Finance and Economics Discussion Series
Divisions of Research & Statistics and Monetary Affairs
Federal Reserve Board, Washington, D.C.

FOMC Responses to Calls for Transparency

Miguel Acosta

2015-060

Please cite this paper as:
Acosta, Miguel (2015). "FOMC Responses to Calls for Transparency," Finance and Economics Discussion Series 2015-060. Washington: Board of Governors of the Federal Reserve System, http://dx.doi.org/10.17016/FEDS.2015.060.

FOMC Responses to Calls for Transparency

Miguel Acosta*

First draft: May 8, 2014
This draft: July 10, 2015

Abstract

I apply latent semantic analysis to Federal Open Market Committee (FOMC) transcripts and minutes from 1976 to 2008 in order to analyze the FOMC's responses to calls for transparency. Using a newly constructed measure of the transparency of deliberations, I study two events that define markedly different periods of transparency over this 32-year period. First, the 1978 Humphrey-Hawkins Act increased the degree to which the FOMC used meeting minutes to convey the content of its meetings. Historical evidence suggests that this increased transparency reflected a response to the Act's requirement that the Fed provide greater detail in reporting with respect to its goals and objectives. Second, the 1993 decision to publish nearly verbatim transcripts also increased transparency. However, the cost was an increasing degree of conformity at each meeting, as evidenced by lower variance in content disagreement at the member level.

Keywords: Federal Open Market Committee, Transparency, Latent Semantic Analysis, Deliberation, Natural Language Processing, Conformity, Central Bank

JEL Codes: E58, H83, D78, D82

*Federal Reserve Board. miguel.acosta@frb.gov. I'm greatly indebted to Professor John Taylor, who advised this paper—originally my honors thesis at Stanford University. This paper would not exist without the help, comments, discussions, and support of José Acosta, Lourdes Acosta, Geoffrey Cox, Lydia Cox, Marcelo Clerici-Arias, Chris Gust, Mary Grebenc, José Gutierrez, Donna Hunter, Stephanie Kalfayan, Mark Lucianovic, Ellen Meade, Allan Meltzer, Elmar Mertens, Kurt von Tish, the Bing Honors College (2013) participants, and seminar participants at the Federal Reserve Board. The views expressed in this paper are solely my responsibility and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any other employee of the Federal Reserve System.

1 Introduction

The statements, minutes, and transcripts of the Federal Open Market Committee (FOMC) are some of the most closely studied and heavily scrutinized government documents. The subject of many academic studies, they provide one of the few windows into the decision making process of this influential institution. Perhaps more significant is the attention paid to these documents by financial market participants and their influence on participants' expectations—and, consequently, on the expectations of agents in the broader economy. Given the importance of the topics contained in these documents, it is not surprising that, over the years, both the public and the FOMC itself have favored greater transparency. The transparency considered here is concerned with document publication, which, in practice, can take two forms: the first is a matter of timing, the second is one of translation. Questions of timing ask which documents should be released, and when. Questions of translation consider the process of distilling detailed information into a concise summary; for example, which words and topics from FOMC meetings should be included in the meeting minutes? Despite the importance of these documents, topics of timing and, to an even lesser extent, translation have received little empirical attention.

My first goal here is to quantify procedural transparency at the FOMC over time. Geraats (2002) defines procedural transparency as the description of how monetary policy decisions are made, which is achieved in part through the publication of records of the deliberative process.[1] The focus of this paper is the extent to which the FOMC uses its meeting minutes—which are released shortly after each meeting—to communicate the content of its meetings—which is contained in transcripts that are not released until years after a meeting has taken place.[2] The first question I address is (i) how has transparency evolved over time? Unlike previous studies that have declared the minutes of a central bank "transparent" if a researcher deems them informative and timely, this paper presents a measure of transparency that objectively quantifies how informative—in the sense of accurately conveying the content of the meetings (that is, transcripts)—the meeting minutes are.[3] This is the first such quantified measure of transparency and is constructed from documents to which both

[1] See appendix A.5 for a discussion about this and other types of transparency. Unless otherwise noted, discussions of "transparency" in this paper refer to procedural transparency.

[2] Currently, transcripts are released five to six years after a meeting has concluded; they were not released at all from 1976–1993. I assume here that the transcripts give the true account of meetings.

[3] Most recently, Dincer and Eichengreen (2014) provide a binary measure of the transparency of minutes, elaborated upon in section 2. The measure here is not a binary variable (for example, 1 if a central bank's minutes are transparent, 0 otherwise) but a variable that can take any real value in the closed interval $[0, 1]$. This, in the words of Kedan and Stuart (2014), addresses the "quality of minutes" and confronts the view that rapid and voluminous reporting of FOMC meetings is, ipso facto, more transparent—a view that ignores the fact that this reporting may be inaccurate or uninformative.

the public, and the FOMC itself, pay a great deal of attention: the meeting minutes and transcripts—the main vehicles through which the public can ascertain the content of FOMC meetings and the nature of the deliberations that underlie monetary policy decisions.[4]

My second goal is to understand what causes the measure of transparency to change. Because the measure is constructed from the minutes and transcripts, these are the only two sources that can explain transparency changes.[5] For example, a call for transparency might result in minutes that, ceteris paribus, more accurately reflect the transcripts. On the other hand, the minutes may more closely resemble the transcripts because meeting discussions become shaped by participants' knowledge that they would be made public—an undesirable result if it lowers the overall quality of discussion. My analysis here is designed to answer the following questions: (ii) what has caused the transparency measure to change and (iii) what effects has this had on the discussion in the Board Room.

To answer questions (i)–(iii), I analyze the minutes and transcripts of FOMC meetings using latent semantic analysis (LSA), along with some other tools from natural language processing (NLP). LSA has proven to be useful in the NLP literature, principally because of its ability to reduce noise in a body of text—a feature that is desirable for this study.[6] Reducing noise in natural language allows for more accurate measurements of document similarities—these similarity measures form the basis of my quantification of transparency and conformity. In particular, to measure procedural transparency, the timeliness of the minutes and detail of the transcripts are exploited. The transcripts give the full account of FOMC meetings but are not released until five years have passed, while the minutes, released with a shorter lag, are a summary of the meetings. I measure procedural transparency by computing the "similarity" between the minutes and transcripts of each meeting using LSA; the idea is that a higher similarity indicates that the Fed is releasing more of the content of its meetings in the more timely minutes.

A key finding is that procedural transparency has generally been increasing since 1976. Additionally, I use the transparency measure to focus on two historical events demarcating notable changes in transparency. First, transparency increased sharply in 1979. Using LSA to compare the behavior of the minutes and transcripts at this time reveals that this was almost entirely a minutes-driven change, with no notable difference in the transcripts. Written evidence from the transcripts suggests that this change was in response to the Full Employment and Balanced Growth Act of 1978, known as the Humphrey-Hawkins Act,

[4]Evidence provided in appendix A.5 confims that the FOMC has spent a considerable amount of time discussing topics related to procedural transparency, and the public pays attention to the publication of minutes and statements.

[5]Again, I take transcript content to be synonymous with meeting discussion.

[6]A feature of LSA not yet used in the economics literature.

which amended Section II of the Federal Reserve Act and mandated that the Fed report how its objectives and plans related to the "short-term [economic] goals" of the President and the Congress (Congress, 1978). The lack of change in the transcripts suggests that the new guidelines, as well as the generally increased scrutiny given to the Federal Reserve, led the Fed to be more transparent in conveying the conversations that they were already having. In other words, the FOMC was already discussing objectives and plans consistent with its mandate and short-term Presidential and Congressional goals, so the only change it made was in conveying these to the public. As such, this appears to have been a successful attempt at increasing transparency at the Fed.

The second change occurred in late 1993 and was also marked by an increase in transparency. From the beginning of the sample period until 1993, very few people knew that transcripts of FOMC meetings existed. In 1993, Congressional inquiries and Fed discussions revealed that the transcripts had indeed been kept, and at that point the Fed decided to publish the old transcripts with a five-year lag. Within 18 months, it also decided to publish new transcripts going forward, with the same lag. This shift from "closed" FOMC meetings to more public meetings has been studied extensively because of the natural experiment that it created. Specifically, this transition allows for a study of FOMC behavior, at the committee and member levels, under different levels of transparency. The analysis here suggests that this decision induced an increase in transparency that was partially driven by a changing transcript. Because the transcripts give detailed accounts of meetings, the move to more-public meetings caused the conversation in the Board Room to change.

While transparency *did* increase, the worry is that there could be some undesirable impact on the policy making process. For example, Meade and Stasavage (2008) cite the 1995 words of Kansas City Fed President Hoenig: "the tape has had some chilling effects on our discussions. I see a lot more people reading their statements." (pp. 704–705.) If increasing transparency caused a response like this, then the "success" of such an increase should certainly be called into question and remembered in the future when considering whether to increase procedural transparency at the Fed. My measure of transparency does not explain how the transcripts changed, but language analysis is still helpful in identifying possible answers. One possibility explored here, which has received attention in other studies, is whether the degree of conformity—an aspect of deliberation—changed once meetings were made more public. If so, this would support Hoenig's claim about the "chilling effects" of transparency. To study the degree of "chilling," a measure of conformity is constructed for each meeting using the statements of individual members.[7] The finding, similar to Meade and Stasavage (2008), Woolley and Gardner (2009), and Hansen et al. (2014), is that the

[7]See section 4 for more detail.

1993 episode resulted in an increase in conformity within the Board Room. In addition, this drop was primarily the result of a decrease in the willingness to offer disagreeing views.

The rest of this paper proceeds as follows. Section 2 reviews the literature in economics, political science and natural language processing that are similar to this paper in content and methodology. Next, section 3 describes the measure of transparency and its behavior over time, then goes on to describe the measure of minute and transcript content and their properties. Section 4 discusses conformity, and section 5 concludes.

2 Previous Studies and Theory

The literature on central bank transparency and communication started at the turn of the century, and many branches have since been pursued—Blinder et al. (2008) provide a thorough survey of this literature through 2008. The branch on which this paper grows concerns the measurement of transparency and has, most recently, been treated by Dincer and Eichengreen (2014), with the first measurements arising in Eijffinger and Geraats (2006).[8] This strand has taken the Geraats (2000, 2002) definitions of different aspects of central bank transparency as a basis for constructing measures of transparency. In this paper, I focus on the component of Geraats' *procedural transparency* that concerns central bank accounts of deliberations, the measurement of which has, thus far, focused primarily on binary and potentially subjective measures of timeliness and informativeness of central bank communications. For example, the relevant measure of procedural transparency in Dincer and Eichengreen (2014) is a binary indicator of whether "the central bank give[s] a comprehensive account of policy deliberations (or explanations in case of a single central banker) within a reasonable amount of time[.]"[9]

Hansen et al. (2014) use the term "transparency" to denote the period after 1993—another binary notion of transparency. The transparency that I study and measure in this paper, however, concerns transparency that (i) is timely and (ii) involves a degree of choice. The question I ask is when external principals express a desire for increased transparency at the Fed, how does the Fed respond? Short of real-time meetings, the minutes and statements are the only vehicles through which the FOMC can *choose* to be more or less transparent about their meetings in a manner that is meaningful to the public in real time—contrast this with the transparency in Hansen et al. (2014) in which, by the time five years have passed,

[8] Dincer and Eichengreen (2014) also thoroughly discuss previous literature.

[9] Taborda (2015) uses language analysis to give descriptive information about the central bank minutes of the central banks of Brazil, Chile, Colombia, Mexico, and Peru, under the premise of studying procedural transparency. Taborda's "interesting trends," however, should not be seen as giving a measure of procedural transparency.

"the main channel through which one expects transparency to operate . . . is career concerns rather than, for example, communication with financial markets to shift expectations about future policy." (pp. 5). The authors thus study the quality of deliberation in FOMC meetings and find results that are corroborated here—namely, that there was an increase in the conformity of views expressed after 1993.[10]

The effect of the 1993 episode on deliberations has been measured in a number of other studies. Meade and Stasavage (2008) coded voiced agreement or disagreement with Greenspan's interest rate proposals in the pre-1993 ('89–'92) and post-1993 ('84–'97) transcripts and found that there was less voiced disapproval after 1993; this finding was predicted by their theoretical "model of deliberation in a committee, where members care both about reaching the correct decision and about convincing an outside audience that they have a high level of expertise" (pp. 196).[11] Woolley and Gardner (2009) found similar results by constructing a measure of "deliberation"—the number of speakers per 100 words of transcript text—that is both highly correlated with the conformity index constructed here ($\rho = 0.79$) and shows a decrease on the order of 10% after 1993, consistent with the finding here (see figure 6). The comprehensive work of Schonhardt-Bailey (2013) studies several aspects of deliberation dynamics of the FOMC using Alceste, a language processing software, and, among other notable findings, concludes that "deliberation" decreased after 1993—noting, as do Woolley and Gardner (2009), that the downward trend appears to have started shortly before 1993.

Bholat et al. (2015) highlight the usefulness of and recent advances in text mining applications concerning central banks. Vector-space models of documents, such as the one presented here, have been applied extensively in finance,[12] and LSA's validity is well documented: see the seminal work by Deerwester et al. (1990) and Dumais et al. (1988), as well as validations by Landauer et al. (1998) and Foltz et al. (1999). At the intersection of the LSA, finance, and central bank communications literatures lie Boukus and Rosenberg (2006) and Hendry and Madeley (2009), both of which find that the LSA "themes" present in the

[10] In particular, there are two aspects of the career concerns literature (formalized in Holmström (1999)) from which the authors draw their conclusions about the overall quality of deliberation: discipline and conformity.

What I don't consider here, and what leads Hansen et al. (2014) to their conclusions, is the question of discipline. (The authors make a strong case for using "themes"—derived from latent Dirichlet allocation—to study discipline.) The authors find evidence of increased discipline and thus conclude that the net effect of transparency is a more informed deliberative process. The conclusion of a "net effect" is determined using their eigenvalue centrality measure of influence.

[11] Lähner (2015) updates and extends Meade and Stasavage (2008). Chappell et al. (2005) also analyze the transcripts to find the interest rate preferences of FOMC members, using them to construct individual reaction functions during the tenures of Burns and (partially) Greenspan.

[12] A small sample of such work includes Hoberg and Phillips (2010), Hanley and Hoberg (2010, 2012), and work by Tim Loughran and Bill McDonald, notably, Loughran and McDonald (2011).

minutes of the Fed and Bank of Canada, respectively, are significantly correlated with different market indicators.[13] However, because of the subjectivity involved in giving meaning to the high dimensional themes produced by LSA, I avoid such interpretations here. Instead, I rely on the proven strengths of LSA for dimension reduction and document comparison. To be clear, Boukus and Rosenberg (2006) and Hendry and Madeley (2009) are similar to the present paper only in that they apply LSA to central bank minutes—the goals of the papers are orthogonal.

3 Transparency

This section presents the measure of transparency and studies why, and how, it has changed over time. First, section 3.1 discusses how vector-space models of documents are used to construct the measure of transparency: the similarity between the minutes and transcripts at each meeting. Section 3.2 shows the results. Because the minute-transcript similarity does not identify the behavior of each of the documents—minutes or transcripts—individually, section 3.3 presents a way of assessing the behavior of these documents and gives the results.

3.1 Measurement: Comparing Documents using Latent Semantic Analysis

Document Comparisons in Vector-Space Models All of the measures constructed in this paper—procedural transparency, content changes, and conformity—rely critically on the ability to compare documents in a holistic and objective way. To illustrate the mechanics of this analysis, this section describes how documents are compared and presents the first application: the measurement of transparency.

Vector-space models of FOMC documents satisfy the goal of holistic and objective analysis. Here, a vector-space model of a single document is a vector whose entries are the number of times that each unique term occurs in that document. If there are multiple documents under consideration, then each document's (empty) vector has as many entries as there are unique terms in the entire body of documents; each entry corresponds to a unique word. The vector for a given document is then populated as follows: for any term occurring in that document, the corresponding entry in the vector is populated with that term's occurrence in that document. Often, the vectors are "preprocessed" in order to reduce

[13] Farka (2011) and Rosa and Verga (2007) also use textual approaches to conclude that Fed and European Central Bank (ECB) communications move markets—the latter ranks ECB president statements from "hawkish" to "dovish." Ehrmann and Fratzscher (2007) use Fed, Bank of England, and ECB communications to study the collegiality of each bank.

semantic noise that arises from using word counts alone; the details of the preprocessing performed here are listed below. For the construction of the procedural transparency measure, the body of documents under consideration—the corpus—consists of all Memoranda of Discussion (MOD, an older version of the transcripts), Records of Policy Actions (ROPA, an older version of the minutes), and minutes and transcripts from meetings physically held in Washington, D.C., between 1967 and 2008, where procedural information (voting records, attendance) is removed.[14]

The similarity of two document vectors, a and b, with m elements a_i and b_i, is

$$\text{sim}(a,b) = \frac{\sum_{i=1}^{m} a_i b_i}{\left(\sqrt{\sum_{i=1}^{m} a_i^2}\right)\left(\sqrt{\sum_{i=1}^{m} b_i^2}\right)}, \tag{1}$$

which is called the *cosine similarity* of a and b because it represents the cosine of the angle that lies between a and b.[15] From this definition, it is clear that two unrelated—orthogonal—documents will have a cosine similarity of 0; this occurs if they share no words. If two documents are scalar multiples of each other, they will have a cosine similarity with absolute value 1. This measure of similarity is desirable, in part, because it is bounded above by 1, unlike a Euclidean distance.[16]

The measure of procedural transparency, for each meeting, is the cosine similarity of that meeting's minutes and transcripts. I construct the variable `MTsim`, which consists of one observation for each FOMC meeting in the corpus described above. For meeting j, call the vector representing the minutes m_j, and call the transcript vector t_j. Each observation in the variable `MTsim` is given by

$$\texttt{MTsim}_j = \text{sim}(m_j, t_j).^{17} \tag{2}$$

[14] See appendix A.1 for a discussion of the various documents released by the FOMC since its inception. The 1967–1976 period is not considered part of the measure, though these documents are included initially as a check of the methodology. Additionally, for the modern-day minutes (1993–present), all words prior to the paragraph that typically begins with "The information reviewed at the x meeting…" (now labeled "Staff Review of the Economic Situation") is removed in order to make these documents look like the ROPA. This also keeps the content of the minutes looking relatively similar over the year, since the first meeting of each year contains discussions of procedural matters (see Meade et al. (2015) for more information on the content of the mintues).

[15] An elementary proof shows that this is equivalent to the cosine in $\mathbf{R}^2$. For $\mathbf{R}^{\geq 2}$, Equation (1) is the primary definition of cosine. Note also that the cosine similarity is nearly equivalent to the correlation coefficient if the vectors have been demeaned.

[16] Controlling for the length of each document allows for a comparison of documents of any length; thus, it is possible to compare one word or phrase to an entire document to see how well that word captures the meaning of the document, as will be shown in figure 11.

[17] The j index is meant to index time, but the letter j is used, instead of t, to remind the reader that each "document" is the j^{th} column in C, described below.

This variable gives a measure for the extent to which the meeting minutes reflect what was said at the meeting—a measure of informativeness or "quality," in the words of Kedan and Stuart (2014). This is precisely the aspect of procedural transparency considered here.

Document Preprocessing and Latent Semantic Analysis As mentioned above, documents (and their vectors) are often preprocessed in order to reduce the semantic noise present in natural language. The goal is to manipulate the document vectors such that cosine similarities between them are closer to what an all-reading human being would compute. The preprocessing performed here, and the parameters chosen, are standard fare in the NLP literature and are meant to improve the accuracy of cosine similarities.[18] Unless otherwise noted, all of the cosine similarities in this paper are computed using document vectors that have been preprocessed as described in the following paragraphs.

First, terms in a "stoplist" are excluded. As is customary, this list contains "common" words that contribute little meaning to the documents, since they are used so often. The excluded words, listed in appendix A.4, are predominantly prepositions, conjunctions, and pronouns. Also excluded are FOMC member last names, months, and Federal Reserve District numbers (*first* through *twelfth*). Numbers are also removed.[19] Additionally, words must contain at least three characters, and fewer than 15. The lower bound should catch any relatively common words not caught by the stop list and any small typographical errors. Similarly, the upper bound should catch some typographical errors or errors in the OCR processing of the original files, such as a conjoining of words—for example, *federalreservesystem*.[20] Next, terms are "stemmed" to their root. For example, the terms *different*, *differ*, and *differing* would all be reduced to *differ*.[21] After that, I use the inverse document frequency to weigh terms globally. This is another tool to minimize, but not eliminate, the weight given to terms that are used too frequently to add much insight into a document's meaning, primarily because they are used frequently in all documents (for example, *monetary*, *policy*, etc.). The first step is to calculate the number of documents in which a term t occurs—also called its document frequency, DF_t. The inverse document frequency for term t, in a corpus of n documents, is given by $IDF_t = \log\left(\frac{n}{DF_t}\right)$, thus yielding a lower weight for terms that appear in many documents. Scaling the word counts for t—also called the term frequencies,

[18]The preprocessing, even without LSA, can also stand alone; Hoberg and Phillips (2010) and Hanley and Hoberg (2010, 2012) select a few of these adjustments and then calculate cosine similarities off of these adjusted document vectors.

[19]The stoplist here also corrects for a few typographical errors in the spelling of member names in the transcripts.

[20]Optical character recognition; many of the files were originally in PDF format.

[21]The tradeoffs here are obvious, but its widespread use in the NLP and LSA literature gives hope that the benefits outweigh the costs. The unconvinced reader might consider the fact that without stemming, the two-word documents [*I play*] and [*she played*] would have a cosine similarity of 0.

TF_t—by IDF_t yields the commonly used $TF - IDF$ weighting scheme that is employed here.

Next, I run the documents through latent semantic analysis, a method first introduced as a way of matching user queries to text results.[22] As Deerwester et al. (1990) explain, word choice is, to a certain extent, random, so there exists a great deal of noise when using words or word-counts alone to capture meaning. Theoretically, however, there is an underlying "latent," or hidden, semantic structure to which any text can be mapped.[23] However, to the extent that language is used to convey an idea in some underlying semantic space, randomness is introduced.

Creating a low-rank approximation of the document vectors is the defining feature of LSA—this reduces the noise alluded to in the previous paragraph. In practice, the first step is to horizontally append the document vectors to form a term-document matrix, C, then compute the singular value decomposition (SVD) of that matrix, given by $C = U\Sigma V^T$. The k-rank approximation of C is computed by zeroing out all but the k largest singular values in the diagonal singular-value matrix Σ to get Σ_k, so that the low-rank approximation of C is given by $C_k = U\Sigma_k V^T$.[24] The elements in C_k are said to be the word counts in the latent semantic, or lower (k) dimension, space. Here, k is set to 200. Figure 1 plots the singular values—the elements $\sigma_i \in \Sigma$—from the singular value decomposition of the matrix associated with the minutes and transcripts since the beginning of the sample, 1967. The 200 dimensions retained account for 64% of the variance,[25] and each of the other 580 dimensions account for less than 1/10 percent of the variance. See appendix A.3 for more discussion on the choice of k.

[22]LSA is particularly concerned with overcoming the problems of polysemy and synonymy. Polysemy occurs when one word can describe many concepts. For example, polysemy would lead the documents [I read a book.] and [I'll book a hotel.] to look more (cosine) similar than an English-speaking human might think. Next, synonymy occurs when any concept can be expressed using many different words. The document [I'll make lodging arrangements.] and the hotel document from above would look (cosine) dissimilar, despite conveying the same idea.

[23]To continue in the language of functions, the functions from language to ideas, f_1, and from ideas to language, f_2, are both non-injective (and, hence, have no inverse). The domain of the first function f_1, from language to ideas, is an infinite set of possible expressions, and the range is the infinite set of "ideas." The other function f_2 has these flipped. The randomness arises from the non-injectivity of the functions; a speaker can express a single idea any multitude of ways (f_1: synonymy), and many ideas can be expressed by a single expression (f_2: polysemy).

[24]As shown by Eckart and Young (1936), this approximation minimizes the Frobenius Norm, $\|C - C_k\|_F$. The Frobenius norm of an $m \times n$ matrix X, with elements x_{ij}, is given by $\|X\|_F := \sqrt{\sum_{i=1}^{m}\sum_{j=1}^{n} x_{ij}^2}$.

[25]One can describe the "percent of variation explained" by the first s singular values by computing $\frac{\sum_{i=1}^{s}\sigma_i}{\sum_{i=1}^{n}\sigma_i}$.

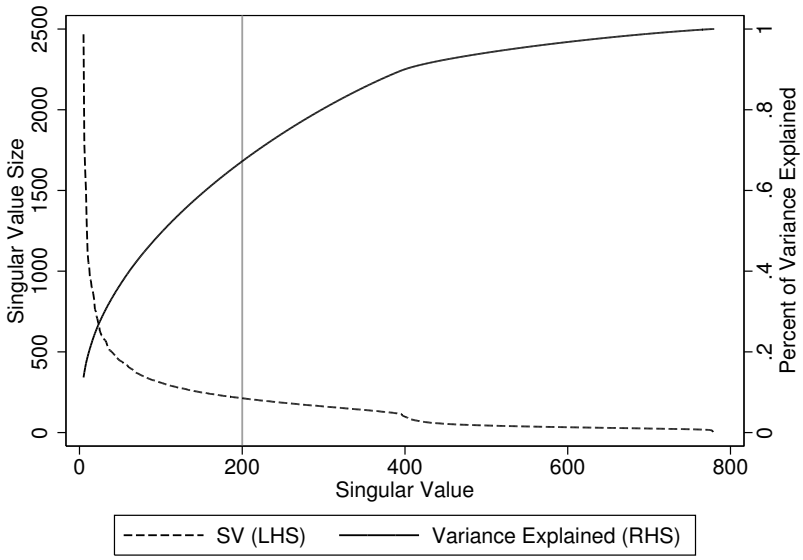

The first four SV (not depicted) are: 7945, 3576, 2863, and 2560. They represent 11% of the variance.

Figure 1: This graph shows the singular values (the elements $\sigma_i \in \Sigma$, $i = 1, \ldots, 780$, from the SVD) and the percentage of the variance explained by these singular values. The plot of the singular values (solid line) excludes the first four SVs.

3.2 How has Transparency Changed?

This section presents a new measure of the component of Geraats's (2002) *procedural transparency* that concerns central bank accounts of policy deliberations. Specifically, the measure assesses the extent to which the FOMC reveals the content of its meetings to the public via summary documents like the minutes. As opposed to previously mentioned measures that group informativeness and timeliness of central bank minutes into a single binary variable, the measure here focuses on the informativeness, or quality, of the minutes. The measure can take any real value in the closed interval $[-1, 1]$ and involves no subjectivity in its measurement.[26] Evidence provided in appendix A.5 supports the notion that FOMC communications—especially the minutes, statements, and transcripts—receive ample attention from the public and Committee and are a crucial aspect of FOMC transparency more broadly. Furthermore, the analysis below is useful in determining whether perceived transparency increases—those that are easy to see, such as more rapid document releases—actually reflect increased informativeness, an assumption that prior research has made in stating that these more conspicuous changes reflect, ipso facto, a more transparent Fed. Finally, understanding changes in the Fed's behavior in response to increased procedural transparency is crucial when crafting policy aimed at increasing transparency. A look at the

[26] All of the similarities in this paper are positive, so the interpretation of negative similarities is not discussed.

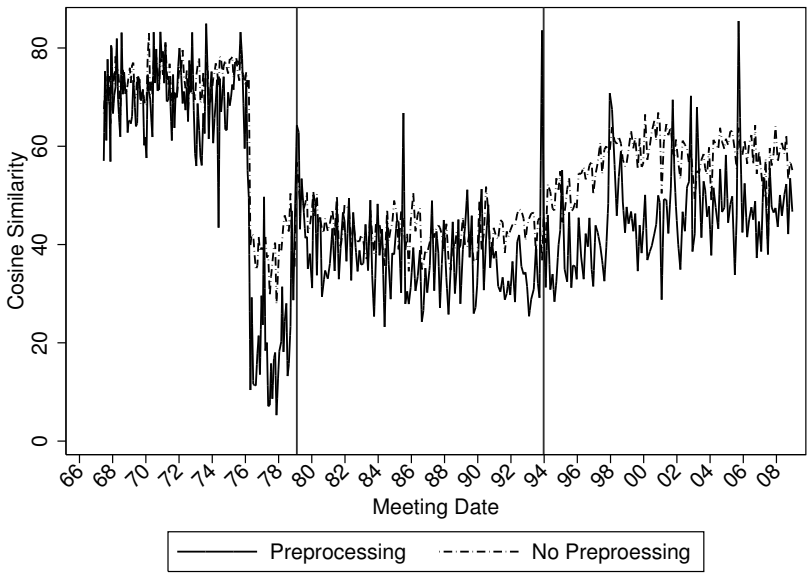

Figure 2: The solid line in this figure gives the time-series representation of `MTsim`, described in section 3.1. The dashed line shows what the measure would look like absent any document preprocessing.

time series of the similarity between FOMC minutes and transcripts in figure 2 reveals two obviously distinct phases of this similarity since 1967.

3.2.1 The Memorandum of Discussion: 1967–1976

The first time period, from 1967 to 1976, reflects the time when the FOMC published the heavily edited MOD and ROPA, referred to here as the transcripts and minutes, respectively.[27] The noticeable drop in minute-transcript similarity at the end of this period reflects the change from the MOD to a true transcript. The preparation of the MOD and ROPA was the responsibility of the FOMC secretary; Arthur Broida was the sitting secretary at the time of the transition, negating the fear that this similarity drop occurred because of a change in the leadership of the writing staff. Woolley (1986) aptly describes the editing process:

> Editorial revisions were made in preparing the [MODs]—which took place in the days immediately following each meeting—and they may not always be regarded as minor. One former governor described this as "toning down" words and making debates sound less emotional. A further and equally important part of the editorial

[27] These documents correspond roughly to the transcripts and minutes of the present day, although the ROPA and present-day minutes are much closer in form than are the MOD and transcripts. Despite these differences, it is convenient to call the ROPA "minutes" and the MOD "transcripts" in order to characterize the change in these *types* of documents over time.

> process consisted of giving more coherence to the statements of some [members] than had actually been true in debate. [A] former governor reported going to [Broida] to congratulate him on his editorial skill: "I didn't say that, but that's what I meant!" (Woolley, 1986, pp. 196)

It is not surprising, then, that this period exhibited exceptionally high similarity between the two documents; they were both highly edited documents, prepared by the staff of a single office in the Fed. Unfortunately, however, because the MOD was not truly a transcript, it is of less use in the present study since it, like the minutes, is a "prepared" document, rather than a "raw" one like the transcripts.

3.2.2 Closed-Door Meetings, Humphrey-Hawkins, and Transcript Release: 1976–2008

The second period consists of the post-1976 section of figure 2. Visually, within this post-1976 region, there are three distinct periods: from 1976 to 1979, from 1979 to 1993, and from 1993 through 2008.[28] The first cutoff, in 1979, coincides with the meeting in which the Humphrey-Hawkins Act first became effective. To more robustly confirm these periods, in table 1 I estimate the following equation:

$$\texttt{MTsim}_t = \alpha + \theta_1 \text{Post93}_t + \theta_2 \text{PreHH}_t + X_t\theta + u_t, \tag{3}$$

where Post93 and PreHH are indicator variables for post-1993 (discussed later) and pre-1979 meetings, respectively. Additionally, I estimate a few other specifications in which the variables in X are changed. In particular, there is a trend variable t; one-quarter-ahead Greenbook forecasts for GDP, unemployment, and GDP-deflator inflation; and indicators for the sitting Chair and Secretary.[29] Because the FOMC secretary is less well known, table 2 lists the secretaries over time. FOMC forecasts for macroeconomic variables from the Fed's Greenbook are used in lieu of realized values because these best reflect the knowledge available to the FOMC concerning forecasts for economic indicators. Because the primary focus of this paper is understanding why the members and FOMC behave a certain way,

[28] I focus on the 1979 and 1993 episodes because they are visually clear in the time series, because there is written evidence in the transcripts of these being important episodes, and because of their political-economic interest. I examined the Quandt likelihood ratio statstic (QLR) (Quandt (1960), with critical values from Andrews (1993)) to test for breaks in the mean between 1976 and 2008 (with 15% trimming, which excludes the 1979 episode), controlling for a linear trend and Greenbook forecasts of macroeconomic indicators. The QLR suggests breaks, at the 5% and 10% levels, in August 1989 and May 1981 respectively, though the QLR was significantly elevated for most of the 1980s. Other (not statistically significant) increases in the QLR occur in late 1993, late 1997, and starting in mid-2002.

[29] The data are from the Philadelphia Fed's Greenbook Data Set, which chooses the forecasts for one meeting each quarter. Forecasts are mapped forward to each meeting until the next forecast is available.

Table 1: Regression results from Equation(3).

	(1) MTsim	(2) MTsim	(3) MTsim	(4) MTsim
Post 93	1.034	4.841*	5.380*	7.546***
	(3.116)	(2.093)	(2.420)	(1.214)
Pre HH	-22.54***	-23.37***	-15.98***	-17.33***
	(5.664)	(5.660)	(2.331)	(1.928)
Constant	31.05*	49.95***	34.25***	37.97***
	(13.03)	(6.152)	(3.694)	(0.857)
X_t	CSMT	CSM	T	
R^2	0.500	0.494	0.391	0.388
Observations	273	273	273	273

Standard errors in parentheses

The row X_t indicates the other variables included in the regression: C and S are indicators for the Chair and Secretary, respectively; M stands for macroeconomic Greenbook forecasts; T is t for each meeting.

The dependent variable (a cosine similarity) has been multiplied by 100 so that it lies between -100 and 100.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

it is best to use the economic information available to them (or, at a minimum, that the Board staff has suggested to them).[30] The most important coefficients here are θ_1 and θ_2, the dummies on post-1993 and pre-1979, respectively. In particular, under each specification, the effect of the Humphrey-Hawkins Act was at least 16 cosine-similarity points. The evidence confirms that the minutes and transcripts became more similar after 1979—that is, transparency increased.[31]

The second discontinuity in the post-MOD era occurs in roughly 1993. The policy change in 1993 has been exploited as a natural experiment in many studies of both the FOMC and deliberative policy making in general. In this year, the Fed agreed to make public, in response to congressional pressure, that tape recordings had been kept of its

[30] As mentioned in Romer and Romer (2008), a significant portion of each meeting, the "economic go-round," is devoted to discussing forecasts and current economic conditions. (Note that Romer and Romer (2008) discuss the member-level forecasts that precede the Fed's biannual testimony before Congress, not the Greenbook forecasts, but the important notion is that these forecasts are relevant to FOMC proceedings.) Furthermore, each member has access to the Fed staff's forecasts before the meetings; thus, within a certain margin of error, the forecasts are informative of what the participants believe to be the economic outlook.

[31] The coefficients are smaller in magnitude but of the same sign (for all but θ_1 in (1)) when both a seven-meeting centered moving average and the measure resulting from un-preprocessed documents are used (sepearately) as the dependent variables. R^2 values are also high under these alternate specifications.

Table 2: FOMC Secretaries and Acting Secretaries Since: 1973–2010

Secretary	Dates in Office	Number of Meetings	Chair
Broida	07/73 – 07/78	28†	Burns, Miller
Altmann	08/78 – 02/83	41	Miller, Volcker
Axilrod	03/83 – 04/86	25	Volcker
Bernard*	05/86 – 02/87	7*	Volcker
Kohn	03/87 – 06/02	123	Volcker, Greenspan
Reinhart	08/02 – 06/07	40	Greenspan, Bernanke
Madigan	08/07 – 06/10	12†	Bernanke

*Bernard was the assistant secretary acting in the interim.

†Number of in-sample meetings.

meetings since 1976. Whether or not the Fed *was made aware* of this fact in 1993 or earlier is still a topic of debate, but in November 1993, the FOMC decided to begin releasing the 1976–1988 transcripts, and by 1995 it had agreed to start releasing transcripts from that point forward, with a five-year lag. The congressional inquiries that started in 1993 spurned much debate in the FOMC over the coming years regarding minutes and transcripts, and the entire episode led to another period in document similarity—again in the upward direction. This again suggests that a call for transparency had the desired effect; namely, the minutes more closely reflected the transcripts. The coefficient θ_1 from Equation (3) (estimated in table 1) confirms this increase under three of the four specifications.

3.3 Minutes and Transcript Content: What's Caused Transparency to Change?

For any change in minute-transcript similarity, a natural step is to also study the underlying structure of the minutes and transcripts to determine in which of these documents (minutes or transcripts) the change occurred. Even at a very high level, analyzing content changes helps build an understanding of the procedural transparency measure.[32] This type of high-level analysis is carried out here by comparing all documents of a document type (for example, all transcripts) to a single document.[33] The choice of this document is, to an extent, arbitrary,

[32]Focusing on the exact content of meetings, or the topics of discussion of each meeting, would certainly be an informative endeavor and would involve looking more closely at the particular elements in U and V from the SVD. The work of Boukus and Rosenberg (2006) did precisely this; they examined the prevalence of particular "themes" over time in FOMC minutes. However, the scope of this research calls only for observing changes in minutes and transcripts over time at a very high level. Note that "themes" are evaluated by looking at the columns of the U matrix from the SVD; each element in the column tells how much each term contributes to that theme. The columns of V give the theme contributions to each document.

[33]The rest of this discussion describes summary *transcripts*, although the discussion for minutes is analogous.

and appendix A.6 provides robustness checks to account for this. Here, for the transcripts, I use a *summary transcript*—which is a concatenation of all transcripts—as the comparison document.[34] This is a terrible summary in terms of length, but the perfect summary in terms of capturing all of the content of all of the transcripts. The variables of interest when studying the content of transcripts and minutes are given by

$$\begin{aligned} \texttt{TSTsim}_j &= \text{sim}(t_j, s_t) \quad \text{and} \\ \texttt{MSMsim}_i &= \text{sim}(m_i, s_m), \end{aligned} \tag{4}$$

respectively, where t_j is the vector representation of the transcript from the j^{th} meeting, s_t is the summary transcript vector, and m_j and s_m are the corresponding vectors for the minutes.

Panels [A] and [B] of figure 3 show the cosine similarities of the minutes and transcripts with the summary minutes and transcripts, respectively. The general pattern of an upside-down parabola exhibited in these graphs is to be expected—the comparison point, the summary document, is an average of all meetings from 1976 to 2008, so it is not surprising that the documents in the middle of the period exhibit the highest similarities. Panels [C] and [D] show the values of $\text{sim}(t_j, t_{j-1})$ and $\text{sim}(m_j, m_{j-1})$—that is, the cosine similarity of the transcripts and minutes with those of the previous meeting—respectively. These final two measures will show any big shifts in document content from meeting to meeting, as well as any anomalous minutes or transcripts.[35]

3.3.1 1979: Humphrey-Hawkins

Focusing first on the early 1979 change in transparency, the measures in figure 3 point to a change in the content of the minutes, with no associated change in transcript content. There is an uptick in the minute–summary minute similarity (Panel [B]), and a large difference between the December 1978 and February 1979 minutes (Panel [D]). The transcripts, on the other hand, remained stable over the beginning of 1979 (Panels [A] and [C]) and only changed once Chairman Volcker entered office. At the beginning of 1979, however, there was no change in leadership (Chairman Miller) or secretary (Altmann), negating the fear that this large change in minute content was caused by a change in leadership. There was also

[34] Formally, for transcripts with vector representations given by $t_1, t_2, \ldots, t_n$, the summary document has vector representation $s_t = \frac{1}{n}(t_1 + \cdots + t_n)$.

[35] These measures also portray the amount of linguistic persistence present in these documents—from meeting to meeting, documents remain remarkably similar; this is especially true of the minutes. In forthcoming work with Ellen Meade, we find persistence in FOMC statements to be of a similar magnitude, with a dramatic increase occurring since the 2008 financial crisis.

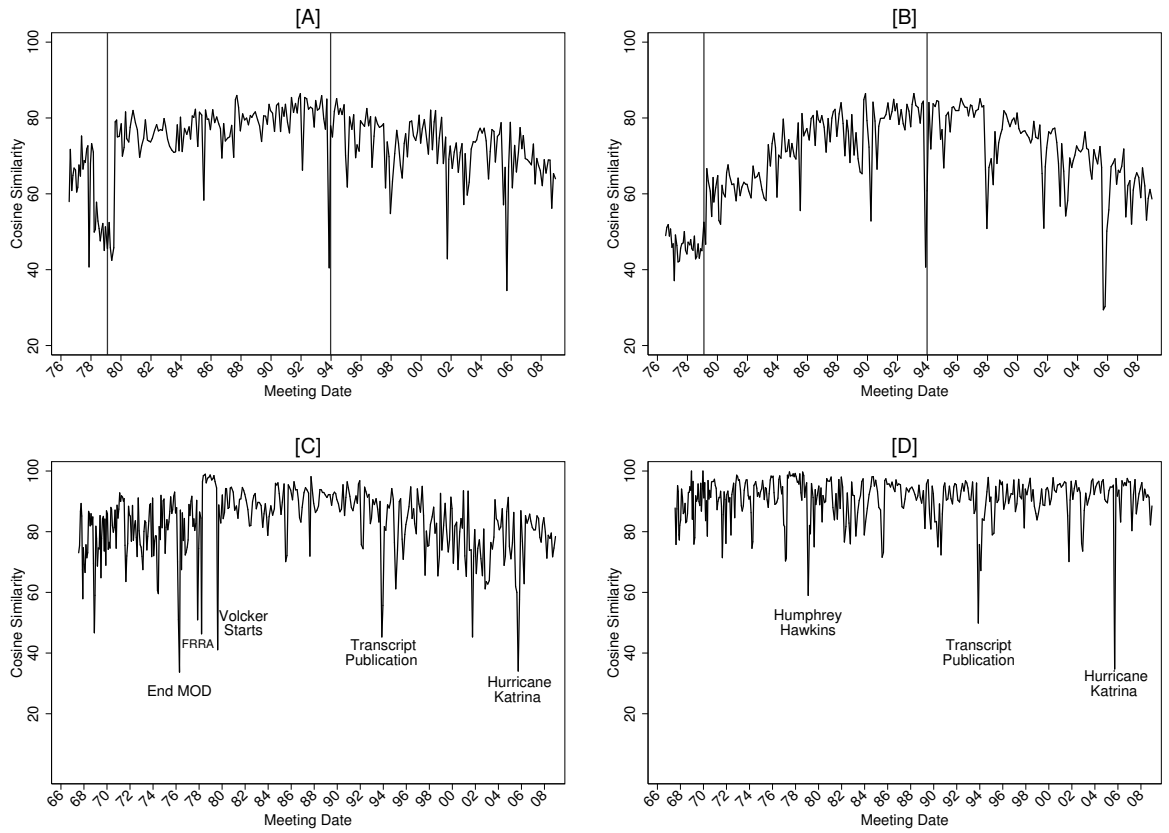

Figure 3: [A] and [B] show the cosine similarities of transcripts and minutes, respectively, with their respective summary documents. [C] and [D] show the similarities of transcripts and minutes with those from the previous meeting.

no official policy made concerning the minutes, or at least not one that was made public. In the absence of any other information, it seems reasonable to conclude that this difference in minute content is the result of the "Humphrey-Hawkins process," as it was called by Chairman Miller. Taken in conjunction with the lack of any noticeable difference in the content of the transcripts, it seems that the effect of the law was not to cause the discussion to change within the Board Room, but, rather, to change the way that the FOMC reported its proceedings to the public. Recall that, at the time, the minutes (ROPA) were thought to be the only records of FOMC proceedings, which helps ensure that the transcript content was not purposefully manipulated in response to the Act.

So, what about the Act would cause such a response—a change in the minutes only? A portion of the Humphrey-Hawkins Act amended Section 2A of the 1913 Federal Reserve Act, which had recently been amended by the Federal Reserve Reform Act of 1977 (FRRA). The FRRA was the act that legally instituted the well-known dual mandate [36] and semi-

[36]"The Board of Governors of the Federal Reserve System and the Federal Open Market Committee shall

annual reports to congressional committees (Congress, 1977). The Humphrey-Hawkins Act emphasized certain points that were to be reported to Congress semiannually (2) and added a new "objective" to be met by the Fed; it states that the Fed should report:

> (2) the objectives and plans of the [Board and FOMC] with respect to the ranges of growth or diminution of the monetary and credit aggregates for the calendar year during which the report is transmitted, *taking account of past and prospective developments in employment, unemployment, production, investment, real income, productivity, international trade and payments, and prices*; and (3) the relationship of the aforesaid objectives and plans to the *short-term goals set forth in the most recent Economic Report of the President* ... and to any *short-term goals approved by the Congress.*
> (Congress, 1978) (Emphasis added.)

While these points are rarely cited when discussing late-'70s Federal Reserve legislation, they are indicative of the atmosphere in which the FOMC was operating; essentially, there was pressure from Congress for the Fed to increase its transparency and more fully explain its actions. This is emphasized by Congress' desire to be informed of *short-term* goals. This, together with an atmosphere in which the Fed was under more scrutiny owing to the economic underperformance of the 1970s, seems a likely explanation of the Fed's decision to increase the transparency of its meetings. That no change in the transcripts was seen indicates that the Fed did not change what it said in its meetings, only what it reported, in this case through the minutes—exactly what Congress wanted.

In order to make better conclusions about the nature of the effect of the Humphrey-Hawkins Act, table 3 shows the estimates of the coefficients in the following equation:

$$\texttt{MSMsim}_t = \alpha + \theta_1 \text{Post93}_t + \theta_2 \text{PreHH}_t + X_t\theta + u_t, \tag{5}$$

where the independent variables are the same as in Equation (3), with the addition of t^2 to control for the "hump shape" present in the minute–summary minute comparisons. Again, the similarity measures are scaled by 100. Regressions (1) and (2) study minute similarity over time with `MSMsim` as the dependent variable, while (3) and (4) use `TSTsim` for studying transcript changes.

Before discussing the Humphrey-Hawkins effects, however, there is one result to be noted about the estimations. The only significant Chair dummy variables are in the transcript regression (3), and the secretary coefficients are larger in magnitude than the Chair

maintain long run growth of the monetary and credit aggregates commensurate with the economy's long run potential to increase production, so as to promote effectively the goals of maximum employment, stable prices and moderate long-term interest rates."

Table 3: Regression Results from Equation(5).

	(1)		(2)		(3)		(4)	
	MSMsim		MSMsim		TSTsim		TSTsim	
Post 93	2.998	(2.490)	-0.150	(1.840)	-1.699	(2.299)	-4.652*	(1.986)
Pre HH	-13.59**	(4.551)	-4.626*	(2.257)	2.509	(4.202)	-5.867*	(2.436)
Miller	-1.431	(3.853)			-10.89**	(3.558)		
Volcker	-3.706	(5.752)			18.19***	(5.310)		
Greenspan	-1.122	(7.766)			25.29***	(7.170)		
Bernanke	5.466	(8.007)			30.70***	(7.392)		
Broida	-5.897	(6.891)			10.42	(6.362)		
Altmann	-9.408	(5.157)			3.756	(4.762)		
Axilrod	-4.026	(3.397)			1.405	(3.137)		
Kohn	-3.092	(4.886)			-1.285	(4.511)		
Reinhart	-6.813	(5.646)			1.455	(5.213)		
Madigan	-9.516	(6.734)			-1.871	(6.217)		
Constant	23.41	(26.21)	-29.91***	(7.433)	37.49	(24.20)	28.39***	(8.024)
X_t	CSMT		T		CSMT		T	
R^2	0.704		0.669		0.573		0.349	
Observations	273		273		273		273	

Standard errors in parentheses

The row X_t indicates the other variables included in the regression: S represents indicators for the Secretary; M stands for macroeconomic Greenbook forecasts; T is t and t^2 for each meeting.

The dependent variables have been multiplied by 100 so that it lies between -100 and 100.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

coefficients in the minutes regression (1) (although none of the Chair or secretary variables are statistically significant in the minutes regressions). This suggests that, as expected, the Chair has a role in the discussion in the Board Room and the secretary might have some impact on the content of the minutes. It is reassuring to see that this type of relationship was picked up using LSA and this equation specification. Finally, the coefficients on the Chairs are large, suggesting a significant role of the Chair in determining Board Room discussion.

Turning back to the Humphrey-Hawkins effect, regressions (1) and (2) confirm that the minutes changed as a result of Humphrey-Hawkins. On the other hand, the effect on the transcripts is small and statistically insignificant. Taken together with the evidence from the lagged minute and lagged transcript similarities, the regression results give more evidence that the effect of Humphrey-Hawkins was to increase transparency through better reporting via the minutes.

3.3.2 1993: Transcript Release

The transparency increase of 1993 implied different behavior—than the 1979 episode—of the minutes and transcripts. Figure 3 suggests that the change this time was in the transcripts,

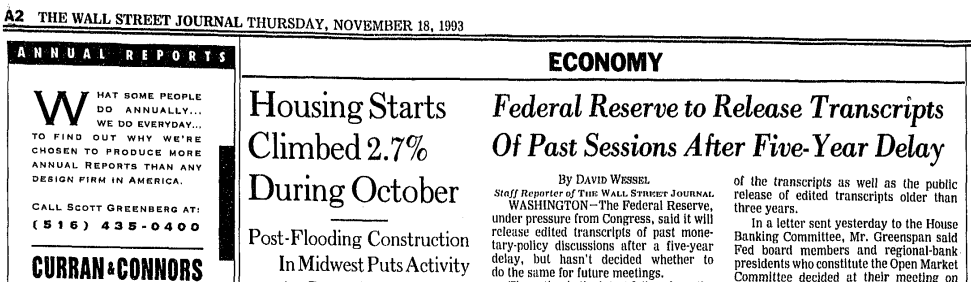

A2 THE WALL STREET JOURNAL THURSDAY, NOVEMBER 18, 1993

ANNUAL REPORTS

WHAT SOME PEOPLE DO ANNUALLY... WE DO EVERYDAY... TO FIND OUT WHY WE'RE CHOSEN TO PRODUCE MORE ANNUAL REPORTS THAN ANY DESIGN FIRM IN AMERICA.

CALL SCOTT GREENBERG AT: (516) 435-0400

CURRAN & CONNORS

ECONOMY

Housing Starts Climbed 2.7% During October

Post-Flooding Construction In Midwest Puts Activity

Federal Reserve to Release Transcripts Of Past Sessions After Five-Year Delay

By DAVID WESSEL
Staff Reporter of THE WALL STREET JOURNAL

WASHINGTON—The Federal Reserve, under pressure from Congress, said it will release edited transcripts of past monetary-policy discussions after a five-year delay, but hasn't decided whether to do the same for future meetings.

of the transcripts as well as the public release of edited transcripts older than three years.

In a letter sent yesterday to the House Banking Committee, Mr. Greenspan said Fed board members and regional-bank presidents who constitute the Open Market Committee decided at their meeting on

Figure 4: The news of the November 1993 decision.

not the minutes. While both the content of the minutes and transcripts began to fluctuate more significantly after 1993 (Panels [A] and [B]), the minutes essentially continued along the expected hump shape, while the transcripts experienced a level decrease in similarity from the summary transcript. The simple regressions, again in table 3, help to verify this difference. The coefficient of interest, θ_1, in these regressions is the one associated with the dummy variable `Post93`, which indicates whether an observation is before or after the November 1993 meeting. This date is used not only because it is the date at which old transcripts were released, but because it was a time when significant attention was paid to the potential publication of the transcripts—presumably, this was a time when changes in behavior would start to be seen, if they were to be seen at all. While it took until 1995 for the decision to be made about current transcripts, the November 1993 decision suggested that a new path was being paved toward greater procedural transparency—a path that was decided in 1995. The coefficient θ_1 on the Post93 dummy in the second transcript regression (4) is both higher (in absolute terms) and more statistically signifcant than the corresponding θ_1 in the minute regression (2). Thus, the effect of the 1993 call for transparency was to decrease the similarity of the transcripts with the summary document; in other words, there was a change in content. This change of about 4.5 cosine similarity points is smaller than the 1979 minutes change, but is still notable.

Additionally, panels [C] and [D] in figure 3 show that both the minutes and transcripts of the November 1993 meeting were significantly different from those of the prior meeting. Relative to the average of these lagged similarity values over the previous eight meetings, the transcript value for November represented a 48% decrease, while the corresponding value for the minutes was 46%. Taken as a whole, the evidence here suggests that there was, indeed, an increase in procedural transparency in 1993 that was caused by some change in the content of the transcripts. Only with this result established is it appropriate to ask questions about what may have changed about the transcripts at this time. As an application of the results discussed here, section 4 attempts to answer this question.

4 Conformity

The decision to publish FOMC transcripts, it has been shown, increased the level of transparency used in this analysis (`MTsim`). However, a serious worry is that instead of FOMC members being more open about the content of their discussions as in the 1979 episode, they opted to alter what they said in meetings to be closer to the carefully worded minutes; this would be consistent with a changing transcript and unchanged minutes. For example, the theory in Meade and Stasavage (2008) predicts that under increased transparency, members would have the incentive to offer fewer dissenting views. Alternatively, they may feel compelled to prepare statements in advance in order to avoid misspeaking, because any error or less-than-intelligent remark would be reported to the public—Hansen et al.'s (2014) finding of increased discipline supports this idea. One possible effect could be on what FOMC members were willing to say in meetings, given the fact that they knew that their individual words and opinions would be published. For example, as documented and confirmed by Schonhardt-Bailey (2013), before 1976 and after 1993, members were far more likely to essentially read off their prepared remarks.

4.1 The Measurement of Conformity

This paper's approach to studying the discussion dynamics of FOMC meetings is to use the text of the transcripts themselves, not as single entities but, rather, as accounts of discussion that can be broken apart by member. Breaking the transcript up into *member documents*, each of which holds the content of one member's utterances at that meeting, allows for a comparison of this *member document* to the words of all other members at the meeting.[37] In other words, each comparison is between what member j said and what everyone else said. In this context, the similarities are taken as showing the extent to which FOMC members "agreed" with other FOMC members at each meeting. With this, a measure of the amount of conformity, or differences in opinion and discussion topics present in the Board Room, can be achieved. Specifically, to measure conformity, the standard deviation of these agreement measures is taken for each meeting; this gives a measure of the *willingness to deviate* from the average level of agreement.[38]

Formally, for each meeting at time t, the transcripts are separated by voting member.

[37] Only voting members are included in this analysis—that is, the portions of the transcripts spoken by non voting participants are excluded.

[38] The term *agreement* is used here in a general sense. It is, literally, how similar the words of one person are to those of a group. If these entities are each speaking about completely orthogonal topics, then they are said to be out of agreement. It is also hoped that if these entities are discussing the same topic but have different opinions about them, their words will be different and, therefore, they will also be said to be out of agreement.

For p meeting members, if a transcript[39] T_t is given by $T_t = m_{1t} + m_{2t} + \cdots + m_{pt}$, where each m_{jt} are the words of FOMC voting member j, then for each member j, a new document, $m_{(-j)t}$, can be created such that $m_{(-j)t} = T_t - m_{jt}$. The variable

$$\texttt{MembSim}_{jt} = \text{sim}(m_{jt}, m_{(-j)t}) \tag{6}$$

can then be computed in order to measure the extent to which that member was in agreement with the others present at the meeting.[40] The variables $\texttt{MeanSim}_t$ and $\texttt{SDsim}_t$ represent the mean and standard deviation of $\texttt{MembSim}_{jt}$, respectively, at meeting t. Again, the latter comprises the measure of conformity.

The corpus for this analysis consists of all the member documents, m_{jt}, and all of the excluded member documents, $m_{(-j)t}$, for all meetings held in Washington, D.C., after March 1976 and all voting members. For the 276 meetings, this results in 6128 documents total, half of which are member documents m_{jt}. The average number of voting members per meeting is 11.1.[41] The choice of k for the SVD rank reduction is, again, 200; see appendix A.3 for discussion of this choice.

4.2 How has Conformity Evolved?

Table 4 show the values of the $\texttt{MembSim}_{jt}$, which are the previously defined values that compose this paper's measure of conformity, $\texttt{SDsim}_t$, where t = October 6, 1986. At this meeting, 12 voting members were present. Of course, the numbers in table 4 are relatively high, given that a cosine similarity of 1 signifies identical documents. However, these figures have a sufficient range and variance in their distribution to allow for a study of their behavior over time.

The results from section 3.2 showed that the transition from "transcriptless" (closed) meetings to recorded and published (open) meetings resulted in some change in the content of FOMC transcripts. The primary finding here is that one change took the form of a

[39] Again, only the portions of the transcript spoken by voting members.

[40] More precisely, by "in agreement," I mean how similar a members words were to those of the other members.

The only assumption that underlies this methodology is that FOMC meetings consist of dialogue and discussion. For if each person came to the meeting solely to "say their piece," and that "piece" was unrelated to every other member's "piece," then these similarity measures would also be low. If each member is essentially discussing the same topics as the rest of the members with varying degrees of agreement, then this measure of conformity is satisfactory, as any differences would come from the varying degrees of agreement.

[41] This average is less than 12 for two reasons: absenteeism and silence. While there are 12 voting positions on the FOMC, there were often meetings where either a spot had not been filled yet, or the person was simply absent. There are also a couple of cases in which a member did not speak for the entire meeting: during the August and September meetings in 1978, the record indicates that Governor David Lilly was silent. This is the only case of this occurring in the data.

Table 4: Measuring FOMC Member Agreement at the October 6, 1981 Meeting

Member	Similarity to Other Members	Member	Similarity to Other Members
Boehne	0.88	Rice	0.54
Boykin	0.60	Schultz	0.89
Corrigan	0.93	Solomon	0.92
Gramley	0.95	Teeters	0.84
Keehn	0.47	Volcker	0.97
Partee	0.90	Wallich	0.83

Figures closer to 1 represent more agreement. The measure is formally defined below. The mean is 0.80 and standard deviation is 0.18.

decrease in member disagreement within meetings; this corroborates the evidence in Hansen et al. (2014) of increased conformity after meetings were made open. Figure 5 gives four visual representations of $\texttt{MembSim}_{jt}$ for each member j and meeting t. Before any statistical analysis of this variable is made, however, there are a few observations that will certainly help guide the thinking behind the statistical/econometric tests. Panel [A] shows all the member-level observations over time, and the second panel ([B]) plots a seven meeting moving average of the range and mean of $\texttt{MembSim}_{jt}$ of each meeting t. Looking at these in conjunction, a visual test suggests that there was a narrowing of the distribution as time progressed; in particular, it appears that there were fewer members who expressed views that conflicted with the rest of the Committee. Statements that were in disagreement would presumably be those below the mean level of agreement, and the increasing minimum shown in Panel [B] confirms that there have been fewer disagreeing views expressed over time.

Panel [C] plots the standard deviation of member disagreement levels for each meeting over time; this is perhaps the best summary measurement of conformity at each meeting, as it, tautologically, gives a measurement for the variance in how much members agreed at each meeting. In other words, it gives a measure of how far members were willing to stray from the mean level of agreement—their aversion to conformity. The standard deviation is also helpful to consider at this stage given that changes in the mean level of agreement/disagreement are not obvious in the graphs, while the standard deviation changes are notable and easy to pinpoint. First, there is a large drop after Chairman Volcker left the Fed, consistent with the commonly held view that "Alan Greenspan's quiet authority was rarely challenged during his 18-year rule. His predecessor, Paul Volcker, clashed with governors appointed by President Ronald Reagan" (Hilsenrath, 2013). A similar, but smaller, drop in standard deviation and increase in mean member-disagreement is also seen when Chairman Bernanke came into office, although this is not given much attention because there are only two in-sample years of data. Another, less striking change in the data, which is more relevant to the changes

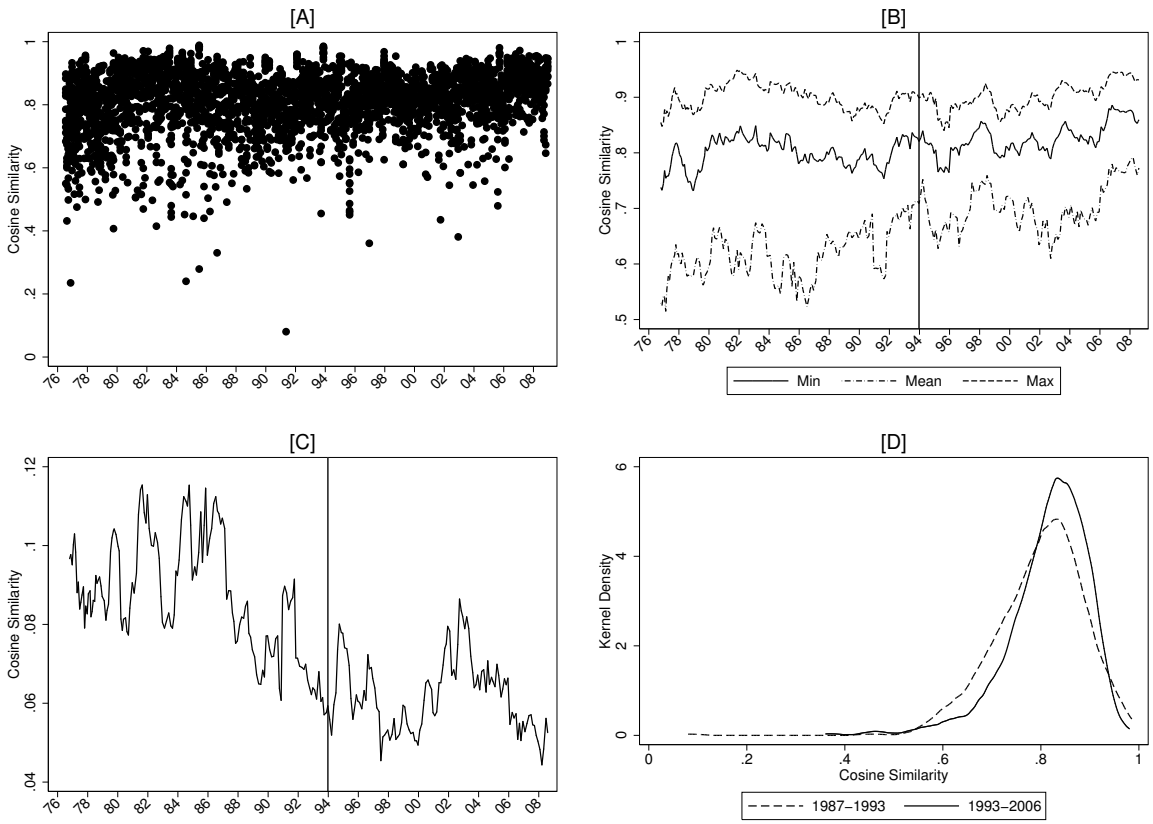

Figure 5: [A] Shows the values of `MembSim`, [B] plots its range and mean (seven-meeting moving average), [C] plots its standard deviation (seven-meeting moving average), and [D] shows its kernel density estimate during Greenspan's term before and after November 1993.

that took place in the early to mid-'90s, is the drop in standard deviation after 1993. This suggests that once members were aware that their words might—and, eventually, would—be made public, the discussion within the Board Room became significantly less varied. This decrease in variance is indicative of the change predicted by theory when meetings are made essentially open to the public (see Hansen et al. (2014), Meade and Stasavage (2008), and Woolley and Gardner (2009) for theory and other results); for example, if a member's words are consistently out of line or contradictory with the FOMC's general discussion, that member might worry that her reputation with the public will suffer—she would be the odd-member out.[42] Finally, panel [D] presents two kernel density estimates of this data under Chairman Greenspan's leadership—one before 1993 and one after. Note that the pre-1993 group had a lower mean, slightly wider distribution, and much longer left tail.

[42]Preliminary results suggest that, controlling for the sitting Chair and secretary, conformity and dissension in voting are positively related at the committee level—an unexpected result. This will be revisited with member-level voting data.

Table 5: Regression results from Equation(7)

	(1) MeanSim	(2) MeanSim	(3) SDsim	(4) SDsim
Pre HH	-2.407*	-3.419***	0.273	0.276
	(1.191)	(0.543)	(0.551)	(0.265)
Post 93	1.570***	2.112***	-1.262***	-2.495***
	(0.450)	(0.330)	(0.208)	(0.161)
Constant	79.66***	80.44***	9.156***	8.685***
	(1.341)	(0.231)	(0.620)	(0.113)
X_t	CSM	M	CSM	CSM
R^2	0.652	0.303	0.780	0.511
Observations	267	267	267	267

Standard errors in parentheses

The row X_t indicates the other variables included in the regression: C and S are indicators for the Chair and Secretary, respectively; M stands for macroeconomic Greenbook forecasts.

The dependent variable (a cosine similarity) has been multiplied by 100 so that it lies between -100 and 100.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

MeanSim Recall that the mean similarity at each meeting is a summary statistic for the general level of agreement at the meeting. In order to find whether the 1993 episode had any effect on agreement at meeting t, the following equation is estimated:

$$\texttt{MeanSim}_t = \alpha + \theta_1 \text{Post93}_t + \theta_2 \text{PreHH}_t + X_t\theta + u_t, \tag{7}$$

where, again, the independent variables are the same as in Equation (3). The results are given in columns (1) and (2) of table 5. The relatively high constant $\alpha \approx 80$ suggests that, in general, FOMC members are discussing the same topics and display a relatively high level of agreement.[43] It is important to see, however, that the magnitude of the coefficient on the post-1993 variable—on average 1.8—is small relative to the mean (≈ 80) of MeanSim. Unfortunately, it is difficult to say how a difference of 1.8 in cosine similarity translates into meaning, especially compared to the relatively high mean cosine similarity of 80. It is likely that because FOMC members generally discuss the same topics it is difficult for the LSA

[43] Of course, this is not a perfect measure, but the belief is that applying LSA to the documents would help distinguish, in a meaningful way, documents that discuss the *same* topics but disagree about them. So to the extent that LSA captures meaning and can make such distinctions, this variable is informative.

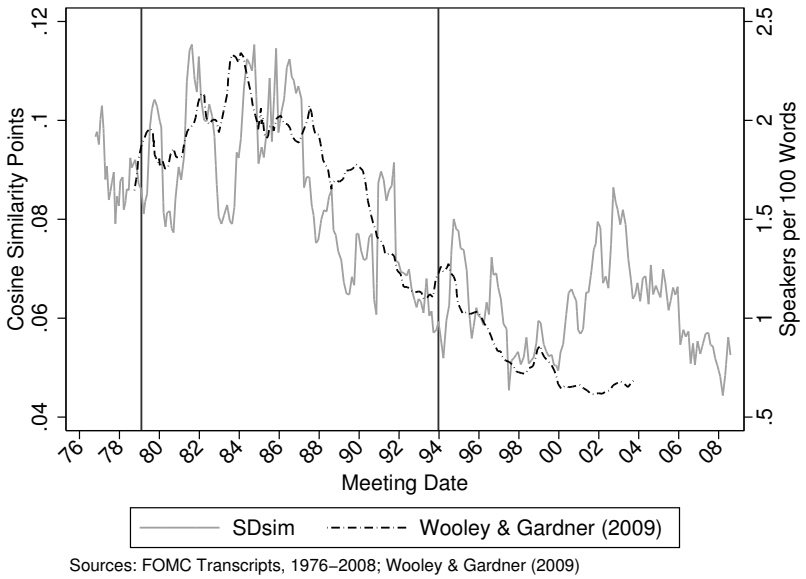

Figure 6: Comparing the `SDsim` measure of FOMC deliberation (from figure 5) with that of Woolley and Gardner (2009). The correlation coefficient $\rho = 0.79$.

algorithm to make much of a distinction.[44]

`SDsim` Regressions (3) and (4) in table 5 study the `SDsim` variable using the same independent variables as in equation (7) and also find significantly ($p < 0.001$) negative effects from the move toward open meetings in 1993, even when controlling for the Chair, secretary, and macroeconomic conditions. The average of the coefficients on the post-'93 variable in (3) and (4) is 1.88 cosine similarity points, which, given the small mean of that variable (about 9), suggests a bigger change than was apparent in the the `MeanSim` regressions. Whereas the mean similarity measured the average level of agreement/disagreement, the standard deviation by definition measures how far away members were willing to depart from this mean level of agreement. Furthermore, that the *difference between the maximum and mean similarity stayed relatively constant* over the period while the *minimum-mean difference noticeably shrank* suggests that the lower standard deviation is due primarily to decreased disagreement (where *disagreement* is similarity below the mean, and *agreement* above)—see panel [B] in figure 5. Thus, the results are consistent with the theory mentioned above; namely, that openness reduces the willingness to offer disharmonious opinions. In the case of the 1993 event, the discussion "dampened" by this measure by roughly 20 percent.[45] Equivalently, conformity increased to the same degree. This is consistent with the results

[44]Broadly speaking, of course, but the expectation is that members discuss issues related to the economy, which makes up a relatively small subspace of all discussable objects.

[45]This is a rough ratio of $\frac{\theta_1}{\alpha}$ in regressions (3) and (4) in table 5.

Table 6: Regression Results from Equation(8)

	(1) MembSim		(2) MembSim	
Pre HH	-0.337	(1.612)	-4.157***	(0.585)
Post 93	3.393***	(0.624)	2.673***	(0.460)
Miller	-0.774	(1.129)		
Volcker	3.531	(1.805)		
Greenspan	3.881*	(1.910)		
Bernanke	8.278***	(2.178)		
BankPres	-1.846***	(0.452)	-1.915***	(0.451)
IsChair	2.498**	(0.793)	2.391**	(0.798)
Constant	72.94***	(2.721)	81.40***	(0.370)
X_t	M			
R^2	0.0688		0.0511	
Observations	2417		2417	

Standard errors in parentheses
The row X_t indicates whether M—macroeconomic Greenbook forecasts—-are included as controls.
The dependent variable (a cosine similarity) has been multiplied by 100 so that it lies between −100 and 100.)
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

found by Woolley and Gardner (2009); figure 6 plots their measure of transparency with SDsim—the measures have a correlation coefficient $\rho = 0.79$.

Member-Level Agreement and Debate The regressions estimated in table 6 study the member-level responses to the variables mentioned above and a few additional controls. This is a less crucial aspect of the analysis of this section, but it serves to tie the data to results found in other studies. Specifically, for member j at meeting t,

$$\texttt{MembSim}_{jt} = \alpha + \theta_1 \text{Post93}_t + \theta_2 \text{PreHH}_t + \delta_1 \text{IsChair}_{jt} + \delta_2 \text{IsBankPres}_{jt} + I_t^{\text{Chair}}\gamma + u_{jt}, \tag{8}$$

where IsChair_{jt} and BankPres_{jt} are dummy variables that indicate whether the member j has the title of Chair or Reserve Bank president during the meeting t, and I_t^{Chair} is a matrix of indicators for the sitting Chair. The estimates show that after 1993, members were more likely to express opinions that were similar to those of the other members. Again, the relevant coefficient, $\theta_1 \approx 3$, is small relative to the mean but suggests that there was at least some change after members could feel the impending move toward open (with a lag) meetings. Without having a feeling for what a "3 cosine similarity change" means for

a discussion, though, especially in a room devoted to discussing a narrow set of topics, it is difficult to know how tangible this difference was. Anecdotal evidence, primarily member accounts, suggests that the change was in fact noticeable; members tended to read their remarks more often in meetings.[46]

Next, the coefficients δ_i on the member characteristic variables (IsChairman and IsBankPres) are both in line with the literature that suggests that Reserve Bank presidents tend to dissent more by voting measures. Here, presidents are more likely to present differing opinions than the rest of the group; this is consistent with the general observation that Reserve Bank presidents tend to be the more outspoken members of the Board (Chappell et al., 2005). Chairs, on the other hand, are less likely to offer dissenting opinions, and Governors (omitted) are in between. That these coefficients are highly significant and of the same magnitude in both specifications suggests that these are important factors in understanding discussion dynamics.

Finally, notice the Chair indicator variables—that is, those that describe who the Chair is at the time of the meeting (notice that Burns is the omitted former Chair for this sample). From these results, it appears that the Chair has played an increasingly important role in the dynamics of FOMC discussion over time. Most notable is Bernanke's effect; at 8 cosine-similarity points, it is the largest effect seen in regression (1). It is again indicative of the collegial style for which he was well known. Of course, such an interpretation is something of a best guess given the general descriptions of his leadership style; it could be that he was so dictatorial that everyone chose to agree with him. However, equally likely is that the discussion was dampened, narrowed, or more tightly focused. Regardless, Chair effects are significant and should certainly be kept in mind when considering ways to affect FOMC transparency and discussion.

5 Conclusion

The analysis presented in this paper has made only a small dent in showing what NLP techniques have to offer to the study of empirical questions in economics. Text offers an alternative perspective from which to approach problems; because text is generally prepared by a human being, the choices made in composing the text, both in content and style, reflect the preferences and mindset of the author(s). For FOMC documents, this quality offers new insights into the decision making process of the institution that has always been associated with a bit of mystique.

After descriptions of previous studies and the general methodology, the main results

[46] See Schonhardt-Bailey (2013) for more discussion on this point.

of the paper were given in sections 3.2, 3.3, and 4. First, the passing of the Humphrey-Hawkins Act (1978) coincided exactly with a large increase in procedural transparency, as measured by the cosine similarity between the minutes and transcripts of each meeting. The textual evidence also shows that this change was primarily driven by the minutes. Historical evidence suggests that the requirement that the Fed report *how* its objectives are related to the short-term goals of the President and Congress, and the increase in public scrutiny of the Fed, caused the FOMC to more transparently report the content of its meetings while making few adjustments to the content of those meetings. In this sense, the transparency increase was relatively uncontaminated by some of the problems that accompanied the 1993 change.

The next major finding is that the 1993–1995 move from completely private meetings to semi public (with a lag) meetings also increased transparency—a shift that was driven by changing transcripts. Of course, LSA cannot say *why* or exactly *how* this change occurred. Given the complexity of these documents, it is difficult to pinpoint what specifically changed about the post-1993 transcripts, so historical accounts again served as a guide. FOMC members expressed concern about potential "chilling" effects—as one bank president[47] put it—of open meetings, suggesting that *deliberation*—specifically, conformity—may have been affected. Using the similarity of each member's words with those of everyone else at each meeting to get a sense of the level of agreement, the evidence suggests that there was a substantial degree of "dampening" in deliberations after 1993. The deliberation case is particularly relevant today because it suggests how the FOMC might respond to increased procedural transparency.

In addition to the historical results about FOMC responses to calls for transparency and corroborating evidence for previous studies of the effect of "openness" on deliberation, this paper has made two contributions. The first, from a methodological standpoint, is the introduction of dimension-reduction-based textual analysis to the economics literature—a procedure that takes advantage of an abundance of NLP research. The second is the refinement of the meaning of FOMC "transparency" beyond a binary notion, coupled with a new measurement for procedural transparency.

Given the richness of the documents released by the FOMC, there is no shortage of avenues for future research. The transcripts, for example, lend themselves to studies of the discussions behind monetary policy decisions—that is, the topics that have been discussed over the years. The more timely documents, such as the minutes and statements, can be exploited further to measure how the public reacts to the words released by the FOMC. These ideas, and many more, can benefit hugely from the application of natural language

[47] See the quote from Meade and Stasavage (2008) in the introduction.

processing to FOMC documents, since the analyses of these documents can be more rapid, accurate, and objective than ever before. □

References

Andrews, D. W. (1993). Tests for parameter instability and structural change with unknown change point. *Econometrica*, 821–856.

Bholat, D., S. Hansen, P. Santos, and C. Schonhardt-Bailey (2015). Text mining for central banks. *Centre for Central Banking Studies Handbook* (33), 1–19.

Blinder, A. S., M. Ehrmann, M. Fratzscher, J. D. Haan, and D.-J. Jansen (2008). Central bank communication and monetary policy: A survey of theory and evidence. *Journal of Economic Literature 46*(4), 910–945.

Boukus, E. and J. Rosenberg (2006). "The information content of FOMC minutes. *Federal Reserve Bank of New York Working Paper*.

Chappell, H. W., R. R. McGregor, and T. Vermilyea (2005). *Committee Decisions on Monetary Policy*. Cambridge: The MIT Press.

Congress, U. S. (1977). *H.R. 9710–95th Congress: Federal Reserve Reform Act*. United States.

Congress, U. S. (1978). *H.R. 50–95th Congress: Full Employment and Balanced Growth Act*. United States.

Danker, D. J. and M. M. Luecke (2005). Background on FOMC meeting minutes. *Federal Reserve Bulletin 91*, 175.

Deerwester, S. C., S. T. Dumais, T. K. Landauer, G. W. Furnas, and R. A. Harshman (1990). Indexing by latent semantic analysis. *Journal of the American Society for Information Science 41*(6), 391–407.

Dincer, N. N. and B. Eichengreen (2014). Central bank transparency and independence: Updates and new measures. *International Journal of Central Banking 10*(1), 189–259.

Dumais, S. T., G. W. Furnas, T. K. Landauer, S. Deerwester, and R. A. Harshman (1988). Using latent semantic analysis to improve access to textual information. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pp. 281–285. ACM.

Eckart, C. and G. Young (1936). The approximation of one matrix by another of lower rank. *Psychometrika 1*(3), 211–218.

Ehrmann, M. and M. Fratzscher (2007). Communication by central bank committee members: Different strategies, same effectiveness? *Journal of Money, Credit and Banking 39*(2-3), 509–541.

Eijffinger, S. C. and P. M. Geraats (2006). How transparent are central banks? *European Journal of Political Economy 22*(1), 1–21.

Farka, M. (2011). The asymmetric impact of "informative" and "uninformative" federal open market committee statements on asset prices. *Contemporary Economic Policy 29*(4), 469–493.

Foltz, P. W., D. Laham, and T. K. Landauer (1999). Automated essay scoring: Applications to educational technology. In *World Conference on Educational Multimedia, Hypermedia & Telecommunications*, Volume 1999, pp. 939–944. EdMedia.

Geraats, P. M. (2000). Why adopt transparency? The publication of central bank forecasts. *Center for International and Development Economics Research, UC Berkeley*.

Geraats, P. M. (2002). Central bank transparency. *The Economic Journal 112*(483), F532–F565.

Geraats, P. M. (2007). The mystique of central bank speak. *International Journal of Central Banking 3*(1), 37–80.

Goodfriend, M. (1986). Monetary mystique: Secrecy and central banking. *Journal of Monetary Economics 17*(1), 63–92.

Hanley, K. W. and G. Hoberg (2010). The information content of IPO prospectuses. *Review of Financial Studies 23*(7), 2821–2864.

Hanley, K. W. and G. Hoberg (2012). Litigation risk, strategic disclosure and the underpricing of initial public offerings. *Journal of Financial Economics 103*(2), 235–254.

Hansen, S., M. McMahon, and A. Prat (2014). Transparency and deliberation within the FOMC: a computational linguistics approach. *Working Paper*.

Hendry, S. and A. Madeley (2009). "Text Mining and the Information Content of Bank of Canada Communications". *Bank of Canada Working Paper*.

Hilsenrath, J. (2013, Sept 22). Yellen would bring tougher tone to Fed. *The Wall Street Journal*.

Hoberg, G. and G. Phillips (2010). Product market synergies and competition in mergers and acquisitions: A text-based analysis. *Review of Financial Studies 23*(10), 3773–3811.

Hoberg, G. and G. M. Phillips (2009). Text-based network industries and endogenous product di fferentiation. Technical report, National Bureau of Economic Research.

Holmström, B. (1999). Managerial incentive problems: A dynamic perspective. *The Review of Economic Studies 66*(1), 169–182.

Kedan, D. and R. Stuart (2014). Central bank minutes. Working paper, Central Bank of Ireland.

Lähner, T. (2015). Inconsistent voting behavior in the FOMC. Technical report, Hannover Economic Papers.

Landauer, T. K., P. W. Foltz, and D. Laham (1998). An introduction to latent semantic analysis. *Discourse processes 25*(2-3), 259–284.

Lindsey, D. E. (2003). A modern history of FOMC communication: 1975-2002. *Memo: Board of Governors of the Federal Reserve System, Washington, DC*.

Loughran, T. and B. McDonald (2011). When is a liability not a liability? Textual analysis, dictionaries, and 10-ks. *The Journal of Finance 66*(1), 35–65.

Meade, E. E., N. A. Burk, and M. Josselyn (2015). The FOMC meeting minutes: An assessment of counting words and the diversity of views. *Federal Reserve Board FEDS Notes*.

Meade, E. E. and D. Stasavage (2008). Publicity of debate and the incentive to dissent: Evidence from the us federal reserve. *The Economic Journal 118*(528), 695–717.

Meltzer, A. H. (2010). *A History of the Federal Reserve, Volume 2*. Chicago: University of Chicago Press.

Quandt, R. E. (1960). Tests of the hypothesis that a linear regression system obeys two separate regimes. *Journal of the American Statistical Association 55*(290), 324–330.

Romer, C. D. and D. H. Romer (2008). The FOMC versus the staff: Where can monetary policymakers add value? *American Economic Review 98*(2), 230–35.

ROPA (1976–1993). *FOMC Records of Policy Actions 1976-1993*. Board of Governors of the Federal Reserve System.

Rosa, C. and G. Verga (2007). On the consistency and effectiveness of central bank communication: Evidence from the ecb. *European Journal of Political Economy 23*(1), 146–175.

Schonhardt-Bailey, C. (2013). *Deliberating American Monetary Policy: A Textual Analysis*. Cambridge, MA: The MIT Press.

Taborda, R. (2015). Procedural transparency in latin american central banks under inflation targeting schemes. a text analysis of the minutes of the boards of directors. *Ensayos sobre Política Económica 33*, 76–92.

Transcripts (1976–2008). *FOMC Transcripts 1976-2008*. Board of Governors of the Federal Reserve System.

Woolley, J. T. (1986). *Monetary Politics: The Federal Reserve and the Politics of Monetary Policy*. Cambridge: Cambridge University Press.

Woolley, J. T. and J. C. Gardner (2009). Does sunshine reduce the quality of deliberation? the case of the federal open market committee. *APSA 2009 Toronto Meeting Paper*.

Zeimpekis, D. and E. Gallopoulos (2006). TMG: A MATLAB toolbox for generating term-document matrices from text collections. In J. Kogan, C. Nicholas, and M. Teboulle (Eds.), *Grouping Multidimensional Data*, pp. 187–210. Springer Berlin Heidelberg.

A Appendix

A.1 Transcripts, Minutes, ROPAs, MOAs, and MODs: The History of FOMC Communications

Understanding the logistical aspects of procedural transparency—which documents are released and when—is a necessary step in the assessment of procedural transparency in the "quality" sense. Since the inception of the modern-day FOMC, it has always communicated in some way the content of its meetings. Table 7 lists the various FOMC publications since

Table 7: FOMC Publications: Banking Act of 1935 to Present (2015). *Release lags in italics.*[48]

Date	Meeting Summaries	Detailed Accounts
1935–1967	Record of Policy Actions *(Annually)*	Minutes *(Confidential)*
1967–1975	Record of Policy Actions *(90 Days)*	Memorandum of Discussion *(5 Years)*
1975–1976	Record of Policy Actions *(45 Days)*	Memorandum of Discussion *(5 Years)*
1976–1993	Record of Policy Actions and Minutes of Actions *(One Meeting)*	Transcripts *(Confidential)*
1993–2005	Minutes *(One Meeting)*	Transcripts *(Five Years)*
2005–	Minutes *(Three Weeks)*	Transcripts *(Five Years)*

1935, along with their release lags.[49] While the nomenclature of the various documents has undergone several changes over the past 80 years, there have, in general, been two types of documents: detailed accounts of FOMC meetings and summaries. In general, the latter were more readily available to the public. Most of the changes in FOMC communications have occurred alongside calls for transparency, and they have formally come from Congressional pressure, legislation, and litigation. The exogeneity of these external pressures permits the study of how the FOMC's procedural transparency responds to these calls. This section provides an overview of the logistical aspects of these responses.

The first significant step toward greater procedural transparency—in the sense of timeliness—came in response to the 1967 Freedom of Information Act (FOIA). Beginning with the April 1967 meeting, the Record of Policy Actions—a summary of the Committee's policy actions and rationales—would be published after a 90-day lag (Danker and Luecke, 2005). And, for the first time, a transcript-like document—the Memorandum of Discussion—was to be released with a five-year lag. This set a precedent for publishing long accounts of FOMC meetings, but because the MOD was a heavily edited account, the 1967–1976 period is not included in the measurement of transparency reported below, for reasons elaborated upon in section 3.2.1.[50]

The 1976 MODs were the last published by the FOMC; after five years of fighting a

[48] Adapted from Danker and Luecke (2005).

[49] Statements, press conferences, and other releases are omitted, since the focus here is on documents whose primary purpose is to convey the meeting discussion.

[50] Though the similarity of minutes and MODs is shown above, this section of the time series should not be taken as measuring procedural transparency, due to the heavily edited nature of the MODs.

claimed FOIA violation, the Committee decided in 1981 to discontinue the MOD,[51] largely at the request of Chairman Burns (Lindsey, 2003). At this point, the Committee decided to release an expanded ROPA shortly after each subsequent meeting; effectively, a 30-day lag. At the time, the reason cited for the discontinuation of the MOD was that "the benefits derived from them did not justify their relatively high costs, particularly in light of the changes made in the [ROPA]" (ROPA, 1993, 05/18/76). However, the more accurate reason seems to be "'fear that Congress would request access' [to the MOD] promptly" (Lindsey, 2003, p. 8) and, as an FOMC subcommittee indicated, "concern about the ability to conduct monetary policy, if the court required prompt release of the memoranda of discussion" (Meltzer, 2010, p. 976). The discontinuation of the MOD started a nearly 20-year period in which the FOMC published no detailed account of its meetings. Most FOMC members were aware that meetings were recorded, but they also believed that these tapes, used only for the production of minutes by Board staff, were recorded over after each meeting.

Contrary to what most members believed, Congressional inquiries (primarily headed by Congressman Henry González) and internal Fed investigations revealed that, in fact, these tapes had been maintained since 1976. In November 1993, the Committee agreed to publish all of the transcripts since 1976; by 1995 the decision was made to reinstate the publication of meeting transcripts after a five-year lag. In addition, the ROPA and MOA were now combined to form the "minutes." In 2005, these minutes began to be released with a three-week lag.

All post-1967 ROPAs, MOAs, MODs, minutes and transcripts were downloaded from `http://federalreserve.gov`, either in PDF format or plain text. Documents in PDF format were converted to plain text using optical character recognition (OCR) software. After procedural information was removed, documents were converted to numerical data using the MATLAB® toolbox Text Matrix Generator (Zeimpekis and Gallopoulos, 2006).

[51] See Goodfriend (1986) for a thorough account.

A.2 A Graphic Representation of Vector-Space Models

Suppose that there are three documents—x, y and z—in a corpus $\mathcal{C}$:

$x =$ Inflation inflation inflation unemployment.

$y =$ Inflation unemployment unemployment unemployment.

$z =$ Unemployment unemployment unemployment.

The size of the vocabulary set, $\mathcal{V} = \{\text{unemployment, inflation}\}$, is two. Note that, at a glance, documents y and z seem relatively *similar* in word usage, whereas x and z seem to be the most dissimilar in this regard. Once a corpus' vocabulary set is created, the next step is to create a *term* x *document* matrix—that is, a matrix with $|\mathcal{C}|$[52] columns and $|\mathcal{V}|$

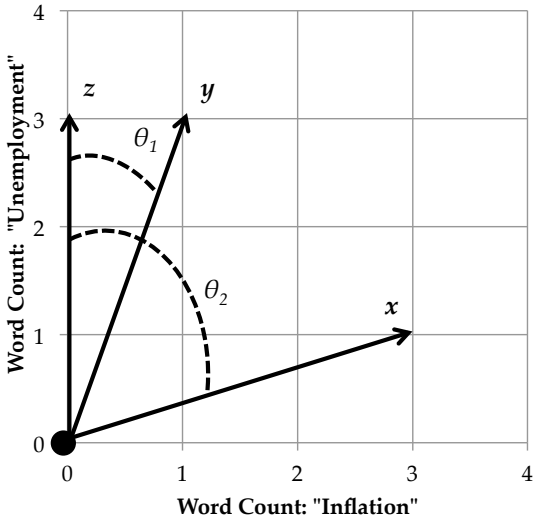

Figure 7: Vector representations of the documents x, y and z.

rows. The entries in the matrix are the number of times that word i appears in document j. An objectionable yet necessary part of vector-space models is that syntactical elements such as capitalization, punctuation, and, most notably, the order of words are ignored. For this example, the matrix is given by

$$\begin{array}{l} \\ \text{Unemployment} \\ \text{Inflation} \end{array} \begin{array}{c} \begin{array}{ccc} x & y & z \end{array} \\ \begin{bmatrix} 1 & 3 & 3 \\ 3 & 1 & 0 \end{bmatrix} \end{array}.$$

Hence, the "vector-space" terminology: each column of the matrix (a vector) corresponds to a document—this is the document's *vector-space representation.* For these documents, consider

[52] $|\cdot|$ is the cardinality operator.

figure 7, in which each document's vector has been plotted in *inflation-unemployment* word count space. In this representation, it is clear that y and z look "closer together" than x and z do, corresponding to the earlier claim about the relative similarities of x, y, and z. In order to quantify this similarity while controlling for the relative magnitude (length) of each vector, the angles, θ_i, that lie between each of the documents can be measured.

For this example, $\text{sim}(z, y) = \cos(\theta_1) = 0.95$ and $\text{sim}(x, z) = \cos(\theta_2) = 0.32$, again confirming the intuitions about the relative similarity of these documents.

A.3 Selection of k in the SVD

Application to the Transparency Corpus Recall figure 1. First, starting from the

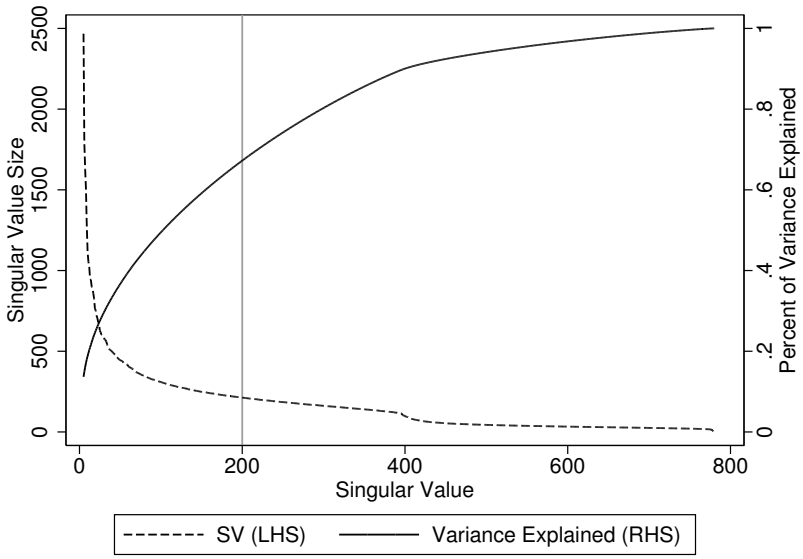

Figure 8: This graph shows the singular values (the elements $\sigma_i \in \Sigma$, $i = 1, \ldots, 780$, from the SVD) and the percentage of the variance explained by these singular values. The plot of the singular values (solid line) excludes the first four SVs.

right of the graph, notice that the data initially lies in 780 dimensions, one for each minutes and one for each transcript for each of the 390 meetings, and the inclusion of all 780 dimensions accounts for all of the variation in the data. The LSA literature has not conclusively determined the "optimal number" of dimensions to retain, although generally a few hundred dimensions are chosen.[53] For this sample, "a few hundred" seems to be a good choice, especially given the relatively sharp drop in singular value significance around the 400th dimension and the stabilization of the singular values between 100 and 300. By the 390th dimension, just under 90% of the variance in the original data has been retained.[54] This drop by itself is indicative of the structure of the documents; because there are 390 meetings in the sample, and because minutes and transcripts are, ideally, accounts of the same meetings, this corpus should lie in roughly 390-dimensional space. This space is not the "latent semantic space," it is simply a reflection of the fact that there were 390 meetings. Entering the "latent semantic space," then, requires retaining fewer than 390 dimensions.[55] Thus, the first $k = 200$ dimensions are retained; the noisiest 580 dimensions have been discarded when recomposing the dimension-reduced, C_k, matrix. Despite the elimination of so

[53] According to (Landauer et al., 1998, p. 269), "the number of dimensions retained in LSA is an empirical issue."

[54] See footnote 25 to see how this is calculated.

[55] The use of "should" and "required," and the discussion in general, is not so much a scientific result, but more of a philosophical or linguistic approach that assumes the existence of such semantic spaces and noise in natural language.

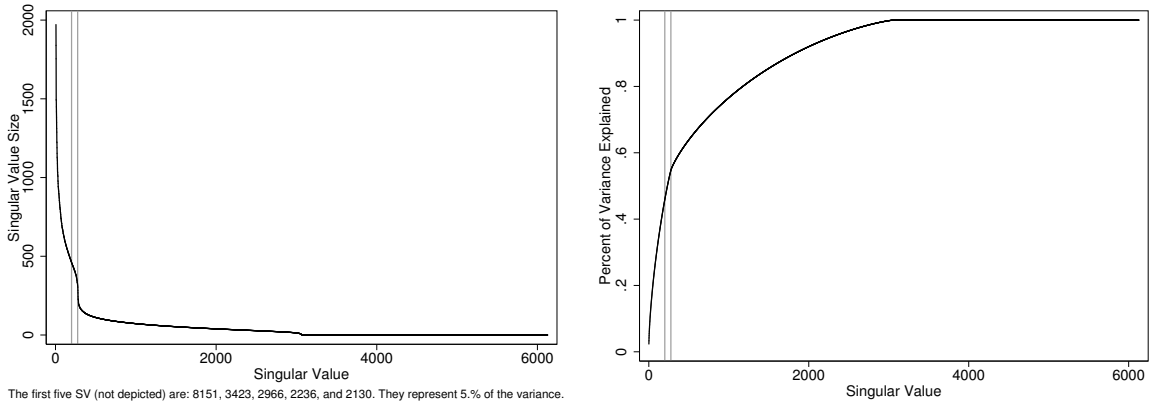

Figure 9: The left-hand side shows all but the first 5 singular values (the elements $\sigma_i \in \Sigma$ from the SVD) from the member corpus. The right-hand side graph show percentage of the variance explained by all 6128 singular values for the member corpus. Vertical lines are drawn at 200 and 276.

many dimensions, the remaining dimensions account for 64 percent of the variance in the data. Each additional dimension accounts for less than $1/10$ percent of the variance.

Application to the Conformity Corpus Figure 9 visually portrays the singular values (SVs) of the "member corpus" used for the computation of the conformity measure. Similar to the findings for the previous corpus, there is a noticeable drop in SV sizes at around 276; this again confirms that the documents are all accounts of the same 276 meetings. Thus, in order to eliminate the noise created by separating transcripts into m_{jt} and $m_{(-j)t}$, no more than 276 factors should be retained; at this point, 55% of variance in the data has been explained. In order to enter the latent semantic space, then, 200 factors are retained. The right-hand-side graph confirms the decreasing explanatory power of the factors after the "kink" at 276. The choice of 200 is in line with the "few hundred" factors that are typically retained; further, the singular values have largely stabilized by this point, and 46% of the variance has been explained. Notice that after 3064, the singular values are zero; since the rank of C is half of the total number of documents. [56]

[56]To be precise, calling the SVs past 3064 "singular values" is incorrect because the number of SVs can be no greater than rank of the matrix. The use of this term serves to ingrain the concept of a singular value and rank. To see why the statement about the rank of C is true, notice that any member document m_j can be written as $m_j = m_o + m_{(-o)} - m_{(-j)}$ for some other member document, m_o, in the same meeting; but this violates the definition of linear independence and thus lowers the rank of the matrix. Eliminating all member documents m_j while leaving $m_{(-j)}$ would leave a full rank matrix that had half as many documents as the original; therefore, the rank of the C matrix is half the number of documents, or 3064.

A.4 Stoplist

Stopwords (Alphabetically L to R)						
a	about	again	all	am	an	and
angell	any	april	are	aren't	as	at
august	balles	baughman	be	because	been	being
bernanke	between	bies	black	blinder	boehne	both
boykin	broaddus	burns	but	by	chairman	coldwell
corrigan	december	did	didn't	do	does	doesn't
doing	don't	down	during	each	eastburn	eighth
eleventh	evans	february	ferguson	few	fifth	first
fisher	for	ford	forrestal	fourth	from	further
gardner	geithner	governor	gramley	gramlich	greenspan	guffey
guynn	had	hadn't	has	hasn't	have	haven't
having	he	he'd	he'll	he's	heller	hendricks
her	here	here's	hers	herself	him	himself
his	hoeing	hoenig	holland	horn	hoskins	how
how's	i	i'd	i'll	i'm	i've	if
in	into	is	isn't	it	it's	its
itself	jackson	january	johnson	jordan	july	june
keehn	kelley	kimbrel	kohn	kroszner	lacker	laware
let's	lilly	lindsey	march	martin	may	mayo
mcdonough	mcteer	me	melzer	meyer	miller	minehan
mishkin	more	morris	moskow	most	mr	mr.
mrs	mrs.	ms	ms.	mullins	mustn't	my
myself	ninth	no	nor	not	november	october
of	off	olson	on	once	only	or
other	ought	our	ours	ourselves	out	over
own	parry	partee	phillips	pianalto	poole	rice
rimbrel	rivlin	roberts	roos	rosengren	same	santomero
schultz	second	secretary	seger	september	seventh	shan't
she	she'd	she'll	she's	should	shouldn't	sixth
so	solomon	some	stern	such	syron	teeters
tenth	than	that	that's	the	their	theirs
them	themselves	then	there	there's	these	they
they'd	they'll	they're	they've	third	this	those
through	timlen	to	too	twelfth	under	until
up	very	vice	volcker	walich	wallich	warsh
was	wasn't	we	we'd	we'll	we're	we've
were	weren't	what	what's	when	when's	where
where's	which	while	who	who's	whom	why
why's	willes	winn	with	won't	would	wouldn't
yellen	you	you'd	you'll	you're	you've	your
yours	yourself	yourselves				

Total: 283

A.5 Evidence of the Relevance of Procedural Transparency

Defining Transparency When used in the context of monetary policy, the word "transparency" can carry different connotations. Thus, to understand how the term is used here, table 8 presents the five forms of transparency relevant to central banks, as defined by Geraats (2002).[57] Procedural transparency is what allows for the main analysis of this paper—the

Table 8: The types of transparency relevant to central banks (Geraats, 2002, pp. F540)

	Type of Transparency **and Description**
1.	*Political Transparency* refers to openness about policy objectives and institutional arrangements that clarify the motives of monetary policy makers. This could include explicit inflation targets, central bank independence and contracts.
2.	*Economic Transparency* focuses on the economic information that is used for monetary policy, including economic data, policy models and central bank forecasts.
3.	*Procedural Transparency* describes the way monetary policy decisions are taken. This includes the monetary policy strategy and an account of policy deliberations, typically through minutes and voting records.
4.	*Policy Transparency* means a prompt announcement and explanation of policy decisions, and an indication of likely future policy actions in the form of a policy inclination.
5.	*Operational Transparency* concerns the implementation of monetary policy actions, including a discussion of control errors for the operating instrument and macroeconomic transmission disturbances.

study of FOMC minutes and transcripts. It encompasses the procedure by which the accounts of FOMC decisions are released to the public via documents. What makes procedural transparency important is that increased procedural transparency presumably leads to increases in the other four types of transparency; it is the mechanism through which the other four are manifested. Specifically, as one of the broadest categories, it demands that the other forms be augmented; openly discussing how decisions are made requires an explanation—this would come in the form of, at the very least, greater political and economic transparency (that is, explicitly stating the data and objectives used). Of course, monetary policymakers are known for their carefully chosen words and strategic lack of specificity at times, but it would seem to be the case that procedural transparency is the channel through which the other, perhaps more economically important forms of transparency are augmented. For example, mandating that the Fed release the theoretical rule it uses to determine policy hinges on effective procedural transparency, though it is a form of economic, policy, and political transparency.

At a more general level, there are a few reasons to believe that procedural transparency is worth studying. First, the public pays attention to the documents released by the Fed. Some anecdotal evidence here is a graph from Google Trends—a service from Google that

[57] Geraats has written much about central bank transparency. See Geraats (2000), where these terms were first defined, or Geraats (2007) for other examples.

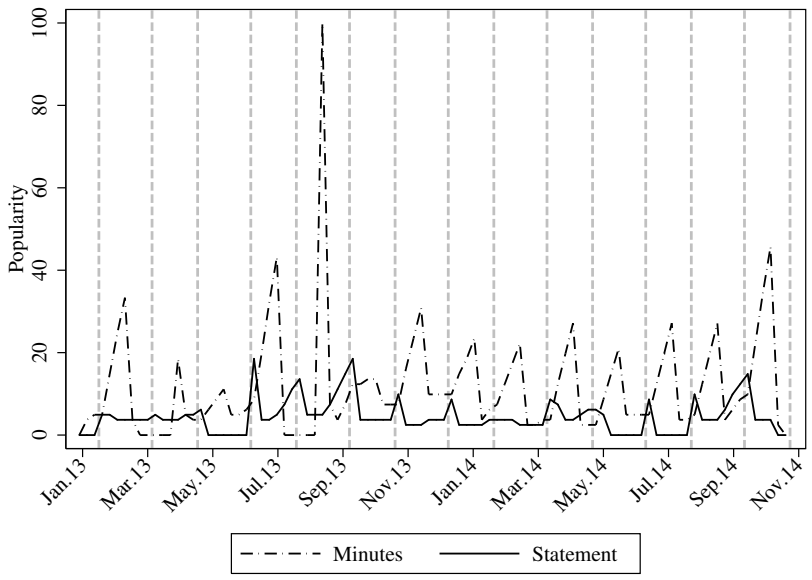

Figure 10: Results from the Google Trend queries "fed statement" and "fed minutes." The y-axis represents the frequency with which a given phrase is searched on Google, and is normalized so that the highest frequency is 100. Thus, this graph does not say how these search terms rank among all other terms, but it does give information about when the terms are searched. Dashed lines are included at FOMC meeting dates.

plots the "interest over time" of any term. The graph, shown in figure 10, shows how popular "fed statement" and "fed minutes" were over a 22-month period between 2013 and 2014. As expected, peaks in interest in the terms is in one-to-one correspondence with FOMC meetings; the three-week lagged release of minutes is also clearly noticeable. So, at the very least, there appears to be public interest in the content of FOMC documents.

Next, while not typically on the FOMC agenda, discussions about its document releases have occurred numerous times in the Board Room. These are meaningful discussions, going beyond the procedural and suggesting that, at a minimum, the FOMC believes that the public cares about its discussions. The following quote from former Vice Chairman Blinder supports this claim:

> Mr. Blinder. I believe...that we now have a situation where the people that speak the least about the Fed's decisions are those at the Fed, and we are interpreted voluminously. There is nothing wrong with that. We will still be interpreted voluminously even if we say things. But our statement is a chance for us to say what we are up to and why. (Transcripts, 2008, 01/31/1995.)

Further, more comprehensive evidence shows that the FOMC was indeed discussing issues of procedural transparency in its meetings. Figure 11 uses the text-analysis techniques described below to infer the extent to which topics of procedural transparency were discussed at each FOMC meeting. This is measured by comparing a query of procedural-transparency-related words (*transcripts, meetings, statement*, etc.) to the transcripts of each meeting. In order to determine the relevance of a query to each transcript, the cosine similarity of that query is taken against each of the transcripts in the low-dimension space—this was the

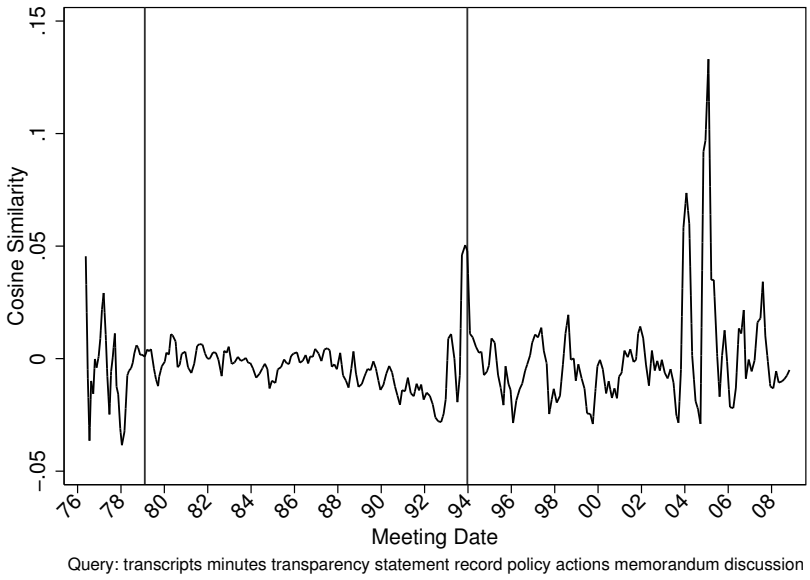

Figure 11: Similarity of the query with the content of each FOMC meeting. The "cosine similarity" says how similar the query is to each document, which in turn describes how much the ideas in the query were discussed in the meetings.

original motivation behind the development of LSA. If the C matrix $C \in \mathbf{R}^{t \times d}$, then the query vector $q \in \mathbf{R}^t$ has components that are mostly 0's except for the terms used in the query. Mapping the query q into its lower-dimensional space representation, q_k, is done by letting $q_k = \Sigma_k^{-1} U^T q$; at this point, it now has the "estimated" term values discussed above—this process is called "folding in."[58] The cosine similarity is then computed against the other documents in the matrix $\Sigma_k V^T$, which is equivalent to computing the similarities in the reconstructed C_k matrix. When interpreting figure 11, recall that when the similarity is around 0, this query is irrelevant to the meeting proceedings (it is orthogonal). Therefore, positive or negative deviations suggest a query whose presence is noticeable in the discussion. It is clear that topics of procedural transparency have, at times, accounted for significant portions of the discussion at FOMC meetings; without smoothing, some of the similarity measures are over 0.2. Furthermore, the discussion coincides with procedural changes in policy—1976 marked the temporary end of transcript publication, and topics of procedural transparency persisted for a few years after that change. Since the early '90s, changes in publication policy have been relatively frequent—in 1993, old transcripts were released; in 1995, new transcripts began to be released; and in 2005, minutes were released with shorter lags. Most of these changes are visible on the graph, indicating that a significant amount of discussion was behind each decision.

In summary, this evidence supports the claim that studying procedural transparency is worthwhile. Both the public and the FOMC pay careful attention to these documents, so

[58] Setting $q_k = \Sigma_k^{-1} U^T q$ is equivalent to taking the familiar least-square approximation of q in k-space: $q_k = (C_k^T C_k)^{-1} C_k^T q = V^T \Sigma_k^{-1} U^T q$, where the leading V^T can be eliminated when taking cosine similarities because the inner product is invariant under multiplication by an orthonormal matrix.

understanding how and why this type of transparency changes is certainly meaningful.

A.6 Summary Document Robustness

As a robustness check to using a summary document to study the behavior of the transcripts (and, analogously, minutes) over time, comparisons were made between all transcripts and the first transcript in the sample. The same was done with the last transcript. In both cases, there was relatively high, but falling, similarity between transcripts within two years; two years away, the similarities became relatively stable at low levels, while continuing to decline as time passed. Panels [A] and [B] in figure 12 show these results. In both panels [A] and [B],

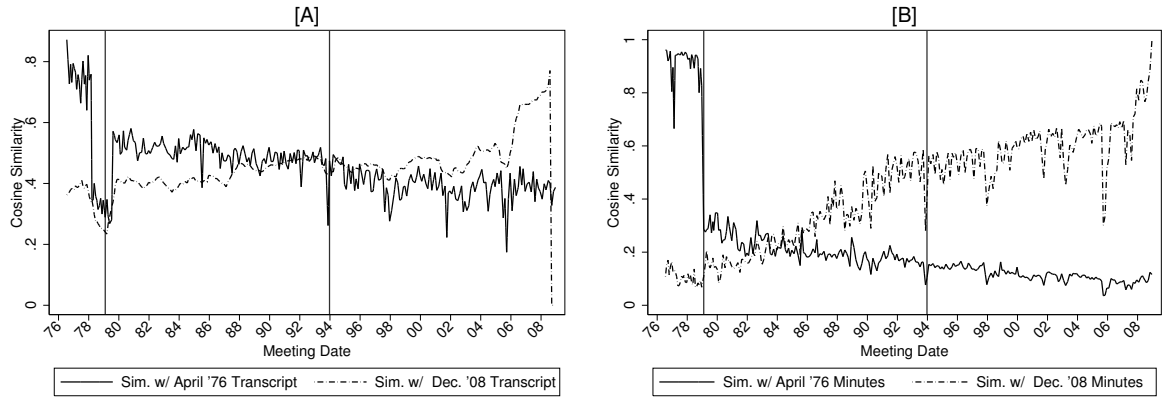

Figure 12: Cosine similarities of transcripts and minutes ([A] & [B], respectively) and first/last documents.

which plot the similarity of the documents with the first and last documents of the sample period, the lines are downward-sloping with respect to the reference document; this was the prediction made about comparing documents pairwise to a single document—namely, that the similarities would decrease as documents became more distant in time from the reference document (first or last).

Table 9: Average Similarities

Document Type	Average Similarity
First Minutes	0.26
Summary Minutes	0.69
Last Minutes	0.42
First Transcripts	0.57
Summary Transcripts	0.81
Last Transcripts	0.57

Average of similarities between the relevant document (that is, the first transcripts) and all other documents of that type (that is, all transcripts).

Table 9 gives a summary of the average similarity measures over time for each of the document types considered for this analysis (first, last, and summary); as predicted, the highest similarities are found when comparing minutes/transcripts to the summary documents.