Deciphering Fedspeak: The Information Content of FOMC Meetings

Narasimhan Jegadeesh

Di Wu*

First Draft: September 2013
Current Draft: March 2015

ABSTRACT

We present a new approach to quantify the economic and policy content of Federal Reserve communications by dissecting the Federal Open Market Committee (FOMC) meeting minutes into distinct economic topics, and simultaneously extract the tone and uncertainty level of each topic. We use market reaction to objectively assess the relative informativeness of each topic, and we find significant incremental informational value from the topic contents, despite that the minutes are released several weeks after the original meetings. Furthermore, we find evidence consistent of the Fed possessing superior information, which is then transmitted to the market through the language of the minutes.

*Narasimhan Jegadeesh is the Dean's Distinguished Professor at the Goizueta Business School, Emory University, and NBER, ph: 404-727-4821, email: Jegadeesh@emory.edu. Di (Andrew) Wu (corresponding author) is a Ph.D. student of Finance at The Wharton School, University of Pennsylvania, ph: 215-898-2543, email: wudi1@wharton.upenn.edu. We would like to thank Itay Goldestin, Jeremy Siegel, Rob Stambaugh, Paul Tetlock and Amir Yaron for helpful comments and suggestions.

1. Introduction

Monetary policies implemented by the Federal Reserve have profound effects on the global economy. Numerous papers in the economics and finance literature examine the determinants and effects of such policies using *quantitative* "Fed proxies" such as the federal funds target rate or the reserve requirement. In addition to these hard data, the Fed routinely releases large amounts of *qualitative* information, such as meeting minutes, transcripts, and speeches, in an effort to foster effective communication with the public and achieve greater operational transparency. While a voluminous literature examines market reactions to quantitative information such as rate changes, very few papers explore the informativeness of these "soft" data conveyed in the language of Fed communications. Do they have incremental information value? How does the market react to these data? Can they be used as alternative predictors of economic and policy outcomes?

Our paper fills the void by presenting an innovative, topic-based approach to determine the informativeness of FOMC meeting minutes, which are detailed summaries of everything discussed during the preceding meeting. Because such discussions encompass a wide range of topics, the proportions of which vary widely from meeting to meeting, we use an automated algorithm based on Bayesian learning to objectively and robustly classify each individual paragraph in the minutes into four distinct economic themes that intuitively correspond to specific Fed mandates and tasks: *Growth*, *Inflation*, *Financial markets*, and *Policy*. We then simultaneously extract contents–the tone and uncertainty level–from the texts of each minute, and by topic. Compared to a manual approach such as Romer and Romer (1989), our objective approach minimizes any potential researcher-induced bias, thereby allowing us to accurately gauge the specific context of each discussion and, for each meeting minute, obtain a granular measure of topic mix that human readers cannot accurately identify. To

further remove any subjectivity, we assess the informativeness of each topic based on financial market's reactions to the release of the minutes.

We find several new results with our approach. First, we demonstrate that the texts of FOMC minutes contain incremental information not incorporated in either rate announcements or the more timely meeting statements, despite the fact that the minutes are released several weeks after the meetings. While Lucca and Moench (2013) find that policy announcement on the day of the FOMC meeting is correlated with significantly higher stock market volatilities, we show that,several weeks after that, the release of the meeting minutes is also correlated with a similar degree of volatility spike in the stock market.

Our next set of tests examine the granular source of this additional informativeness from individual topics. We first demonstrate that, when treated as a single unit, each document as a whole does not yield informative results: neither whole-document tone nor uncertainty is significantly related to market reaction. However, the market do find the discussion on individual topics informative, and assign different informational value to different topics. The market finds traditional "dual mandate" themes, such as *Inflation*, most informative. Interestingly, the market also reacts strongly to the content of the relatively new topic of *Financial markets*, reflecting the Fed's increasingly important role of maintaining systemic stability, particularly during and after the recent financial crisis. Furthermore, we find that the *Policy* topic is not only deemed informative by the market, but its discussion is also orthogonal to existing economic conditions, indicating that the FOMC members do not necessarily follow fixed guidelines such as the Taylor rule when setting the monetary policy. Our topic-content Scores also hold significant predictive power for real economic activities, which we explore in a related research.

The results above suggest that the Fed possesses superior information than other market participants. Our next tests examine whether such superior information is transmitted to

the market through "soft channels" conveyed by language of the minutes. We show that the price jumps at the release of the minutes do not revert, and market volatility is greatly reduced after the release of the minutes. This is consistent with information transmission into the market at the time of the minutes' release.

Our paper contributes to the literature on three fronts. First, our paper is the first in finance to use a topic-based textual analysis approach on the FOMC minutes, and our approach provides collection of intuitive indicators on multiple facets of the economy and monetary policy, which are also orthogonal from existing economic variables. Alternative text-based economic indicators also exist, such as Baker, Bloom, and Davis (2013), which is based on counting the frequency of uncertainty-related words in news reports. By contrast, our policy indicators are derived directly from the language of policy makers themselves. Unlike news reports, each FOMC minute is likely to be painstakingly scrutinized by the market, and the usage of every word from the minutes thus matters. This is evidenced by the significant market reaction to our measures. As such, our economic and policy indicators are likely to contain more policy-relevant information and less noise.

Second, our paper furthers the burgeoning literature of financial textual analysis by being the first to employ a paragraph-level information retrieval system that moves beyond the traditional word-based approach employed in current literature such as Tetlock (2007), Hanley and Hoberg (2010), Loughran and McDonald (2011), and Jegadeesh and Wu (2013). This paper is the first in finance to employ on FOMC minutes the Latent Dirichlet Allocation (LDA) model of automated topic retrieval, which has been successfully employed to characterize topics of a wide variety of document sources, from journal articles in Nature to patient-discharge reports.[1] LDA is ideally suited to our collection of FOMC minutes for the following reasons: first, the topic mix and content of FOMC minutes are sufficiently varied, which leads to both robust and intuitively appealing classification results that are

[1]For a list of LDA applications and an evaluation of their effectiveness, see Blei, Ng, and Jordan (2003).

on par with or exceeds manual classification by researchers.[2] Second, compared to manual approaches, our approach is entirely objective, relying only on the structure of the provided texts, and does not require subjective input from researchers. Third, many paragraphs in the FOMC minutes exhibit several topics without a dominant topic. In this case researchers would have difficulty manually identifying the proportion of each topic, while our algorithm outputs the proportion directly, enabling us to compute a unified topic-content score for each minute.

Furthermore, we provide a model-free alternative of time-varying monetary policy. Structural models such as Ang, Boivin, Dong, and Loo-Kung (2002), Campbell, Pflueger, and Viceira (2013) and Sims and Zha (2006) usually posit the existence of latent policy "regimes" beyond the observable data, and estimate such regimes in a structural VAR setting. However, the specific mechanism from which policies are generated depends on the underlying model supplied by the researcher, which can be subjective. By contrast, our approach directly outputs the economic and policy contents from the texts of FOMC minutes. Our Policy Score series can be interacted with any identifiable economic variables, thereby explicitly generating "latent" states such as policy *tone*, *aggressiveness*, or *uncertainty*, etc. Therefore, our text-based measure nicely complements the interest-rate-based structural models by providing additional rich data moments.

The rest of the paper is organized as follows. Section 2 describes our sample and data sources. Section 3 introduces our automated, topic-based content analysis methodology. Section 4 reports the results of our empirical tests and explores the sources of predictive power of our measures. Section 5 concludes.

[2]We manually select 50 paragraphs and employ 10 research assistants to classify them manually into our topic collection and to identify the topic mixture. On average the algorithm agrees with human researchers in 46 out of 50 cases. See Section 3.2 for details.

2. Data

2.1. Introduction of FOMC meetings and minutes

This subsection provides a brief overview of the logistic details of FOMC meetings and the release of the meeting minutes. From the early 1980s, the FOMC holds eight regularly scheduled meetings per year, during which members discuss the economic outlook and formulate monetary policy. Any policy change decided at the meeting is immediately implemented through open market operations. Prior to 1994, no public announcement about policy was made and the market inferred any policy change through the size and direction of the open market operations on the next day. Starting from January 1994, specific policy changes were made public in a short statement released immediately after the meeting.

Moreover, during each meeting, detailed records of the discussions are kept, then summarized in the form of Meeting Minutes, which are released to the public after a delay.[3] The minutes contain no new information received between the meeting date and the release date, and instead serve as an overview of the members' internal discussions on their economic outlook, as well as a nuanced explanation of the rationale for any policy change.

The meeting minutes follow a highly structured writing style. They are routinely consisted of four major sections. The first section outlines the administrative detail of the meeting and reviews previous open market operations. The second section provides the staff's review and outlook of the economic and financial situation, prepared in advance of the meeting. The next two sections provide the bulk of the economic content: the third section details the FOMC members' discussion of the current economic and financial situation, as well as their own economic outlook and projections. The last section is mostly related to policy and discusses the rationale for current policy and outlook for future policies. We

[3]The delay ranges between three and eight weeks. The Fed implemented a series of accelerated release schedule during the 1990s and 2000s, which shortened the lag from eight from three weeks.

remove the first section prior to processing the documents since it is unlikely to contain any economically meaningful content.

2.2. The FOMC minutes sample

We download all FOMC meeting minutes between June 1991 and December 2012 from FOMC's web site using a customized web crawling algorithm, and obtain all textual data from the PDF documents using a text extraction engine.[4] We record the date of the meeting and the date and time for the release of each minute. Our sample consists of 176 meeting minutes (thereafter referred to as Documents).

For each Document, we use a proprietary parsing algorithm to simultaneously achieve the following: 1) remove the introductory section of the Document that lists participant names and administrative matters, and remove the section on specific open market operations (e.g. amount of securities purchased); 2) break the document into individual paragraphs; 3) record the document section that the paragraph is located, and, 5) obtain paragraph length in the number of words. This procedure produces 24,793 unique sentences and 4,784 paragraphs. The average sentence length is 28 words.

2.3. Market reaction data

In many of our tests, we use high-frequency trading data from both equity and bond markets in order to measure market reactions to the contents of the minutes. For the equity market, we use tick-by-tick trading data from the SPDR exchange-traded fund by State Street to proxy for the overall level of stock market response. The SPDR, launched in 1993, follows the S&P 500 index with negligible tracking error. Trading volume has increased dramatically since 2000, making SPDR one of the most liquid stocks. Since volume prior to 2000 is low, we restrict our sample period from 2000 to 2012.

[4]Minutes downloaded in PDF at http://www.federalreserve.gov/monetarypolicy/fomccalendars.htm

Next, we construct our event window around the time when the meeting minutes are released. We then calculate return volatility during the event window. After 1997, the official release time for the meeting minutes is 2:00pm Eastern Standard Time. However, it is possible that some minutes are released early or late. As such, for each scheduled minutes release day, we use an automated algorithm to simultaneously search the FOMC's official web site, Bloomberg, and Factiva, and record the earliest time that the minutes (or news about the minutes) are reported among the three sources. The actual release time ranges between 2:00pm and 2:11pm. Therefore, we construct our event window as the 15-minute window between 2:00pm and 2:15pm each day, such that it fully encompasses the reported release times. Our result is robust to alternative event window specifications. The results are also similar using windows ranging from 15 minutes to two hours.

We calculate event-window return, and then compute raw return volatility as the squared event-window return. Specifically, for each minute t in our sample,

$$R_t = \frac{P_{t,2:15pm} - P_{t,2:00pm}}{P_{t,2:00pm}} \tag{1a}$$

$$V_t = R_t^2 = \left(\frac{P_{t,2:15pm} - P_{t,2:00pm}}{P_{t,2:00pm}} \right)^2 \tag{1b}$$

Because we use a very short, 15-minute window in constructing the market volatility measure, confounding effects from other macroeconomic variables are negligible, as no other important economic data is announced during window, and very rarely during the same day. To further ensure that any volatility change during our short event window is solely a contemporaneous response to the minutes' release, rather than a delayed response to other macroeconomic events, we separate the event window volatility into an expected and unexpected part. Specifically, we compute the unexpected volatility on the release day as

the difference between the raw volatility and the average event window volatility, computed per Equation (1b), in the past k trading days:

$$UV_{t,k} = V_t - \sum_{j=1}^{k} \frac{V_{t-j}}{k} \tag{2}$$

In general we set k between 5 and 30 trading days. Most results in our Tables are reported using $k = 20$ days. The results are little changed when k is set to 5, 10, or 30 trading days. We therefore omit the k-subscript and instead use the notation in UV_t subsequent discussions.

3. Methodology

Because each Minute is a summary of everything that is discussed during the preceding meeting, it is a mixture of a wide range of topics: The actual Minute examples in Appendix B demonstrate that, while one paragraph discusses the latest developments on inflation, another paragraph might provide outlook on financial markets. Another paragraph might discuss both. Several complications arise from these multi-faceted texts: First, which discussions are informative and which are not? Second, many words have different connotations under different contexts. For example, *increase* is considered a positive word in the economic growth context, but has negative connotations in the inflation context. How do we separate these contexts? Third, the proportion and content of discussions on each topic are likely to vary from meeting to meeting. How should one accurately measure these proportions?

These are our motivation for using a topic-based approach that isolates the content of each topic prior to content extraction. This approach allows us to address the above concerns simultaneously by 1) on the paragraph level, accurately gauging the context of each paragraph, and 2) on the document level, obtaining a granular measure of time-varying

content proportions that human readers cannot accurately identify. Overall, our approach adds another dimension that enhances traditional content analysis. This section describes our methodology to separate the FOMC minutes into individual topics and extract the content from each topic.

3.1. The Latent Dirichlet Allocation (LDA) algorithm

We first classify each Minute into distinct topics with the Latent Dirichlet Allocation (LDA) algorithm first developed by Blei et al. (2003), which belongs to a broader class of probabilistic topic models that use hierarchical Bayesian analysis to uncover the underlying semantic structure of textual documents. The common intuition behind such topic models can be summarized by two statistical distributions, which constitute the latent data generating process: The base unit of our analysis is a paragraph. Each paragraph is sufficiently summarized as a distribution over a collection of topics, each of which is, in turn, a distribution over the collection of English words used in the sample texts. For example, a paragraph that discusses inflation should be represented by a distribution that places a high weight on a topic that places high weights on words such as *prices, CPI, inflation*, etc. By contrast, a topic that places high weights on *foreign trade* and *imports* should receive a low weight in this paragraph distribution.

However, the two distributions are unobservable from the point of the researcher. The advantage of probabilistic topic models is that, using Bayesian techniques, such models efficiently infer the hidden distributional properties from the observable data (i.e. the collection of documents). LDA represents one particular parameterization of the model: We assume that these two latent distributions belong to the Dirichlet family. Then, armed with this functional form and the observed words in each paragraph, we compute the posterior (i.e. empirical) paragraph and topic distributions using the standard Bayes Theorem. These empirical distributions are the main outputs of the model. The only inputs in LDA are the

document texts and the number of topics. As such, compared to a manual classification approach, researcher-induced subjectivity and bias are minimized.

We illustrate our approach with a simple example. Suppose that the full set of revelant FOMC vocabulary consists of only $V = 4$ words (ignore common words such as *we*, *the*, etc): {*employment*, *layoff*, *imports*, *trade*}. We are given $D = 3$ paragraphs:

1. *Employment situation is good and layoff has declined.*
2. *Imports have increased and the outlook for trade is good.*
3. *Imports look good, and employment situation is also good.*

A human reader would intuitively recognize that the first paragraph is about employment and the second is about foreign trade. The third paragraph is a mixture of both. Suppose we fit the LDA model with $N = 2$ topics. If the model performs satisfactorily, then first, the posterior topic distributions should clearly and intuitively identify the topics and thus be something similar to:

- $\hat{\beta}_1 \equiv \{\hat{P}_{topic1}(employment), \hat{P}_{topic1}(layoff), \hat{P}_{topic1}(imports), \hat{P}_{topic1}(trade)\}$
  $= \{0.55, 0.43, 0.01, 0.01\}$
- $\hat{\beta}_2 \equiv \{\hat{P}_{topic2}(employment), \hat{P}_{topic2}(layoff), \hat{P}_{topic2}(imports), \hat{P}_{topic2}(trade)\}$
  $= \{0.01, 0.01, 0.60, 0.48\}$

Next, the posterior topic mixture in each paragraph should correspond to the human reader's intuition:

- $\hat{\theta}_1 \equiv \{\hat{P}_{paragraph1}(Topic1), \hat{P}_{paragraph1}(Topic2)\} = \{0.99, 0.01\}$
- $\hat{\theta}_2 \equiv \{\hat{P}_{paragraph2}(Topic1), \hat{P}_{paragraph2}(Topic2)\} = \{0.01, 0.99\}$
- $\hat{\theta}_3 \equiv \{\hat{P}_{paragraph3}(Topic1), \hat{P}_{paragraph3}(Topic2)\} = \{0.51, 0.49\}$

We proceed with our LDA classification of the FOMC minutes simply by generalizing this example to our sample of $D = 4,784$ unique paragraphs. This set of paragraphs becomes our document collection and our input to the LDA algorithm. Stop words, such as *a, the*, etc., and the most common 100 words from the documents, are removed prior to processing. Consistent with Jegadeesh and Wu (2013), we reduce each word in the document back to its root form. This procedure results in $V = 3,565$ unique words.

Next, we hypothesize that there are $N = 8$ unique topics in the document. Our results are robust to alternative specifications from $N = 5$ to $N = 10$.[5] This is the only manual step in the entire process. Here, each of the N topics represents a distribution over the V words in the FOMC vocabulary, and each paragraph is a mixture of the N topics. We assume that the observable data, i.e. words in each document, is generated from a probabilistic data generating process parameterized as follows:

1. Each of paragraph $d = 1, \ldots, D$ contains a mixture of N topics. Let the proportion of topic n in paragraph d be $\theta_{d,n}$ and let the vector $\theta_d = [\theta_{d,1}, \ldots, \theta_{d,N}]'$ represent the true topic mixture of paragraph d. For each d, we assume that this mixture follows an order-N Dirichlet distribution over the N topics, governed by the latent, parameter vector μ of size N:

$$\theta_d \sim Dirichlet_N(\mu)$$

2. Given paragraph d's topic mixture θ_d, let the assignment of each word i in document d into topics be $Z_{d,i}$, where $Z_{d,i} \in \{1, \ldots, N\}$. We assume that this assignment follows the multinomial distribution governed by the document-specific topic vector θ_d described in the previous step:

$$Z_{d,i}|\theta_d \sim Multinomial(\theta_d) \tag{3}$$

[5]Because each FOMC minute contains at least four sections, it is likely that $N \geq 5$. When the number of topics increase, some topics become redundant. However, the algorithm results in a similar number of major topics after grouping similar topics as discussed below.

Suppose there are I_d unique words in document d. Let the vector Z_d denote the collection of the topic assignment of all words within d, i.e. $Z_d = \{Z_{d,i}\}_{i=1}^{I_d}$

3. The N topic distributions (applied universally to all paragraphs) are in the collection $\beta = \{\beta_1, \ldots, \beta_N\}$. Each topic β_n also follows an order-V Dirichlet distribution over the V words, governed by the latent scalar parameter ϕ:

$$\beta_n \sim Dirichlet_V(\phi) \tag{4}$$

4. For each word i in document d, there are V choices to choose from based on our FOMC vocabulary. Conditional on the chosen topic for word i in Step 2 above (i.e. a draw from Distribution (3)), and on the structure of the topic distribution from Step 3 (i.e. a draw from Distribution (4)), we assume that actual choice of the word, $W_{d,i}$, follows a multinomial distribution governed by the resulting word-topic assignment $\beta_{Z_{d,i}}$:

$$W_{d,i} | \left(\{\beta_n\}_{n=1}^{N}, Z_{d,i} \right) \sim Multinomial(\beta_{Z_{d,i}})$$

Similarly, let the W_d denote the collection of the vocabulary choice of all words within document d: $W_d = \{W_{d,i}\}_{i=1}^{I_d}$

The above four distributions constitute the latent data generating process that results in our observable document collection $\{W_d\}_{d=1}^{D}$. Recall that they are not directly observable to the researcher. Instead, the only observable data is the occurrence of the actual words i in each document d, i.e. W_d. We can then write the overall data generating process as the

joint distribution of latent variables $\{\beta_n\}_{n=1}^N, \{\theta_d\}_{d=1}^D, \{Z_d\}_{d=1}^D$ and the observable variable $\{W_d\}_{d=1}^D$:

$$
\begin{aligned}
&P\left(\{\beta_n\}_{n=1}^N, \{\theta_d\}_{d=1}^D, \{Z_d\}_{d=1}^D, \{W_d\}_{d=1}^D\right) \\
&= \prod_{n=1}^{N} P(\beta_n) \prod_{d=1}^{D} P(\theta_d) \left[\prod_{i=1}^{I_d} P(Z_{d,i}|\theta_d) P\left(W_{d,i}|\{\beta_n\}_{n=1}^N, Z_{d,i}\right)\right]
\end{aligned}
$$

where $P(\cdot)$ are the respective (Dirichlet or multinomial) density functions specified above.

Now that we observe our FOMC document collection $\{W_d\}_{d=1}^D$, we can compute the posterior distribution of the document-topic structure given the observed documents using Bayes' Rule:

$$
P\left(\{\beta_n\}_{n=1}^N, \{\theta_d\}_{d=1}^D, \{Z_d\}_{d=1}^D | \{W_d\}_{d=1}^D\right) = \frac{P\left(\{\beta_n\}_{n=1}^N, \{\theta_d\}_{d=1}^D, \{Z_d\}_{d=1}^D, \{W_d\}_{d=1}^D\right)}{P\left(\{W_d\}_{d=1}^D\right)}. \quad (5)
$$

Similar to other Bayesian inference methods, the numerator in Equation (5) and can be easily computed. The denominator is by construction a double integral and therefore cannot be feasibly computed. However, it can be efficiently approximated using a Gibbs sampler. We use a customized Gibbs sampler for fast implementation and defer the technical aspects of the Bayesian inference to the Online Appendix.

3.2. Results from the LDA inference

Once the posterior probabilities are computed, we compute the posterior expectations of two key latent variables, which represent the main output from the LDA algorithm:

1. Posterior vocabulary distribution for each topic: $\{\hat{\beta}_1, \ldots, \hat{\beta}_N\}$
2. Posterior topic mixture for each paragraph in our collection: $\{\hat{\theta}_1, \ldots, \hat{\theta}_D\}$

The first set of output from our LDA procedure identifies the topics. For each topic k, $\hat{\beta}_k = [\hat{\beta}_{k,1}, \ldots, \hat{\beta}_{k,V}]'$, and each entry $\hat{\beta}_{k,j}$ represents the *probability that the word j characterizes topic k*. Our FOMC document collection has $V = 3,565$ unique terms. As a result, each $\hat{\beta}_k$ contains 3,565 entries. Panel A of Table 1 reports the top 20 words for each topic. Panel B of the same table reports the top 20 N-grams (phrases) for each topic, which is a result of a post-estimation likelihood testing procedure that we discuss in the Online Appendix. The results show that the topics are clearly identified. First, the top words from each topic reported Panel A are mostly distinct and identify their respective topics with almost no ambiguity. For example, the topic 1) consists of keywords such as *consumer*, *sales*, *spending*, etc., indicating that this topic is about domestic consumption. Topic 4) consists of keywords such as *energy*, *core*, etc., suggesting that this topic is about inflation. Similarly, the rest of the topics can be easily identified as related to 2) employment, 3) investment, 5) policy, 6) foreign trade, 7) financial markets, and 8) economic growth.

Because several of the identified topics are individual components of economic growth, to facilitate economic interpretation and to better map the topics to specific Fed mandates, we perform one last post-estimation task and group all growth-related topics into one topic. Specifically, we group the Consumption, Investment, Foreign Trade, Employment and Growth topics into one *Growth Mandate* topic. This corresponds to the Fed's mandate of optimal employment (or optimal growth). The rest of the topics can also be pinpointed to specific Fed mandates and tasks. The *Inflation Mandate* topic corresponds to the Fed's second mandate of stable prices. The *Financial Market Mandate* topic corresponds to the Fed's increasingly important role as a regulator of financial markets. Finally, the *Policy* topic is relatively orthogonal to other topics and specifically addresses the plan, performance, and outlook of monetary policies. We use this topic more extensively in a follow-up research.

The second set of output is the set of paragraph-level topic mixture vectors, $\{\hat{\theta}_1, \ldots, \hat{\theta}_D\}$. From this collection, each paragraph d has one mixture, $\hat{\theta}_d = [\hat{\theta}_{d,1}, \ldots, \hat{\theta}_{d,N}]'$. Our grouping

procedure in the previous paragraph results in 4 major topics. As a result, each vector $\hat{\theta}_d$ has 4 entries, where each $\hat{\theta}_{d,n}$ corresponds to the *proportion* of paragraph d that is devoted to topic n. The four entries sum up to one for each paragraph. We plot the time series of the proportion of each topic in Figure 1. The shaded area in Figure 1 corresponds to NBER-designated recession periods.

Interestingly, Figure 1 shows significant time variation in the proportion of the FOMC minutes devoted to each topic. For example, from 1991 to 2011, an increasingly smaller proportion of the minutes has been devoted to the *Growth Mandate* topic. This is particularly true during recessions, when the proportion drops significantly. This decrease has been offset by increases in the proportions of the other three topics, particularly the *Financial Markets Mandate* topic: during the latest recession, its proportion more than tripled. This likely reflects the Fed's increasingly important regulatory role in maintaining the stability of the financial markets, such as negotiating the rescue of systematically important banks and the subsequent TARP initiatives. Overall, this table demonstrates that the FOMC minutes are not uniformly-written documents that always address one particular issue, but a compendium of discussions on various issues, whose proportion change continuously over time. This highlights the importance of our topic-based approach.

Finally, as an additional robustness check, we randomly select 50 paragraphs from each of the two groups that satisfy the following properties:

1. Paragraphs classified as containing $\geq 99\%$ of a single topic.
2. Paragraphs classified as containing a mixture of two, three and four topics. (each topic having a proportion of at least $\geq 10\%$)

A selection of the texts are presented verbatim in Appendix A. We then ask a team of 10 human readers to identify the topic mixtures of these 100 paragraphs, without revealing the LDA classification result. For paragraghs that are identified by the LDA as containing

only a single topic, human readers and the LDA agree in 49 of the 50 paragraphs (e.g. they both identify a paragraph into the Growth Mandate topic). For multiple-topic paragraphs, human readers agree with the LDA in 46 of the 50 paragraphs about the number and type of the topic. However, they often have difficulties pinning down the exact proportions of each topic, especially when the number of topics is higher than three. For example, many readers identify the last paragraph of Appendix A as containing 25% of each topic, whereas the LDA offers a different proportion mix that is potentially accurate. This test demonstrates two advantages of our automated topic classification approach. First, it offers an accurate topic classification that is consistent with common intuition. Second, for each document, it offers a granular, time-varying topic mixture that is more accurate than manual reading, therefore minimizing researcher-induced bias.

3.3. Extraction of contents

Having obtained the estimate of topic proportions, we now proceed to extract the contents for each paragraph-topic combination, using a bag-of-words approach similar to Tetlock (2007) and Jegadeesh and Wu (2013). Specifically, for each paragraph, we simultaneously extract the tone and uncertainty of each topic by tabulating the frequency of keywords in the respective tone and uncertainty lexicons. The Tone content consists of the frequency of positive words (*positive tone*), negative words (*negative tone*), and the difference in frequency between positive and negative tonal words (*net tone*) in a comprehensive tonal lexicon that merges the Harvard IV-4 Psychosociological Dictionary[6] and the financial tonal lists developed by Loughran and McDonald (2011). The *uncertainty* content is the frequency of keywords in the "uncertain words" lexicon developed by Loughran and McDonald (2011).[7]

[6] Available at http://www.wjh.harvard.edu/∼inquirer/homecat.htm.
[7] Available at http://www3.nd.edu/∼mcdonald/Word_Lists.html.

Since there are four topics for each paragraph, the topic contents can be summarized in four paragraph-level content *Scores*. Specifically, for meeting t, paragraph d, topic $n = 1, \ldots, 4$, and content $c \in \{positivetone, negativetone, nettone, uncertainty\}$:

$$Score^t_{d,n,c} = \hat{\theta}^t_{d,n} F^t_{d,c}$$

where $\hat{\theta}^t_{d,n}$ is the topic-n proportion estimate for paragraph d from LDA, and $F^t_{d,c}$ is the number of occurrences of content words from the respective tonal or uncertainty lexicons in paragraph d. In addition, we isolate a list of words that are associated with quantity increases and decreases.[8] We reverse the connotation for these words when they are used in the Inflation Mandate topic. For example, *gain* is considered as a positive word by both lexicons. However, it should be treated as a negative word when used in the Inflation Mandate context.[9] In general, a higher Tone Score indicates a more positive or less negative tone, and a higher uncertainty Score reflects a higher degree of uncertainty in the paragraph.

Next, similar to Jegadeesh and Wu (2013), we aggregate the Scores to the document level as the sum of individual paragraph scores, weighted by the inverse of paragraph length in number of words:

$$Score^t_{n,c} = \sum_{d=1}^{D_t} Score^t_{d,n,c} \left(\frac{1}{T^t_d} \right), \tag{6}$$

where T^t_d is the total number of words in paragraph d, and D_t is the total number of paragraphs in Document t. The term $\left(\frac{1}{T^t_d} \right)$ reflects the intuitive notion that the strength of the topic tone is negatively related to overall paragraph length. Longer paragraphs are more difficult to read and process, and are therefore downweighted.

[8] Available at http://fnce.wharton.upenn.edu/profile/1661/.

[9] Because a *gain* in inflation is considered negative by the Fed and thus increases the likelihood of tightening actions.

Figure 2 plots the 10-period moving average of the document-level Net Tone and Uncertainty Scores of each of the four topics over time. For ease of comparison, the Scores are standardized to a mean of zero and standard deviation of one. This figure demonstrate the large difference between the topic contents over time. For example, the Growth Mandate Score is procyclical, becoming more positive during boom periods and turning sharply negative during recessions. This is probably not surprising because FOMC members' discussions on this topic is likely based on their review and outlook of the underlying economic conditions, which are likely to be quite persistent. On the other hand, the Financial Market Mandate Score is clearly countercyclical: it is more negative during boom periods and turns sharply positive about half way into recessions, especially during the latest recession from 2008-09. This pattern probably indicates the Fed's increasing level of concern about the stability of financial markets during boom periods: the level of speculative trading and exotic financial products and instruments (e.g. dot-com stocks, subprime mortgages and derivative contracts) usually increase toward the end of booms, and as a result, the level of systematic risk in the financial markets increases. The Fed, as a market regulator, likely foresees this risk and conveys its caution via Minute languages. Similarly, during busts, the Fed is instrumental in maintaining financial stability, and thus voices its commitment via more positive languages related to financial markets. This finding further suggests that the discussions on some topics are probably more informative than others. Our topic-based approach therefore can highlight the informative topics and construct content indices that are deemed important by the market, and therefore useful for predicting future economic conditions and policies. The rest of our paper is focused on assessing this ability.

4. Empirical tests and results

This section discusses our empirical tests and reviews the test results. We first assess the informativeness of the minutes as a whole, and then measure the relative informativeness

of each individual topic and of its content using market reaction. We then relate our topic-specific content Scores to macroeconomic variables and explore the determinants of the Scores. Finally, we explore the source of informativeness from the minutes by analyzing whether the price reactions to the minutes are permanent.

4.1. Informativeness of FOMC minutes as a whole

We first measure the amount of incremental information contained in the texts of the Minutes in addition to the "hard data" released immediately after the original meetings, such as interest rate changes. Instead of subjectively assess the Minutes' informativeness, we use market reaction as a natural and objective measure. Specifically, to the extent that the release of the Minute is deemed informative by the market, it should "move the market", i.e. positively correlated with market volatilities around the time when they are released. We specify the following relation between minutes' release and market reaction:

$$V_t = a + bL_t + e_t \tag{7a}$$

$$V_t = \alpha + \beta L_t + \sum_{j=1}^{k} \gamma_j V_{t-j} + \epsilon_t \tag{7b}$$

where V_t is the 15-minute event window market volatility (on both release and non-release days) computed per Equation (1b), L_t is a dummy variable that equals to one if a Minute is released at date t. Each regression uses approximately 3,700 days of observation.

We also use a more robust specification in Regression (7b). Because we use a very short window in constructing the market volatility measure, confounding effects from other macroeconomic variables are negligible, as no other important economic data is likely announced during this window. However, it is still possible for the coefficient estimates to be positively biased. Consider the following scenario: suppose a Minute is completely

uninformative, but the market expectation about the content of the minutes can be erroneously distorted between the meeting date and the release date, e.g. by interim speeches from other Fed officials. Therefore, when the Minute is released, the market corrects its wrong expectation, thereby registering a higher than normal volatility. This produces a positive bias on the coefficient estimates of b in Regression (7a). However, because interim changes in expectation are also associated with changes in market volatility, we can use the volatility of the k-days between the meeting date and the release date to control for the effect of changing expectations. The estimate for β therefore measures the true level of informativeness of the minutes, conditional on all prior expectations.

We fit Regression (7b) above using k=0, 5, 10, and 20 trading days. As an additional robustness check, we also extend the release window to 20 minutes from 1:55pm to 2:15pm. To facilitate interpretation, we scale all regression coefficients by the unconditional mean of V_t across all observations. The coefficient estimates $\hat{b}$ and $\hat{\beta}$ can thus be interpreted as the *incremental* volatility introduced by the release of the minutes as a percentage of the average volatility in the event window across both release and non-release days.

Table 2 reports the coefficient estimates. The estimate for the release dummy, L_t, is significantly positive for all specifications, ranging between 0.692 and 0.925. This indicates that the minutes as a whole are deemed to be highly informative by the market, because the release days experience volatilities that are between 69.2% and 92.5% higher than normal. Both Fleming and Piazzesi (2005) and Lucca and Moench (2013) find that on the actual FOMC meeting days, market volatilities are significantly higher. By contrast, the minutes are released several weeks after the original meetings. Our finding that market volatility is also significantly higher on release days thus suggests that the language of the FOMC minutes does contain additional, "soft" information not incorporated in the quantitative policy announcements. Such information can be viewed as corresponding to the latent policy regimes identified in structural models such as Sims and Zha (2006).

Next we examine the informativeness of the specific contents of the Minutes. We compute the Tone and Uncertainty Scores on the document level, but without LDA processing, and relate them to unexpected market volatilities in the following regression:

$$UV_t = \alpha + \beta_c Score_c^t + \epsilon_t \tag{8}$$

where UV_t is the unexpected volatility around the event window on release date t, computed per Equation (2), adjusted using volatility in the previous $k = 20$ trading days. $Score_c^t \in \{Score_{nettone}^t, Score_{uncertainty}^t\}$ are the document-level net tone and uncertainty Scores, computed per Equation (6) while setting all $\hat{\theta}$'s equal to one. Each regression uses 176 observations corresponding to the minutes' release dates.

The above regression explore the relation between the overall document tone/uncertainty and market reaction. If, for example, the market interprets a more positive overall tone as more informative, then it would respond more to a more positively-written document, resulting in higher volatility and a positive estimate for $\beta_{nettone}$. Similarly, more ambiguous documents have higher uncertainty Scores, which might not elicit a pronounced market response because they do not contain as much incremental information.

Table 3 reports the regression estimates. None of the coefficient estimates for the document-level content Scores are statistically different from zero, with t-stats ranging between -1.52 and 0.80. This suggests that, on the document level, the market does not perceive the tone Scores as useful, as neither more positive nor more negative tone Scores are associated with higher market volatility. This finding further indicates the usefulness of our granular, topic-based approach: if some individual topics are consistently more informative than others, then variations in individual topic contents can be associated with changes in volatility not detected on the document level.

4.2. Relative informativeness of individual FOMC topics

The results in the previous subsection indicate that, although the Minutes do contain incremental information, the content Scores that we compute for the minutes as a whole are not as informative. This is likely because the individual topics within the FOMC minutes are not *equally informative*: the market might deem some topics more valuable than others. For example, our discussions with industry practitioners reveal that they consistently find the discussion on inflation to be more informative than that on economic growth, which they find containing more "stale" information. As such, some topics in the FOMC minutes should receive more market attention than others. Therefore, similar to the argument in the previous subsection, the informativeness of discussions on each topic can be measured by market reaction. We first evaluate the relative informativeness of the topics themselves using our proportion measure $\hat{\theta}$. We then assess the informativeness of each topic's contents using Tone and Uncertainty Scores. We specify the following relations between topic proportion, content Scores, and event window return volatility for each $c \in \{positive\,tone, negative\,tone, net\,tone, uncertainty\}$:

$$UV_t = a + \sum_{n=1}^{4} b_n \hat{\theta}_{n,t} + rX_t + e_t \tag{9a}$$

$$UV_t = \alpha + \sum_{n=1}^{4} \beta_n Score_{c,n}^t + \gamma X_t + \epsilon_t \tag{9b}$$

where UV_t is the unexpected volatility around the event window on release date t, computed per Equation (2), and X_t is the vector of macro controls variables that include:

- $IntRate$: the latest daily closing yield of 10-year Treasury notes obtained from the Department of Treasury.

- *UnEmp*: latest monthly rate of unemployment obtained from the Bureau of Labor Statistics.
- *Recession*: a dummy variable which equal to one if meeting date t falls within a NBER-designated recession period.

Each regression uses 118 daily observations from 2000 to 2012. In this setting, an estimate of b_n or $\beta_{c,n}$ that is statistically different from zero indicates informativeness of a topic, or its content Score: a significantly positive $\hat{\beta}$ suggests that the market respond more to a more positive topic tone while a significantly negative $\hat{\beta}$ suggests that the market find more information in a more negative topic tone. Similarly, a significantly positive $\hat{b}$ for topic n indicates that the market finds the discussion of this topic informative regardless of the tone.

Table 4 displays the coefficient estimates from the proportion Regression (9a). The Financial Market Mandate topic is omitted from the regression to prevent multicolinearity. Tables 5 to 7 report the coefficient estimates for the tone Scores in Regression (9b). All independent variables in the regressions are standardized to mean zero and unit standar ddeviation. First, perhaps not surprisingly, the coefficient estimates for the Growth, Inflation, and Policy proportions are all statistically significant and positive. This is consistent with the findings in Table 2 that the minutes do contain additional information, and this is true for the topic level as well: as the discussion on each topic becomes more detailed (thus higher proportions), more information is transcribed in the texts, and therefore the market responds more to the topic.

More interesting are the results from the content Score regressions, as they suggest that the market indeed views different topics differently and assigns different informational value to them. First, the coefficient estimates for the Growth Mandate's Net tone Score and Positive Score are significantly positive, indicating that the market deems more positive growth-related discussions as more informative. For the Inflation Mandate topic, the market

finds both positive and negative discussions informative, as the estimates for both tone Scores are significantly positive. This is intuitive because, compared to the broader Growth topic, the more specialized Inflation topic is more relevant to upcoming Fed actions: a positive discussion might suggest upcoming tightening actions, while a more negative tone indicates bad current economic conditions but perhaps more easing actions in the future.

Next, the coefficient estimates for the Net and Positive tone Scores of the Financial Market Mandate topic is significantly negative while the estimate for its Negative Score is significantly positive. This suggests that the market finds more negative discussions of financial markets as more informative. This result is also consistent with our intuition because the Financial Market topic is countercyclical. If during good times, the market sees more negative market-related discussions (e.g. cautions on speculative trading and exotic products), then it might view this as a signal of impending market top. During bad times, the Fed is probably already expected to help the financial markets, therefore a more positive discussion might be priced in. For the Policy topic, the only statistically significant estimate is the Negative Score. This is consistent with the notion that the market is more afraid of unanticipated tightening actions than an easing policy stance, therefore it assigns a higher information value to such discussions. The finding is also consistent with the practitioners' practice of keeping a more vigilant lookout for surprising rate hikes, as those are more unpredictable, and treating rate reductions during bad times as more or less given.

Our next set of tests explore the *directional* nature of our topic Scores. For example, does the market really think a negative discussion on inflation as a negative thing? Not necessarily. After all, such discussion could indicate bad current conditions, but also signal future easing actions. As such, the informativeness of the topics is reflected by not only by

market volatility, but also from the relation between the Scores and raw returns during the event window. We therefore fit the following regression:

$$R_t = \alpha + \sum_{n=1}^{4} \beta_n Score_{c,n}^t + \gamma X_t + \epsilon_t, \tag{10}$$

where R_t is the 15-minute event-window market returns constructed according to Equation (1a) and X_t is the vector of controls used in the previous regression.

This regression explores the micro relation between the topic tones and directional returns. If the market indeed thinks that a particular tone for a particular topic is good/bad news, then it should respond accordingly, resulting in a positive/negative estimate for β. Panel B of Table 5 reports the coefficient estimates for the Net Tone Scores. First, surprisingly, the estimates for the Inflation and Financial Market Mandate topics are both significantly *negative*: The market interprets a more positive tone of these topics as bad news. The market return is about 0.1% lower with a one-standard-deviation increase in the tones' net positivity. This suggests that, particularly during bad times, inflationary actions (more positive inflation discussions) or more positive discussions on financial markets might be viewed as additional signs of weakness in the economy, thus resulting in a negative market reaction. For example, the market turned sharply lower following the news of Fed-initiated bailouts of both Bear Stearns and AIG. This results in a negative estimate for $\beta_{inflation}$ and $\beta_{financial\,market}$.

More significantly, the estimate for the Policy topic is strongly positive: market return is on average 0.1% higher during the 15-minute window with a one-standard-deviation increase in the Policy topic tone Score. Here, in a fashion similar to Jegadeesh and Wu (2013) who demonstrate the incremental information value of 10-K discussions, this result suggests the additional information value of the Policy Score beyond traditional quantative indicators such as interest rates. Because all policies in the current meeting has already been

implemented many weeks ago, the fact that there is still significant reaction to the policy discussions suggests that this Score captures 1) latent *quanlitative* regimes of monetary policy not observable from interest rate decisions alone, such as policy "aggressiveness", and 2) indication of future policy moves. Therefore, our Policy Score can be viewed as a direct quantification of the latent policy regimes in structural models such as Sims and Zha (2006). This measure of the qualitative aspects of policy is therefore likely to hold significant predictive powers for both future policy and real economic variables. We explore this issue in detail in an upcoming paper.

4.3. Determinants of topic proportion and tone

How "orthogonal" are our text-based measures from existing economic indicators? After all, the Fed is likely to take into account current economic conditions when formulating monetary policies. In addition, a whole section of the Minutes is devoted to reviewing current economic conditions and providing an outlook for future conditions. Many theoretical and structural frameworks of monetary policy making, for example the Taylor rule, also stiputates that monetary policy, usually proxied by the nominal interest rate, is related to changes in economic variables such as output, inflation and unemployment. Does the Taylor rule matter when the FOMC members are in the meeting room? This subsection specifically examines the relation between the proportion and content Scores of each topic and current economic conditions. From our discussion in Subsection 3.3, we expect that the Growth and Inflation Mandate Scores to be procyclical and follow the Taylor rule. The Financial Market Mandate Score is countercyclical. Moreover, as the FOMC members have the most latitude in their policy discussions, the effect of macro variables on the Policy Score is likely to be ambiguous: if the Fed correctly anticipates economic cycles and changes policy before the cycle changes, then we might not see a significant relation between policy proportion/content

and contemporaneous macro variables. The Policy Score is therefore likely to be the most orthogonal among the topics.

We examine the determinants of the topic proportion and content Scores via the following regressions:

$$\hat{\theta}_{n,t} = a + bIntRate_t + rUnEmp_t + dRecession_t + e_t \tag{11a}$$

$$Score_n^t = \alpha + \beta IntRate_t + \gamma UnEmp_t + \delta Recession_t + \epsilon_t \tag{11b}$$

where $Score_n^t$ is the Net Tone Content Score for topic n and Minute t, computed per Eq. (6). $\hat{\theta}_{n,t}$ is the topic-n proportion in Minute t estimated using Eq. (5). We fit the regression using all 176 minute Documents from 1991 to 2012.

Table 8 presents the coefficient estimates for the proportion regressions. The proportion of the Growth Mandate topic is positively related to interest rates and negatively related to unemployment. The exact opposite relations can be found for the Inflation Mandate and Financial Market Mandate. This suggests that during bad times, the Fed is more concerned about inflation (or deflation) and the health of the financial markets, leaving economic growth as a secondary concern. Table 9 reports the coefficient estimates for the Content Score regressions. Both the Growth and Inflation Mandate topics are procyclical: their content Scores are positively related to interest rates and both topics becomes significantly negative during recessions. The Financial Market Mandate topic is highly countercyclical. Again this stresses the role of the Fed as a market regulator: during recessions, it seems to take an aggressive stance on helping the financial markets and maintaining systematic stability. Interestingly, the Policy topic is not significantly related to existing economic conditions: the coefficient estimates for none of the macro variables are statistically significant. This again highlights the fact that the policy discussions during the meetings probably incorporate factors beyond current economic conditions, and therefore, the Policy topic can serve as a leading indicator of the economy, which is corroborated by Fig. 1, where the Policy Score

usually "flips" half way into the economic cycle. We explore the predictive power of the topic Scores in an upcoming paper.

4.4. Sources of predictive power and information transmission

This section explores the reason why the minutes are informative. On the surface, this seems puzzling: the minutes are released a long time after the meetings, why would they contain any incremental information at all? However, there is an important distinction: the staff economists at the Fed and the FOMC members have access to a wide variety of confidential economic data, such as detailed records of interbank lending, that are not observable to other researchers. It is likely, therefore, that their information set is superior to that of other market participants. Due to the confidential nature of the data, they cannot disclose any quantitative facts in the minutes. However, it possible that such "inside" information influences the tone and uncertainty level of the minutes' language.

We can jointly test the above hypothesis and whether such "soft" information is transmitted to the market by observing the market reactions to the minutes: if there is new information about the future economy *and* the information is transmitted to the market through the minutes, then the effect on prices should be permanent rather than temporary, and the price "jump" on the minutes release day that we document in the last section should be persistent and not revert quickly. In other words, following the initial spike, market volatility should decline, since the newly incorporated information reduces the overall level of uncertainty.

Figure 3 plots the average minute-by-minute return volatility in 15-minute bins for both release and non-release days. First, we confirm the same pattern found in the treasury market: market volatility spikes dramatically to about 1.6 times the normal levels on days where FOMC minutes are released. More importantly, volatility quickly declines after the

minutes are released to about 20% lower than non-release days. As a result, the initial price jump at the release do not on average revert back, and prices on average stay at the new levels. This permanent "shift" in prices indicates that the overall level of uncertainty in the market is lower after the release of the minutes, and supports our hypothesis that information is indeed transmitted from the Fed to the markets through the text of the minutes.

5. Conclusion

We present a novel approach in financial content analysis to determine the informativeness of FOMC meeting minutes. This automated approach is based on the Latent Dirichlet Allocation (LDA) algorithm, which enables us to dissect minutes into distinct topics and simultaneously extract the tone and uncertainty level of each topic. In an event study setting, we use market reaction to assess the relative informativeness of each topic and find a significant relation between topic contents and market volatility. Furthermore, we find evidence consistent of the Fed possessing superior information, which is transmitted to the market through the text of the minutes.

The upcoming version of our paper adds several important extensions. First, we greatly expand the coverage of our high-frequency trading data sample, using futures data from the Chicago Mercantile Exchange in addition to the equities data. Furthermore, we show that tone and uncertainty level of the minutes have significantly better predictive power for economic and policy variables than professional forecasts.

Our measures of economic and policy outlook/uncertainty are both model-independent and robust, and can be readily applied to structural macroeconomic models, as well as reduced-form predictive models where policy uncertainty serves as an input. We are currently exploring several of such these issues.

References

Ang, A., Boivin, J., Dong, S., Loo-Kung, R., 2002. Monetary policy shifts and the term structure. Review of Economic Studies 78, 429–457.

Baker, S. R., Bloom, N., Davis, S. J., 2013. Measuring economic policy uncertainty. Working paper, University of Chicago, Chicago, IL.

Blei, D. M., Ng, A. Y., Jordan, M. I., 2003. Latent dirichlet allocation. the Journal of machine Learning research 3, 993–1022.

Campbell, J. Y., Pflueger, C., Viceira, L. M., 2013. Monetary policy drivers of bond and equity risks. Working paper, Harvard University, Cambridge, MA.

Fleming, M. J., Piazzesi, M., 2005. Monetary policy tick-by-tick. Working paper, University of Chicago, Chicago, IL.

Hanley, K. W., Hoberg, G., 2010. The information content of IPO prospectuses. Review of Financial Studies 23, 2821–2864.

Jegadeesh, N., Wu, D., 2013. Word power: A new approach for content analysis. Journal of Financial Economics 110, 712–729.

Loughran, T., McDonald, B., 2011. When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. Journal of Finance 66, 35–65.

Lucca, D. O., Moench, E., 2013. The Pre-FOMC announcement drift. Journal of Finance, Forthcoming.

Romer, C. D., Romer, D. H., 1989. Does monetary policy matter? a new test in the spirit of friedman and schwartz. In: *NBER Macroeconomics Annual 1989, Volume 4*, MIT Press, pp. 121–184.

Sims, C. A., Zha, T., 2006. Were there regime switches in U.S. monetary policy? American Economic Review 96, 54–81.

Tetlock, P. C., 2007. Giving content to investor sentiment: The role of media in the stock market. Journal of Finance 62, 1139–1168.

Appendix A Select FOMC Paragraphs with LDA Classification Results

Part 1. Single Topic Examples

Example 1. (99% growth mandate, other topics negligible)

With regard to developments and prospects in key sectors of the economy, members noted that despite further survey indications of eroding consumer confidence, consumer expenditures had strengthened in recent months after a pause earlier in the year. The pickup had featured rising sales of motor vehicles, and while the latter had slipped recently, a number of special factors such as shortages of popular models at the end of the model year and the effects of flooding in some parts of the Midwest suggested the need to withhold judgment on any downward shift in the underlying demand for motor vehicles. Tourism was reported to have strengthened considerably in many areas this summer, though there were major exceptions. As had been true for an extended period, consumer attitudes continued to be inhibited by concerns about employment opportunities, especially given further reductions in defense spending, the ongoing restructuring and related downsizing of many business operations, and the continuing efforts by business firms to limit the number of their permanent employees in order to hold down the rising costs of health care and other nonwage worker benefits. Members noted, however, that the growth in employment thus far this year, while tending to involve many low paying jobs, had greatly exceeded the rate of expansion in 1992. In the view of at least some members, appreciable further growth was likely as business firms found it increasingly difficult in an expanding economy to meet growing demands through outsourcing, temporary workers, and overtime work. Some members also noted that the newly legislated taxes on higher incomes would tend to curtail some consumer spending. The timing of that effect was uncertain; tax liabilities had already risen, but some payments on the added tax liabilities were not due until April of 1994 and 1995.

Example 2. Inflation Mandate (99% inflation mandate, other topics negligible)

The core consumer price index advanced at a faster rate in the first quarter than it had in the fourth quarter, reflecting the pass-through of higher energy prices and a leveling off of goods prices after sizable declines last year. The higher goods price inflation owed, in part, to the recent run-up in the prices of non-oil imports, energy, and other commodities. The price index for core personal consumption expenditures also rose at a faster rate in the first quarter than it had late last year. Despite the rise in inflation this year, however, the cumulative increase in the overall consumer price index for the year ending in March was somewhat less than the advance for the twelve months ending in March 2003. In the year ending in March, the increase in the price index for total personal consumption expenditures was similar to that of a year earlier. Survey measures of near-term inflation expectations edged up somewhat in March and April, but measures of longer-term expectations decreased. With regard to labor costs, average hourly earnings of production or nonsupervisory workers on private nonfarm payrolls rose notably less for the twelve months ending in March than they had in the year-earlier period. The overall increase in the employment cost index for private industry for the twelve months ending in March was about the same as that for the twelve-month period ending a year earlier, as wages and salaries decelerated and benefits accelerated.

Example 3. Financial Market Mandate (99% market mandate, other topics negligible)

Participants noted that financial markets were volatile over the intermeeting period, as investors responded to news on the European fiscal situation and the negotiations regarding the debt ceiling in the United States. However, the broad declines in stock prices and interest rates over the intermeeting period were seen as mostly reflecting the incoming data pointing to a weaker outlook for growth both in the United States and globally as well as a reduced willingness of investors to bear risk in light of the greater uncertainty about the outlook. While conditions in funding markets had tightened, it was noted that the condition of U.S. banks had strengthened in recent quarters and that the credit quality of both businesses and households had continued to improve.

Example 4. Policy Mandate (99% policy mandate, other topics negligible)

Participants discussed a number of policy tools that the Committee might employ if it decided to provide additional monetary accommodation to support a stronger economic recovery in a context of price stability. One of the policy options discussed was an extension of the period over which the Committee expected to maintain its target range for the federal funds rate at 0 to 1/4 percent. It was noted that such an extension might be particularly effective if done in conjunction with a statement indicating that a highly accommodative stance of monetary policy was likely to be maintained even as the recovery progressed. Given the uncertainty attending the economic outlook, a few participants questioned whether the conditionality of the forward guidance was sufficiently clear, and they suggested that the Committee should consider replacing the calendar date with guidance that was linked more directly to the economic factors that the Committee would consider in deciding to raise its target for the federal funds rate, or omit the forward guidance language entirely.

Part 2. Multiple Topic Examples

Example 5. (56% growth, 43% inflation)

The information reviewed at this meeting suggested that economic activity had weakened further in the opening months of the year. Production cutbacks were evident in a wide range of industries, and private payrolls had fallen markedly, especially in the goods producing sector. On the positive side, consumer confidence had rebounded sharply since the cease-fire in the Persian Gulf, retail sales and housing starts had strengthened recently, and exports had continued to expand. Broad measures of prices had slowed or contracted in January and February, but excluding energy and food prices, increases in those measures were higher than in previous months. Wage increases had moderated over the past several months.

Example 6. (83% financial market, 17% policy)

Committee members and Board members agreed that, with few exceptions, the functioning of most financial markets, including interbank markets, no longer showed significant impairment. Accordingly they agreed that the statement to be released following the meeting would indicate that the Federal Reserve would be closing the Asset-Backed Commercial Paper Money Market Mutual Fund Liquidity Facility, the Commercial Paper Funding Facility, the Primary Dealer Credit Facility, and the Term Securities Lending Facility on February 1, 2010. Committee members also agreed to announce that temporary liquidity swap arrangements between the Federal Reserve and other central banks would expire on February 1. In addition, the statement would say that amounts available through the Term Auction Facility would be scaled back further, with $50 billion of 28-day credit to be offered on February 8 and $25 billion of 28-day credit to be offered at the final auction of March 8. The statement also would note that the anticipated expiration dates for the Term Asset-Backed Securities Loan Facility remained June 30,

2010, for loans backed by new-issue commercial mortgage-backed securities, and March 31, 2010, for loans backed by all other types of collateral. Members emphasized that they were prepared to modify these plans if necessary to support financial stability and economic growth.

Example 7. (34% growth, 31% financial market, 35% policy)

Open market operations during the intermeeting period continued to be directed toward maintaining the existing degree of pressure on reserve positions. The federal funds rate rose briefly in response to year-end pressures, but it otherwise tended to remain close to the 5-1/4 percent level expected with an unchanged policy stance. Other short-term interest rates generally were unchanged to slightly higher over the intermeeting period. Rates on intermediate- and long-term securities edged higher on balance in reaction to incoming data on economic activity that were on the firm side of market expectations; the increases in such rates appeared to be tempered, however, by favorable market reactions to new data on wages and prices. The generally positive news on economic growth and inflation along with favorable reports on earnings appeared to reinforce the optimism of equity market investors, and major indexes of stock prices increased markedly further over the intermeeting period.

Example 8. (39% growth, 13% inflation, 20% financial market, 26% policy)

In their discussion of the economic situation and outlook, FOMC meeting participants indicated that the worsening financial situation, the slowdown in growth abroad, and incoming information on economic activity had led them to mark down significantly their outlook for growth. While economic activity had evidently already been slowing over the summer, the turmoil in recent weeks had apparently resulted in tighter financial conditions and greater uncertainty among businesses and households about economic prospects, further limiting their ability and willingness to make significant spending commitments. Recent measures of business and consumer sentiment had fallen to historical lows. Participants generally expected the economy to contract moderately in the second half of 2008 and the first half of 2009, and agreed that the downside risks to growth had increased. While some expected an improving financial situation to contribute to a recovery in growth by mid-2009, others judged that the period of economic weakness could persist for some time. Several participants indicated that they expected some fiscal stimulus in coming quarters, but they were uncertain about the extent and duration of the resulting support to economic activity. Participants agreed that in coming quarters inflation was likely to move down to levels consistent with price stability, reflecting the recent declines in the prices of energy and other commodities, the appreciation of the dollar, and the expected widening of margins of resource slack. Indeed, some saw a risk that over time inflation could fall below levels consistent with the Federal Reserve's dual objectives of price stability and maximum employment.

Figure 1. FOMC Topic Proportions Over Time

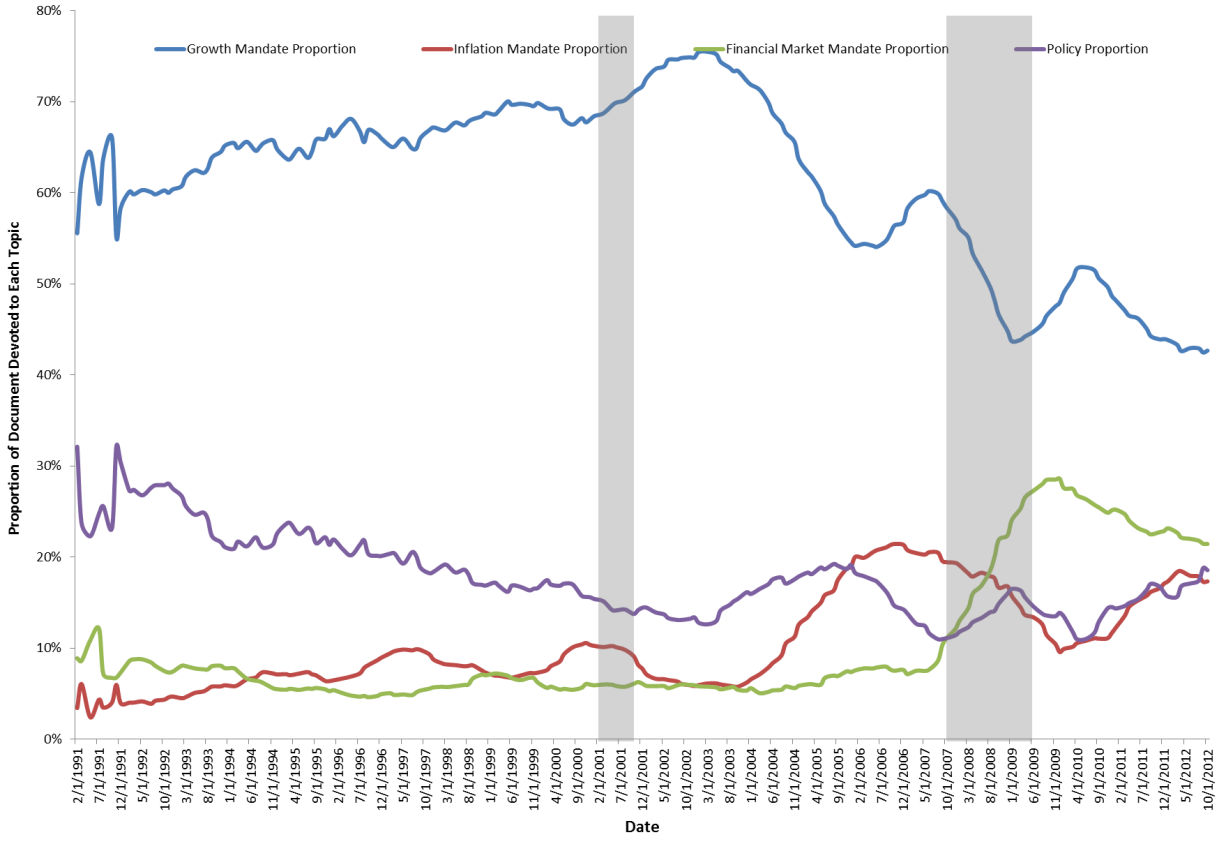

Figure 2. FOMC Topic Content Scores Over Time

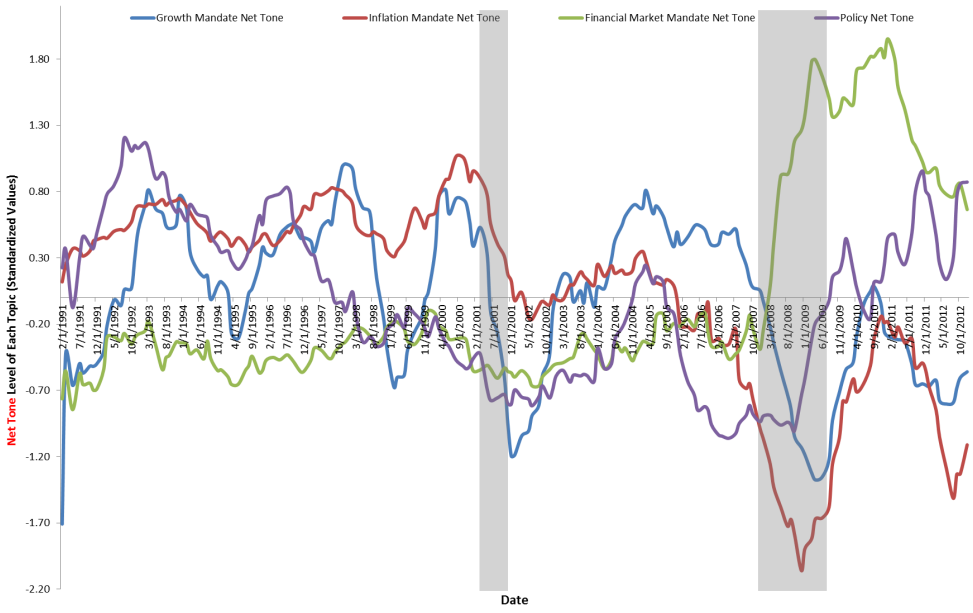

(a) Topic Net Tone Scores Over Time

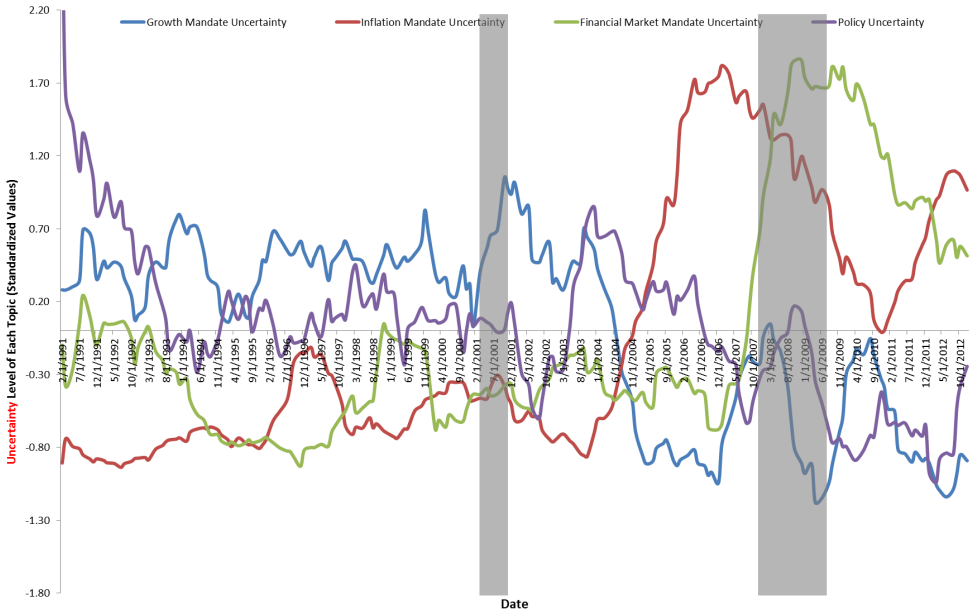

(b) Topic Uncertainty Scores Over Time

Figure 3. Market reaction to the release of FOMC minutes

This figure plots the average return volatility in 15-minute bins for both release and non-release days. The first chart shows the raw level of volatility for both series, and the second chart shows the ratio of volatility on minute release days over that on nonrelease days. Return volatility is calculated as the standard deviation of minute-by-minute returns in each 15-minute bins. The sample period is 2000-2012.

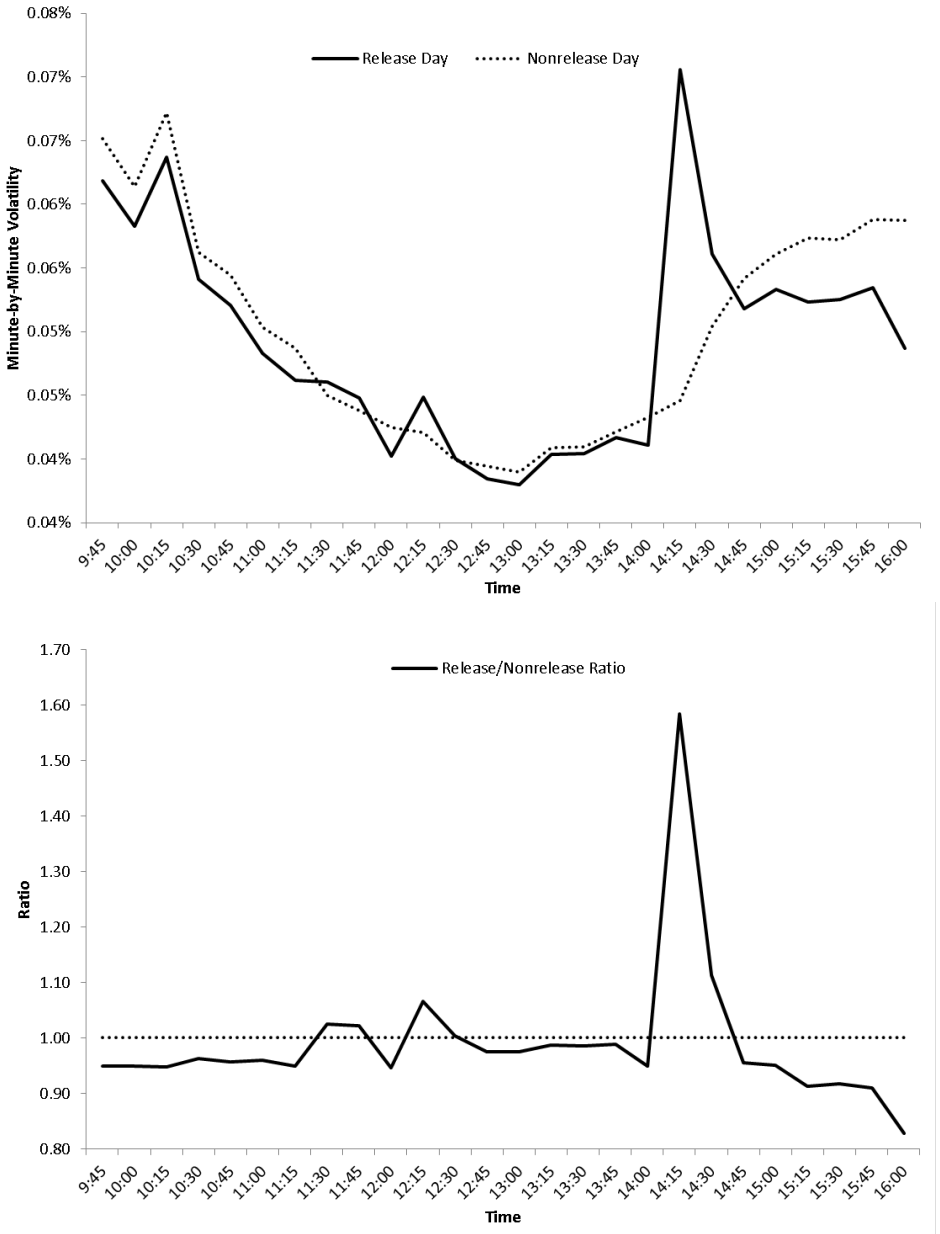

Table 1
Distribution of Top LDA Topic Keywords

Topic1		Topic2		Topic3		Topic4	
Weight	Word	Weight	Word	Weight	Word	Weight	Word
0.0503	consumer	0.0479	unemployment	0.0333	business	0.0408	energy
0.0410	sales	0.0474	employment	0.0309	equipment	0.0379	consumer
0.0362	housing	0.0242	labor	0.0263	production	0.0312	core
0.0173	homes	0.0184	job	0.0221	sales	0.0232	labor
0.0152	mortgage	0.0164	workers	0.0218	inventories	0.0191	expectations
0.0148	starts	0.0163	projections	0.0206	inventory	0.0158	food
0.0146	home	0.0156	payroll	0.0204	manufacturing	0.0140	costs
0.0146	real	0.0155	nonfarm	0.0193	investment	0.0118	compensation
0.0141	income	0.0141	private	0.0183	construction	0.0118	index
0.0139	construction	0.0141	manufacturing	0.0160	motor	0.0109	producer
0.0120	household	0.0134	hiring	0.0159	output	0.0108	moderate
0.0107	family	0.0129	services	0.0152	industrial	0.0093	previous
0.0107	single	0.0117	construction	0.0149	capital	0.0092	hourly
0.0104	gains	0.0110	gains	0.0128	nonresidential	0.0091	subdued
0.0101	retail	0.0109	industries	0.0111	shipments	0.0082	reflecting
0.0100	motor	0.0107	civilian	0.0103	vehicles	0.0080	gdp
0.0092	expenditures	0.0096	forecasts	0.0089	stocks	0.0079	elevated
0.0085	vehicles	0.0090	payrolls	0.0085	ratio	0.0079	real
0.0082	existing	0.0089	jobs	0.0084	software	0.0077	oil
Topic5		Topic6		Topic7		Topic8	
Weight	Word	Weight	Word	Weight	Word	Weight	Word
0.0484	policy	0.0355	foreign	0.0203	credit	0.0191	business
0.0208	monetary	0.0339	exports	0.0160	yields	0.0117	demand
0.0168	reserve	0.0334	dollar	0.0144	financial	0.0097	labor
0.0136	risks	0.0249	imports	0.0139	securities	0.0094	financial
0.0118	funds	0.0235	trade	0.0137	loans	0.0087	effects
0.0113	directive	0.0209	economies	0.0132	debt	0.0077	anticipated
0.0106	agreed	0.0197	countries	0.0123	spreads	0.0077	pressures
0.0101	financial	0.0177	major	0.0113	funds	0.0069	fiscal
0.0098	easing	0.0148	industrial	0.0109	equity	0.0067	potential
0.0088	consistent	0.0142	currencies	0.0108	banks	0.0063	productivity
0.0081	statement	0.0133	deficit	0.0108	bank	0.0060	investment
0.0078	tightening	0.0128	euro	0.0107	commercial	0.0057	firms
0.0077	stability	0.0128	japan	0.0078	nonfinancial	0.0057	costs
0.0077	expectations	0.0121	united	0.0076	corporate	0.0057	anecdotal
0.0068	stance	0.0108	exchange	0.0072	issuance	0.0055	favorable
0.0068	action	0.0104	area	0.0071	investors	0.0053	extent
0.0065	view	0.0099	services	0.0070	interest	0.0051	consumer
0.0060	pressures	0.0095	products	0.0069	bonds	0.0049	continuing
0.0059	move	0.0093	emerging	0.0065	net	0.0048	risks

Table 2
Market Reaction to the Release of FOMC Minutes

Panel A. 15-Minute Volatility				
Indep. Var	Number of Lags in Control			
	(None)	(5)	(10)	(20)
Release Dummy	0.6927**	0.9255***	0.8508**	0.8623**
	(3.42)	(3.41)	(3.18)	(3.24)
No. Obs	3703	3698	3693	3683
adj. R-sq	0.019	0.103	0.133	0.149
Panel B. 20-Minute Volatility				
Indep. Var	Number of Lags in Control			
	(None)	(5)	(10)	(20)
Release Dummy	0.4592**	0.5011*	0.5451*	0.5472*
	(2.75)	(1.97)	(2.17)	(2.17)
No. Obs	3703	3698	3693	3683
adj. R-sq	0.011	0.064	0.091	0.097

Table 3
Market Reaction to the Overall Content of FOMC Minutes

	Document Level Tone		
	Net	Positive	Negative
Tone	-0.00000157	-0.00000139	-0.000000113
	(-1.71)	(-1.52)	(-0.11)
Interest Rate	-0.000000781	-0.00000128	-0.00000105
	(-0.79)	(-1.29)	(-1.02)
Unemployment	-0.000000728	-0.000000580	-0.000000948
	(-1.20)	(-0.90)	(-1.54)
Recession	-0.00000459	-0.00000155	-0.00000214
	(-1.91)	(-0.75)	(-0.84)
N	118	118	118
adj. R-sq	0.057	0.052	0.033

Table 4
FOMC Topic Proportion and Market Reaction

Panel A. Unexpected Volatility		
	Models	
	(1)	(2)
Growth Mandate	0.00000222**	0.00000509**
	(2.76)	(3.18)
Inflation Mandate	0.00000466***	0.00000652***
	(5.35)	(4.90)
Policy	0.00000215*	0.00000344**
	(2.20)	(2.70)
Interest Rate		-0.000000220
		(-0.20)
Unemployment		0.00000136
		(1.67)
Recession		0.000000762
		(0.33)
N	118	118
adj. R-sq	0.192	0.212
Panel B. Directional Price Change		
	Models	
	(1)	(2)
Growth Mandate	0.0000364	0.000490
	(0.15)	(1.05)
Inflation Mandate	-0.00000902	0.000486
	(-0.04)	(1.26)
Policy	0.000649*	0.000834*
	(2.22)	(2.25)
Interest Rate		0.000672*
		(2.12)
Unemployment		0.000566*
		(2.40)
Recession		0.000271
		(0.41)
N	118	118
adj. R-sq	0.012	0.022

Table 5
FOMC Topic Net Tone Score and Market Reaction

Panel A. Unexpected Volatility		
	Models	
	(1)	(2)
Growth Mandate	0.00000195*	0.00000178*
	(2.06)	(2.01)
Inflation Mandate	-0.000000738	-0.000000536
	(-1.08)	(-0.73)
Financial Mkt Mandate	-0.00000176**	-0.00000153*
	(-2.64)	(-2.08)
Policy	0.000000117	-0.000000605
	(0.16)	(-0.71)
Interest Rate		-0.00000244*
		(-2.13)
Unemployment		-0.000000954
		(-1.52)
Recession		-0.00000281
		(-1.17)
No. Obs	118	118
adj. R-sq	0.051	0.076
Panel B. Directional Price Change		
	Models	
	(1)	(2)
Growth Mandate	0.000109	0.0000211
	(0.52)	(0.09)
Inflation Mandate	-0.000853*	-0.001097**
	(-2.03)	(-2.65)
Financial Mkt Mandate	-0.000881*	-0.000942*
	(-2.01)	(-2.40)
Policy	0.000726*	0.000970*
	(2.06)	(2.16)
Interest Rate		0.000840**
		(2.81)
Unemployment		0.000468*
		(2.62)
Recession		-0.000226
		(-0.33)
No. Obs	118	118
adj. R-sq	0.037	0.074

Table 6
FOMC Topic Negative Tone Score and Market Reaction

Panel A. Unexpected Volatility		
	Models	
	(1)	(2)
Growth Mandate	-0.00000100	-0.000000612
	(-1.29)	(-0.75)
Inflation Mandate	0.00000214**	0.00000268**
	(3.03)	(3.35)
Financial Mkt Mandate	-0.00000185**	-0.00000250**
	(-3.03)	(-2.98)
Policy	0.00000169*	0.00000184*
	(2.13)	(2.22)
Interest Rate		0.000000655
		(0.63)
Unemployment		0.00000112
		(1.66)
Recession		-0.000000945
		(-0.42)
No. Obs	118	118
adj. R-sq	0.132	0.141

Table 7
FOMC Topic Positive Tone Score and Market Reaction

Panel A. Unexpected Volatility		
	Models	
	(1)	(2)
Growth Mandate	0.00000144**	-0.000000251
	(2.34)	(-0.26)
Inflation Mandate	0.00000191**	0.00000197**
	(2.69)	(2.73)
Financial Mkt Mandate	-0.00000165*	-0.00000221*
	(-2.29)	(-2.28)
Policy	0.00000105	0.000000686
	(1.42)	(0.81)
Interest Rate		-0.000000181
		(-0.16)
Unemployment		0.000000596
		(0.89)
Recession		0.000000302
		(0.12)
No. Obs	118	118
adj. R-sq	0.102	0.094

Table 8
Topic Proportion and Macroeconomic Variables

	(1) **Growth Mandate**	(2) **Inflation Mandate**	(3) **Fin. Mkt Mandate**	(4) **Policy**
Interest Rate	0.228***	-0.417***	-0.246***	0.335***
	(5.79)	(-9.57)	(-8.69)	(7.13)
Unemployment	-0.250***	0.164***	0.319***	0.187***
	(-6.76)	(4.01)	(12.00)	(4.25)
Recession	-0.485**	-0.248	0.905***	-0.0999
	(-2.68)	(-1.24)	(6.95)	(-0.46)
N	176	176	176	176
adj. R-sq	0.459	0.336	0.720	0.230

Table 9
Topic Tone Scores and Macroeconomic Variables

Panel A: Net Tone				
	Growth Mandate	**Inflation Mandate**	**Fin. Mkt. Mandate**	**Policy**
Interest Rate	0.116*	0.306***	-0.190***	0.140
	(2.54)	(7.33)	(-4.55)	(1.96)
Unemployment	-0.0610	-0.0230	0.220***	0.307
	(-1.42)	(-0.59)	(5.62)	(1.92)
Recession	-1.400***	-1.030***	0.786***	-0.390
	(-6.63)	(-5.36)	(4.09)	(-1.79)
No. Obs	176	176	176	176
adj. R-sq	0.265	0.391	0.392	0.219
Panel B: Positive Tone				
	Growth Mandate	**Inflation Mandate**	**Fin. Mkt. Mandate**	**Policy**
Interest Rate	0.223***	-0.345***	-0.201***	0.178***
	(4.97)	(-7.39)	(-5.74)	(3.65)
Unemployment	-0.166***	-0.203***	0.261***	0.282***
	(-3.95)	(-4.64)	(7.96)	(6.17)
Recession	-0.392	-0.129	1.155***	0.115
	(-1.90)	(-0.60)	(7.18)	(0.52)
No. Obs	176	176	176	176
adj. R-sq	0.298	0.238	0.572	0.173
Panel C: Negative Tone				
	Growth Mandate	**Inflation Mandate**	**Fin. Mkt. Mandate**	**Policy**
Interest Rate	0.130**	-0.396***	-0.170***	0.115
	(2.64)	(-9.17)	(-4.43)	(1.20)
Unemployment	-0.121**	-0.138***	0.239***	0.0747
	(-2.62)	(-3.40)	(6.66)	(1.52)
Recession	-0.879***	0.389	1.188***	0.677
	(-3.89)	(1.96)	(6.74)	(1.82)
No. Obs	176	176	176	176
adj. R-sq	0.154	0.347	0.488	0.048