A New Measure of Central Bank Transparency and Implications for the Effectiveness of Monetary Policy

Miguel Acosta*
Columbia University

January 15, 2022

Abstract

Transparency has been posited as a channel through which monetary policy is made more effective. However, empirical studies of this question and other questions concerning the role of transparency have lacked access to a time-varying high-frequency measure of transparency. This paper presents a new measure of the transparency of Federal Reserve deliberations, derived from the documents that the Fed uses to record and summarize each of its meetings. The measure—the similarity of the minutes and transcripts of each Federal Open Market Committee (FOMC) meeting—is largely, though not entirely, shaped by FOMC leadership. Monetary policy shocks have about a 40 percent larger effect on nominal and real interest rates when the prevailing level of transparency is high, suggesting an important role for transparency in determining the efficacy of monetary policy. These effects are primarily driven by transparency about monetary policy strategies conditional on the state of the economy.

Keywords: Federal Open Market Committee, Transparency, Natural Language Processing, Central Bank, Monetary Policy

JEL Codes: E58, H83, D78, D82

*Columbia University, Department of Economics. m.acosta@columbia.edu. This paper previously circulated under the title "FOMC Responses to Calls for Transparency" (Acosta, 2015). That incarnation of the paper benefited greatly from discussions with my undergraduate honors thesis advisor, John Taylor; the support of my parents and sister: José Acosta, Mary Grebenc, and Lourdes Acosta; and the helpful comments of Lydia Cox, Geoffrey Cox, Marcelo Clerici-Arias, Chris Gust, José Gutierrez, Donna Hunter, Stephanie Kalfayan, Mark Lucianovic, Ellen Meade, the late Allan Meltzer, Elmar Mertens, Kurt von Tish, the Bing Honors College (2013) participants, and seminar participants at the Federal Reserve Board. This version was further improved by participants in the Monetary Economics Colloquium at Columbia, as well as by discussions with Pierpaolo Benigno (editor), Lydia Cox, Juan Herreño, Jennifer La'O, Frederic Mishkin, José Montiel Olea, Emi Nakamura, Andrea Prat, Ricardo Reis, Martin Rotemberg, Joseph Saia, Jón Steinsson, and several anonymous referees. A huge thanks goes to Stephen Hansen and Lucas Husted, who graciously provided me with data, and Blaine Helleloid, who provided outstanding research assistance. This work was supported financially by the Alfred P. Sloan Foundation Pre-Doctoral Fellowship in Behavioral Macroeconomics awarded through the NBER.

1 Introduction

Over the years, the world's major central banks have become more transparent in many respects—explicit inflation targets are well-established, policy actions are announced, and forecasts are provided to the public. An aspect of transparency that is more difficult to quantify, however, is the extent to which the publicly-provided rationales for policies reflect the reaction function of the monetary policy committee. As stressed by Woodford (2001, 2005) and numerous speeches by Federal Reserve Chairs,[1] communicating these rationales helps the public to better understand how policymakers might react to different future states of the economy, which in turn gives monetary policy greater control over longer-term interest rates. In this paper, I provide the first high-frequency measure of this "procedural transparency." With a quantitative measure in hand, we can then tackle the question of whether transparency make monetary policy more effective. The second part of this paper provides an answer in the affirmative.

I construct the measure of procedural transparency—henceforth referred to as "transparency" unless otherwise noted—using the documents that Federal Open Market Committee (FOMC) releases to the public as records of its policymaking meetings. Since 1976, the FOMC has consistently recorded nearly-verbatim records of its meetings in documents called the *transcripts* and shorter summaries called the *minutes*. The transcripts, while detailed accounts of how the Committee comes to decisions, are not released until at least five years after a meeting has taken place. (Before 1993, they were not expected to be released at all.) The lag between meetings and the release of the transcripts renders much of the information they contain stale in terms of understanding the Committee's current thinking. The minutes, on the other hand, are released with a much shorter lag—today, three weeks. My measure of transparency takes advantage of this timing: it is the similarity of the minutes and transcripts of each meeting, computed using natural language processing (NLP) techniques.

Minute-transcript similarity and the distribution of the topics in the minutes and transcripts are most-strongly associated with FOMC leadership. Minute-transcript similarity was at its highest during the late Greenspan years (early 2000s), and increased noticeably after 1993, the year in which

[1] See, for example, Bernanke (2010, 2013, and Yellen (2013). The aspects of transparency described by these papers is not *exactly* "procedural transparency" as used here, but it is argued below that procedural transparency is the relevant form of transparency for explaining the rationale of policy, which is what these papers and speeches broadly consider.

the FOMC began publishing the records of its meetings. This is consistent with predictions from earlier studies of this event. The distribution of topics in the transcripts also changed significantly in 1993. Minute transcript similarity is not predictable by macroeconomic variables, though changes in the distributions of the minutes and transcripts are. Higher transparency is also weakly associated with monetary shocks—as measured by several authors—that are smaller in magnitude, lending credence to the notion that the measure helps to inform the public about the Fed's policy. In sections 2 and 3, I describe the construction of the measure and its properties, respectively, to establish that it provides a meaningful measure of Fed transparency.

In its most literal interpretation, minute-transcript similarity captures the overlap between the distributions of topics discussed in the transcripts, and those discussed in the minutes. My measure should therefore be interpreted as a way to understand the divergence between what receives the Committee's attention when a decision is being made, and how that thinking is described.[2] Given a higher level of transparency, then, the public and financial market participants should be able to better-predict the Fed's policies, as elevated transparency implies that the Fed's communications are providing clearer insights into policymakers' thinking. This motivates the main empirical question of the paper: how does transparency affect the effectiveness of monetary policy, as measured by the pass-through from short-term nominal rates to long-term nominal *and* real rates?

Section 4 turns to the role that transparency plays in determining the effectiveness of monetary policy. Specifically, I show that the monetary policy shocks of Nakamura and Steinsson (2018) have larger effects on real interest rates when transparency is elevated. In sections 5.1 and 5.2, I show that these results are primarily driven by transparency about monetary policy strategies conditional on the economic outlook. Additionally, the effects of monetary policy estimated by Nakamura and Steinsson (2018) are shown to be slightly downwardly biased. This arises because some of the larger monetary shocks have been delivered at times when transparency is low. But, it is precisely when transparency is low that monetary policy shocks have smaller effects on real interest rates, possibly

[2]Note that a clear, concise, and informative summary of a long discussion—a discussion possibly filled with tangents, misunderstandings, and other banter—should not be expected to cover the same topics in the same proportion, though my proposed measure would penalize such deviations. Variation in minute-transcript similarity could arise from the noise present in natural language/conversation, variation in how difficult a meeting is to describe (say, for example, because of a complicated policymaking environment), or strategic considerations regarding the transmission of information. That said, to my knowledge, there is no evidence that the minutes are anything but a forthright effort to summarize the transcripts (though, such evidence might be hard to find). Additionally, I find in section 5.2 that my results are robust to using measures of transparency about obviously-meaningful topics—i.e., measures that should be less-affected by these concerns.

because the public cannot make as much sense of the short-run surprises as they relate to the path of future rates. These findings are robust to concerns that transparency may be proxying for some other variable—forward guidance, the state of the economy, and public uncertainty about monetary policy.

2 Measuring Procedural Transparency

2.1 Previous Measurements

The literature on central bank transparency and communication started in earnest at the turn of the century—Blinder et al. (2008) provide a thorough survey of this literature through 2008. This paper contributes to a branch of this literature concerning the measurement of transparency—most recently treated by Dincer and Eichengreen (2014), preceded by Eijffinger and Geraats (2006),[3] who based their measures on the Geraats (2001, 2002) definitions of different aspects of central bank transparency. This paper centers around a particular component of Geraats' *procedural transparency* that concerns central bank accounts of deliberations, the measurement of which has, thus far, focused primarily on fairly aggregated and slow-moving measures of the timeliness and informativeness of central bank communications.[4] 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[.]" In this paper, the use of natural language processing techniques allows these shortfalls to be circumvented by using the text of each FOMC meeting. Because the measure is constructed from text, it reflects fairly detailed changes in communications and transparency (it is real valued, as opposed to taking a discrete number of values). Because it changes at every FOMC meeting, it gives a high frequency measure of transparency that, in practice, changes much more often than previous measures that primarily capture large regime changes in communications policies.

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

[4] See appendix B for a discussion about the different types of transparency relevant for a central bank. Procedural transparency, as defined by Geraats (2002) is the description of how monetary policy decisions are made, which is achieved in part through the publication of records of the deliberative process.

2.2 How to Measure Transparency

My proposed measure of procedural transparency is the similarity between the minutes and transcripts of each FOMC meeting, which I will refer to as "minute-transcript similarity" or simply "transparency" for the rest of the paper. The transcripts contain a nearly-verbatim record of each FOMC meeting, yet are not released until at least five years after an FOMC meeting has taken place—before 1993, they were not released at all.[5] The minutes are shorter summaries—typically 10-20 pages to summarize 100-200 page transcripts—released three weeks after each meeting has taken place.[6] Evidence suggests that the minutes are intended to be accurate portrayals of what was discussed at each meeting, and are not intended to obfuscate the content of the discussion. Chair Yellen, questioned about this in her June 2016 press conference, responded that "[T]he minutes are always—have to be an accurate discussion of what happened at the meeting."[7] In addition, the FOMC has to vote on the minutes, presumably reducing the possibility of systematic obfuscation.

The first step in computing minute-transcript similarity is to represent each transcript and minutes as a distribution over a finite number of topics, using Latent Dirichlet Allocation. The similarity between the minutes and the transcripts for a particular meeting is the Kullback-Leibler similarity of the distributions of the two documents—a measure that lies in the interval $[0, 1]$. The measure bears a striking resemblance to Chair Greenspan's interpretation of how the public understood FOMC communications, which he voiced in the September 2003 FOMC meeting: "[A] number of those in the market don't listen to the subtleties; they just take note of how much time we are spending talking about a particular subject" (Transcripts, 1976–2008).

Before detailing the procedure used to compute minute-transcript similarity, it is important

[5]"Nearly verbatim" refers to the fact that the exact words are lightly edited. From the Federal Reserve Board's website: *The most detailed record of FOMC meeting proceedings is the transcript. Beginning with the 1994 meetings, the FOMC Secretariat has produced the transcripts shortly after each meeting from an audio recording of the proceedings, lightly editing the speakers' original words, where necessary, to facilitate the reader's understanding. Meeting participants are given an opportunity within the subsequent several weeks to review the transcript for accuracy.*

For the meetings before 1994, the transcripts were produced from the original, raw transcripts in the FOMC Secretariat's files. These records have also been lightly edited by the Secretariat to facilitate the reader's understanding. In addition, where one or more words were missed or garbled in the transcription, the notation "unintelligible" has been inserted. In some instances, words have been added in brackets to complete a speaker's apparent thought or to correct an obvious transcription error or misstatement.

[6]The release lag has changed over the years–see A for more details.

[7]This is not a recent sentiment, in 2003, Kansas City Fed President stated in an FOMC meeting that "the minutes are our representation of the discussions that occurred at the meeting" (Transcripts, 1976–2008, Sept. 2003). Several other pieces of anecdotal evidence are available, including discussions I have had with Federal Reserve officials and a former Fed governor.

to understand what this measure is, and its potential shortcomings. The measure represents the extent to which the content of the transcripts is reflected proportionately in the minutes.[8] Because LDA represents each document as a distribution over topics, the mass placed on each topic describes the amount of the document devoted to that topic. Thus, only when the minutes devote exactly the same amount of space to each topic as the transcripts will the measure equal unity. The view taken in this paper is that it is not the job of the minutes writers to editorialize the FOMC's discussion—e.g. to eliminate the side of a debate that does not ultimately "win"—but instead to convey the discussion accurately. Put differently, minutes that fully communicate FOMC discussions are taken as transparent.

That said, the inclusion of obviously-irrelevant discussions—e.g. "when should we break for lunch"—should not be a necessary condition for minutes to be transparent. Of course some divergence between the two documents should be expected—while a conversation might be centered around a topic, the actual words used or topics discussed might only-noisily represent that topic—owing to, for example, digressions or misunderstandings. The underlying assumption I make is that this noise is fairly constant over time, only affecting the level of my measure and not its changes. In section 5.2, I show that my empirical findings are robust to more-narrowly defined measures of transparency that are constructed using economically important topics (and should thus be relatively free from this type of noise). This is a benefit of using the fairly complex language model described in the next section—it allows for documents to be analyzed about detailed topics and thus separate out the types of discussions that can add noise to my measure.

2.3 Language Model

Several steps are involved in computing the topic distribution for each FOMC transcript and minutes, with the ultimate goal being a representation of a document into well-understood topics in a high-frequency, holistic, and interpretable way. A drawback of the approach used here—along with the vast majority of NLP techniques—is that documents are represented as "bags of words," i.e., the order of words does not matter. The only features of a document that are retained are counts of the number of times that each unique word appears in that document. This is a necessary evil

[8]Note that the measure does not consider the similarity of documents *at different meetings*, only the similarity of documents related to the same meeting.

when working with a large body of text—the 600 documents used in the analysis here contain 13 million individual, and 170,000 unique, words.[9]

After converting the documents into computer readable formats, these raw text files are "preprocessed" using several techniques that are standard in working with natural language. This preprocessing achieves three goals. The first goal is to reduce the effect of errors that might arise from working with a data-source that may contain typographical errors or other errors arising from the fact that several of the documents had to be converted from typewritten documents. To that end, only the letters of the alphabet are retained, and every unique word must appear at least three times over the entire corpus, otherwise it is dropped. Words shorter than three characters or longer than fifteen are also dropped—the intention of the latter being to remove words that may have been accidentally concatenated.

The second goal of the preprocessing is to reduce the noise that arises from grammatical constraints: the words *increase, Increase!, increasing* and *increased* all convey essentially the same meaning, yet a simple numerical representation of the words in the document might treat them as completely different words, since it knows no better. To that end, words are stemmed to their lexical root, so that in the example above, every occurrence of the "increase" words are stemmed to *increas*.[10] Terms in a "stoplist" are also 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 are the "generic", "dates and numbers," and "geographic" lists from Loughran and McDonald (2011), who carefully constructed these lists to be relevant in a context of finance.

The final goal of preprocessing is to reduce the noise that arises when ideas need to be mapped into words, and vice versa. This goal is addressed via an application of Latent Dirichlet Allocation, developed by Blei et al. (2003).[11] Because LDA enjoys widespread use in the NLP community, and even within the Economics literature, the treatment here is brief. First, the observed corpus contains words, with $w_{d,n}$ being n^{th} word in the d^{th} document ($d \in \{1 \ldots, D\}$), where each document has N_d (so that $n \in \{1, \ldots, N_d\}, \ \forall d$). This is all that is observed. LDA posits each document as a

[9] This is based off of simply splitting the documents by whitespace, i.e., with no preprocessing.

[10] Stemming is performed using the Lancaster Stemmer as implemented in Python's Natural Language Toolkit.

[11] At issue here are the problems of synonymy and polysemy. 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 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 dissimilar, despite conveying the same idea.

distribution over a fixed number, K, of topics—K is chosen by the researcher. Topics are in turn distributions over the V unique terms in the corpus. More precisely, each document δ_d is a draw from a Dirichlet distribution, a distribution over vectors that lie in the K-simplex. The distribution, δ_d, that is drawn from the Dirichlet is a latent variable. The same is true of topics: each topic, ϕ_k is a draw from a V-dimensional Dirichlet. With a topic distribution in hand, each observed document (which has a fixed length, N_d) is populated one word at a time. For the n^{th} slot of document d, a topic is drawn from a multinomial distribution, with parameter δ_d. Thus, if δ_d is heavily concentrated on topic 1, then several words will be drawn from topic 1. The drawn topic, $z_{d,n} \in \{1, \ldots, K\}$, is then used to draw a word from a multinomial distribution with parameter $\phi_{z_{d,n}}$. So, if a topic has a distribution that places heavy weight on "whale," then "whale" will come up often when that topic is drawn, and will thus show up often in documents that have a high probability placed on that topic. The implementation here—including the choice of the priors for the Dirichlet distribution—follows very closely that of Hansen et al. (2017), who estimate LDA on a subset of the FOMC transcripts that I consider using a Gibbs-sampler outlined in Griffiths and Steyvers (2004).[12] The number of topics, K, is set to 50, chosen using a five-fold cross-validation technique similar to that outlined in Hansen et al. (2017) and described in C.

Rather than estimating the topic model over the complete documents in my corpus, I instead begin by splitting these documents (the minutes and transcripts) into sentences using a grammatical sentence parser.[13] I then estimate the word and topic distributions over every sentence in the FOMC minutes and transcripts between 1976 and 2014. In the notation established above, then, d indexes sentences so that each sentence has an estimated distribution over topics. With these estimated distributions, I then estimate the topic distributions for each transcript as a whole ($\theta_t \in \mathbb{R}^K$), and each of the minutes as a whole ($\mu_t \in \mathbb{R}^K$). This follows the approach of Hansen et al. (2017), the purpose being to have each document focused on a small number of topics, in hopes that the latent topics can be determined more easily.[14]

[12] And very graciously provided by the authors at https://github.com/sekhansen/text-mining-tutorial.

[13] I parse the minutes and transcripts into sentences using the English probabilistic context-free grammar developed by Klein and Manning (2003) as implemented in the Stanford Parser Java package. This sentence parser uses rules of English grammar to split sentences, as opposed to simple rules based solely on punctuation. This, for example, avoids erroneously splitting sentences at decimal points or after abbreviations(e.g. "Ms.").

[14] This substantially increases D to about 260,000, though there are 300 minutes and 300 transcripts. This approach of using already-estimated topics to estimate the topic distribution of an excluded document underlies the right panel of figure 10, with the excluded document reading "transcripts minutes record policy actions memorandum discussion communications."

Table 1: Topic Descriptions

Topic	Top 10 Words										Dissent	FFR	GDP	Inflation
1	level	unemploy	percent	rate	employ	low	averag	high	declin	remain				
2	credit	bank	spread	loan	debt	investor	mortgag	bond	larg	money			✓	
3	hous	sale	sector	inventori	activ	declin	construct	home	start	weak			✓	
4	meet	statement	altern	committe	chang	languag	fomc	paragraph	direct	word	✓		✓	
5	think	make	commun	public	view	way	help	decis	use	forward	✓			
6	market	financi	labor	particip	improv	condit	forc	gener	substanti	note	✓			
7	purchas	asset	treasuri	end	secur	program	hold	back	reduc	size			✓	
8	spend	busi	consum	invest	incom	capit	household	confid	increas	continu	✓			
9	risk	balanc	concern	view	side	possibl	sheet	downsid	outlook	signific				
10	time	take	need	think	begin	adjust	process	place	step	normal				
11	chang	littl	move	think	bit	probabl	reason	case	quit	fact		✓	✓	
12	like	look	see	come	go	time	forward	back	number	reason				
13	polici	monetari	action	fiscal	accommod	tighten	eas	appropri	addit	support	✓	✓	✓	
14	line	dollar	right	panel	lrb	rrb	shown	foreign	show	left		✓		
15	term	long	expect	run	longer	short	consist	near	time	object	✓			
16	growth	continu	economi	pace	slow	moder	product	expans	trend	econom		✓	✓	
17	reserv	bank	oper	feder	system	account	open	central	facil	direct	✓		✓	
18	recent	data	month	suggest	year	past	indic	survey	seen	evid				
19	peopl	lot	get	work	think	way	talk	just	say	thing	✓			
20	price	increas	cost	higher	pressur	rise	energi	oil	wage	declin		✓	✓	✓
21	economi	uncertainti	problem	situat	world	face	shock	deal	event	creat				
22	forecast	project	growth	real	gdp	greenbook	staff	point	half	revis	✓			
23	effect	demand	import	factor	reflect	part	posit	suppli	export	state		✓	✓	
24	inflat	expect	percent	measur	core	remain	pce	rang	higher	low				✓
25	rate	fund	point	interest	basi	feder	target	lower	fed	rais	✓			
26	report	good	product	industri	firm	district	manufactur	order	contact	nation	✓			
27	differ	use	model	base	output	gap	rule	structur	estim	actual				
28	question	presid	issu	chairman	discuss	comment	said	governor	mention	ye			✓	
29	think	go	want	say	know	just	dont	get	realli	way			✓	
30	econom	period	committe	outlook	member	develop	stabil	inform	intermeet	consider				✓

The "top 10 words" column of 1 contains the ten words that most-prominently contribute to the composition of each topic. Formally, these are the words corresponding to the ten largest elements of ϕ_k for each topic k. For purposes of interpretation I show the estimated topics for a smaller model with $K = 30$ (my empirical results in section 4 are nearly unchanged quantitatively using the larger model) that I estimated over the sample beginning in 1995 since that forms the basis of my empirical estimates in section 4. In general the topics seem intuitive–the first four topics might be subjectively called topics about labor, credit markets, housing markets, and policy statement language. Not every topic is directly related to an economic concept—the fifth topic contains words that might be used in a debate regarding policy communications.

To better aid in categorizing the estimated topics, the last four columns of table 1 indicate which of the topics are useful for "predicting" a few external observable variables, denoted by e_t. Specifically, I gather the number of dissenting votes at each FOMC meeting from Thornton and Wheelock (2014); and the change in the target of the Federal Funds rate, GDP growth, and inflation from FRED.[15] I then estimate which topics are useful predictors of these variables by estimating

[15]In terms of FRED mnemonics, the FFR target is `DFEDTAR` when it is available and the midpoint of `DFEDTARL` and `DFEDTARU` when it is not; GDP growth is the four-quarter difference in the log of `GDPC1`; and inflation is the twelve-month difference in the log of `PCEPI`.

the following regression using the LASSO objective function

$$e_t = \beta_0 \mathbf{1}\{\text{ELB}\}_t + \sum_{k=1}^{K} \beta_k \delta_t^k + \text{error}_t$$

where δ_t^k is the estimated presence of each topic k in document d. I estimate this regression twice: once using the presence of topics in the minutes (i.e. replacing δ_t^k with μ_t^k) and once using their presence in the transcripts (i.e. replacing δ_t^k with θ_t^k). I take the union of the selected topics between these two regressions. I select the LASSO regularization parameter using ten-fold cross validation.[16] This again helps to shed light on the estimated topics. The selected topics for inflation are the most intuitive, with words like "price, inflat, econom" being the top words for the selected topics (20, 24, and 30). Instead topics about appropriate communication and policy (4, 5, 25), and the general outlook for the economy (most of the other selected topics) can predict the number of dissents at each meeting.

2.4 Transparency Index Definition

Ultimately, the object of interest in this paper is not the topics themselves, but rather the relative entropy of the minutes for the transcripts (i.e., the Kullback-Leibler similarity) and the entropies of the minutes and transcripts on their own. Specifically, given the topic distribution of the transcripts for the FOMC meeting occurring at time t, θ_t, and the topic distribution of the minutes for the same meeting, μ_t, define the three quantities:

$$\text{Minute-Transcript Similarity:} \qquad \tau_t \equiv \exp\left[-\sum_{k=1}^{K} \mu_t^k \ln\left(\frac{\theta_t^k}{\mu_t^k}\right)\right]$$

$$\text{Entropy of the Minutes:} \qquad H(\mu_t) \equiv -\sum_{k=1}^{K} \mu_t^k \ln\left(\mu_t^k\right)$$

$$\text{Entropy of the Transcripts:} \qquad H(\theta_t) \equiv -\sum_{k=1}^{K} \theta_t^k \ln\left(\theta_t^k\right)$$

with θ_t^k being the k^{th} element of θ_t, and analogously for μ_t. The first measure is the negative exponential of the Kullback-Leibler divergence, a distance function for distributions—intuitively it

[16] I include a time-varying constant since the FFR is one of my target variables. I implement the estimation and cross-validation using the LassoCV module in Python's sklearn package; the sample is split into ten disjoint subsets that are the same for each e_t.

Figure 1: Procedural Transparency

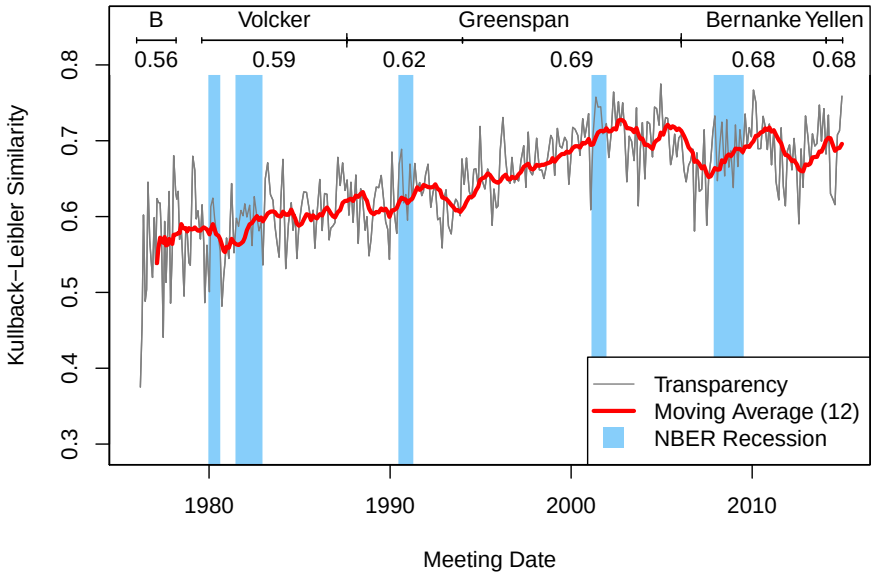

describes the information loss from assuming the truth is the minutes when it is really the transcripts. The other two measures, the entropies of the minutes and transcripts, are the expected values of the information content of a random variable that is distributed according to μ_t and θ_t. Entropy in this case achieves a maximum when both distributions place equal mass on each topic, i.e., $\mu_t^k = 1/K, \forall k$, and decreases as mass moves away from certain topics and concentrates on others.[17] Thus, entropy in this context can be cast intuitively as a measure of how dispersed a conversation/document is—the lower the entropy, the more concentrated the discussion.

For the construction of the procedural transparency measure, the body of documents under consideration consists of all Records of Policy Actions (ROPA, an older version of the minutes), and minutes and transcripts from meetings physically held in Washington, D.C., between April 1976 and December 2014, where procedural information (voting records, attendance) is removed.[18]

[17] Information about a variable drawn from a uniform distribution is more valuable than information about a variable drawn from a point mass–one already know where the point mass is, but has no idea where the uniform variable is.

[18] See A for a discussion of the various documents released by the FOMC since its inception. 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).

Figure 2: Entropy of the Minutes (Left) and Transcripts (Right)

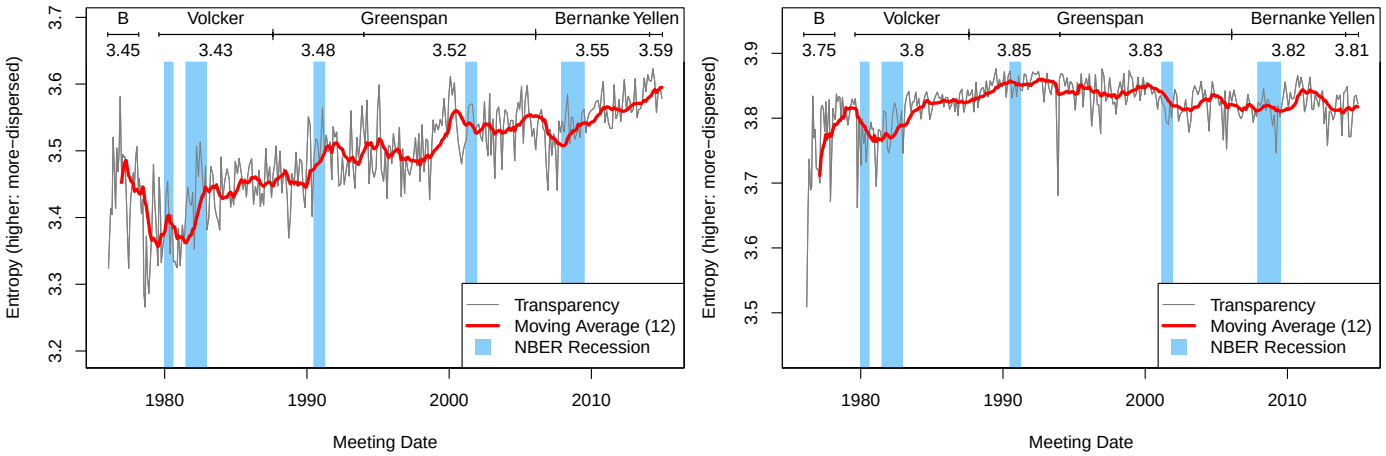

3 Notable Features of the Minutes, Transcripts, and Procedural Transparency

This section presents my measure of procedural transparency and three exercises meant to better understand it. First, in section 3.1 I compare the measure to several external variables, in order to understand its systematic components. My primary finding is that fluctuations in the measure, while correlated with some of these external variables, is not a proxy for something simpler. Next, in section 3.2 I highlight, anecdotally, that the measure captures meaningful differences between the minutes and transcripts. Finally, in section 3.3 I argue that my measure—though not directly observable in real-time—might be roughly observable to the public. This is done using a newspaper-based measure of central bank transparency.

3.1 Correlations with External Variables

Figure 1 presents the time series of minute-transcript similarity, and figure 2 contains the entropies of the minutes and transcripts. All three measures are rather noisily distributed around slower moving trends, shown in the images as twelve-meeting trailing moving averages. Owing to concerns—mentioned in section 2.2—that the raw measures are likely influenced by idiosyncratic noise arising from the noise inherent in natural language, the moving averages of these series are the main measures considered for the rest of the paper.

In order to analyze the series more carefully, tables 2 and 3 present the results of regressing

Table 2: Regression Coefficients for Communications Variables

	Transparency		Minutes Entropy		Transcripts Entropy	
		MA		MA		MA
t	**2.50**	**3.24**	**2.63**	**3.23**	1.07	1.41
Transparency MA	**0.33**	**0.43**	0.09	0.14	0.18	0.19
Minutes Ent. MA	0.13	0.17	0.26	0.43	0.17	0.16
Transcript Ent. MA	0.04	0.09	0.08	0.07	**0.61**	**0.92**
TT State	-0.14	-0.11	-0.10	-0.07	0.12	0.30
RR Shocks$_{t+1}$	-0.03	-0.04	-0.03	-0.02	0.07	0.06
| RR Shocks$_{t+1}$|	-0.08	**-0.16**	**-0.13**	**-0.22**	**-0.13**	**-0.21**
NS Policy Shock$_{t+1}$	0.06	0.13	0.06	0.11	0.06	**0.22**
| NS Policy Shock$_{t+1}$|	0.13	-0.01	-0.05	**-0.32**	0.08	-0.03
FFR	**-0.55**	**-0.77**	**-0.70**	**-0.88**	-0.20	**-0.37**
Change in FFR	-0.07	-0.04	-0.09	-0.04	0.01	-0.04
| Change in FFR |	**-0.17**	**-0.23**	**-0.27**	**-0.29**	**-0.24**	**-0.28**
Δy_t	-0.20	-0.16	-0.19	-0.14	0.06	0.16
$\mathbb{E}_{\text{GB}}[\Delta y_{t+1}]$	0.02	0.13	0.16	**0.30**	0.07	0.08
Unemployment (u)	-0.23	-0.24	-0.10	-0.13	-0.20	**-0.35**
$\mathbb{E}_{\text{GB}}[u_{t+1}]$	-0.26	-0.28	-0.18	-0.24	-0.22	**-0.39**
π_t^{PCE}	**-0.61**	**-0.75**	**-0.68**	**-0.80**	**-0.41**	**-0.60**
$\mathbb{E}_{\text{GB}}[\pi_{t+1}]$	**-0.56**	**-0.70**	**-0.67**	**-0.78**	**-0.46**	**-0.57**

Dependent Variable: This table reports coefficients of univariate regressions of the three communications measures (transparency, and the entropies of the minutes and transcripts) on the variables in the rows of the table. The regressions are shown for the moving average of each measure, and the level of the measure. In all rows the relevant communication measures has been linearly detrended and standardized.

Independent Variables: The variable t is a linear trend. *TT State* is the "state of the economy" variable of Tenreyro and Thwaites (2016); "RR" are the monetary policy shocks of Romer and Romer (2004), updated through the sample; NS Policy Shock and NS Fed Funds Shock are the monetary and Fed-Funds futures shocks of Nakamura and Steinsson (2018); FFR is the federal funds target (or the midpoint of its target range, or the actual value when neither is available); Δy_t is the annualized quarterly growth rate of real GDP from FRED; u_t is the civilian unemployment rate from FRED; π_t^{PCE} is the annualized quarterly growth rate of the PCE price level; the rows $\mathbb{E}_{\text{GB}}[x_{t+1}]$ correspond to the Greenbook forecast of x in the quarter following the FOMC meeting at time t (these correspond to the macroeconomic series from above, except that CPI inflation is used instead of PCE in order to have a longer sample from the Greenbook). The monetary policy shocks are timed such that the regression corresponds to the transparency of the minutes prevailing immediately before the shock is emitted. All variables have been standardized over the regression sample so that the coefficients reflect correlation coefficients.

Sample: The sample sizes [for the moving averages] are as follows: 74 for the NS policy shocks (Jan. 2004–Mar. 2014) to match the results from section 4; 276 [265] for the RR shocks (Mar. 1977–Oct. 2008); 261 for $\mathbb{E}_{\text{GB}}[\pi_{t+1}]$ (Oct. 1979–Dec. 2014); and 326 [315] for all other variables (Mar. 1977–Dec.2014).

Bolded estimates are statistically significant at at least the 5% level, calculated using heteroskedasticity and autocorrelation-consistent asymptotic standard errors with the automatic lag selection method of Newey and West (1994), as implemented by Baum et al. (2010).

Table 3: Jointly Testing Correlates of Communications Measures

	Transparency		Minutes Entropy		Transcripts Entropy	
t	0.01^{**}	(0.01)	0.01^{***}	(0.00)	0.00	(0.01)
Burns (70–78)	2.24^{**}	(1.06)	0.08	(0.41)	−0.78	(1.00)
Miller (78–79)	2.19^{**}	(0.86)	−0.38	(0.32)	0.48	(1.01)
Volcker (79–87)	1.52^{**}	(0.61)	-0.52^{**}	(0.25)	−0.38	(0.70)
Greenspan (87–06)	1.05^{**}	(0.42)	-0.42^{**}	(0.20)	0.57	(0.43)
Bernanke (06–14)	0.05	(0.23)	-0.37^{***}	(0.12)	0.48^{*}	(0.27)
Yellen (14-18)	.	.	.	.	.	.
Broida (73–78)	.	.	.	.	.	.
Altmann (78–83)	0.07	(0.21)	-1.01^{***}	(0.20)	0.76^{**}	(0.37)
Axilrod (83–86)	0.03	(0.23)	0.61^{**}	(0.25)	1.14^{***}	(0.44)
Bernard (86–87)	0.22^{*}	(0.12)	−0.06	(0.09)	0.30	(0.20)
Kohn (87–02)	0.13^{**}	(0.06)	-0.12^{**}	(0.06)	−0.11	(0.09)
Reinhart (02–07)	0.15	(0.19)	-0.30^{*}	(0.16)	-0.97^{***}	(0.35)
Madigan (07–10)	-0.48^{***}	(0.16)	-0.58^{***}	(0.13)	0.13	(0.30)
English (10–15)	−0.10	(0.19)	0.34^{**}	(0.13)	−0.07	(0.28)
FFR	0.02	(0.04)	−0.03	(0.02)	−0.02	(0.03)
ΔFFR	−0.03	(0.04)	0.01	(0.02)	−0.04	(0.03)
|ΔFFR|	−0.05	(0.04)	0.00	(0.00)	−0.01	(0.01)
Post-1993	0.33	(0.30)	−0.07	(0.23)	-0.58^{**}	(0.27)
TT State	−0.39	(0.30)	−0.06	(0.20)	1.04^{***}	(0.34)
Δy_t	0.13	(0.09)	−0.03	(0.02)	−0.03	(0.03)
Unemployment	−0.07	(0.07)	0.05^{*}	(0.03)	0.09	(0.07)
π_t^{PCE}	0.01	(0.08)	0.02	(0.04)	0.08	(0.06)
Cons.	-3.00^{***}	(1.13)	−0.46	(0.75)	-3.31^{*}	(1.87)
Chair	0.000		0.000		0.000	
Secretary	0.001		0.000		0.000	
Fed Funds	0.632		0.050		0.577	
Macro Vars.	0.447		0.531		0.274	
R^2	0.935		0.868		0.954	
N	315		315		315	

Note: The table shows results for regressions of the three communications variables (in the columns) on dummies of the sitting chairs of the FOMC (Burns–Yellen); dummies for the secretaries (Broida–English); and several other variables defined in table 2. Chairs and secretaries have their years in office in parentheses. The rows with italicized labels contain p-values for tests of joint significance of groups of variables in the regression: *Chair* jointly tests the chair dummies; *Secretary* tests the secretary dummies; *Fed Funds* tests the three FFR variables; and *Macro Vars.* tests output growth, unemployment, and inflation. These p-values, as well as the standard errors in parentheses, are computed using asymptotic standard errors with the automatic lag selection method of Newey and West (1994), as implemented by Baum et al. (2010). The communications variables are standardized to be mean-zero with unit standard deviation over the regression sample, which is each regularly-scheduled FOMC meeting between March 15, 1977 and December 13, 2014 (the meetings between April 1976 and March 1977 are dropped when the moving average is formed).

the standardized measures on several variables. The first, table 2, shows the regression of these communications variables on the other variables one at a time. The lessons here are largely consistent with the results in table 3, which estimates the coefficients jointly for variables that span the entire 1976–2014 sample. In table 2 the variables are standardized so that the coefficients represent correlation coefficients; in table 3, measures are scaled so that they can be interpreted as the number of standard deviations by which the measure moves when the variable in the row increases by one unit.[19] The regressions highlight some notable features of the series.

First, the three measures all contain positive linear trends to different degrees of statistical significance—between the two tables, the minutes and transcripts have generally increased in their breadth of coverage (higher entropy), and transparency has generally increased. One notable jump—albeit statistically insignificant—is that procedural transparency increased after the FOMC became aware that its transcripts would be released to the public in 1993, and the transcripts also became more focused (lower entropy). The fact that this event may have had an effect on the content of the transcripts—through its effect on the Committee's deliberation—is the subject of a literature that has largely concluded (with the exception of the final paper) that the event caused a move towards less debate and more formal discussions (like what might be found in the minutes); see Meade (2005), Meade and Stasavage (2008), Acosta (2015), Egesdal et al. (2015), Hansen et al. (2017), and Woolley and Gardner (2017).

Next, the sitting Secretary of the FOMC is strongly correlated with the level of all three measures—the p-value for a test that these variables are jointly non-zero is negligible in all cases. This is encouraging, for it suggests that the person in charge of overseeing the creation of each document has a role in determining its properties. The sitting Chair is also correlated with all three measures. The Madigan dummy—a proxy for the Great Recession—in the minutes regression shows that the minutes became much more focused during the crisis and its aftermath.

By-and-large the measures are not predicted by current macroeconomic events and monetary policy once other controls are included, as can be seen in table 3. Unconditionally the variables tend to be countercyclical: transparency is lower when the Federal Funds rate and inflation are high, and both the minutes and transcripts becoming more focused during these times. The similarity between the minutes and the transcripts also tends to decrease in meetings preceding large (in absolute value)

[19] Over the sample period used for that regression.

monetary policy shocks. In the case of the shocks of Nakamura and Steinsson (2018), this implies that markets are more surprised at time $t+1$ when the minutes from time t were less informative about the transcripts.

3.2 Anecdotal Understanding of the Measure

A less-systematic approach to understanding what the transparency measure is capturing is to read the underlying documents. Natural places to start are the points where the measure was particularly low and high. The lowest post-Volcker observation occurred in the December 1989 meeting, though the story starts in October 1989 with a meeting of the G7. In that meeting, the parties involved decided that the dollar was overvalued, and agreed to a coordinated action to flood the market with dollars in order to remedy this. In the October 1989 FOMC meeting, shortly after this coordinated action began, there was a debate about whether this action interfered with the Fed's statutory mandate to achieve price stability. Some parties were concerned that the Fed would be "implicated in talking out of one side of [its] mouth about price stability goals and yet agreeing to constantly flooding the market with dollars" (Transcripts, 1976–2008, Governor Johnson). Others, like President Guffey, felt differently, stating "I'm not terribly concerned about the price stability issue in the sense that with sterilized intervention I think for some long period in the future we can go about a price stability objective without much problem" (Transcripts, 1976–2008, President Guffey). At the end of the discussion, Vice Chair Corrigan suggested that the staff prepare a "presentation for the Committee where [it] would take a look at this question of price stability in five years in some systematic way." The December 1989 meeting contained that presentation and ensuing discussion, though there was no mention of it in the minutes of that meeting, this perhaps owing to President Guffey's October concern about "bringing this issue to a confrontational stage outside the confines of this Committee and the Treasury." This episode contrasts with the meeting of December 2004—the meeting with the highest level of transparency—in which the Committee undertook a lengthy discussion concerning communications policy. Specifically, they discussed the possibility of accelerating the release of the minutes from six to three weeks—a policy they subsequently implemented with the minutes of that meeting—and also made a record of this discussion in the minutes. These episodes highlight the fact that the measure captures meaningful discrepancies between the minutes and the transcripts.

Figure 3: Monetary Policy Transparency in the News

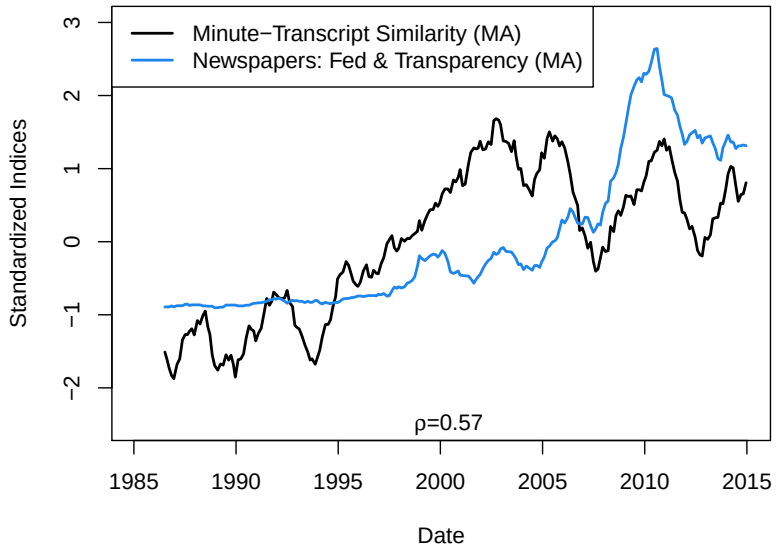

3.3 Observability of the Measure

Lastly, figure 3 shows the transparency measure alongside a measure of transparency derived from newspapers. Specifically, in the spirit of Baker et al. (2016), Husted et al. (2020) construct a measure of monetary policy uncertainty by counting the number of articles in a given period of time that appear in major newspapers containing the terms "uncertainty," "monetary policy," and "Federal Reserve."[20] The measure is divided by the number of articles that contain "Federal Reserve" for each newspaper in each period, in order to control for the volume of articles over time and the different focuses of each newspaper. After scaling the normalized counts by newspaper to have unit variance, the resulting series are summed to form the monetary policy uncertainty index. The blue line in figure 3 is constructed in nearly the same way, with the exclusion of "monetary policy" and "uncertainty" and the inclusion of "transparent" or "transparency."[21] The resulting series is positively and significantly correlated with the transparency measure derived from the minutes and transcripts. Additionally, the moving average of the minute-transcript similarity is more-highly correlated with the newspaper-based measure (and its moving average) than the raw measure. This suggests that minute-transcript similarity, and its moving average, despite being based on a document that is not visible to the public, is something that is in some way observable, given its positive relationship with

[20] They also include close synonyms of these three terms.
[21] A huge thanks to Lucas Husted, who constructed the new index.

this clearly-observable newspaper-based measure. Perhaps this arises because much of the variation in transparency is driven by the sitting Chair, an easily observable variable.

4 Effectiveness of Monetary Policy

Views on the role of transparency in monetary policy making have evolved greatly over the last 50 years. Today, transparency is often touted as a means through which monetary policy is made more effective—in 2013, then-Chair Ben Bernanke stated in a speech that "transparency about the framework of policy has aided the public in forming policy expectations, reduced uncertainty, and made policy more effective." In this section I address this question—whether transparency makes monetary policy more effective—and provide evidence suggesting that it does.

4.1 Defining Effectiveness

The first question that arises when seeking an empirical answer to this question is how to define monetary policy effectiveness. One answer is a policy that allows a central bank to achieve its objective, such as price stability. Blinder et al. (2008) discuss the literature that has taken this approach. A general problem that arises is that establishing causal inference is challenging. Another possibility is that effective monetary policy is able to affect market expectations—Blinder et al. discuss the empirical literature that largely supports this proposition. In this vein, using a high-frequency identification strategy, Nakamura and Steinsson (2018) establish a causal link between monetary policy and real interest rates: the more-surprising the monetary policy announcement, the greater the movement in real interest rates.

While these empirical studies discussed in Blinder et al., Nakamura and Steinsson, and their predecessors have brought a deepened understanding of the effects of monetary policy and its communication, they say little about the role that transparency plays in determining these effects. As it pertains to transparency about the decisionmaking process, the answer is not obvious. An important characteristic of each meeting for understanding its transparency is the context in which it occurred—the state of the economy and the committee, for example. In a more-complicated policy environment, when the content of the minutes are more-heavily scrutinized,[22] high trans-

[22] A simple Google Trends search confirms this—"Fed Meeting" and "Fed Minutes" were the most popular around late 2007, mid 2013, and late 2015; corresponding the beginning of a monetary easing, the months the "Taper

parency might increase or decrease uncertainty and, thus, the ability of policy to have any effects. Using cross-sectional variation in transparency Naszodi et al. (2016) highlight that transparency reduces forecast uncertainty, while increased volume of communication can have the opposite effect, as Lustenbergera and Rossib (2020) find. That said, previous work has postulated that increased transparency should enhance policymakers' abilities to affect the real economy through longer term interest rates (e.g. Woodford, 2005), and the following exercises provide evidence supporting this hypothesis.

4.2 Empirical Strategy

My empirical strategy builds off of the work of Nakamura and Steinsson, who estimate the following equation:

$$\Delta i_t = a + b\,\varepsilon_t + \text{error}_t \tag{1}$$

where i_t stands for a multitude of nominal and real interest rate forwards and yields. I take estimates of daily nominal rates from Gürkaynak et al. (2007), and real rates from Gürkaynak et al. (2010).[23] I start my analysis in 2004 since that is when data on all interest rates under consideration are available, as discussed in Gürkaynak et al. (2010). The variable ε_t is a high-frequency monetary policy shock identified as the first principal component of the change in Fed-Funds and Eurodollar futures out to four quarters.[24] This change is taken over a narrow window around FOMC statement releases. Nakamura and Steinsson find that, in response to a monetary shock, nominal and real interest rates move by a similar amount several years into the term structure. The fact that *real* interest rates move is taken as evidence of monetary non-neutrality.

4.3 Results and Robustness

I slightly augment equation (1) in order to answer the question of whether monetary policy is more or less effective when the Fed is more-accurately representing the content of its meetings:

$$\Delta i_t = \alpha + \beta\varepsilon_t + \gamma(\varepsilon_t \times \bar{\tau}_{t-1}) + \phi\bar{\tau}_{t-1} + \text{error}_t, \tag{2}$$

Tantrum," and the departure from the zero lower bound.

[23] These are available for download from the FEDS working-paper version of these papers: Gürkaynak et al. (2006) and Gürkaynak et al. (2008) respectively.

[24] I take this from the replication materials of Nakamura and Steinsson.

Figure 4: Monetary Policy Effectiveness: Regression Results

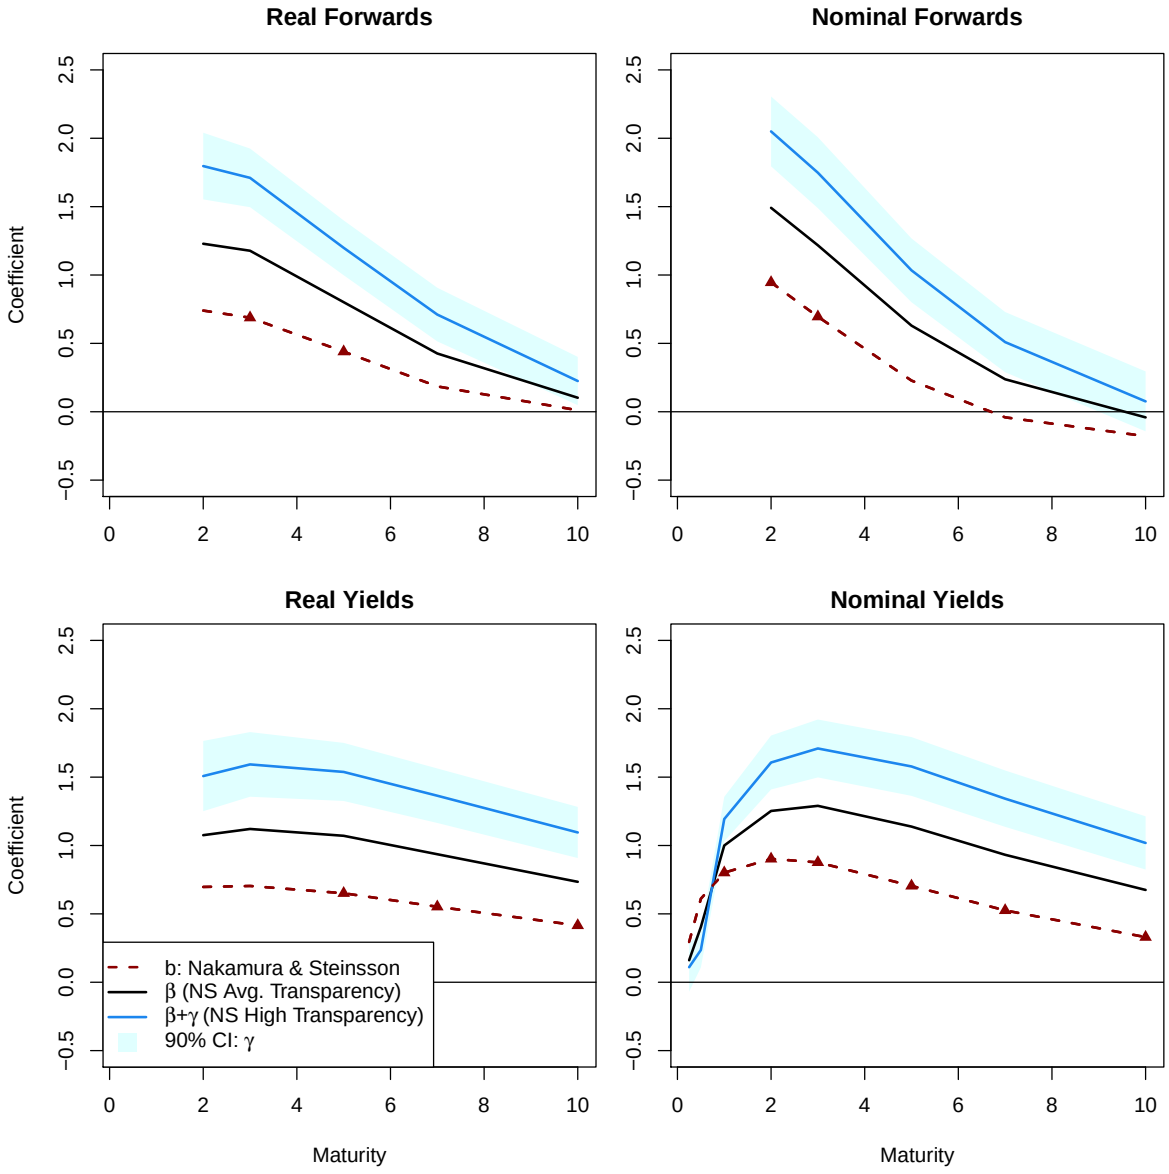

Note: These figures show the results of estimating equations 1 and 2. The four panels correspond to whether i_t is a real or nominal interest rate yield or forward. The x-axis in each plot refers to the relevant maturity for each rate (3- and 6-month, and 1-, 2-, 3-, 5-, 7- and 10-year rates). The blue lines show $\beta + \gamma$ surrounded by the 90% confidence interval of γ computed using heteroskedasticity-robust standard errors. The black lines show β, and the red-dashed lines show b—b will be different than the corresponding values in Nakamura and Steinsson (2018), since my sample runs from January 2004–March 2014 (in contrast to their sample of 2000– March 2014) though I follow Nakamura and Steinsson in dropping the July 2008 – July 2009 period (which has little impact on the results). This forms 74 observations. The shocks are scaled so that the effect of the shock on one-year nominal yields is unity when transparency is at its average level. Transparency is standardized to have unit variance and zero-mean over the sample, so that γ corresponds to the interaction effect when transparency is one standard deviation above its mean. The red triangles are drawn whenever the estimated value of b is statistically significantly below β at at-least the 10% level, again using robust standard errors.

where $\bar{\tau}_t$ is the twelve-meeting moving average of minute-transcript similarity, which has been standardized for interpretability.[25] Because the minutes of meeting $t-1$ are released between $t-1$ and t, the value of $\bar{\tau}$ at time $t-1$ reflects the prevailing level of procedural transparency at time t. The coefficient of interest is γ. With $\gamma > 0$, monetary policy has larger effects on interest rates when transparency is above its mean—a standardization of $\bar{\tau}_t$ means that γ can be interpreted as the interaction effect when procedural transparency is one standard deviation above its mean.

Figure 4 shows the estimated coefficients.[26] The estimated values of γ are positive for nearly every interest rate under consideration, with several of these being statistically significantly different than zero. Averaging the estimates of $\beta + \gamma$ for real yields shows that the effect of monetary policy shocks on real yields is 43 percent higher when transparency is one standard deviation above its mean than when it is at its average level. The estimates of β and b are consistent with the estimates and conclusion of Nakamura and Steinsson (2018)—that nominal and real rates move together far out into the term structure—even using a slightly different sample period.

The discussion in section 3 leaned against an interpretation of minute-transcript similarity as a proxy for something unrelated to procedural transparency. That said, one might be worried that the correlation of transparency with the entropy of the transcripts (which was shown to be correlated with economic conditions) would lead the estimates of equation (2) to simply replicate the conclusions of Tenreyro and Thwaites (2016)—namely, that monetary policy is less effective during recessions. One might also worry that transparency—having a slightly positive trend over time—is serving as a proxy for a Fed that has increasingly relied on longer-term forward guidance. Another concern is that this measure could be proxying for uncertainty regarding monetary policy. In order to alleviate these worries, figure 5 shows the estimates of γ when different controls—and

[25] The reason for using a moving average has been discussed above. Twelve meetings is one-year worth of meetings at the beginning of the sample—this switches to eight meetings in 1981. The results are robust to using several other lag lengths—see D. Section 3 pointed to noise and observability as reasons motivating the use of a moving average—the latter is important here. If this procedural transparency is not perceptible by the public then it is difficult to imagine reasons for which it might have effects. Given that transcripts are not released for several years, it seems unlikely that the public could realize the level of procedural transparency in real time. However, given that the Committee members give speeches and other public commentary, the public should have a sense of what is on the mind of Committee members and be able to compare this to the concerns enumerated in the minutes.

[26] Appendix E contains a table of the estimates.

Figure 5: Monetary Policy Effectiveness: Robustness of Regression Results

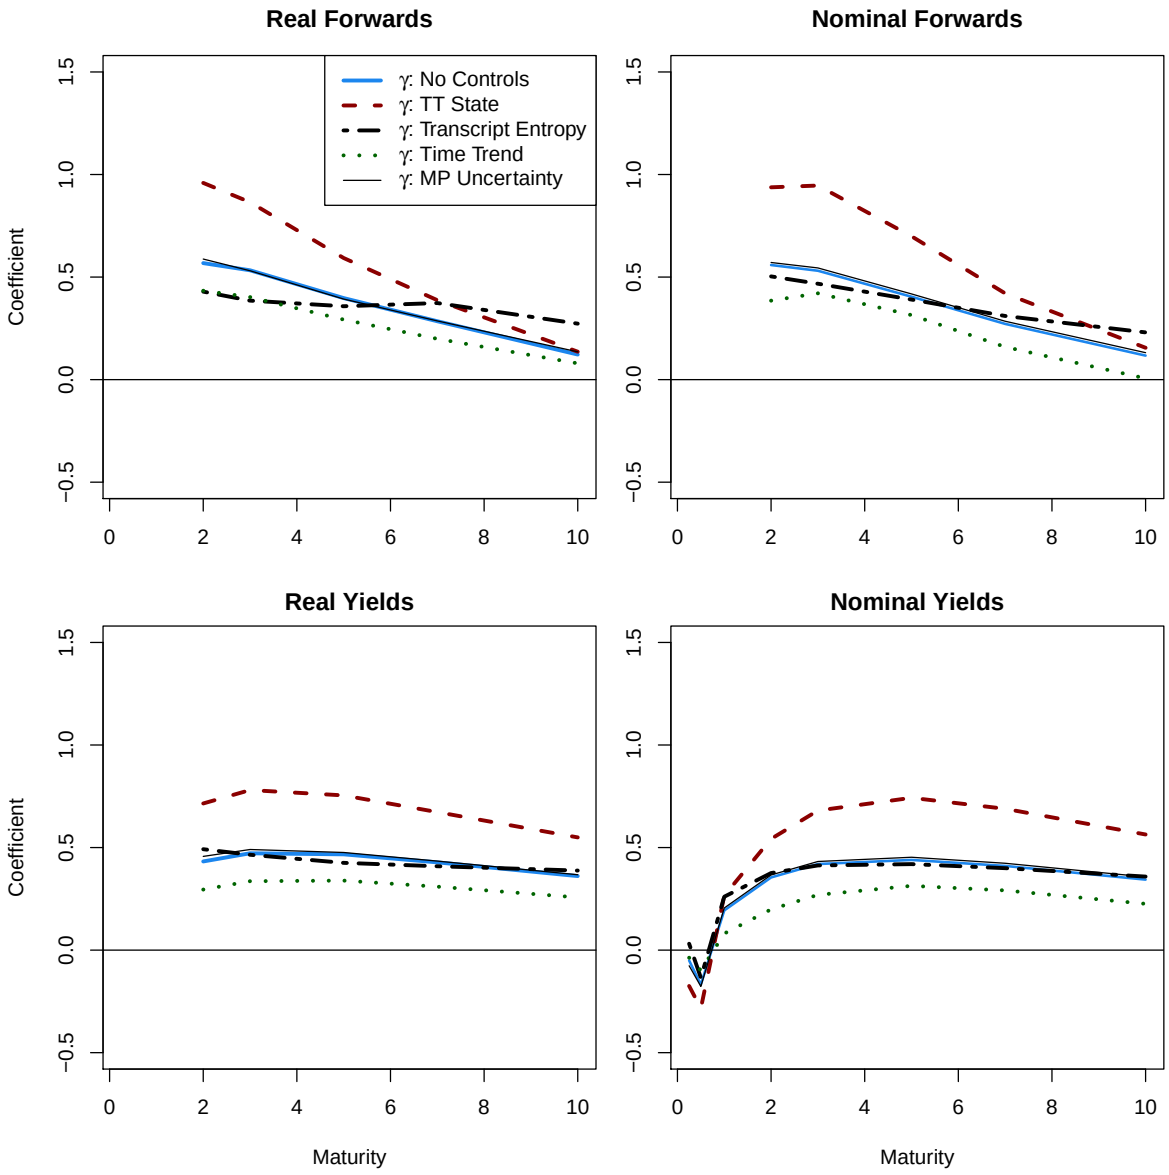

Note: The graphs show the coefficients γ estimated based on equation (3). See the note to figure 4 for details about the sample. The lines labeled "TT State" include the "state of the economy" variable of Tenreyro and Thwaites (2016) as x_t, where t refers to the quarter in which the FOMC meeting took place. The lines labeled "Transcript Entropy" have $x_t = H(\theta_{t-1})$—the lagged entropy of the transcripts.The "Time Trend" label refers to the case in which $x_t = t$. Finally, "MP Uncertainty" is the level of uncertainty about monetary policy—as measured by Husted et al. (2020)—as it stood at the end of the previous FOMC meeting. Again, the monetary policy shocks are normalized so that $\beta = 1$ for the one-year nominal yields regression over the sample period.

their interactions with the monetary shock—have been included in the estimating equation:

$$\begin{aligned}\Delta i_t = \alpha + \beta\varepsilon_t + \gamma(\varepsilon_t \times \bar{\tau}_{t-1}) + \phi\bar{\tau}_{t-1} \\ + \omega(\varepsilon_t \times x_t \quad) + \pi x_t + \text{error}_t,\end{aligned} \tag{3}$$

where x_t is the entropy of the transcripts, the state variable of Tenreyro and Thwaites (for the first concern), a time trend (for the second concern), or the monetary policy uncertainty index of Husted et al. (2020). The results are consistent with the earlier findings—γ is positive for nearly every interest rate under consideration.

4.4 Omitted Variables Bias

Finally, figure 4 also highlights a slight downward bias when equation (1) is estimated without controlling for the role that transparency plays. It is typically the case that $b < \beta$, and in a few cases this difference is statistically significant—red triangles are shown whenever $b < \beta$ is statistically significant at at least the 10% level.[27] Why is this the case? Mechanically, as in any omitted variables bias problem, one has to consider the following relationship:

$$\varepsilon_t \times \bar{\tau}_{t-1} = \psi_0 + \psi\ \varepsilon_t + \text{error}_t.$$

The estimate of ψ will have the sign of $\text{cov}(\varepsilon_t \times \bar{\tau}_{t-1}, \varepsilon_t) \approx \mathbb{E}[\varepsilon_t^2 \bar{\tau}_{t-1}]$, which was shown to be slightly negative in table 2—that is, monetary policy shocks tend to be larger when transparency is lower. With this relationship, however, the estimate of b will not be b but instead $\hat{b} = \beta + \gamma \cdot \psi$. With $\psi < 0$ and $\gamma > 0$, this implies that b underestimates the true effect of monetary policy shocks, β. The intuition for this result is as follows. Consider a large positive shock. The largest monetary shocks tend to occur when transparency is low ($\psi < 0$). However, low transparency also means that the effect of these large shocks on interest rates will be lower (the interaction effect, $\gamma\varepsilon_t\bar{\tau}_{t-1}$, is small or negative on average with small or negative $\bar{\tau}_{t-1}$). Theoretical work may help to clarify this chain of events, though it does suggest that transparency may be a double-edged sword for the effectiveness of monetary policy, if moving interest rates is the definition of effectiveness. Larger

[27] This is computed using a seemingly-unrelated regressions model.

monetary shocks are emitted when the Fed is being less transparent about its discussions. But, possibly because the public cannot make as much sense of these short-run surprises, the shocks are less-easily transmitted to longer term interest rates. The next section provides empirical evidence in support of this interpretation.

5 Interpretation: Transparency about What?

The main transparency measure used in the estimation of equation (2) is an aggregate measure of transparency, though the richness of the underlying data allows for transparency to be measured along different dimensions. In section 5.1, I revisit my estimates using two new measures of transparency: transparency about discussion of the economic outlook, and transparency about monetary policy strategies. I find that my empirical estimates are driven by the latter. To provide additional color to these results I estimate in section 5.2 transparency regarding topics that can be used to "predict" inflation and FOMC dissents, described at end of section 2.3. Consistent with the results in section 5.1, I find that transparency about monetary policy discussions—presumably contentions discussions if they predict FOMC dissents—allow interest rate shocks to pass through more fully to longer term rates. That is not robustly the case for transparency about topics that primarily reveal information about the state of the economy.

5.1 Economic Outlook and Policy Strategy Transparency

The structure of the transcripts suggests a natural first-step towards understanding what minute-transcript similarity is picking up, and which aspect of it plays an important role in determining the effectiveness of monetary policy. As documented by Hansen et al. (2017), "FOMC meetings have two major parts related to the monetary policy decision: the economic situation discussion... followed by the monetary policy strategy discussion." The authors treat these sections as separate in their analysis, and in this section I follow their lead—in so doing, I refer to the first section as ECSIT, and the latter as MPS. I use the breakdowns of the transcript and minutes in order to create an ECSIT transparency index, and an MPS transparency index.[28] The ECSIT index reflects

[28]Stephen Hansen graciously supplied the breakdown of the transcripts through 2011, and I updated this breakdown through 2014. I also performed the corresponding split for the minutes manually since 1995. To have a sense, the ECSIT portion of the minutes in 2014 included the sections titled *Staff Review of the Economic Situation, Staff Review of the Financial Situation, Staff Economic Outlook,* and *Participants' Views on Current Conditions and*

Figure 6: Transparency by Part of Transcript

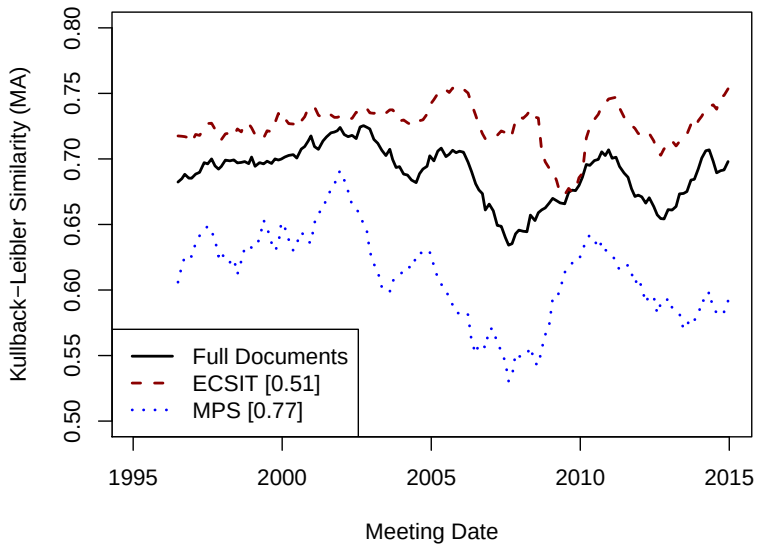

Note: This graph shows the 12-meeting moving-average of transparency using the full minutes and transcripts (solid black line); the ECSIT portions of the minutes and transcripts (red-dashed line); and the MPS portion of the minutes and transcripts (blue-dotted lines). The bracketed numbers represent the correlation coefficient of the of the (moving average of the) ECSIT and MPS transparency measures with the full-document transparency measure. For reference, the ECSIT and MPS discussions take up on average 73 and 12 percent of the minutes, and 51 and 25 percent of the transcripts. These ratios have remained essentially constant since 1995.

the extent to which the Committee provides details about its reading of the state of the economy to the public through the minutes. The MPS index, on the other hand, captures discussions about what this reading implies for monetary policy. The MPS therefore include discussions of the models and targets preferred by policymakers, and their policy preferences more generally.

In order to create these indexes, I estimate the document distributions for the ECSIT and MPS sections of the minutes and transcripts of each meeting.[29] This gives, for every FOMC meeting, a measurement of the transparency of the ECSIT portion of the meeting, and of the MPS portion, displayed in figure 6.

Figures 7 and 8 repeat the analysis regarding the role that transparency plays in determining the effectiveness of monetary policy. Specifically, they present the results of estimating equations 1 and 2, where $\bar{\tau}_{t-1}$ is replaced with the twelve-meeting moving-average of MPS and ECSIT similarity, respectively. The results for MPS similarity are quite similar to those that included minute-

the Economic Outlook, and the MPS section was the section titled *Committee Policy Action*. More details can be provided on request.

[29] Using the same estimated LDA topics as above.

Figure 7: Monetary Policy Effectiveness: Regression Results using MPS Similarity

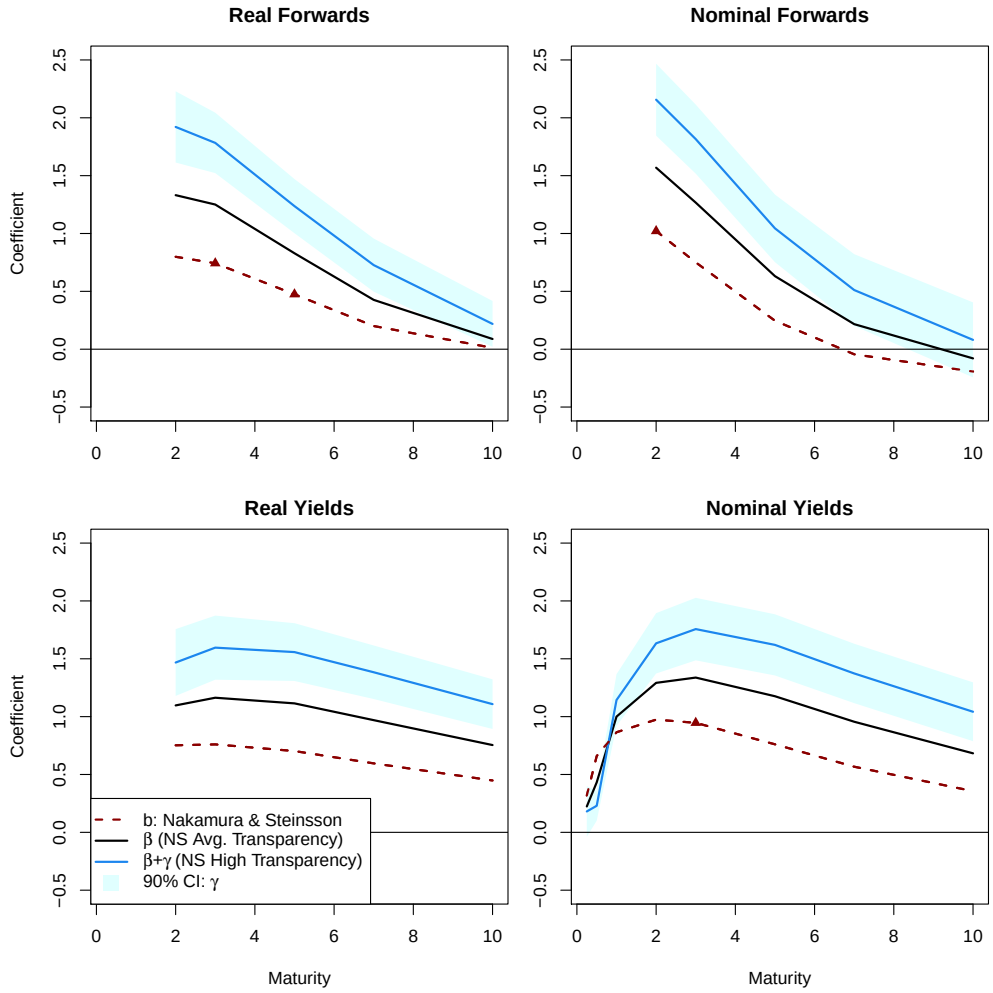

Note: These graphs present the results of estimating equations 1 and 2, where $\bar{\tau}_{t-1}$ is replaced with the twelve-meeting moving-average of MPS similarity. Everything else is the same as in figure 4, so its note can be referenced for further details.

transcript similarity (figure 7). This is not true for ECSIT similarity, despite the fact that ECSIT takes up a much-larger portion of the FOMC discussion (roughly 40% vs. 20 % for MPS). Thus, the fact that monetary policy is more effective when transparency is elevated owes more to transparency regarding policymakers' views about the appropriate monetary policy—conditional on their reading of the state of the economy—than to transparency about the readings themselves.

Figure 8: Monetary Policy Effectiveness: Regression Results using ECSIT Similarity

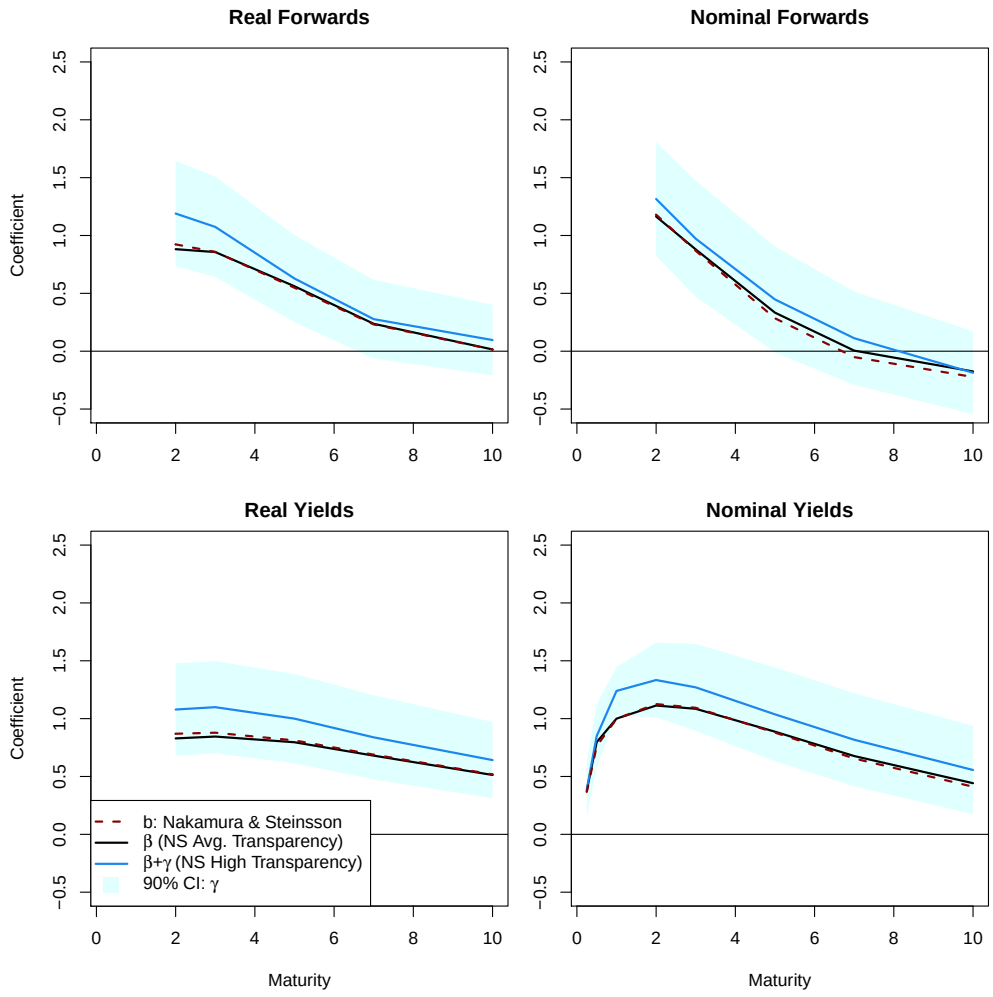

Note: These graphs present the results of estimating equations 1 and 2, where $\bar{\tau}_{t-1}$ is replaced with the twelve-meeting moving-average of ECSIT similarity. Everything else is the same as in figure 4, so its note can be referenced for further details.

5.2 Topic-Specific Transparency

An alternative approach for measuring transparency along different dimensions is to focus on the transparency of specific topics estimated by the LDA language model. In this section I study the role of transparency about topics that are similar to those presented in the previous section: transparency about monetary policy, and transparency about economic fundamentals. To assess the first, I measure the transparency about topics whose presence in the minutes or transcripts predict the number of dissents at each FOMC meeting in a LASSO regression. Of the 50 topics estimated by LDA, I denote the topics selected to predict dissents by $\mathcal{K}^D \subseteq \{1, \ldots, 50\}$. For transparency

Figure 9: Monetary Policy Effectiveness: Regression Results using Topics-Specific Similarities

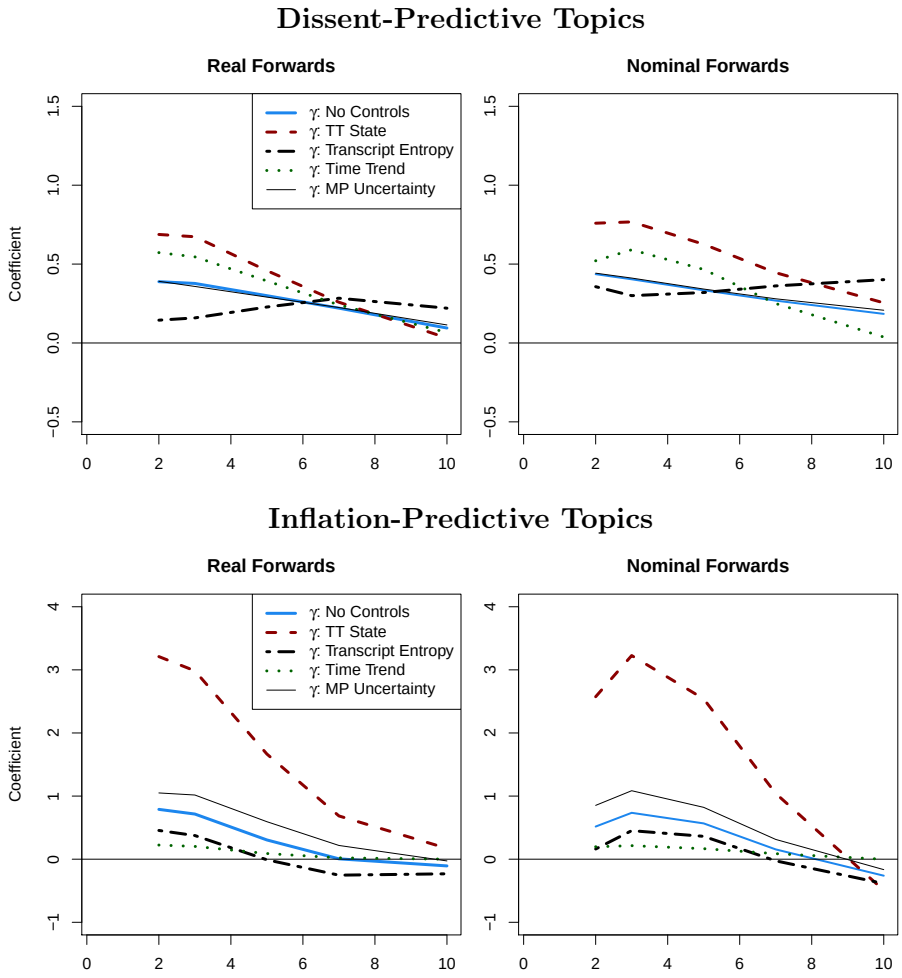

Note: The graphs show the coefficients γ estimated based on equation (3), where $\bar{\tau}_{t-1}$ is replaced with the twelve-meeting moving-average of the similarities of the topics in the minutes and transcripts that are predictive of the number of FOMC dissents and inflation. Everything else is the same as in figure 5, so its note can be referenced for further details. Notice the wider axis for the Inflation panel.

regarding economic fundamentals, I measure the transparency about topics whose presence in the minutes or transcripts predict the level of CPI inflation in the month corresponding to each FOMC meeting. I denote these topics by $\mathcal{K}^{\pi}$. More details regarding the selection of these topics can be found at the end of section 2.3.

I define the transparency for dissent-related topics τ_t^D and inflation-related topics τ_t^{π} as the cosine similarity between the (truncated) distributions of the minutes and transcripts over the selected topics. Recalling from section 2.4 that the topic distributions of the minutes and transcripts

are given by μ_t and θ_t, the transparency measures are given by

$$\tau_t^D = \frac{\sum_{k\in\mathcal{K}^D} \theta_t^k \mu_t^k}{\sqrt{\left(\sum_{k\in\mathcal{K}^D}(\theta_t^k)^2\right)\left(\sum_{k\in\mathcal{K}^D}(\theta_t^k)^2\right)}}$$
$$\tau_t^\pi = \frac{\sum_{k\in\mathcal{K}^\pi} \theta_t^k \mu_t^k}{\sqrt{\left(\sum_{k\in\mathcal{K}^\pi}(\theta_t^k)^2\right)\left(\sum_{k\in\mathcal{K}^\pi}(\theta_t^k)^2\right)}}.$$

Cosine similarities are a commonly used measure of document similarity in the NLP literature, and are simply the uncentered correlation coefficient between the topics used in each document.[30]

Figure 9 shows the estimates of γ for the two transparency measures for specifications that include the various controls considered in the construction of figure 5. The estimates of γ from the specification that uses τ^π are not robustly positive, in contrast to estimates of γ estimated using dissent-predictive topics. This further suggests that it is topics relevant the setting of monetary policy—conditional on an economic outlook—that causes changes in short-term interest rate to affect longer term nominal and real interest rates.

6 Conclusion

A central bank has a plethora of channels through which it can be transparent. Whether it is the publication of inflation reports, timely summaries of policy decisions, or post-meeting press conferences, the objective typically is to explain to the public the rationale behind policy decisions. Nowhere can these rationales be better captured than in the actual meetings in which these decisions are considered and made. The measure of transparency I propose in this paper—the similarity between the minutes and transcripts of each FOMC meeting—captures the extent to which the content of these meetings is described to the public. While this minute-transcript similarity has fluctuated over time, generally the Fed has become more transparent about its reasoning over the last 40 years. My measure is only weakly correlated with economic conditions and the policymaking environment more generally, further supporting the case that minute-transcript similarity captures something more than simply fluctuations in the discussions of FOMC meetings. Anecdotal evidence shows that this measure does indeed capture meaningful discrepancies between the two documents,

[30] I use cosine similarity instead of the Kullback-Leibler divergence because the distributions of the minutes and transcripts over the selected topics do not sum to unity.

and there is little evidence to suggest that these discrepancies are purposeful—they are more likely to do with the fact that the writers of the minutes face a difficult task.

Evidence suggests that monetary policy shocks have larger effects on interest rates when minute-transcript similarity is high. Additionally, when the role of transparency is neglected, these shocks seem to have smaller effects because the largest of these shocks tend to be delivered at times when transparency is low, which is also when the shocks have smaller effects on interest rates. These results suggest that high transparency allows the public to better understand what monetary policy communications and short-term interest rate movements imply for the path of future policy, captured by longer-term interest rates.

References

Acosta, Miguel, "FOMC Responses to Calls for Transparency," *Board of Governors of the Federal Reserve System, Finance and Economics Discussion Series*, 2015, *60.*

Baker, Scott R., Nicholas Bloom, and Steven J. Davis, "Measuring Economic Policy Uncertainty," *The Quarterly Journal of Economics*, 2016, *131* (4), 1593–1636.

Baum, Christopher F, Mark E Schaffer, and Steven Stillman, "ivreg2: Stata module for extended instrumental variables/2SLS, GMM and AC/HAC, LIML and k-class regression.," 2010.

Bernanke, Ben S., "Central Bank Independence, Transparency, and Accountability," in "Speech at the Institute for Monetary and Economic Studies International Conference, Bank of Japan, Tokyo, Japan" 2010.

_ , "Communication and monetary policy," in "Speech at the National Economists Club Annual Dinner, Herbert Stein Memorial Lecture, Washington, DC" 2013.

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

Blinder, Alan S., Michael Ehrmann, Marcel Fratzscher, Jakob De Haan, and David Jan Jansen, "Central Bank Communication and Monetary Policy: A Survey of Theory and Evidence," *Journal of Economic Literature*, 2008, *46* (4), 910–945.

Danker, Deborah J. and Matthew M. Luecke, "Background on FOMC Meeting Minutes," *Federal Reserve Bulletin*, 2005, *91*, 175.

Dincer, N. Nergiz and Barry Eichengreen, "Central Bank Transparency and Independence: Updates and New Measures," *International Journal of Central Banking*, 2014, *10* (1), 189–259.

Egesdal, Michael, Michael Gill, and Martin Rotemberg, "How Federal Reserve Discussions Respond to Increased Transparency," *Available at SSRN*, 2015.

Eijffinger, Sylvester C.W. and Petra M. Geraats, "How Transparent are Central Banks?," *European Journal of Political Economy*, 2006, *22* (1), 1–21.

Geraats, Petra M, "Why adopt transparency? The publication of central bank forecasts," 2001.

_ , "Central Bank Transparency," *The Economic Journal*, 2002, *112* (483), F532–F565.

_ , "The Mystique of Central Bank Speak," *International Journal of Central Banking*, 2007, *3* (1), 37–80.

Goodfriend, Marvin, "Monetary Mystique: Secrecy and Central Banking," *Journal of Monetary Economics*, 1986, *17* (1), 63–92.

Griffiths, Thomas L and Mark Steyvers, "Finding scientific topics," *Proceedings of the National Academy of Sciences*, 2004, *101* (suppl 1), 5228–5235.

Gürkaynak, Refet S, Brian P Sack, and Eric T Swanson, "Market-based measures of monetary policy expectations," *Journal of Business & Economic Statistics*, 2007, *25* (2), 201–212.

Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright, "The U.S. Treasury Yield Curve: 1961 to the Present," *Board of Governors of the Federal Reserve System, Finance and Economics Discussion Series*, 2006, *28*.

_ , _ , and _ , "The TIPS Yield Curve and Inflation Compensation," *Board of Governors of the Federal Reserve System, Finance and Economics Discussion Series*, 2008, *5*.

Gürkaynak, Refet S, Brian Sack, and Jonathan H Wright, "The TIPS yield curve and inflation compensation," *American Economic Journal: Macroeconomics*, 2010, *2* (1), 70–92.

Hansen, Stephen, Michael McMahon, and Andrea Prat, "Transparency and Deliberation within the FOMC: a Computational Linguistics Approach," *Quarterly Journal of Economics*, 2017.

Husted, Lucas, John Rogers, and Bo Sun, "Monetary Policy Uncertainty," *Journal of Monetary Economics*, 2020, *115*, 20–36.

Klein, Dan and Christopher D Manning, "Accurate unlexicalized parsing," in "Proceedings of the 41st annual meeting of the association for computational linguistics" 2003, pp. 423–430.

Lindsey, David E, "A Modern History of FOMC Communication: 1975-2002," *Board of Governors of the Federal Reserve System, Washington, DC*, 2003.

Loughran, Tim and Bill McDonald, "When is a Liability not a Liability? Textual Analysis, Dictionaries, and 10-Ks," *The Journal of Finance*, 2011, *66* (1), 35–65.

Lustenbergera, Thomas and Enzo Rossib, "Does Central Bank Transparency and Communication Affect Financial and Macroeconomic Forecasts?," *International Journal of Central Banking*, 2020.

Meade, Ellen E., "The FOMC: Preferences, Voting, and Consensus," *Federal Reserve Bank of St. Louis Review*, 2005, *87*, 93–101.

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

_ , Nicholas A. Burk, and Melanie Josselyn, "The FOMC Meeting Minutes: An Assessment of Counting Words and the Diversity of Views," *Federal Reserve Board FEDS Notes*, 2015.

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

Nakamura, Emi and Jón Steinsson, "High Frequency Identification of Monetary Non-Neutrality," *Quarterly Journal of Economics*, 2018, *133* (3), 1283–1330.

Naszodi, Anna, Csaba Csavas, , Szilard Erhart, and Daniel Felcser, "Which Aspects of Central Bank Transparency Matter? A Coomprehensive Analysis of the Effect of Transparency of Survey Forecasts," *International Journal of Central Banking*, 2016.

Newey, Whitney K and Kenneth D West, "Automatic lag selection in covariance matrix estimation," *The Review of Economic Studies*, 1994, *61* (4), 631–653.

Romer, Christina D. and David H. Romer, "A New Measure of Monetary Shocks: Derivation and Implications," *American Economic Review*, 2004, *94* (4), 1055–1084.

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

Tenreyro, Silvana and Gregory Thwaites, "Pushing on a string: US monetary policy is less powerful in recessions," *American Economic Journal: Macroeconomics*, 2016, *8* (4), 43–74.

Thornton, Daniel L and David C Wheelock, "Making sense of dissents: a history of FOMC dissents," *Federal Reserve Bank of St. Louis Review*, 2014, *96* (3), 213–227.

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

Woodford, Michael, "Monetary Policy in the Information Economy," in "Symposium on Economic Policy for the Information Economy" Federal Research Bank of Kansas City 2001, pp. 297–370.

_ , "Central bank communication and policy effectiveness," in "The Greenspan Era: Lessons for the Future" Federal Reserve Bank of Kansas City 2005, pp. 399–474.

Woolley, John T and Joseph Gardner, "The effect of "sunshine" on policy deliberation: The case of the Federal Open Market Committee," *The Social Science Journal*, 2017, *54* (1), 13–29.

Yellen, Janet, "Communication in monetary policy," in "Speech to the Society of American Business Editors and Writers 50th Anniversary Conference, Washington, DC" 2013.

A 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 4 lists the various FOMC publications since 1935, along with their release

Table 4: FOMC Publications: Banking Act of 1935 to Present (2015). *Release lags in italics.*[2]

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)*

lags.[3] 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.[4]

The 1976 MODs were the last published by the FOMC; after five years of fighting a claimed FOIA violation, the Committee decided in 1981 to discontinue the MOD,[5] 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, 1976–1993, 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

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

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

[4] 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.

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

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 ROPAs, MOAs, 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.

B Definition and Relevance of Procedural Transparency

When used in the context of monetary policy, the word "transparency" can carry different connotations. To understand how the term is used here, table 5 presents the five forms of transparency relevant to central banks, as defined by Geraats (2002).[6] Procedural transparency—my focus in this

Table 5: 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.

paper—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, and it is the mechanism through which the other four are manifested. For example, mandating that the Fed release the theoretical rule it uses to determine policy is a form of economic, policy, and political transparency. However, effective implementation of this policy hinges on effective procedural transparency because both the rule and deviations from it require detailed explanations. The same

[6] Geraats has written much about central bank transparency. See Geraats (2001), where these terms were first defined, or Geraats (2007) for other examples.

goes for the practice of establishing and explaining an inflation target—especially a "medium-term" target as is done in current practice. Simply put, the communications of the Fed are "a chance for [the FOMC] to say what [they] are up to and why" (Transcripts, 1976–2008, Alan Blinder, Jan. 1995).

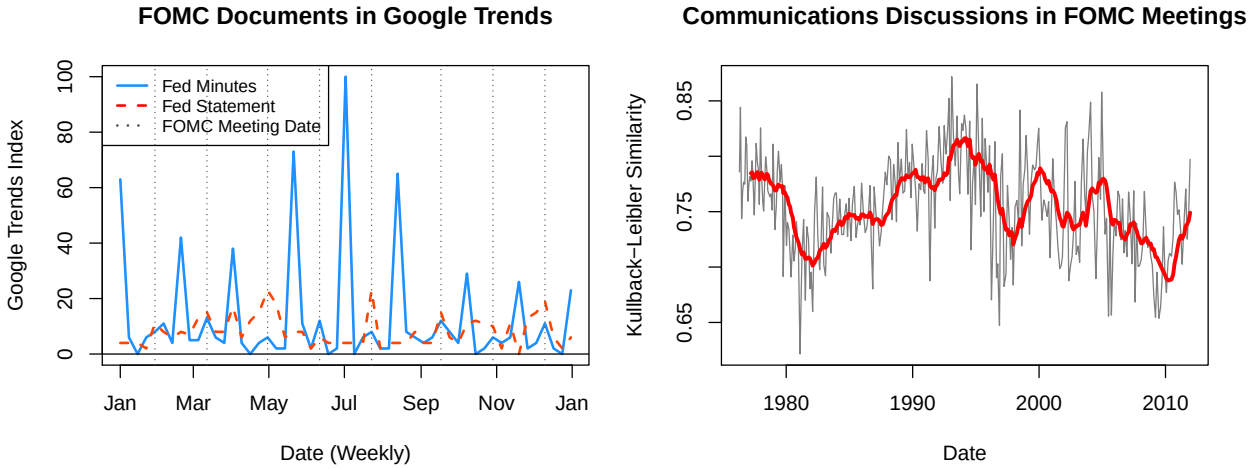

Figure 10: The left panel presents results from the Google Trend queries "fed statement" and "fed minutes." The y-axis represents the frequency with which a given phrase is searched on Google, and is normalized so that the highest frequency is 100. Thus, this graph does not say how these search terms rank among all other terms, but it does give information about when the terms are searched. Dashed lines are included at FOMC meeting dates. The right panel shows the Kullback-Leibler similarity, defined below, of the transcripts at each point in time with a query that contains the names of FOMC minutes and transcripts over time, and the word *communication*. The full list is "transcripts minutes record policy actions memorandum discussion communications."

Fed communications also receive a considerable amount of attention from the public at large, and the FOMC itself. Figure 10 supports this claim. The left panel is a graph from Google Trends—a service from Google that plots the "interest over time" of any search term—that shows how popular "fed statement" and "fed minutes" were over the year 2017. 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.

The right panel of figure 10 provides evidence that the FOMC discusses issues of procedural transparency in its meetings. Using the text-analysis techniques described below, it portrays the extent to which topics of procedural transparency were discussed at each FOMC meeting. This is measured by computing the similarity between the transcripts and a list of procedural-transparency-related words (*transcripts, minutes, communication*, etc.). The similarity of this topic with the transcripts fluctuates meaningfully, coinciding with changes in communications 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 1990s, changes in publication policy have been relatively frequent—in 1993, the FOMC decided to start making its transcripts public; and in 2005, the lag between the meeting and the release of the minutes was reduced. These changes are visible on the graph, indicating that a significant amount of discussion was behind each decision. Thus, given the role of procedural transparency in monetary policymaking generally, and the fact that both the

public and the Fed pay close attention to the documents used here to measure it, the rest of this paper is devoted to its study.

C Number of Topics: Cross-Validation and Robustness

A common way to select topics for LDA is to estimate the model using a fraction of documents in the corpus, then compute how "perplexed" the model is by the held-out documents that were not used for estimation. In figure 11, I show the results of performing this cross-validation for different values of K (from 10–100 by 10, and from 120–200 by 20). For each value of K, the sample is split into thirds—two thirds is used as a training sample, the term distributions are estimated, then using these distributions the topic distributions for the held-out documents are computed. The perplexity of these held-out documents is then computed. This is done five times for each value of K, shown by the black lines in the figure. The blue line shows the average of the black dots for each value of K. The red-dashed line shows the sum of squared of residuals from fitting two lines to the blue curve, where the two lines are split at a given value of K. This is called finding the "knee-point" of the blue line, or the point at which the perplexity drops off most-sharply. The knee of perplexity here is at $K = 50$, so 50 topics are used for the analysis in the paper.

Figure 11: Cross-Validation Results

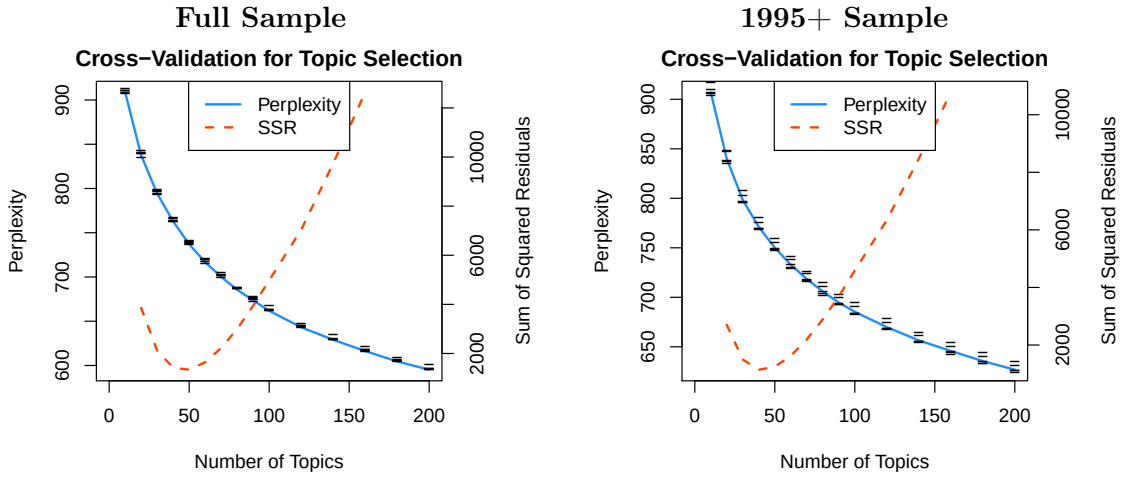

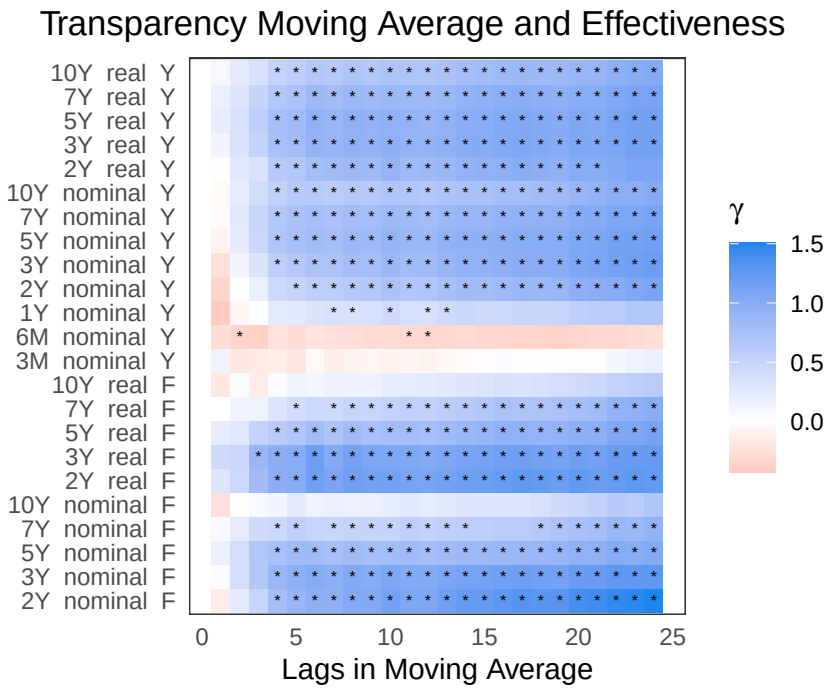

Figure 12: Estimates of γ for Different Moving Averages

D Robustness for Moving Averages

Figure 12 shows the results of γ from equation (2), where the moving average of $\bar{\tau}$ uses lag lengths from zero to 25. Stars are drawn whenever the level of statistical significance is at least 10%. All standardizations and scalings are performed as described in table 4. Notably, the coefficients are heavily skewed towards positive. Also, given how unlikely it is that the minute-transcript similarity for the most-recent meeting is observed (given that the transcript is not released for five years), the first few moving-averages, while negative, should be given little weight.

E Tabular Representation of Figures 4 and 5

	Baseline		TT State		Trans. Entropy		Time Trend		MPU	
	β	γ	β	γ	β	γ	β	γ	β	γ
Real Forwards										
2-Year	1.23	0.57	2.39	0.96	1.19	0.43	-2.00	0.43	1.28	0.59
	(0.27)	(0.15)	(1.36)	(0.27)	(0.26)	(0.18)	(1.32)	(0.14)	(0.46)	(0.14)
3-Year	1.18	0.53	2.00	0.87	1.14	0.39	-1.72	0.40	1.48	0.53
	(0.24)	(0.13)	(1.16)	(0.24)	(0.24)	(0.18)	(1.41)	(0.14)	(0.41)	(0.13)
5-Year	0.80	0.40	0.75	0.59	0.79	0.36	-0.81	0.29	1.33	0.39
	(0.22)	(0.12)	(0.88)	(0.20)	(0.21)	(0.20)	(1.26)	(0.12)	(0.34)	(0.12)
7-Year	0.43	0.28	-0.22	0.39	0.45	0.37	-0.03	0.20	0.92	0.29
	(0.21)	(0.12)	(0.80)	(0.18)	(0.20)	(0.19)	(0.95)	(0.10)	(0.32)	(0.12)
Nominal Forwards										
2-Year	1.49	0.56	2.34	0.94	1.45	0.50	-0.56	0.38	1.98	0.57
	(0.24)	(0.16)	(1.43)	(0.24)	(0.24)	(0.27)	(0.93)	(0.12)	(0.46)	(0.14)
3-Year	1.22	0.53	2.35	0.95	1.19	0.47	-1.60	0.42	1.63	0.54
	(0.27)	(0.16)	(1.41)	(0.26)	(0.26)	(0.25)	(1.39)	(0.15)	(0.48)	(0.15)
5-Year	0.63	0.40	1.13	0.70	0.63	0.39	-1.20	0.31	0.89	0.42
	(0.27)	(0.14)	(1.04)	(0.25)	(0.25)	(0.22)	(1.59)	(0.15)	(0.38)	(0.15)
7-Year	0.24	0.27	-0.01	0.42	0.25	0.31	0.24	0.16	0.48	0.28
	(0.27)	(0.13)	(0.89)	(0.24)	(0.23)	(0.23)	(1.34)	(0.13)	(0.33)	(0.15)
Real Yields										
2-Year	1.08	0.43	1.58	0.72	1.08	0.49	-0.57	0.30	1.16	0.46
	(0.27)	(0.16)	(1.16)	(0.25)	(0.25)	(0.20)	(0.89)	(0.11)	(0.51)	(0.13)
3-Year	1.12	0.47	1.78	0.78	1.12	0.46	-1.00	0.34	1.24	0.49
	(0.25)	(0.14)	(1.16)	(0.24)	(0.24)	(0.19)	(0.98)	(0.11)	(0.47)	(0.13)
5-Year	1.07	0.47	1.61	0.75	1.06	0.43	-1.10	0.34	1.32	0.48
	(0.22)	(0.13)	(1.05)	(0.22)	(0.21)	(0.17)	(1.09)	(0.12)	(0.41)	(0.12)
7-Year	0.94	0.43	1.21	0.67	0.93	0.41	-0.88	0.31	1.26	0.43
	(0.21)	(0.12)	(0.92)	(0.20)	(0.20)	(0.17)	(1.06)	(0.11)	(0.37)	(0.11)
Nominal Yields										
3-Month	0.16	-0.05	-0.72	-0.17	0.18	0.03	0.52	-0.04	0.53	-0.08
	(0.11)	(0.11)	(0.72)	(0.14)	(0.11)	(0.12)	(0.21)	(0.06)	(0.25)	(0.08)
6-Month	0.40	-0.17	0.17	-0.28	0.40	-0.13	0.61	-0.10	0.52	-0.18
	(0.10)	(0.08)	(0.26)	(0.11)	(0.10)	(0.10)	(0.24)	(0.05)	(0.18)	(0.07)
1-Year	1.00	0.19	1.00	0.27	1.00	0.26	1.00	0.08	1.00	0.21
	(0.16)	(0.10)	(0.76)	(0.14)	(0.16)	(0.16)	(0.29)	(0.06)	(0.28)	(0.09)
2-Year	1.25	0.35	1.50	0.54	1.24	0.38	0.62	0.20	1.47	0.37
	(0.19)	(0.12)	(1.00)	(0.17)	(0.18)	(0.20)	(0.42)	(0.08)	(0.35)	(0.11)
3-Year	1.29	0.42	1.80	0.68	1.27	0.41	0.02	0.27	1.59	0.43
	(0.20)	(0.13)	(1.12)	(0.19)	(0.20)	(0.22)	(0.63)	(0.09)	(0.38)	(0.12)
5-Year	1.14	0.44	1.80	0.74	1.12	0.42	-0.63	0.31	1.45	0.45
	(0.22)	(0.13)	(1.11)	(0.21)	(0.21)	(0.22)	(0.97)	(0.11)	(0.38)	(0.12)
7-Year	0.93	0.41	1.43	0.69	0.92	0.40	-0.58	0.29	1.22	0.42
	(0.22)	(0.13)	(0.98)	(0.20)	(0.21)	(0.21)	(1.08)	(0.11)	(0.34)	(0.12)
10-Year	0.68	0.34	0.87	0.56	0.67	0.36	-0.10	0.23	0.96	0.36
	(0.21)	(0.12)	(0.84)	(0.20)	(0.20)	(0.20)	(1.05)	(0.11)	(0.30)	(0.12)

Note: The table shows the coefficients β and γ estimated based on equations (2) (for columns labeled "Baseline") and (3) (for the other columns). See the notes to figures 4 and 5 for details about the sample, column labels, and normalizations. Robust standard errors are in parentheses.

Figure 13: Robustness of Transparency Measure to K and Sample Period

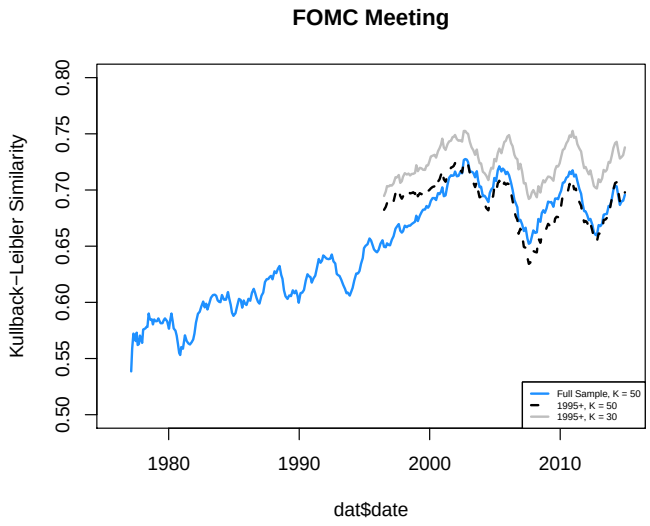

F Robustness to Text-Analysis Features

Figure shows the moving-average of the transparency measure using LDA topics estimated over the 1995+ sample (black-dashed line), which has a correlation of 0.86 with the full-sample measure. In addition, the gray-solid line shows the measure with $K = 30$, which has a correlation of 0.90 with the 1995+ sample measure with $K = 50$.

Figure 14 shows the results of estimating (2) for interest rate forwards when the LDA model is estimated on the full sample

Figure 14: Robustness of Regression Results to K and Sample Period

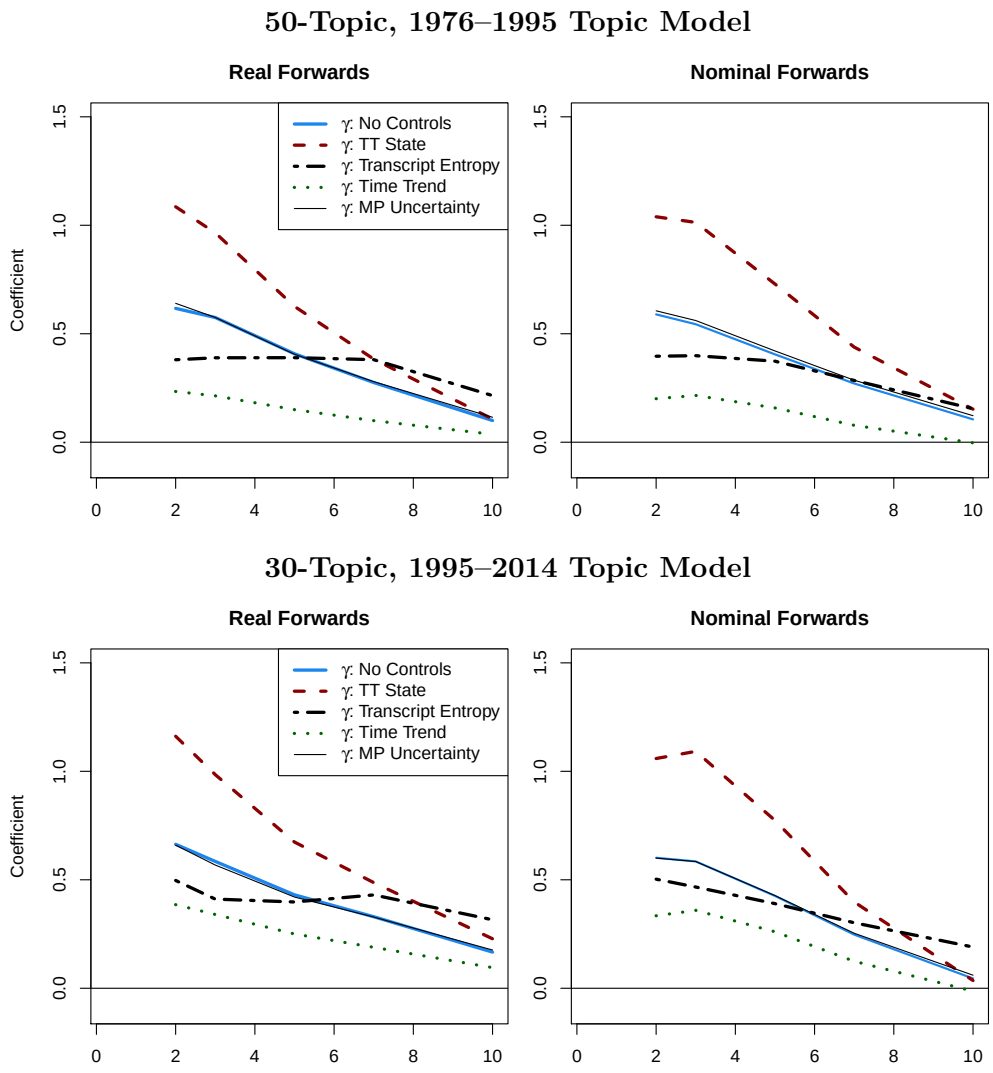

Note: The graphs show the coefficients γ estimated based on equation (3), where $\bar{\tau}_{t-1}$ is estimated using the 1976–2014 LDA model (top panel) and the 30-topic 1995+ LDA model (bottom panel). Everything else is the same as in figure 5, so its note can be referenced for further details.