard. *The Greenspan Era: Lessons for the future 1196.*

Carroll, Royce, Jeffrey B Lewis, James Lo, Keith T Poole, and Howard Rosenthal. 2009. Comparing NOMINATE and IDEAL: Points of difference and Monte Carlo tests. *Legislative Studies Quarterly 34*(4), 555–591.

Chang, Jonathan, Jordan Boyd-Graber, Sean Gerrish, Chong Wang, and David M. Blei. 2009. Reading Tea Leaves: How Humans Interpret Topic Models. *Proceedings of the 23rd Annual Conference on Neural Information Processing Systems*, 288–296.

Chang, Kelly H. 2003. *Appointing Central Bankers: The Politics of Monetary Policy in the United States and the European Monetary Union.* New York: Cambridge University Press.

Chappell, Henry W, Thomas M Havrilesky, and Rob Roy McGregor. 1995. Policymakers, institutions, and central bank decisions. *Journal of Economics and Business 47*(2), 113–136.

Chappell, Henry W., Thomas M. Havrilesky, and Rob Roy McGregor. 2000. Monetary Policy Preferences of Individual FOMC Members: A Content Analysis of the Memoranda of Discussion. *The Review of Economics and Statistics 79*(3), 454–460.

Chappell Jr, Henry W, Thomas M Havrilesky, and Rob Roy McGregor. 1993. Partisan monetary policies: Presidential influence through the power of appointment. *The Quarterly Journal of Economics 108*(1), 185–218.

Clinton, Joshua, Simon Jackman, and Douglas Rivers. 2004. The Statistical Analysis of Roll Call Data. *American Political Science Review 98*(02), 1–16.

Crawford, Vincent P and Hans Haller. 1990. Learning how to Cooperate: Optimal play in Repeated Coordination Games. *Econometrica*, 571–595.

Diouf, Ibrahima and Dominique Pépin. 2017. Gender and central banking. *Economic Modelling 61*, 193–206.

Eichler, Stefan and Tom Lähner. 2014. Forecast dispersion, dissenting votes, and monetary policy preferences of fomc members: the role of individual career characteristics and political aspects. *Public Choice 160*(3-4), 429–453.

Eijffinger, Sylvester, Ronald Mahieu, and Louis Raes. 2017. Inferring hawks and doves from voting records. *European Journal of Political Economy*.

El-Shagi, Makram and Alexander Jung. 2015. Does the Greenspan era provide evidence on leadership in the FOMC? *Journal of Macroeconomics 43*, 173–190.

Gerlach-Kristen, Petra and Ellen E. Meade. 2010. Is There a Limit on FOMC Dissents? Evidence from the Greenspan Era. Working Paper 2010-16.

Goodman, Leo A. 1985. The Analysis of Cross-Classified Data Having Ordered and/or Unordered Categories: Association Models, Correlation Models, and Asymmetry Models for Contingency Tables with or without Missing Entries. *The Annals of Statistics 13*(1), 10–69.

Gordon, David B and Eric M Leeper. 1994. The dynamic impacts of monetary policy: an exercise in tentative identification. *Journal of Political Economy 102*(6), 1228–1247.

Hallerberg, Mark and Joachim Wehner. 2018. When do you get economists as policy makers? *British Journal of Political Science*, 1–13.

Havrilesky, Thomas and John Gildea. 1991. The Policy Preferences of FOMC Members as Revealed by Dissenting Votes: Comment. *Journal of Money, Credit, and Banking* (1), 130–138.

Havrilesky, Thomas and John Gildea. 1995. The biases of federal reserve bank presidents. *Economic Inquiry 33*(2), 274–284.

Hayo, Bernd and Pierre-Guillaume Méon. 2013. Behind closed doors: Revealing the ECB's decision rule. *Journal of international money and finance 37*, 135–160.

Hix, Simon, Bjørn Høyland, and Nick Vivyan. 2010. From doves to hawks: A spatial analysis of voting in the Monetary Policy Committee of the Bank of England. *European Journal of Political Research 49*(6), 731–758.

Johnson, Juliet. 2016. *Priests of Prosperity: How Central Bankers Transformed the Post-communist World.* Cornell University Press.

Kaplan, Stephen B. 2017. Partisan technocratic cycles in latin america. *Electoral Studies 45*, 219–229.

Lauderdale, B. E. and T. S. Clark. 2014. Scaling Politically Meaningful Dimensions Using Texts and Votes. *American Journal of Political Science* (3), 754—-771.

Lowe, Will. 2017. Putting it all on the line: Some unified theory for text scaling. *Working Paper*.

Lowe, Will, Ken Benoit, Michael Laver, and Slava Mikhaylov. 2011. Scaling Policy Positions from Coded Units of Political Texts. *Legislative Studies Quarterly* (1), 123–155.

Lucas, Christopher, Richard A Nielsen, Margaret E Roberts, Brandon M Stewart, Alex Storer, and Dustin Tingley. 2015. Computer-assisted text analysis for comparative politics. *Political Analysis*.

Martin, Andrew D and Kevin M Quinn. 2002. Dynamic Ideal Point Estimation via Markov Chain Monte Carlo for the US Supreme Court, 1953–1999. *Political Analysis 10*(2), 134–153.

Meade, Ellen E. 2005. The FOMC: Preferences, Voting, and Consensus. *Federal Reserve Bank of St. Louis Review* (2), 93–101.

Meade, Ellen E. and D. N. Sheets. 2005a. Regional Influences on FOMC Voting Patterns. *Journal of Money, Credit, and Banking* (4), 661–677.

Meade, Ellen E. and D. Nathan Sheets. 2005b. Regional Influences on FOMC Voting Patterns. *Journal of Money, Credit and Banking 37*, 661–677.

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

Plosser, Charles I. 2014. Monetary rules: Theory and practice. *Journal of Economic Dynamics and Control 49*, 144–146.

Poole, Keith T. and Howard Rosenthal. 1997. *Congress: A Political-Economic History of Roll Call Voting*. Oxford University Press.

Proksch, Sven-Oliver and Jonathan B. Slapin. 2012. Institutional Foundations of Legislative Speech. *American Journal of Political Science 56*, 520–537.

Riboni, Alessandro and Francisco J. Ruge-Murcia. 2010. Monetary Policy by Committee: Consensus, Chairman Dominance, or Simple Majority? *The Quarterly Journal of Economics 125*(1), 363–416.

Roberts, Margaret E., Brandon M. Stewart, Dustin Tingley, Christopher Lucas, Jetson Leder-Luis, Shana Kushner Gadarian, Bethany Albertson, and David G. Rand. 2014. Structural topic models for open-ended survey responses. *American Journal of Political Science 58*(4), 1064–1082.

Rosa, Carlo and Giovanni Verga. 2007. On the Consistency and Effectiveness of Central Bank Communication: Evidence from the ECB. *European Journal of Political Economy 23*, 146–175.

Schelling, Thomas. 1978. *Micromotives and Macrobehavior*. Toronto, Canada: George J. McLeod Ltd.

Schonhardt-Bailey, Cheryl. 2013. *Deliberating American Monetary Policy: A Textual Analysis.* MIT Press.

Schwarz, Daniel, Denise Traber, and Kenneth Benoit. 2017. Estimating intra-party preferences: comparing speeches to votes. *Political Science Research and Methods 5*(2), 379–396.

Shapiro, Carl and Hal R Varian. 2013. *Information Rules: A Strategic Guide to the Network Economy.* Harvard Business Press.

Slapin, Jonathan B. and Sven-Oliver Proksch. 2008. A Scaling Model for Estimating Time-Series Party Positions from Texts. *American Journal of Political Science 52*(3), 705–722.

Taylor, John. 1993. Discretion vs Policy Rules in Practice. *Carnegie Rochester Conference Series on Public Policy*, 195–214.

Thornton, Daniel L. and David C. Wheelock. 2014. Making Sense of the Dissents: A History of FOMC Dissents. *Federal Reserve Bank of St. Louis Review 96*(3), 213–227.

Tootell, Geoffrey MB. 1996. Appointment procedures and fomc voting behavior. *Southern Economic Journal*, 191–204.

Wang, Xuerui, Andrew McCallum, and Xing Wei. 2007. Topical n-grams: Phrase and Topic Discovery, with an Application to Information Retrieval. *In Proceedings of the 7th IEEE International Conference on Data Mining*.

Warsh, Kevin. 2014. Transparency and the Bank of England's Monetary Policy Committee. A review. Available from `www.bankofengland.co.uk/publications/Documents/news/2014/warsh.pdf`.

Woodford, Michael. 2005. Central Bank Communication and Policy Effectiveness. NBER working paper no. 11898, National Bureau of Economic Research.