

Bertsch, Christoph; Hull, Isaiah; Lumsdaine, Robin L.; Zhang, Xin

Working Paper

Central bank mandates and monetary policy stances: Through the lens of Federal Reserve speeches

Sveriges Riksbank Working Paper Series, No. 417

Provided in Cooperation with:

Central Bank of Sweden, Stockholm

Suggested Citation: Bertsch, Christoph; Hull, Isaiah; Lumsdaine, Robin L.; Zhang, Xin (2022) :
Central bank mandates and monetary policy stances: Through the lens of Federal Reserve
speeches, Sveriges Riksbank Working Paper Series, No. 417, Sveriges Riksbank, Stockholm

This Version is available at:

<https://hdl.handle.net/10419/272874>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

SVERIGES RIKSBANK
WORKING PAPER SERIES

417



Central Bank Mandates and Monetary Policy Stances: through the Lens of Federal Reserve Speeches*

Christoph Bertsch, Isaiah Hull, Robin L. Lumsdaine, and Xin Zhang

October 2022

WORKING PAPERS ARE OBTAINABLE FROM

www.riksbank.se/en/research

Sveriges Riksbank • SE-103 37 Stockholm

Fax international: +46 8 21 05 31

Telephone international: +46 8 787 00 00

The Working Paper series presents reports on matters in the sphere of activities of the Riksbank that are considered to be of interest to a wider public.

The papers are to be regarded as reports on ongoing studies and the authors will be pleased to receive comments.

The opinions expressed in this article are the sole responsibility of the author(s) and should not be interpreted as reflecting the views of Sveriges Riksbank.

Central Bank Mandates and Monetary Policy Stances: through the Lens of Federal Reserve Speeches*

Christoph Bertsch¹, Isaiah Hull^{†2}, Robin L. Lumsdaine³, and Xin Zhang^{1,4}

¹Sveriges Riksbank

²BI Norwegian Business School; CogniFrame

³Kogod School of Business, American University; Erasmus University
Rotterdam; National Bureau of Economic Research (NBER); Tinbergen
Institute; Center for Financial Stability

⁴BIS Innovation Hub Nordic Centre

Sveriges Riksbank Working Paper Series

No. 417

October 2022

*We are grateful for comments from Francesco Bianchi and seminar participants at Sveriges Riksbank and Universidad Carlos III and participants at the 2022 Society for Financial Econometrics Annual Meeting. Nicole Brynjolfsson provided excellent assistance in collecting and organizing some of the speeches used in this paper as part of the Center for Financial Stability's Summer Internship Program. This work was conducted while Lumsdaine was on partial secondment to the Office of Financial Research, US Department of the Treasury, under an Intergovernmental Personnel Agreement (IPA). We thank Edvin Ahlander for excellent research assistance. The opinions expressed in this article are the sole responsibility of the authors and should not be interpreted as reflecting the views of the BIS, Sveriges Riksbank, the US Department of the Treasury, or any of the other institutions they represent. The authors have no relevant conflicts of interest to declare. All remaining errors are our own.

[†]Corresponding author. Email: isaiah.hull@bi.no. Tel: +46 076 589 06 61. Finance Department, BI Norwegian Business School, Nydalsveien 37, 0484 Oslo, Norway.

Central Bank Mandates and Monetary Policy Stances: through the Lens of Federal Reserve Speeches

Abstract

When does the Federal Reserve deviate from its dual mandate of pursuing the economic goals of maximum employment and price stability and what are the consequences? We assemble the most comprehensive collection of Federal Reserve speeches to-date and apply state-of-the-art natural language processing methods to extract a variety of textual features from each paragraph of each speech. We find that the periodic emergence of non-dual mandate related discussions is an important determinant of time-variations in the historical conduct of monetary policy with implications for asset returns. The period from mid-1996 to late-2010 stands out as the time with the narrowest focus on balancing the dual mandate. Prior to the 1980s there was a outsized attention to employment and output growth considerations, while non dual-mandate discussions centered around financial stability considerations emerged after the Great Financial Crisis. Forward-looking financial stability concerns are a particularly important driver of a less accommodative monetary policy stance when Fed officials link these concerns to monetary policy, rather than changes in banking regulation. Conversely, discussions about current financial crises and monetary policy in the context of inflation-employment themes are associated with a more accommodative policy stance.

Keywords: Natural Language Processing, Machine Learning, Central Bank Communication, Financial Stability, Zero Shot Classification, Extractive Question Answering, Semantic Textual Similarity.

JEL classification: C63, D84, E32, E7.

The mandate of the Federal Reserve (Fed) and the institution itself have evolved over time, which is also reflected in the historical variation in the conduct of monetary policy. Today's dual mandate of pursuing the economic goals of maximum (sustainable) employment and price stability was only codified in the Federal Reserve Reform Act of 1977 and the Full Employment and Balanced Growth Act of 1978. The dual mandate was shaped as the consequence of increased recognition for the Fed's role for price stability and paved the way for an embrace of an active monetary policy. Yet it is not always clear how the central bank balances its dual goals. As a result, researchers and market participants alike often turn to the communication of Federal Reserve officials to glean information regarding the tradeoff. Interestingly, an anecdotal read of past policy debate in central banking identifies emphasis on other potential goals, for instance financial stability, and more recently climate risk.

Federal Reserve speeches give valuable insights into the Fed's interpretation of its mandate and the reasoning behind the historic conduct of monetary policy. They provide a direct way to measure policy makers' objective function, beyond the clear dual mandate. In this paper we dissect the information in speeches of Fed officials to identify whether the discussion of other policy goals can drive deviations from the Fed's dual mandate. Finding this to be the case, we then consider the consequences of these deviations for the conduct of monetary policy and for the relationship between asset prices and monetary policy regimes.

There is a growing literature analyzing historical monetary policy (Clarida et al., 2000; Orphanides, 2003; Sims and Zha, 2006; Bianchi, 2012; Bianchi et al., 2022). A popular way to assess changes in monetary policy is the Taylor-rule framework (Taylor, 1993, 1999). Conceptually, adjustments to (or deviations from) a monetary policy rule may be driven by factors that are inside or outside the Federal Reserve's official mandate and the speeches given by the members of the Federal Open Market Committee (FOMC), that is, the Federal Reserve Board of Governors and the Federal Reserve Bank presidents, provide us with a unique lens through which to analyze how monetary policy is shaped by the concerns of policy makers. We contribute to the existing literature in two main ways.

First, we analyze the concerns expressed in Fed speeches that are associated with salient changes in the historical conduct of monetary policy. Second, we document the consequences of Fed communication about these concerns on both, changes in monetary policy and asset prices. Moreover, we unearth how these consequences are shaped by the nature of the concerns. Our methodology draws from state-of-the-art natural language processing (NLP) methods to extract a variety of text features from speeches by Fed officials at the paragraph level, including their most significant concerns. This allows us to understand whether or not the concerns expressed in speeches are related to the Federal Reserve’s dual mandate of pursuing the economic goals of maximum employment and price stability, or to other potentially important aspects such as financial stability. We also measure the temporal focus of the concerns, a key source of information that has been largely absent from existing work on central bank communication.¹

To analyze the Fed speeches we draw on recent advances in NLP and use state-of-the-art transformer models (Vaswani et al., 2017). Specifically, we use the BERT (Bidirectional Encoder Representations from Transformers) language model introduced by Devlin et al. (2019) and the modified RoBERTa (Robustly Optimized BERT pretraining Approach) NLP model by Liu et al. (2019), which were trained to perform extractive question answering. In contrast to the dictionary-based methods that are commonly used in economics,² transformer models allow us to classify text based on a remarkable capacity to understand language, taking left-right dependencies, negations and modifiers into account. Based on this, we can perform extractive question answering to understand whether certain content is present (the classification task) and to analyze semantic textual similarity by measuring the cosine similarity between contextualized embeddings. For the classification task, we generate text features using a variation of the zero-shot classification algorithm by Pushp and Srivastava

¹A notable recent exception is Byrne et al. (2022), who quantify the temporal dimension with a rules-based temporal tagger to study the communication transmission mechanism and the information deficit.

²Examples of these include the financial dictionaries introduced in Loughran and McDonald (2011) to measure tone and sentiment of corporate reports and dictionaries tailored to measure financial stability sentiment by Correa et al. (2021) or the hawkishness of monetary policy by Apel and Grimaldi (2014).

(2017), which entails the labelling of sequence embeddings. For the measurement of semantic textual similarity we employ a sentence BERT model with a refined pretraining process introduced by Reimers and Gurevych (2019) and Wang et al. (2021).

Our first finding is that non dual-mandate related concerns, primarily centered around financial stability, constitute an important determinant of changes in the historical conduct of monetary policy, explaining deviations towards a less accommodative monetary policy, especially when forward-looking in nature and linked to monetary policy by Fed officials in their speeches. The airing of financial stability concerns and the acceptance of using monetary policy to achieve financial stability objectives increased during the Great Moderation. Unsurprisingly, the emergence of financial stability concerns can be traced to measures of credit and the experience of financial turmoil. Distinct from forward-looking financial stability concerns and backward-looking crisis experiences, discussion of a current financial crisis and monetary policy in the context of dual-mandate-related inflation-employment themes is associated with a more accommodative policy stance.

We also find that the concerns discussed in speeches by Fed officials are associated with changes in asset valuations, with a prominent role played by the non-dual mandated related text features that are associated with salient changes in the conduct of monetary policy. Not surprisingly, the consequences of discussions about financial crises are negative for asset valuations. The picture is, however, more nuanced for financial stability discussions, which are only negatively associated with valuations for equity and risky assets, but not for bonds and safe assets. This result is consistent with potential concerns of Fed officials about the risk taking channel of monetary policy (Jiménez et al., 2014; Dell’Ariccia et al., 2017). Moreover, we corroborate the result and its interpretation by evidence that shows that financial stability concerns are an important determinant of asset valuation-monetary policy regimes, a topic recently discussed in Bianchi et al. (2022).

Also, a higher academic focus of Fed speeches is negatively associated with valuations for equity and risky assets, while the valuations of bonds and safe assets are positively associated

with a past or present temporal focus. This result stands against the backdrop of a positive association between the discussion of financial stability concerns and a heightened academic focus. Nevertheless, we also find that Fed speeches referencing the academic literature tend to oppose the use of both monetary policy and banking regulation to achieve financial stability. From an institutional viewpoint, an interesting observation is that the financial stability and academic focus scores are highest at the Federal Reserve Banks of New York and Richmond, as well as at the Federal Reserve Board.

Taking a historical perspective, we find that prior to the 1980s Fed speeches display an outsized attention to output and employment considerations in their speeches, which is consistent with an over-weighting of output and employment considerations and resonates with the view that there was an insufficient emphasis on price stability leading up to the Great Inflation in the mid 1960s and 1970s (Meltzer, 2009). After the Great Financial Crisis non-dual mandate related discussions centered around financial stability considerations emerged, which we show to be important. This observation motivates our study and allows us to complement an existing literature discussing dual mandate related factors that can help to explain the historical conduct of monetary policy.

The existing literature documents a low frequency time variation in the historical conduct of monetary policy (Clarida et al., 2000; Boivin, 2006), as well as in the monetary policy rule perceived by financial forecasters (Bauer et al., 2022). Most prominently, the high inflation period in the late 1970s and early 1980s stands out as discretionary episode with large deviations from the Taylor rule (Nikolsko-Rzhevskyy et al., 2014), which is a popular framework to assess changes in monetary policy (Taylor, 1993, 1999).

There have been attempts to rationalize variations in the historical conduct of monetary policy. In relation to financial stability there may be benefits from aggressive inflation targeting (Bernanke and Gertler, 2001a) and from moderate leaning against the wind (Caballero and Simsek, 2020), but macroprudential policy is generally seen as superior Schularick et al. (2021). On an informational level, Orphanides (2001, 2003) documents the importance of

misinformation at the time when the monetary policy decisions are made, while (Meltzer, 2009) highlights the role of misconceptions about economic theory that led the Fed to chase the Philips curve in an attempt to lower unemployment during the 1970s. The subsequent shift in tack under the chairmanship of Paul Volcker (August 1979 - August 1987) followed the aforementioned formalization of today’s dual mandate (Federal Reserve Reform Act of 1977 and the Full Employment and the Balanced Growth Act of 1978), that made it possible to commit to a hawkish monetary policy regime, which, according to Bianchi (2012), the Fed entered into in the second half of 1980.³ In a recent paper Shapiro and Wilson (2022) analyze the 2000-2011 period. Using textual analysis and a Taylor framework, the authors argue that the Fed’s implied inflation target was below 2%.

Unlike the previous studies, our focus is on how monetary policy is shaped by the discussion of non-dual mandate related concerns in Fed speeches. An analysis of the textual features identifies several structural breaks at dates coinciding with important events and monetary policy transitions, such as the end of the Bretton Woods international monetary system, peaks in the fed funds rate and inflation, the passage of banking regulation, changes in Fed Chair, and the global financial crisis. We find that the hawkish policy regime is significantly associated with non-dual mandate related discussion centering around the use of monetary policy to achieve financial stability objectives. There is a recent literature (see Peek et al. (2016) and Istrefi et al. (2021)) emphasizing the role of financial stability related factors that contribute to the interpretation of communication and monetary policy of the Federal Reserve.⁴ Our study differs by starting with a generic model of textual features without pre-specifying the scope of text analysis to financial stability. In addition, we are

³Before the late 1970s, the Fed’s role was to provide a flexible supply of currency and central bank reserves to prevent banking panics and “to promote maximum employment, production, and purchasing power,” as prescribed by the Employment Act of 1946. Moreover, the institutional view under Chairman Martin (April 1951- January 1970) was that the Fed had an obligation to contain the interest rates for government debt by financing deficits (Meltzer, 2009). Within the boundaries of the Gold Standard, inflation was controlled automatically, but this changed after its abandonment in 1971. The experience of the Great Inflation in the 1970s led to an increased appreciation of price stability by Congress, paving the way for the Fed’s “dual mandate” to promote maximum sustainable employment and price stability.

⁴van Dieijen and Lumsdaine (2019) consider whether financial stability has become a de facto ternary mandate but do not consider the implications for monetary policy.

able to capture more sophisticated text features, for instance the policy discussion regarding whether one can use monetary policy to achieve a financial stability goal. Moreover, we argue that the different layers of communications unveiled by our model are the key to understanding monetary policy stances and the policy regimes.

The paper is organized as follows. Section 1 discusses the data we use and Section 2 introduces the NLP methods. Thereafter, Section 3 introduces the text features extracted using NLP and presents a set of descriptive findings. Section 4 discusses the econometric results. Finally, Section 5 concludes.

1 Data

We use three different types of data. The first is a collection of speeches given by Federal Reserve Bank presidents and members of the Board of Governors, which we use to measure Fed communication about financial stability and its institutional mandate. The second is a corpus of journal articles and working papers related to central banking, which we use to refine the natural language processing (NLP) models, allowing us to extract higher quality features from the Fed speeches. The third is a set of macroeconomic and financial variables, which are used as both controls and the dependent variable in different regression exercises. For the sake of consistency and to ensure coverage over the long sample period, we take most of these variables from the Macrohistory Database, introduced by Jordà et al. (2016).

1.1 Federal Reserve Speeches

Our primary source of text data is a novel collection of speeches given by presidents of Federal Reserve Banks and members of the Board of Governors of the Federal Reserve System. It includes 5,742 speeches, given by 94 speakers, and spans the period between 1914 and 2020.⁵ While the coverage of speeches given is not complete over the sample period, it is, as far

⁵We build on a sample by van Diejen and Lumsdaine (2019), which we extend over a longer time period and augment with speeches of the Federal Reserve Bank presidents.

as we are aware, the most comprehensive collection assembled. Table 1 shows the speakers, districts, and speech counts for the 20 individuals who appear most frequently. Notably, the list contains at least one president from all Federal Reserve Banks other than Cleveland, as well as members of the Board of Governors.

Most Common Speakers in Fed Speech Corpus

Speaker	District	Count
Robert Forrestal	ATL	312
Michael Moskow	CHI	254
Ben Bernanke	FRB	205
Robert Parry	SF	169
Alan Greenspan	FRB	167
Paul Volcker	NY, FRB	166
William Dudley	NY	161
Richard Fisher	DAL	160
Jeffrey Lacker	RIC	140
Eric Rosengren	BOS	137
Darryl Francis	STL	134
John Williams	SF, NY	134
William Poole	STL	132
Dennis Lockhart	ATL	130
Charles Evans	CHI	116
Roger Ferguson	FRB	115
Narayana Kocherlakota	MIN	113
Charles Plosser	PHL	110
Thomas Melzer	STL	99
Hugh Galusha	MIN	93

Table 1: Our corpus contains 5,742 speeches from 94 Federal Reserve officials, spanning the period between 1914 and 2020. The table above shows the speaker name, Federal Reserve Bank district, and speech count for the 20 individuals with the highest speech counts in the sample.

Speech frequency increases over the sample period. This is a consequence of two factors. First, the Fed increased the frequency of its public communication over the sample. And second, works that were produced more recently are more likely to have been digitized and incorporated into an online collection. As shown in Figure 1, there are fewer than 20 speeches

in the corpus for most years between 1914 and the early 1960s. However, after 1960, the number of speeches given and available annually tends to rise over time. The number also jumps again in the 1990s.⁶ We will focus primarily on the period between 1960 and 2020.

Federal Reserve Speech Count

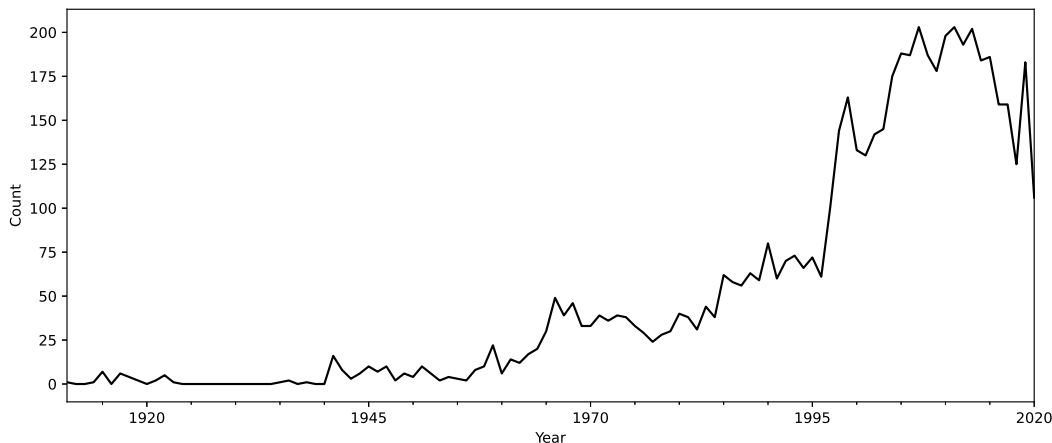


Figure 1: The figure above plots the annual count of Federal Reserve speeches in the corpus. While the corpus spans the period between 1914 and 2020, both coverage and speech frequency increase considerably in the 1960s and again in the 1990s.

1.2 Journal Articles and Working Papers

The second text corpus consists of journal articles and working papers that were drawn from the Semantic Scholar Open Corpus (S2ORC), introduced in Lo et al. (2020). The latest version of the dataset contains metadata for 136M papers, including titles, abstracts, and citations. Within this corpus, we focus on the 2.3M journal articles and working papers that were identified by S2ORC as being from the field of economics.

We next identify the subset of articles that discuss topics related to macroeconomics, monetary economics, and financial markets. We then filter those articles using two additional

⁶In the late 1990s under Chairman Greenspan, there was a gradual shift to more transparency, which is not only reflected in a large increase in the frequency of speeches by presidents of Federal Reserve Banks and members of the Board of Governors, but also in other communication, such as the introduction in 1994 of FOMC statements in conjunction with all rate changes and later on in 2000 for all FOMC meetings.

criteria. First, the article must have an abstract available, because abstracts will be used in the training process. And second, it must be published in a journal or working paper series that has at least 500 entries in the S2ORC database. The second criterion is intended to filter out articles from obscure journals. The final sample comprises 328,370 articles.

From this final sample, we construct two datasets. The first consists of the full abstract text for the 328,370 articles. This is used to extend and fine-tune the pretraining of the NLP models we use in this paper, which are already pretrained on large text corpora, including the full English language Wikipedia.⁷

To construct the second dataset, we start by selecting abstracts that reference central banks, which reduces the sample to 29,781 articles. We then tokenize (divide) each abstract into sentences and identify two types of sentence pairs: 1) sentences in the same abstract; and 2) sentences in different abstracts. We randomly select an equal number of both types of pairs, yielding a total sample size of 194,227. This is intended to emulate the dataset construction process for the next sentence prediction (NSP) task in BERT models (Vaswani et al., 2017).⁸ We do not, however, train on the NSP task, but instead fine-tune models using the approach in Reimers and Gurevych (2019), because our objective is to measure the similarity between of sentences.

Table 2 provides article counts for a subset of the 10 most common journals and working paper series in the sentence pair corpus. The most common is the Social Science Research Network (SSRN), which accounts for almost 14% of the 29,781 articles. In total, there are 283 journals and working paper series included in the corpus.

1.3 Macroeconomic and Financial Data

In addition to the text data, we also use macroeconomic and financial data in regression exercises. All variables used are listed in Table 3, along with a brief description of the variable and its source. With the exception of the output gap, all variables are taken from the

⁷See Section 2.2.1 for a description of the models and pretraining process.

⁸This task involves using the current sentence in a sequence to predict the next sentence.

Most Common Journals in Sentence Pair Corpus

Journal	Article Count
SSRN Electronic Journal	4,164
Journal of Banking and Finance	1,388
Journal of Money, Credit and Banking	618
IMF Working Papers	590
Journal of Finance	585
National Bureau of Economic Research	427
Applied Economics	369
Econometric Reviews	357
Economic Modelling	356
Journal of International Money and Finance	348
Other journals	20,579
Total number of articles	29,781

Table 2: This table provides article counts for the 10 most common journals and working paper series in the corpus we construct to train our NLP models. In total, the final corpus includes 283 journals and 29,781 articles. It is a subset of the 2.3M economics articles in the S2ORC corpus (Lo et al., 2020).

Macrohistory Database (Jordà et al., 2016). The output gap is measured as the percentage difference between actual and potential (real) GDP. Actual GDP is taken from the Bureau of Economic Analysis (BEA) and potential GDP is measured by the Congressional Budget Office (CBO). All variables are measured at an annual frequency and span the period between 1960 and 2020.

2 Method: FedBERT

Our objective is to examine the Fed’s interpretation of its own mandate. We do this by extracting text features that are informative about this objective from Federal Reserve speeches and examining their variation across time, district, and speaker. Many of the features that are most informative about our research question involve sequences of words, such as sentences and paragraphs, rather than individual words. As such, we will use sequence-to-sequence

Macroeconomic and Financial Series

Variable	Description	Source
bond return	total return on government bonds	Jordà et al.
cpi inflation	consumer price index (1990 = 100)	Jordà et al.
crisis	dummy for financial crisis	Jordà et al.
debt-to-gdp ratio	public debt-to-GDP ratio	Jordà et al.
equity return	total return to equity	Jordà et al.
house prices	nominal house prices	Jordà et al.
ltd ratio	loan-to-deposit rate	Jordà et al.
output gap	percentage difference between actual and potential (real) GDP	CBO, BEA, FRED
risky return	total return on risky assets	Jordà et al.
safe return	total return on safe assets	Jordà et al.
interest rate	short term interest rate	Jordà et al.
total loans	total loans originated to non-financial sector	Jordà et al. Jordà et al.

Table 3: The output gap variable is taken from the St. Louis Fed’s FRED database Federal Reserve Bank of St. Louis (n.d.). It is computed as the percentage difference between actual and potential (real) GDP. The underlying measures of actual and potential GDP are computed by the Bureau of Economic Analysis (BEA) and Congressional Budget Office (CBO), respectively. The remaining macroeconomic and financial variables are taken from the Macrohistory Database (Jordà et al., 2016).

(S2S) modeling for NLP tasks, building on recently introduced variants of the Transformer model (Vaswani et al., 2017). This section will provide a brief overview of the NLP tasks we perform, the models we use, and the types of features we extract. We will also discuss the pretraining and fine-tuning process for these models, but will relegate the technical details about the NLP models to Section A.3 of the Appendix.

2.1 The NLP Task

Many NLP tasks can only be performed using an S2S model. Others do not require an S2S model, but achieve better performance when one is employed. Machine translation,

for instance, yields low quality results when the translation task is performed at the word level. An entire sentence typically needs to be processed and interpreted before a suitable translation can be generated. This is, in part, because words have different meanings in different contexts. See A.2 in the Appendix for an overview of S2S modeling.

Another example of an inherently sequential NLP task is extractive question answering. This involves finding a subsequence of text that contains an answer to a question. An NLP model would receive a “context” sequence (the original text) and a query as inputs. It would then yield a subsequence of the context as an output. Both the inputs and output are sequences. State-of-the-art S2S models, such the transformer models we use in this paper, are naturally suited to such problems.

For our purposes, S2S modeling will mostly involve the transformation of a sequence of words in a paragraph into a sequence of “contextualized” words. The words themselves will be represented using dense vectors called embeddings.⁹ The contextualized embeddings will encode information about the meaning of the word in the context in which it was used. For example, the word “run” will be encoded differently in the following two sequences:

Sequence 1: “If depositor concerns are not addressed, there could be a bank run.”

Sequence 2: “If stock prices increase again tomorrow, it will be the longest bull run in history.”

The output of transformer models – sequences of contextualized embeddings – can be used to perform sentiment analysis, zero shot classification (classification without training), textual similarity measurement, machine translation, contextual embedding generation, and extractive text summarization – all with state-of-the-art performance. Transformers can also be extended and trained to perform supervised learning tasks, such as sentiment classification. Furthermore, such extensions can be fine-tuned at a low computational cost to yield

⁹See Gentzkow et al. (2019) for an overview of embeddings.

further improved performance in a specific domain, such as central banking communication.

We will use transformer models to extract features from Fed speeches, focusing on elements related to the Fed’s mandate. Using transformer models that were fine-tuned for different language tasks will enable us to measure subtle features of a speech, such as the speaker’s primary concern in a given paragraph or the extent to which a statement expresses approval of monetary policy as a means of achieving financial stability. It will also allow us to identify whether certain content – such as concern about bank liquidity – is present in a paragraph.

2.2 The NLP Models

Prior to the introduction of the Transformer model (Vaswani et al., 2017), most S2S models treated language as an autoregressive process, where the meaning of a word was assumed to depend only on the words that immediately preceded it. The Transformer model replaced the autoregressive elements in S2S modeling with the attention mechanism, which enabled the model to learn the relationships between words, irrespective of where they appeared in a sequence. Figure 2 illustrates the application of the attention mechanism to sequences taken from two Federal Reserve speeches.

One benefit of the removal of the autoregressive assumption is that it enabled the construction of models that had a greater capacity for parallelization. The Transformer model also introduced several other innovations that have proven useful for natural language processing, including multi-headed attention and positional encodings. We discuss the Transformer model’s architecture and contributions to S2S modeling in Section A.3 of the Appendix.

In addition to the improvements in model architecture, transformer models also made use of extensive pretraining. This entails training the model on a large, auxiliary dataset that is unrelated to the NLP task of interest. The purpose of pretraining is to create a baseline or “foundation” model that understands a given language using large amounts of

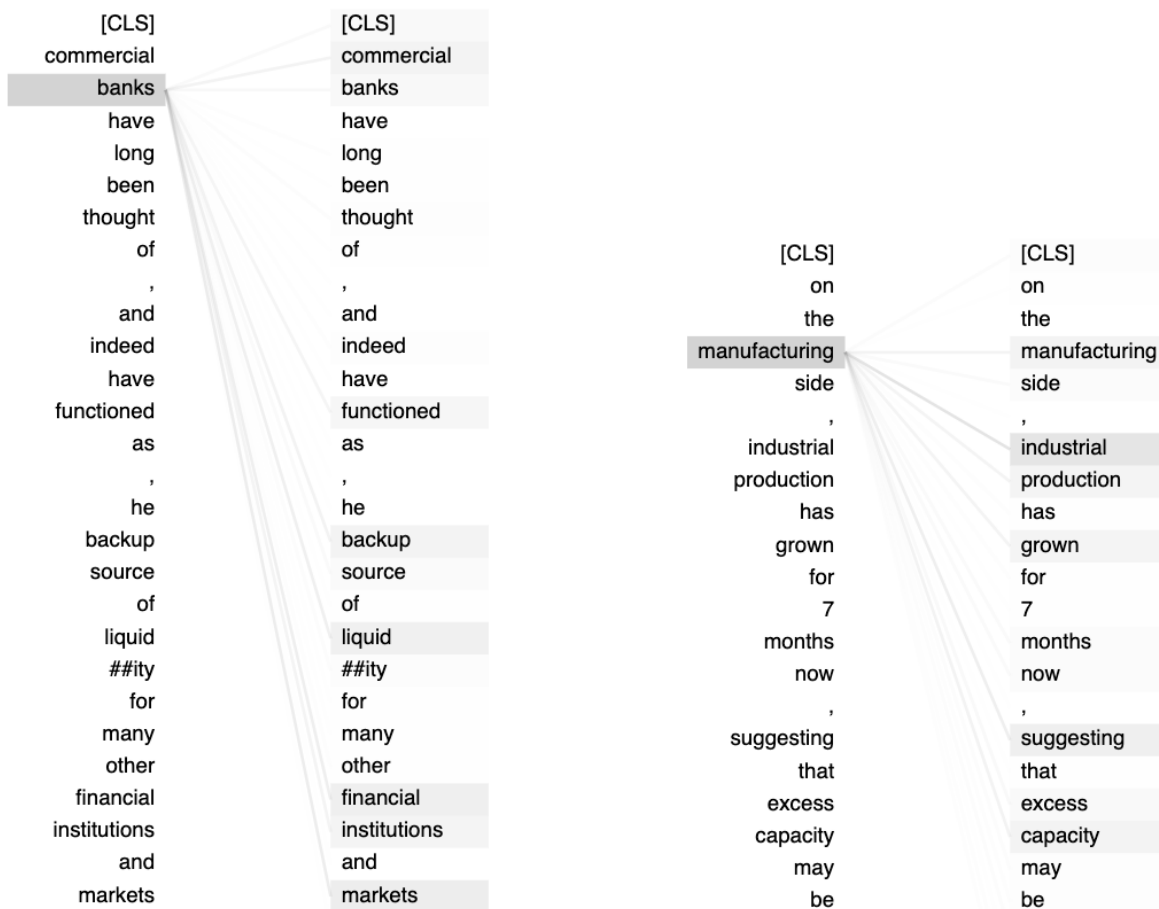


Figure 2: The left panel illustrates self-attention applied to the word “banks” in a sequence from a speech given by Gary Stern in January 2009. The words “liquid(ity),” “financial institutions,” and “markets” are not close to “banks” in the sequence, but the attention mechanism determines that they are highly relevant for contextualizing banks. The right panel illustrates self-attention applied to the word “manufacturing” in a sequence taken from a speech given in September 2002 given by Cathy Minehan, then-President of the Federal Reserve Bank of Boston. The attention mechanism identifies the importance of “industrial,” “production,” and “capacity” for contextualizing “manufacturing” in the sequence. Note that the [CLS] token indicates the start of a sequence and ## indicates that a token is outside of the corpus.

text. This model can then be fine-tuned on a small amount of text that is closely related to the language task of interest. In our case, for example, the model can be pretrained on all English language Wikipedia articles, newspaper articles, and journal articles. It can then be fine-tuned on a much smaller corpus of Federal Reserve speeches. Since we use a variety of pretrained and fine-tuned models in this paper – and also extend the pretraining and fine-tuning process – we discuss both topics in the remainder of this subsection.

2.2.1 Pretraining

The NLP literature has shown that pretraining a language model can yield substantial performance benefits on subsequent (downstream) tasks (Dai and Le, 2015; Peters et al., 2018; Salimans and Sutskever, 2018; Howard and Ruder, 2018). Indeed, some of the most substantial gains in the development of language models have come from changes in the pretraining procedure and data, rather than changes in model architecture. The BERT and RoBERTa models, for instance, achieved state-of-the-art performance on NLP benchmarks using the same Transformer architecture introduced in Vaswani et al. (2017), but with modified training processes.

BERT. The BERT model, introduced in Devlin et al. (2019), proposed several modifications to the training process used in Vaswani et al. (2017). Most notably, the pretraining process concentrated on two tasks: 1) masked language modeling (MLM); and 2) next sentence prediction (NSP). Both tasks automatically generate labels, allowing for pretraining on arbitrarily large text datasets that have not been labeled by a human.

For concreteness, we will consider how the BERT model would use the passage given in the quote below in the training process. The text is taken from a speech given by Gary Stern, then-President of the Minneapolis Federal Reserve Bank, in January 2009.

“Commercial banks have long been thought of, and indeed have functioned as, the backup source of liquidity for many other financial institutions and markets.

Banks continue to play this role, but it has become more challenging today to do so because some lenders find themselves capital-constrained as a result of recent losses and/or sizable, unanticipated additions to their balance sheets of formerly off-balance-sheet instruments.”

The MLM task involves masking randomly-selected words in a sequence and then training the model to predict them. In the quote above, MLM might generate the following sequence and labels.

Sequence: “Commercial [MASK]₁ have long been thought of, and indeed have functioned as, the backup source of [MASK]₂ for many other financial institutions and markets.”

Labels: [MASK]₁ = banks, [MASK]₂ = liquidity.

Another benefit of pretraining with MLM is that it allows for bidirectionality in the interpretation of sequences. Rather than training BERT to predict the next word in a sequence, conditional on the preceding words, it is instead trained to predict a missing word in a sequence, conditional on all words before and after it. This results in a pretrained model that has a substantially expanded capacity to understand language.

The other training task, next sentence prediction (NSP), also allows for the automatic generation of labels. This task presents the model with a sequence of two sentences drawn from the corpus. The model must determine whether the second sentence follows the first or whether it is drawn from a different place in the document. Again, returning to the speech from Gary Stern, consider the following three sentences.

Sentence A: *“Commercial banks have long been thought of, and indeed have functioned as, the backup source of liquidity for many other financial institutions and markets.*

Sentence B: *“Banks continue to play this role, but it has become more challenging today to do so because some lenders find themselves capital-constrained as a result of recent losses and/or sizable, unanticipated additions to their balance sheets of formerly off-balance-sheet instruments.”*

Sentence C: *“On the positive side, term funding is more readily available than at the height of the crisis, and risk premia have diminished through much of the financial sector.”*

In the speech, Sentence B follows Sentence A; whereas Sentence C is selected from a random location in the text. The MLM process might yield sequence (A,B), which could be passed to BERT with the training label `IsNextSentence`. Alternatively, it could yield the sequence (A,C), which could be passed to BERT with the training label `IsNotNextSentence`.

Because neither the MLM nor NSP tasks require labeled data, Devlin et al. (2019) were able to train BERT on two large text corpora: the 800M word Book Corpus from Zhu et al. (2015) and a 2500M-word corpus constructed from English language Wikipedia articles.¹⁰ This increase in the size of the training corpus allows for a corresponding increase in model scale.

Devlin et al. (2019) show that the pretrained BERT model can then be fine-tuned to achieve state-of-the-art performance on downstream tasks, such as question answering and language inference. This entails training a model to take the contextualized embedding output from BERT to use as an input to a supervised learning task. This can be done with a few hours of GPU training and does not require modifications to the model’s architecture that are specific to the task. The authors used this approach to achieve state-of-the-art performance on 11 NLP benchmarks.

The fine-tuned versions of BERT introduced in Devlin et al. (2019) remain near the state-of-the-art on the GLUE (Wang et al., 2019) and SQuAD (Rajpurkar et al., 2018) benchmarks,

¹⁰Devlin et al. (2019) introduced two pretrained versions of BERT: 1) `BERTBASE`, which has 12 transformer blocks, a hidden dimension size of 768, 12 attention heads, and 110M parameters; and 2) `BERTLARGE`, which has 24 transformer blocks, a hidden dimension size of 1024, 16 attention heads, and 340M parameters.

which are used to assess the performance of NLP models on specific language tasks. The models that have since surpassed them use either modifications of BERT or alternative Transformer model architectures. Fine-tuned BERT models have also demonstrated better performance on domain-specific NLP tasks than models that were trained exclusively using text from that domain.

RoBERTa. Liu et al. (2019) introduced the RoBERTa model, arguing that the gains from BERT are primarily attributable to the modification of the pretraining process, which allows for an increase in model size by several orders of magnitude. They also argue that this is true of most landmark language models, such as ELMo (Peters et al., 2018), GPT (Salimans and Sutskever, 2018), XLM (Conneau and Lample, 2019), and XLNet (Yang et al., 2019).

The RoBERTa model attempts to provide further improvements over the pretraining process in BERT. Specifically, it employs larger datasets, removes the NSP task, trains on larger sequence lengths, employs a dynamic masking pattern, and uses a new dataset (CC-News).

Liu et al. (2019) find that the resulting modifications to the pretraining process boost the “robustly optimized” version of BERT (RoBERTa) to match or outperform the models that were introduced after BERT. We will use versions of the RoBERTa model for many of the NLP-related tasks in this paper.

2.2.2 Fine-tuning

The success of supervised transfer learning – that is, training on a corpus in one domain and performing prediction in a different domain – was one of the motivations for the construction of BERT.¹¹ Since BERT does not require human labeling, it opened up the possibility of performing transfer learning with much larger text corpora. This first step is typically described as “pretraining” in the context of BERT.

¹¹See Conneau et al. (2017) and McCann et al. (2017) for examples of transfer learning.

Once a model has been pretrained, we can fine-tune it to perform a task in our domain of interest. We typically accomplish this by *freezing* all of the layers, except the output layer and a few of the layers that directly precede it. We then perform supervised learning on a domain-specific task and with domain-specific data. Note that freezing a layer prevents its parameters from being updated. This means that in `BERTBASE`, for instance, fewer than 100K of the 340M parameters will need to be trained. The rest of the model will serve as a state-of-the-art text feature extractor.

2.3 Text Feature Generation

We have now discussed the NLP problem we encounter in extracting information from Federal Reserve speeches, the class of models we use to solve it, and the features that make such models appealing. In this subsection, we will discuss three different types of text features that these class of models can extract. We will start with zero shot learning, which we will use to classify text into categories without training a model. We will then discuss extractive question answering, which we will use to identify areas of concern for the speeches in our corpus. Finally, we will examine semantic textual similarity (STS) measurement, which will allow us to evaluate whether two statements are closely related. In one application, we will use this to determine whether speakers discuss monetary policy as a tool for achieving financial stability.

2.3.1 Zero Shot Classification

For some exercises in this paper, we will need to perform text classification. This normally involves training a supervised learning model with labelled text categories. However, in the absence of an appropriate dataset or labels, it will not be possible to perform supervised classification. An alternative approach, introduced in Pushp and Srivastava (2017), trains a general model that can perform classification on arbitrarily chosen categories without first

labelling the data.¹²

Algorithm 1 below provides an outline of the approach to zero shot classification taken in Pushp and Srivastava (2017). We start by constructing a corpus using labelled text data that is outside of our domain of interest. We then embed the text sequences and the labels in the same embedding space. Pushp and Srivastava (2017) do this by averaging word embeddings for the sequence; however, we can also do this using a sentence transformer, such as SBERT, which will generate a contextualized embedding for the sequence. We then concatenate pairs of sequence and label embeddings and train a supervised model to predict whether or not they match.

We can then perform zero shot classification by embedding candidate labels and sequences of interest in the same space and then using the model to perform classification. This will yield the probability that the sequence and candidate label match.

Algorithm 1: Zero Shot Text Classification

- 1 Construct a corpus of text sequences and labels.
 - 2 Embed the sequences and labels in the same space.
 - 3 Concatenate pairs of sequence and candidate label embeddings.
 - 4 Train supervised model to predict whether pair matches.
 - 5 Embed labels and sequences of interest in same space.
 - 6 Perform classification.
-

Below, we provide an example of zero shot classification on our text corpus using the approach outlined in Algorithm 1 and with a fine-tuned BERT model. We attempt to classify whether the statement given is about one of the following four categories: 1) financial stability; 2) output; 3) inflation; and 4) the labor market. The scores can be interpreted as a probability distribution over the categories.

Sequence: *“Banks continue to play this role but it has become more challenging today to do so because some lenders find themselves capital constrained as a result*

¹²See Yogatama et al. (2017), Zhang et al. (2019), and Yin et al. (2019) for alternative approaches to zero shot text classification.

of recent losses and or sizable unanticipated additions to their balance sheets of formerly off balance sheet instruments.”

Candidate Classes: [‘financial stability’, ‘output’, ‘inflation’, ‘labor market’]

Scores: [0.718, 0.203, 0.048, 0.031]

The model identifies “financial stability” as the most probable label, even though it was not trained on a central bank communication corpus. As we will see in later exercises, transformer models perform well on simple zero shot classification tasks on Fed texts.

2.3.2 Extractive Question Answering

The original BERT model (Devlin et al., 2019) was trained to perform extractive question answering on the Stanford Question Answering Dataset (Rajpurkar et al., 2016), which consists of 100,000 human-labelled questions, answers, and context passages. The context is taken from a Wikipedia article. The model must correctly predict the subsegment of the text that contains the correct answer to the question. BERT exceeded human-level performance on both benchmarks for that task.

We use extractive question answering with pretrained BERT models in several exercises in this paper. In most cases, we attempt to discern the speaker’s most pressing concern in a passage. The two examples given below provide the query, context, and model output for passages in a speech given by then-President of the Federal Reserve Bank of St. Louis, Darryl Francis, in February 1972.

Query 1: What is the most significant concern in the passage?

Context 1: “The suspension of the convertibility of the dollar into gold and the imposition of a 10 percent import surcharge last summer ran the risk of mass foreign retaliation in the form of destructive trade barriers.”

Output 1: mass foreign retaliation

Query 2: What is the most significant concern in the passage?

Context 2: “Another significant aspect of the President’s new policies announced August 15 are the measures taken to reverse the deteriorating US balance of payments.”

Output 2: deteriorating US balance of payments

Once the most significant concerns have been extracted from a speech, they can then be converted to a set of numerical features, such as contextualized word embeddings via BERT or a sentence embedding via SBERT.

2.3.3 Semantic Textual Similarity

For several exercises, we will need to measure the similarity between pairs of passages from speeches using a measure called semantic textual similarity (STS). The BERT and RoBERTa models, which we use throughout the paper, achieve state-of-the-art performance on the STS task; however, they require each pair of passages to be input into the model simultaneously to produce an STS score and hence are computationally intensive.

Sentence BERT. An alternative to this approach is to construct sentence embeddings, which we can compute individually for each passage, and then measure the cosine similarity between pairs of embeddings. See Equation (1) for the construction of cosine similarity and note that S is a sentence embedding. Note that 10,000 sentences would require around 50M STS pair computations. As such, it is necessary to use this more computationally efficient means of performing the comparison.

$$\text{sim}(S_i, S_j) = \frac{S_i \cdot S_j}{\|S_i\| \|S_j\|} \quad (1)$$

Sentence embeddings have been explored in Kiros et al. (2015), Conneau et al. (2017), and Cer et al. (2018). We will make use of the approach in Reimers and Gurevych (2019), which

modifies the pretrained BERT and RoBERTa models to produce contextualized sentence embeddings, rather than contextualized word embeddings. This approach uses Siamese and triplet networks (Schroff et al., 2015) to train the model, which have objective functions that are comparable to cosine similarity. It is trained using the SNLI (Bowman et al., 2015) and NLI (Williams et al., 2018) datasets.

SBERT Pretraining. The sentence BERT (SBERT) models employed in the paper were pretrained first by Devlin et al. (2019) and then Reimers and Gurevych (2019). We refine the pretraining using an unsupervised learning process called a Transformer-based Sequential Denoising Auto-Encoder (TSDAE). Specifically, we use the training process in Wang et al. (2021), along with the data described in Section 1, which is compiled from the S2ORC corpus (Lo et al., 2020).

The TSDAE approach to training sentence embeddings was based on earlier work by Vincent et al. (2010) and Hill et al. (2016). The training task entails injecting noise into the input embeddings and then training the model to recover the denoised embeddings. Much like the MLM and TSP tasks for BERT, TSDAE does not require labels, making it an attractive choice for refining the pretrained model on domain-specific text. It also achieves state-of-the-art results, which approach the performance of supervised methods on domain-specific texts.

SBERT Fine-tuning. In addition to pretraining the SBERT model, we also fine-tune it on the semantic textual similarity (STS) task using pairs of sentences drawn from paper abstracts in the S2ORC corpus (Lo et al., 2020). See Section 1 for an overview of the construction of the dataset.

3 Interpretation of Text Features

In the previous section, we discussed three methods for extracting text features using transformer models: zero shot classification, extractive question answering, and semantic textual similarity measurement. In this section, we will describe a selection of the features we extracted from Federal Reserve speeches using those methods. Our objective is to examine whether they adequately represent the concept we intended to measure in the text. We describe how some of the more informative features evolved over time and across district.

Most of the features discussed in this section are constructed from paragraph-length sequences. For relatively simple features, we use zero shot classification with the RoBERTa model to determine whether a paragraph discusses a certain concept, such as financial stability. For more subtle concepts, we extend the pretraining and fine-tuning of the model to improve its capacity to understand central bank texts. We then measure cosine similarity between paragraphs of speech text and the statement we want to evaluate. Finally, we standardize the raw scores by subtracting the sample mean and dividing by the sample standard deviation.

Using S2S models – and, in particular, the RoBERTa model – has at least three advantages. First, in contrast to dictionary-based methods, RoBERTa’s classifications are based on the entire paragraph, taking long-run dependencies, negations, and modifiers into account. Second, RoBERTa automatically identifies related terms and, thus, does not rely on the ex-ante identification of all relevant terms. And third, unlike average word embeddings, RoBERTa accounts for the context in which words are used.

3.1 Time Variation in Text Features

Since we are interested in measuring changes in the Federal Reserve System’s interpretation of its mandate over time, we will start by examining the time variation in the text features. In this subsection, we discuss what features we measure, how they appear to perform, and

whether they provide any insights into the Federal Reserve’s interpretation of its mandate.

3.1.1 Dual Mandate Content

We first partition the text into content that is related to the Fed’s “dual mandate” and content that is not. Because the sample starts prior to the introduction of the dual mandate, we measure whether paragraphs discuss (at least one of) inflation, employment, or output growth, rather than attempting to identify references to the dual mandate itself. The evolution of the series is plotted in Figure 3.

Text Feature: Dual Mandate Content

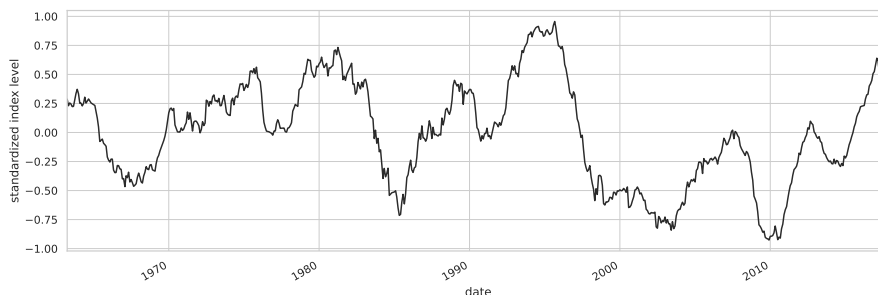
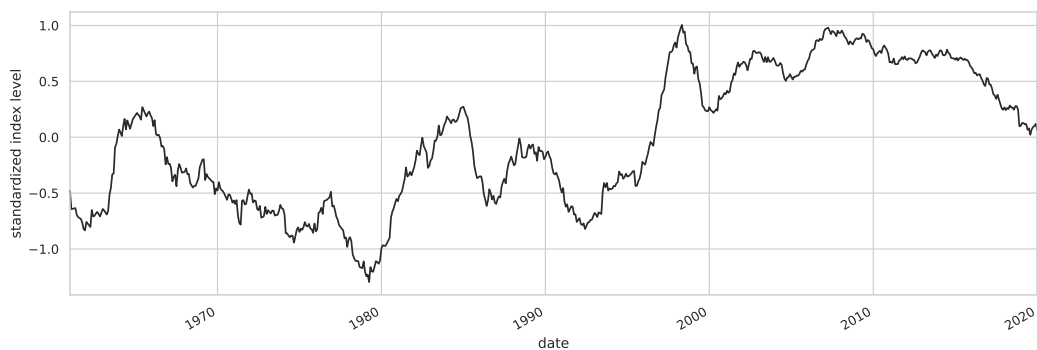


Figure 3: The figure shows the 2-year rolling mean of the standardized cosine similarity score from BERT for “inflation, employment, and output growth” and each paragraph in the text.

The plot indicates that there is short- and long-term variation in the dual mandate content of speeches. We examine the content of paragraphs that have low dual mandate content and find that they are largely related to the financial sector, financial crises, and banking regulation. To demonstrate this informally, we use extractive question answering with the RoBERTa model to identify the speaker’s concern in each paragraph that has a low dual mandate content score. We then construct word clouds for the pre-Great Moderation period and the Great Moderation period (Figure 4). In both cases, the discussion is dominated by topics related to finance and banking regulation.

Latin American debt, amounting to 176.5 percent of their capital (Sachs, 1987). Thereafter, the financial stability index declined during the first half of the Great Moderation.

Text Feature: Financial Stability



Text Feature: Financial Crisis

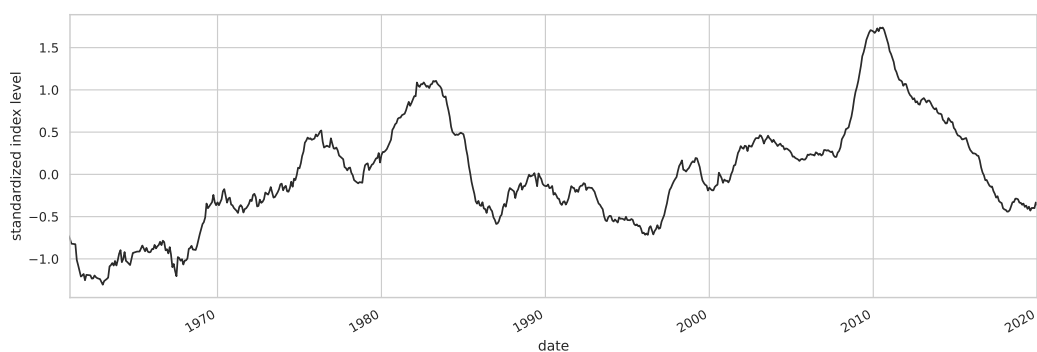


Figure 5: The figures above shows the standardized classification score from BERT for the terms “financial stability” and ”financial crisis.” This is computed by classifying each paragraph as describing “financial stability” / “financial crisis” or not. We then compute the average score for each month, standardize it, and then plot the 2-year rolling mean.

In addition to the financial stability index, we also construct a financial crisis index. This allows us to differentiate between discussions about the Fed’s role in achieving financial stability and deliberations about specific financial crises. The bottom panel of Figure 5 plots the 24-month rolling mean of the financial crisis feature. In contrast to the financial stability discussion, there is a general upward trend in financial crisis content prior to the

start of the Great Moderation in the mid-1980s. There is also a lull between 2000 and 2007. Interestingly, the Asian and Russian financial crises, as well as the collapse of LTCM do not generate a spike in the financial crisis index. This may be explained by the effectiveness of the “Greenspan put” in combination with the much lower exposure of the major U.S. banks to the 1998 external crisis events than the 1982 external crisis events. Also interesting is the peak of the financial crisis index in 2010, the time of the European sovereign debt crisis, and the way the index declines rapidly after that peak. In contrast, the financial stability index remains high. Taken together, these contrasting patterns illustrate that while discussion of the global financial crisis itself quickly subsided, discussions of financial stability in its aftermath continued.

3.1.3 Mandate Description

During the period we examine, the Federal Reserve tended to talk about either the components of its dual mandate or financial stability. In Figure 6, we plot the series for inflation and output growth together with the financial stability index to examine how the Fed’s focus shifted from dual mandate content to financial stability over time and across cycles. Consistent with inflation data, textual content about inflation in speeches increases during the period in the late 1970s and early 1980s, an era that has been classified as discretionary with large deviations from the Taylor rule (see, e.g., Nikolsko-Rzhevskyy et al., 2014). Another notable feature of the plot is that financial stability and output growth content diverge prior to the Great Moderation, but then positively comove thereafter. This suggests that the Great Moderation could have been a turning point with respect to Fed communication. For this reason, we will include sample splits in our regression exercises in Section 4, so that we can evaluate the pre-Great Moderation and Great Moderation periods separately. It is also interesting that the inflation and output growth features appear to comove in the early 1970s and again during the Volcker years.

Text Features: Inflation, Output Growth, and Financial Stability

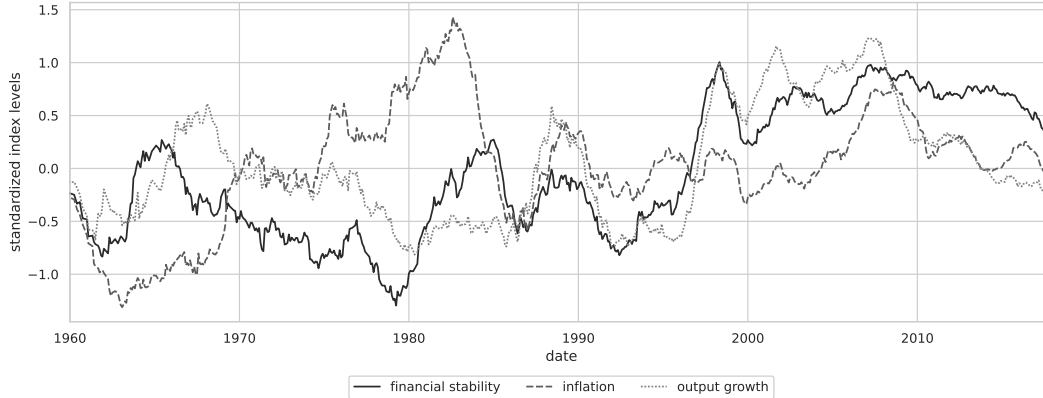


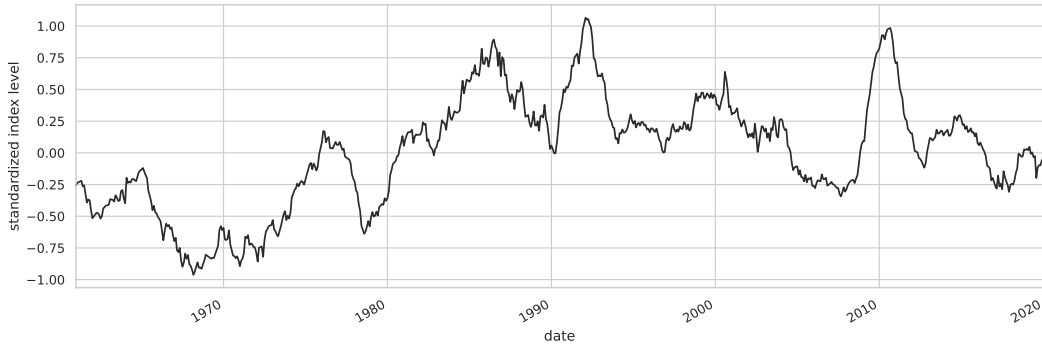
Figure 6: The figure shows the 2-year rolling mean of the standardized classification score from BERT for the terms “financial stability,” “output growth,” and “inflation.”

3.1.4 Financial Stability Position

In addition to measuring broad text features related to financial stability, we also construct two features that provide a direct description of the Fed’s evolving perception of its mandate. The first feature, shown in the top panel of Figure 7, measures the cosine similarity between each paragraph and the statement “banking regulation should be used to achieve financial stability.” Peaks for this series are typically closely related to crisis events and some of the peaks in this figure resemble those seen in the financial crisis text feature in the bottom panel of Figure 5. Yet there are some important differences; in particular, the peak in the late 1980s seems to follow a similar peak in Figure 5, while the peak in 1992, the year the Basel I capital regulations were implemented, is only apparent in the cosine similarity. In contrast, the bottom panel of Figure 7 measures the cosine similarity between each paragraph and the statement “monetary policy should be used to achieve financial stability,” providing us with a measure of how inclined the Fed is to use monetary policy as a means to achieve financial stability.

We can see a clear difference between the evolution of this series and the one for banking regulation, suggesting that the Fed places emphasis on this nuance. In particular, these

Cosine Similarity: Banking Regulation and Financial Stability



Cosine Similarity: Monetary Policy and Financial Stability

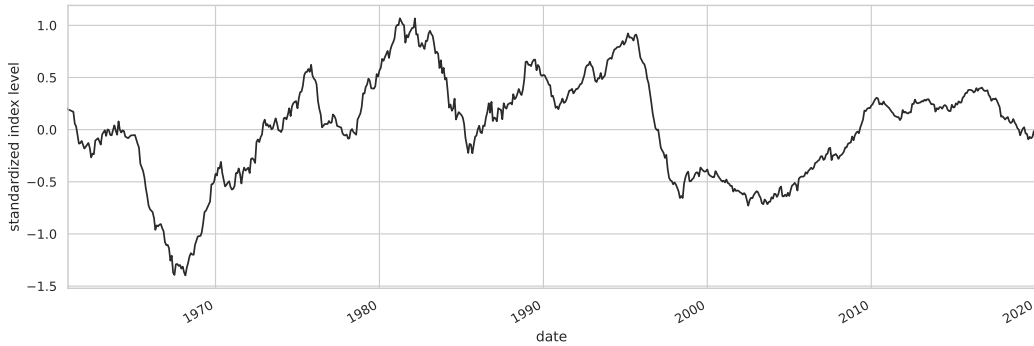


Figure 7: Both panels plot standardized cosine similarity scores. The top panel makes use of a cosine similarity score computed between the statement “banking regulation should be used to achieve financial stability” and the contextualized embeddings for paragraphs with high classification scores for “financial stability.” Cosine similarity scores are repeated for the same exercise in the bottom panel, but using the statement “monetary policy should be used to achieve financial stability” instead of “banking regulation.”

figures indicate less appetite (or alternatively, less ability, given the zero lower bound) for using monetary policy to achieve financial stability in the past two decades than in the three decades prior. In addition, the more jagged nature of the banking regulation plot suggests a more episodic role for banking regulation to achieve financial stability, rather than being a sustained tool.

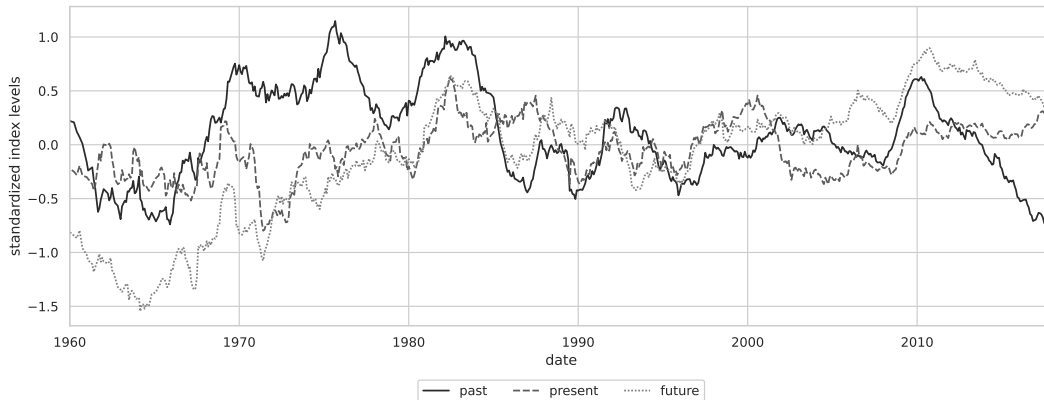
Looking further at the cosine similarity depicted in the bottom panel of Figure 7, it is also apparent that support for the idea that monetary policy should be used to achieve financial stability varies according to the level of accommodation in the Fed’s stance, and in particular, the cosine similarity appears to increase as the central bank tightens and decreases (or, in the last two decades remains low or constant) during easing periods. This pattern is consistent with an endorsement of a forward looking “leaning against the wind” view, rather than a “financial instability is caused by monetary tightening” view. We will explore this aspect in our regressions in Section 4.

3.1.5 Speech Temporal Focus

In addition to distinguishing between the content of financial stability concerns, it may also be useful to identify the tense or the temporal focus of a statement. This could indicate whether a paragraph is discussing a past crisis, an unfolding crisis, or the prospect of a future crisis and may indicate the extent to which Federal Reserve officials rely on historical precedent, are reactive, or are proactive in making their decisions. We use zero shot learning to classify each tense separately, allowing for the possibility that no clear tense is established in a given statement.

The standardized indices for past, present, and future focus are plotted in the top panel of Figure 8. The bottom panel shows the difference between future and past focus, which has been rising since the start of the Great Moderation. The increased use of future tense coincides with the Fed’s emphasis on increased transparency and forward guidance. Prior to the Great Moderation, there is evidence that an increased use of future tense coincides with

Text Features: Past Focus, Present Focus, Future Focus



Text Features: Difference Between Future and Past Focus

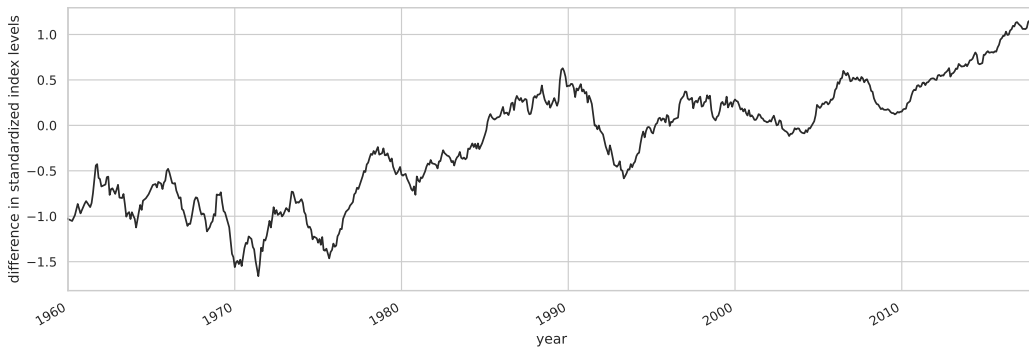


Figure 8: The figure in the top panel shows the standardized classification score for the speech tense for a focus on the “past”, “present” and “future.” This is computed by classifying each paragraph and then computing the average score for each month, standardizing it, and then plotting the 2-year rolling mean. The figure in the bottom panel shows the 2-year rolling mean of the difference between “future” and “past” focus.

tightening episodes, whereas the shift towards a focus on the past coincided with the easing that began in April 1989.

3.1.6 Concern Type and Level

Not all content contained in Fed speeches is equally informative. In order to identify statements and terms that reveal internal concerns, we make use of extractive question answer-

ing, as described in Section 2.3.2. Specifically, for each paragraph, we query the model with “What is the most significant concern in the passage?” This yields a concern, which is extracted from the text, along with a score that indicates the model’s uncertainty.

In Figure 9, we visualize the evolution of “most significant” concerns as they relate to financial stability over time. We do this by first identifying paragraphs about financial stability using zero shot classification. We then extract the concerns from these paragraphs using the approach described above. These concerns are then visualized as a word cloud for three separate periods: 1960-1983, 1984-2006, and 2007-2020. What is telling in the first period is that the most significant concerns align with the Fed’s dual mandate. “Price” and “inflation” are dominant, perhaps reflecting the high inflation episode in the early 70s, but words related to employment, the economy, and monetary policy are also apparent. Overall, most words appear to be related to economic rather than financial terms. In the second episode, more financial words become apparent. Both “banking” and “financial” are larger, while “inflation” has faded somewhat. “stability” is very dominant during this middle period, and “price” also remains important; in the latter part of this episode, a number of the Fed’s speeches focused on the importance of price stability. In the final time period, encompassing the global financial crisis and its aftermath, the words “financial” and “stability” have near-equal prominence, with “price” shrinking in size to be similar to “inflation”, reflecting the muted inflationary period that has existed throughout this time period. Two other items of note: (1) while “financial” and “stability” are concerning terms, “instability” is not used nearly as much, and (2) the concerning words related to financial stability appear to focus more on the “price stability” part of the dual mandate than on the “maximum employment” part.

3.2 Institutional Variation in Text Features

Another source of variation in text features arises from differences in the institutional arrangements and concerns of different Federal Reserve Banks and the Board of Governors.

For instance, the Cleveland and New York Federal Reserve Banks may have different concerns, since there are differences in the compositions of their respective regional economies and the bank holding companies they supervise. Furthermore, given the local appointment processes, some district banks may favor presidents with academic profiles, while others may favor those with private sector experience. In this subsection, we will attempt to document the variation we were able to measure across districts and the Board of Governors.

Financial Stability Focus. First, we find that the Federal Reserve Banks of New York and Richmond have the highest average financial stability scores, consistent with the fact that most of the systemically important financial institutions (SIFIs) are in those two districts. In addition, many of the speakers with high scores (e.g., Bernanke, Dudley, Evans, Geithner, George, Kocherlakota, Kohn, Kroszner, Lacker, Minehan, Plosser, Stern, Tarullo, Warsh) were officials in the period surrounding the GFC (2007-2009), a time when the overall financial stability index was near its highest levels. Many of the other speakers with high scores were members of the Board of Governors; in general Reserve Bank Presidents' speeches have less financial stability content, although the three highest index scores are associated with Reserve Bank Presidents.¹³ The most negative scores are also associated with Reserve Bank Presidents, many of whom served near the start of the Great Moderation (e.g., Eastburn, Kimbrel, Parry, Roos, Willes). Overall, the results align well with our intuition and we refer the interested reader to Tables 15 - 17 in Appendix Section A.4 for the average scores for financial stability for each speaker and institution.

Future Focus. Second, we find that the Federal Reserve Banks of New York and Richmond also have the highest score on future focus (see Table 15). This is consistent with the previous finding, given the strong connection between future focus and financial stability documented in the cosine similarity results. As noted previously, these districts supervise the largest,

¹³In addition to Minehan and Lacker, Loretta Mester, President of the Cleveland Fed has a very high score; the Cleveland Fed co-hosts an annual financial stability conference.

most systemically important financial institutions in the US. The NY Fed also has the highest scores for past and present focus; taken together, these results indicate that tense usage is clearer/more distinct in speeches by its presidents than in speeches by presidents of other districts. Looking across the individual speakers' scores given in Tables 16 and 17, there is a strong positive correlation between the future focus and financial stability scores (0.65), suggesting that discussions of financial stability involve use of more forward-looking language.¹⁴ As a result, many of the speakers with high future focus and financial stability scores are the same ones, encompassing many of those that were on the FOMC during the GFC.

Academic Inclination. There is a long-standing debate in the literature over the extent to which academic discussions influence central bank deliberations and policy-making. Some, such as Mankiw (2006), diverge from the majority position and argue that policy making is largely uninformed by the academic literature. Reviewing a memoir written by a former member of the Board of Governors, Mankiw states:

[His] analysis of economic fluctuations and monetary policy is intelligent and nuanced, but it shows no traces of modern macroeconomic theory. It would seem almost completely familiar to someone who was schooled in the neoclassical- Keynesian synthesis that prevailed around 1970 and has ignored the scholarly literature ever since. ... It is typical of economists who have held top positions in the world's central banks.

In addition, Howitt (2012) argues that central bankers often face crises that have not been adequately studied by the academic literature and, thus, typically lead the literature. As such, their lack of reliance on the literature may be a consequence of the lack of availability of useful work.

¹⁴In contrast, the correlation between past or present focus and financial stability is -0.14 and 0.34, respectively.

With respect to financial stability, the academic literature has largely argued that macroprudential policy – not monetary policy – should be used to achieve financial stability (Vollmer, 2021). As such, the Fed’s interpretation of its mandate and its belief about what may be achievable with monetary policy may also depend on its openness to dialogue with academics. To capture this, we use zero shot learning to identify whether a paragraph refers to an “academic debate or the academic literature.”

Just as the NY and Richmond districts have the highest future focus and financial stability scores, so also do they have a high academic focus score. In the case of Richmond, this no doubt reflects the academic backgrounds of the two main presidents in our corpus (Broadus and Lacker), while for the NY Fed the presidents’ backgrounds are more mixed but still reflect a high reliance on academic debate and literature. The Federal Reserve Board also has a high academic focus score, unsurprisingly given the more than 200 staff economists that support its work.

Communication Clusters. Finally, we examine clustering in the financial stability content of Federal Reserve Bank speeches. We cluster on the content of speeches, which we represent using contextualized sequence embeddings produced using the RoBERTa model. For each paragraph of each speech, this yields a 768-dimensional vector, where each dimension corresponds to a text feature. We then compute the average vector for each district-month.

To prepare the data for visualization, we first perform dimensionality reduction using the t-distributed stochastic neighbor embedding (t-SNE) introduced in van der Maaten and Hinton (2008). The t-SNE approach is a nonlinear dimensionality reduction technique that preserves both local and global structure and is intended to produce representations that can be visualized in 2 or 3 dimensions, which help to differentiate and visualize clusters, but where the distances across clusters are not readily interpretable.

We first apply t-SNE to the full sample, which spans the period between 1960 and 2020, and visualize it in Figure 10. The markers in the visualization represent the different Federal

Reserve Banks, as well as the Board of Governors (FRB). The Federal Reserve Bank of New York (FRBNY – denoted NY in the Figure) appears to have a nearly self-contained cluster in the bottom of the figure, suggesting that the content of their financial stability discussions is closer to their own past discussions than it is to contemporaneous communication at other institutions of the Federal Reserve System.

In contrast to the tight and distinct cluster for the FRBNY, the Federal Reserve Bank of St. Louis's communication forms a long cluster that overlaps with the communication of many other Federal Reserve banks. Other regional banks have tight clusters (e.g., ATL, DAL), similar to the FRBNY, but are considerably closer to the clusters of other districts, indicating that their communication is dominated more by contemporaneous views than by past discussions.

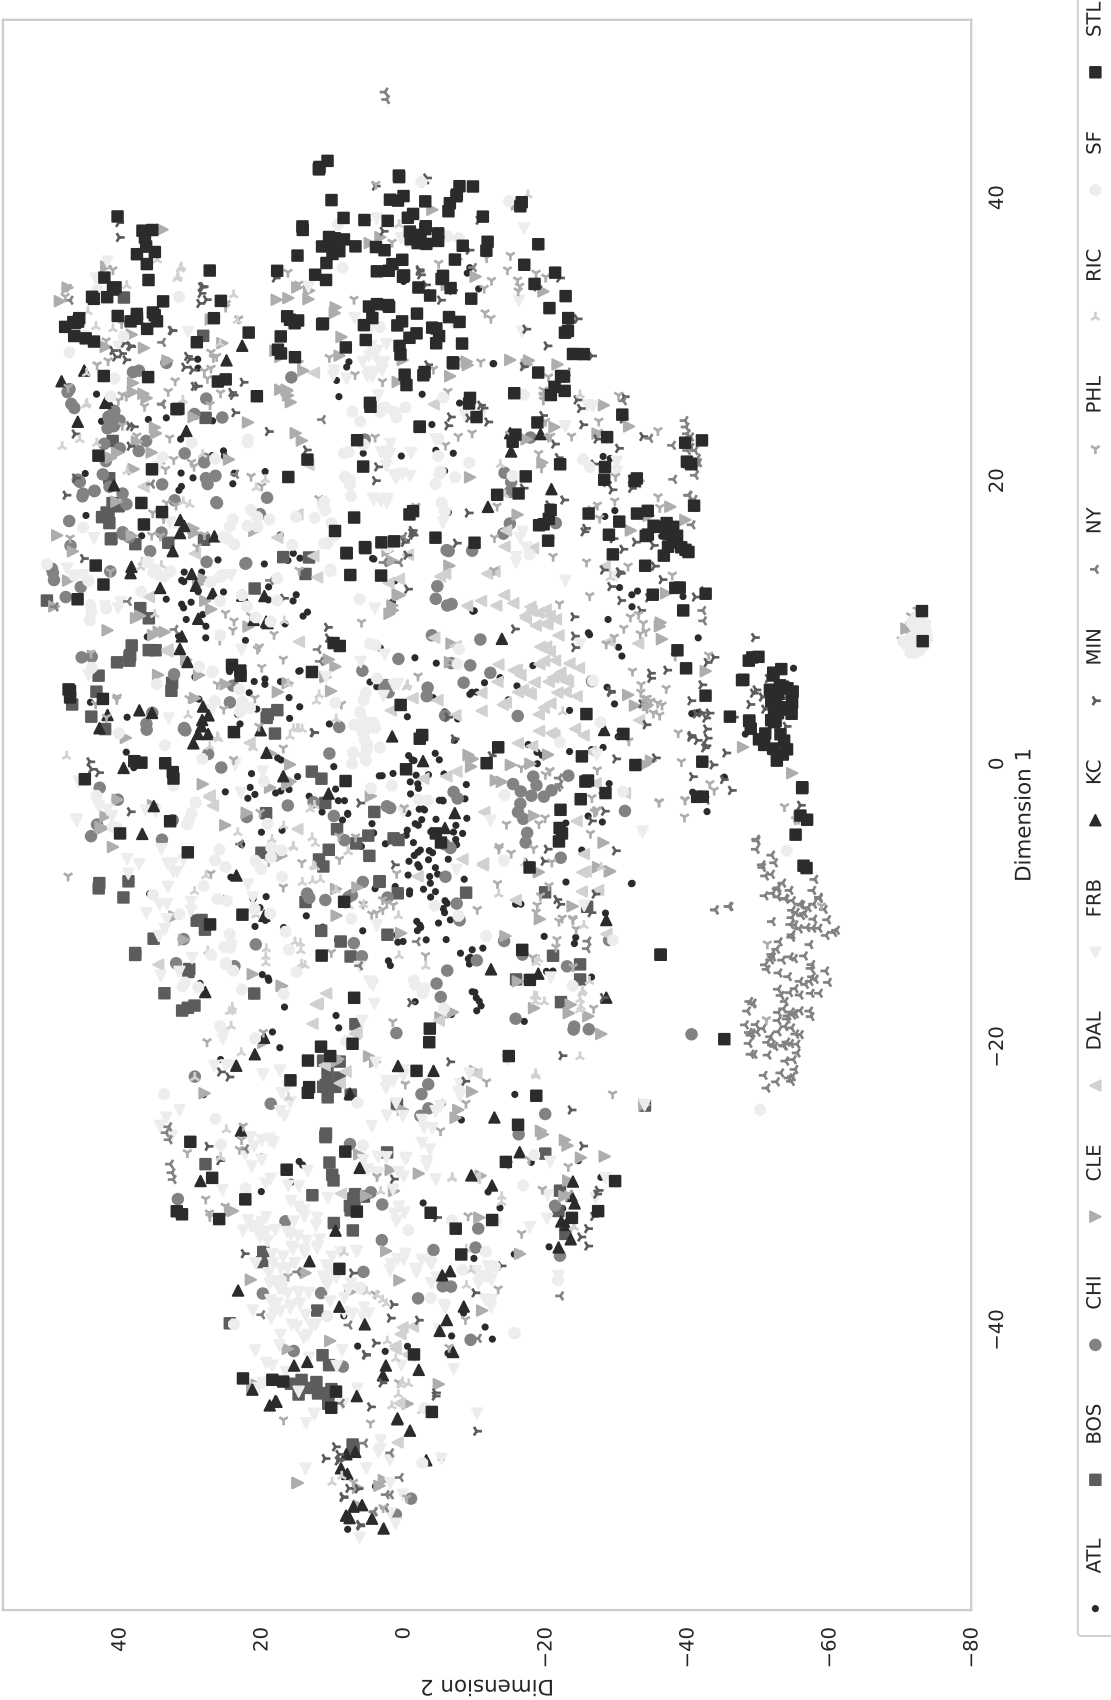


Figure 10: The figure visualizes the output of the t-stochastic nearest neighbors (t-SNE) algorithm applied to sequence embeddings for paragraphs that discuss financial stability. We compute embeddings for all documents in the corpus over the period between 1960 and 2020. The sequence embeddings are produced for each paragraph using a RoBERTa Sentence Transformer model with extended pretraining on abstracts from the S2ORC corpus, as well as fine-tuning on sentence pairs from the S2ORC corpus. Each dot corresponds to a 2-dimensional representation of the average sequence embedding for a given month and institution.

4 Econometric Results

We next move to the empirical tests. In Section 4.1 we analyze the determinants of financial stability concerns. Section 4.2 explores the association between the text features and financial ratios. Next, Section 4.3 analyzes structural breaks in text features, which we relate to important events. Thereafter, we investigate in Section 4.4 whether our semantic variables can shed light on the time variation in the historical conduct of monetary policy by conducting extended Taylor rule regressions. Finally, Section 4.5 links our text features to asset valuation–monetary policy regimes.

4.1 Determinants of Financial Stability Concerns

For our first set of regressions we employ a regression framework, where we concentrate on identifying the states of the economy and how the discussion of topics that are most closely associated with concern about financial stability might vary with these states. Our main regression exercises attempt to determine the underlying drivers and content of financial stability discussions at the Fed.

4.1.1 Empirical Specification

On a conceptual level, financial stability concerns suggest being closely connected to the experience of a financial crisis, as well as to discussions about bank liquidity and bank capital. Per its dual mandate, which was codified in the Federal Reserve Reform Act of 1977 and the Full Employment and Balanced Growth Act of 1978, the key objectives in the conduct of monetary policy are to promote maximum employment and price stability. Consequently, we would expect that concerns about inflation and employment may feature prominently when Fed officials express concerns about financial stability. Hence, we include these text features as explanatory variables in our regression, alongside a battery of macroeconomic and financial controls.

The primary specification is given in Equation (2). We also consider a separate set of specifications that lag each of the macroeconomic and financial variables.

$$y_t = \beta_0 + \beta_1\tau_{jt}^{fc} + \beta_2\tau_{jt}^{\pi} + \beta_3\tau_{jt}^e + \beta_4\tau_t^{bc} + \beta_5\tau_t^{bl} + \beta_6\nu^m + \beta_7\nu^f + \zeta_k + \gamma_t + e_{jt} \quad (2)$$

We use y_t to represent the dependent variable. In the first set of regressions, y_t is a measure of the financial stability content of a particular paragraph. The explanatory variables include various text features indicated by the variable τ . Such features are computed at the paragraph level and, thus, have both paragraph (j) and time (t) variation. The superscripts on τ indicate what variable or group of controls it encompasses. Additionally, we include district fixed effects (ζ_k) and time fixed effects (γ_t) in some specifications. The following text variables are included in regressions: financial crisis (fc), inflation (π), employment (e), bank capital (bc), and bank liquidity (bl).

The macroeconomic and financial controls are indicated by ν and have time variation (t) only. The macro variables (ν^m) include inflation, the output gap, house prices, and the debt-to-gdp ratio. Financial variables (ν^f) include the short term interest rate, a financial crisis indicator, the loan-to-deposit rate, and the natural log of the total loan volume to the non-financial sector. We refer the interested reader to Table 3 in Section 1.3 for details about the macro and financial data definitions and sources.

All regressions use either Newey-West standard errors or cluster at the institution level. Additionally, we provide the results with both contemporaneous regressions and regressions with some or all of the controls lagged. We also show fixed effect specifications, where non-text controls are dropped and both year-month and institution fixed effects are employed. Finally, the baseline regression spans the full sample (1960-2020), but we also include two sample splits: 1960-1983 (pre-Great Moderation) and 1984-2020 (Great Moderation).

4.1.2 Regression results

The empirical results are presented in Tables 4-6. In Table 4, the dependent variable y_t is the standardized financial stability index, computed using zero shot learning. In Table 5, it is the measure of cosine similarity between the statement “monetary policy should be used to achieve financial stability” and the content of a paragraph. Table 6 also uses a cosine similarity measure as the dependent variable – the relationship between “banking regulation should be used to achieve financial stability” and the content of a paragraph.

Financial Crises and Financial Stability. Our first finding, given in Table 4, columns (1) to (4), is that a one standard deviation increase in discussion of financial crises is associated with a 0.09 standard deviation increase in discussion of financial stability. Importantly, the result holds even in the presence of macroeconomic and financial controls (column 2), and even when year-month fixed effects are used (column 4). Thus, irrespective of the state of the economy and financial system, an uptick in interest in financial crises appears to be strongly and statistically significantly associated with discussion of financial stability. This also holds when we lag the macroeconomic controls, as shown in Table 11 in the Appendix.

Dual Mandate and Financial Stability. Another important question about financial stability is whether it is typically discussed in the context of the Fed’s dual mandate or whether it is treated as a separate concern or objective. In the first four columns of Table 4, the coefficients on the inflation and employment text features are positive and strongly statistically significant in the full sample, indicating that financial stability discussion is positively associated with dual mandate related concerns over the full sample. Furthermore, we can see that the association with financial stability is stronger for employment, where the effect size is roughly three times as large as that for inflation.

At first glance this result may seem at odds with a literature suggesting that aggressive inflation-targeting is conducive to financial stability (Bernanke and Gertler, 2001b). How-

Table 4: Federal Reserve Speech Focus: Financial Stability

	(1)	(2)	(3)	(4)	(5)	(6)
inflation _{jt} [text]	0.0434*** (0.0029)	0.0429*** (0.0029)	0.0429*** (0.0101)	0.0434*** (0.0029)	-0.0352*** (0.0058)	0.0629*** (0.0033)
employment _{jt} [text]	0.1207*** (0.0021)	0.1205*** (0.0021)	0.1205*** (0.0065)	0.1206*** (0.0021)	0.1578*** (0.0050)	0.1114*** (0.0023)
financial crisis _{jt} [text]	0.0913*** (0.0029)	0.0890*** (0.0029)	0.0890*** (0.0149)	0.0882*** (0.0029)	0.0670*** (0.0066)	0.0950*** (0.0032)
bank liquidity _{jt} [text]	0.1466*** (0.0040)	0.1429*** (0.0040)	0.1429*** (0.0127)	0.1408*** (0.0040)	0.1773*** (0.0096)	0.1347*** (0.0044)
bank capital _{jt} [text]	0.2957*** (0.0043)	0.2995*** (0.0043)	0.2995*** (0.0152)	0.3008*** (0.0043)	0.3463*** (0.0104)	0.2901*** (0.0047)
past focus _{jt} [text]	-0.1221*** (0.0020)	-0.1178*** (0.0020)	-0.1178*** (0.0080)	-0.1162*** (0.0020)	-0.1058*** (0.0043)	-0.1198*** (0.0022)
present focus _{jt} [text]	0.0723*** (0.0017)	0.0738*** (0.0017)	0.0738*** (0.0042)	0.0744*** (0.0017)	0.0513*** (0.0035)	0.0799*** (0.0019)
future focus _{jt} [text]	0.0915*** (0.0021)	0.0900*** (0.0021)	0.0900*** (0.0072)	0.0899*** (0.0021)	0.0633*** (0.0046)	0.0945*** (0.0023)
academic focus _{jt} [text]	0.1639*** (0.0021)	0.1615*** (0.0021)	0.1615*** (0.0076)	0.1603*** (0.0021)	0.1563*** (0.0053)	0.1608*** (0.0023)
debt-to-gdp ratio _t		0.2435*** (0.0287)	0.2435*** (0.0737)		0.5241** (0.2177)	0.3065*** (0.0375)
loan-to-deposit ratio _t		0.0031*** (0.0006)	0.0031*** (0.0017)		0.0058* (0.0035)	0.0042*** (0.0007)
Sample Period	1960-2020	1960-2020	1960-2020	1960-2020	1960-1983	1984-2020
Macro Controls	NO	YES	YES	NO	YES	YES
Financial Controls	NO	YES	YES	NO	YES	YES
Housing Controls	NO	YES	YES	NO	YES	YES
Month x Year FEs	NO	NO	NO	YES	NO	NO
District FEs	YES	YES	YES	YES	YES	YES
Standard Errors	NW	NW	CL	NW	NW	NW
Adj. R-squared	0.5229	0.5244	0.5244	0.5273	0.5087	0.5287
N	263,613	263,613	263,613	263,613	49,262	214,351

Notes: The dependent variable in all regressions is the zero shot classification score for the term “financial stability.” A higher score indicates that the text in a given sequence of words – such as a paragraph or a sentence – describes a topic related to financial stability. All controls that include [text] indicate that they are text features measured using zero shot classification. For instance “inflation [text]” is the classification score for whether a sequence describes inflation. The features “past focus,” “present focus,” and “future focus” classify the tense of a sequence to determine whether the speaker was discussing the past, present, or future. “Academic focus” classifies whether a speaker was discussing an academic debate or the academic literature. Macro controls include the output gap and consumer price inflation. Financial controls include the log of total loans to the non-financial sector, the short term interest rate, and an indicator for whether a financial crisis occurred in a given year. The financial controls, housing controls, debt-to-gdp ratio, and loan-to-deposit ratio are taken from the *Macrobistory Database*, introduced by Jordà et al. (2016). Note that j indexes paragraph and t indexes time. Standard errors are either Newey-West (NW) or clustered (CL). * $p < .1$, ** $p < .05$, *** $p < .01$. Table 11 in the Appendix presents the results with lagged macroeconomic and financial variables.

Table 5: Federal Reserve Speech Advocacy for the Use of Monetary Policy to Achieve Financial Stability

	(1)	(2)	(3)	(4)	(5)
inflation _{jt} [text]	0.2086*** (0.0024)	0.2086*** (0.0134)	0.1941*** (0.0024)	0.2027*** (0.0055)	0.2077*** (0.0026)
employment _{jt} [text]	-0.2035*** (0.0027)	-0.2035*** (0.0123)	-0.1991*** (0.0026)	-0.1844*** (0.0060)	-0.2111*** (0.0029)
financial crisis _{jt} [text]	0.1139*** (0.0031)	0.1139*** (0.0098)	0.1154*** (0.0030)	0.1481*** (0.0073)	0.1110*** (0.0035)
bank liquidity _{jt} [text]	-0.1482*** (0.0045)	-0.1482*** (0.0167)	-0.1386*** (0.0044)	-0.1773*** (0.0113)	-0.1344*** (0.0049)
bank capital _{jt} [text]	0.0492*** (0.0044)	0.0492*** (0.0175)	0.0478*** (0.0043)	0.0798*** (0.0114)	0.0380*** (0.0047)
past focus _{jt} [text]	0.0368*** (0.0026)	0.0368*** (0.0083)	0.0317*** (0.0025)	0.0113** (0.0055)	0.0396*** (0.0029)
present focus _{jt} [text]	-0.0014 (0.0023)	-0.0014 (0.0093)	-0.0009 (0.0022)	0.0357*** (0.0048)	-0.0134*** (0.0025)
future focus _{jt} [text]	0.0629*** (0.0026)	0.0629*** (0.0086)	0.0606*** (0.0025)	0.0205*** (0.0058)	0.0717*** (0.0028)
academic focus _{jt} [text]	-0.0665*** (0.0027)	-0.0665*** (0.0117)	-0.0609*** (0.0026)	-0.0920*** (0.0069)	-0.0573*** (0.0029)
debt-to-gdp ratio _t	0.2605*** (0.0554)	0.2605 (0.4653)	0.2605 (0.4653)	0.4878 (0.4147)	0.3316*** (0.0745)
loan-to-deposit ratio _t	-0.0073*** (0.0011)	-0.0073 (0.0081)		0.0174** (0.0068)	-0.0130*** (0.0014)
Sample Period	1960-2020	1960-2020	1960-2020	1960-1983	1984-2020
Macro Controls	YES	YES	NO	YES	YES
Financial Controls	YES	YES	NO	YES	YES
Housing Controls	YES	YES	NO	YES	YES
Month x Year FEs	NO	NO	YES	NO	NO
District FEs	YES	YES	YES	YES	YES
Standard Errors	NW	CL	NW	NW	NW
Adj. R-squared	0.0962	0.0962	0.1499	0.1400	0.0990
N	263,613	263,613	263,613	49,262	214,351

Notes: The dependent variable in all regressions is the cosine similarity between the contextualized embedding for a given sequence of words and the contextualized embedding that corresponds to the sequence “monetary policy should be used to achieve financial stability.” A higher score indicates that a sequence is more likely to be advocating for the use of monetary policy to achieve financial stability. All controls that include [text] indicate that they are text features measured using zero shot classification. For instance “inflation [text]” is the classification score for whether a sequence describes inflation. The features “past focus,” “present focus,” and “future focus” classify the tense of a sequence to determine whether the speaker was discussing the past, present, or future. “Academic focus” classifies whether a speaker was discussing an academic debate or the academic literature. Macro controls include the output gap and consumer price inflation. Financial controls include the log of total loans to the non-financial sector, the short term interest rate, and an indicator for whether a financial crisis occurred in a given year. The financial controls, housing controls, debt-to-gdp ratio, and loan-to-deposit ratio are taken from the *Macroeconomic History Database*, introduced by Jordà et al. (2016). Note that j indexes paragraph and t indexes time. Standard errors are either Newey-West (NW) or clustered (CL) at the district level. * $p < .1$, ** $p < .05$, *** $p < .01$. Table 12 in the Appendix presents the results with lagged macroeconomic and financial variables.

Table 6: Federal Reserve Speech Advocacy for the Use of Banking Regulation to Achieve Financial Stability

	(1)	(2)	(3)	(4)	(5)
inflation _{jt} [text]	-0.2994*** (0.0023)	-0.2994*** (0.0260)	-0.2891*** (0.0023)	-0.1905*** (0.0045)	-0.3200*** (0.0027)
employment _{jt} [text]	-0.2020*** (0.0023)	-0.2020*** (0.0127)	-0.1884*** (0.0023)	-0.1557*** (0.0047)	-0.2110*** (0.0026)
financial crisis _{jt} [text]	0.2212*** (0.0030)	0.2212*** (0.0192)	0.2130*** (0.0030)	0.1106*** (0.0060)	0.2437*** (0.0034)
bank liquidity _{jt} [text]	-0.1550*** (0.0053)	-0.1550*** (0.0222)	-0.1469*** (0.0051)	-0.1670*** (0.0115)	-0.1455*** (0.0058)
bank capital _{jt} [text]	0.3765*** (0.0053)	0.3765*** (0.0216)	0.3577*** (0.0051)	0.2945*** (0.0115)	0.3842*** (0.0057)
past focus _{jt} [text]	-0.0013 (0.0024)	-0.0013 (0.0060)	-0.0038 (0.0023)	-0.0170*** (0.0047)	-0.0023 (0.0027)
present focus _{jt} [text]	0.0743*** (0.0021)	0.0743*** (0.0050)	0.0678*** (0.0020)	0.0824*** (0.0040)	0.0712*** (0.0024)
future focus _{jt} [text]	-0.0334*** (0.0024)	-0.0334*** (0.0050)	-0.0291*** (0.0023)	-0.0286*** (0.0048)	-0.0334*** (0.0027)
academic focus _{jt} [text]	-0.0252*** (0.0026)	-0.0252 (0.0168)	-0.0229*** (0.0025)	-0.0184*** (0.0056)	-0.0235*** (0.0028)
debt-to-gdp ratio _t	0.2173*** (0.0492)	0.2173 (0.2851)		2.0998*** (0.3215)	0.4425*** (0.0654)
loan-to-deposit ratio _t	0.0052*** (0.0010)	0.0052 (0.0087)		-0.0044 (0.0050)	-0.0017 (0.0013)
Sample Period	1960-2020	1960-2020	1960-2020	1960-1983	1984-2020
Macro Controls	YES	YES	NO	YES	YES
Financial Controls	YES	YES	NO	YES	YES
Housing Controls	YES	YES	NO	YES	YES
Month x Year FEs	NO	NO	YES	NO	NO
District FEs	YES	YES	YES	YES	YES
Standard Errors	NW	CL	NW	NW	NW
Adj. R-squared	0.1894	0.1894	0.2243	0.1378	0.2031
N	263,613	263,613	263,613	49,262	214,351

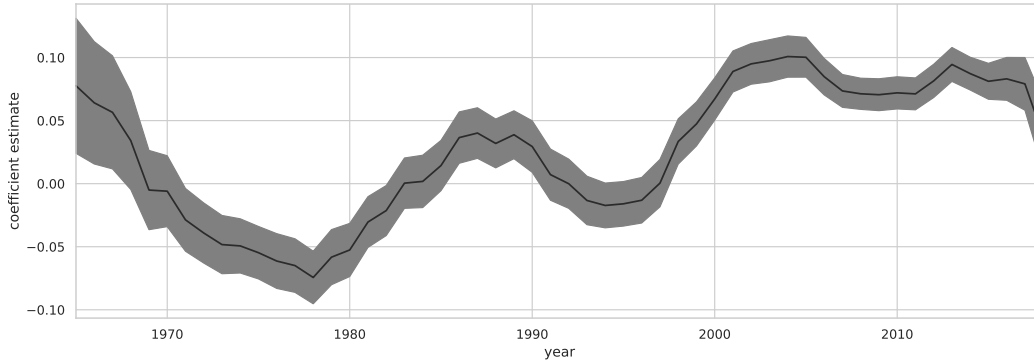
Notes: The dependent variable in all regressions is the cosine similarity between the contextualized embedding for a given sequence of words and the contextualized embedding that corresponds to the sequence “bank regulation should be used to achieve financial stability.” A higher score indicates that a sequence is more likely to be advocating for the use of monetary policy to achieve financial stability. All controls that include [text] indicate that they are text features measured using zero shot classification. For instance “inflation [text]” is the classification score for whether a sequence describes inflation. The features “past focus,” “present focus,” and “future focus” classify the tense of a sequence to determine whether the speaker was discussing the past, present, or future. “Academic focus” classifies whether a speaker was discussing an academic debate or the academic literature. Macro controls include the output gap and consumer price inflation. Financial controls include the log of total loans to the non-financial sector, the short term interest rate, and an indicator for whether a financial crisis occurred in a given year. The financial controls, housing controls, debt-to-gdp ratio, and loan-to-deposit ratio are taken from the *Macroeconomic Database*, introduced by Jordà et al. (2016). Note that j indexes paragraph and t indexes time. Standard errors are either Newey-West (NW) or clustered (CL) at the district level. * $p < .1$, ** $p < .05$, *** $p < .01$. Table 13 in the Appendix presents the results with lagged macroeconomic and financial variables.

ever, when the sample is split into two sub-periods (columns (5) and (6)), the relationship between inflation and financial stability discussion appears to be negative prior to the Great Moderation, when there were episodes of high inflation and output growth was less stable. In addition, in the first subsample (1960-1983), the employment effect is about 25% stronger than over the full sample. In contrast, during the Great Moderation period, the relationship between inflation and financial stability becomes strongly positive. In this sub-period inflation targeting also became more prominent at the Federal Reserve, suggesting that, in line with the literature, stabilizing prices and the financial sector are not treated as separate concerns or as being in conflict. In addition, during the Great Moderation, the effect size for employment declines while that for inflation increases, so that during this sub-period the employment effect is actually only about twice that for inflation.

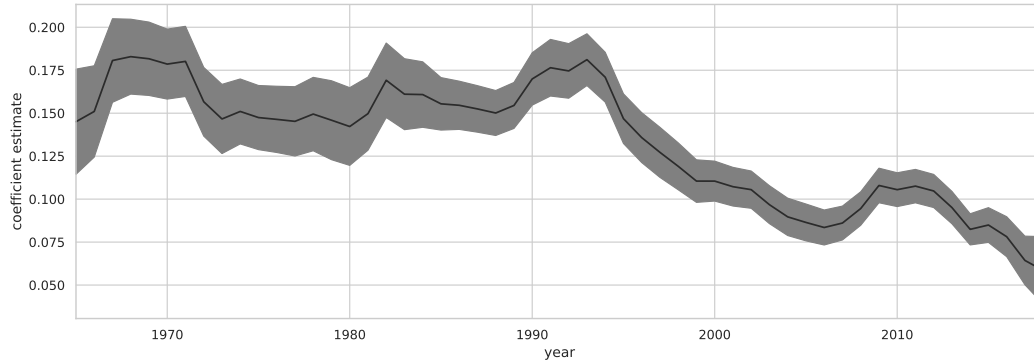
These findings can also be seen in the first two panels of Figure 11, which show the coefficient estimates corresponding to the inflation and employment text features from a series of five-year rolling regressions of financial stability discussion on the full set of text, macroeconomic, and financial controls. Consistent with the sample split findings, the Great Moderation appears to have changed how the dual mandate is discussed in conjunction with financial stability. Whereas financial stability and inflation are more likely to be discussed jointly; financial stability and employment are less likely to be discussed together in the same paragraph.

Academic Focus. Next, we find that a focus on academic debates and the academic literature is strongly associated with increased discussion of financial stability in Table 4. In all specifications, a one standard deviation increase in academic focus is associated with about a 0.16 standard deviation increase in financial stability content, an effect magnitude that exceeds those for either of the aspects of the Fed's dual mandate. This suggests that, among the topics that the Federal Reserve discusses, the discussion of financial stability appears to have an above average concentration on the academic discussion, possibly reflecting the

Rolling Regression: Inflation Text Feature Coefficient



Rolling Regression: Employment Text Feature Coefficient



Rolling Regression: Academic Focus Text Feature Coefficient

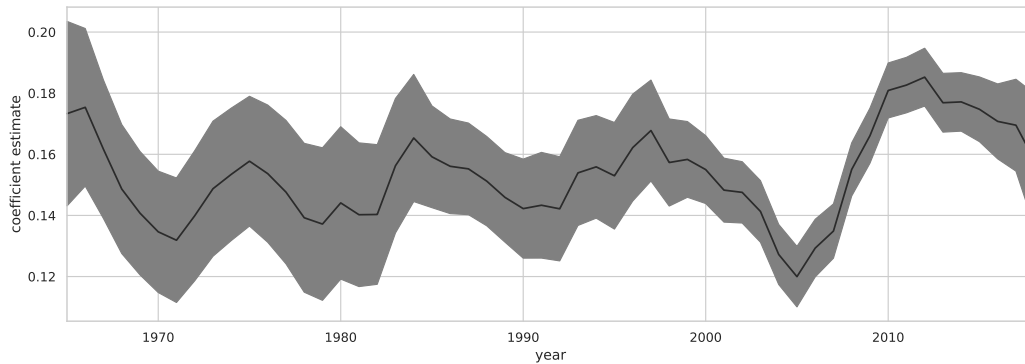


Figure 11: The plots above show yearly coefficient estimates from rolling regressions of the financial stability text feature on the inflation, employment, financial crisis, bank liquidity, bank capital, tense focus, and academic focus text features, as well as month-year and district fixed effects. Along with the point estimates, we show 95% confidence intervals. We use a 5-year window for the rolling regressions.

need for a conceptual framework.

Table 5 considers the extent to which various text features are discussed in the context of advocating that monetary policy be used to achieve financial stability. The significant negative coefficient on the academic focus variable indicates that Fed speeches that reference the academic literature tend to oppose the use of monetary policy to achieve financial stability (Vollmer, 2021). Table 6 presents similar results in the context of advocating that banking regulation be used to achieve financial stability. Fed speeches that reference the academic literature also tend to oppose the use of banking regulation for this purpose, although the effect is about one-third the magnitude of the effect size for the use of monetary policy.

Comparing the two subsamples across all three tables (columns (5) and (6)) suggests that, although there is little change in the association between academic focus and financial stability overall, the use of academic literature to voice opposition to the use of monetary policy to achieve financial stability appears to have declined during the Great Moderation (relative to the period before the Great Moderation) while the use of it to oppose the use of banking regulation for this purpose has increased. A closer look of how this relationship evolved over time reveals that the influence of academic content over the financial stability discussion increased during the Great Moderation, but declined during Greenspan's tenure, rising again during the Bernanke era. This is illustrated in the third panel of Figure 11, which plots the coefficient estimates corresponding to the academic text feature from a five-year rolling regression.

Speech Tense. Table 4 suggests that statements about financial stability tend to be framed in terms of the present and future, perhaps suggesting that they are in response to ongoing or anticipated events. This tendency increased during the Great Moderation. Advocacy for the use of monetary policy to achieve financial stability typically hinges on the use of past examples and future hypotheticals, as indicated by Table 5; this advocacy more than tripled during the Great Moderation. In contrast, advocacy for the use of bank

regulation is typically focused on events in the present, as indicated by Table 6. Moreover, there is slightly less use of speech tense for such advocacy in the second subsample compared to period before the Great Moderation.

Financial Variables and Discussion of Financial Stability. Turning to financial variables, we find in Table 4 the expected result that an increase in the debt-to-gdp ratio or an increase in total lending appears to be associated with a considerable increase in discussion of financial stability; however, this increased discussion does not necessarily lead to advocacy for a particular approach to policy. Tables 5 and 6 show that a higher debt-to-gdp ratio tends to be associated with a increase in the advocacy of monetary policy and banking regulation as a means of achieving financial stability. The results for the loan-to-deposit ratio, however, are more mixed.

4.2 Financial Markets

We next demonstrate that the concerns articulated in Federal Reserve speeches matter for financial ratios and returns. We do this by estimating a modified version of Equation (2), where the dependent variable y_t is a measure of annual asset returns and the text feature financial stability is introduced as an additional explanatory variable (τ_{jt}^{fs}). Table 7 reports the association between the variables of interest and returns on equity, bonds, risky assets, and safe assets. For sources and return definitions, see Table 3 in Section 1.3.

From a conceptual viewpoint we would expect that the concerns articulated in Federal Reserve speeches matter not only for asset returns, but also that there may be a differential effect when comparing riskier and safer asset classes. This is because the impact of Fed communication is likely to occur via the risk-taking channel of monetary policy risk perceptions and risk attitudes of market participants, as well as through expectations about monetary conditions, which affect the riskiness of bank lending, valuations and risk measures (Jiménez et al., 2014; Dell’Ariccia et al., 2017).

Table 7: Federal Reserve Speech Impact on Asset Returns

	(1)	(2)	(3)	(4)	(5)
Return type	Equity	Equity	Bond	Risky	Safe
inflation _{<i>jt</i>} [text]	0.0012** (0.0005)	0.0012** (0.0005)	0.0015** (0.0004)	0.0004 (0.0002)	0.0008*** (0.0002)
employment _{<i>jt</i>} [text]	0.0029*** (0.0004)	0.0029*** (0.0004)	0.0015*** (0.0004)	0.0021*** (0.0002)	0.0007*** (0.0002)
financial stability _{<i>jt</i>} [text]	-0.0008* (0.0004)	-0.0009** (0.0004)	0.0001 (0.0004)	-0.0004** (0.0002)	0.0000 (0.0002)
financial crisis _{<i>jt</i>} [text]	-0.0019*** (0.0005)	-0.0019*** (0.0005)	-0.0053*** (0.0004)	-0.0016*** (0.0002)	-0.0026*** (0.0002)
bank liquidity _{<i>jt</i>} [text]	-0.0033*** (0.0007)	-0.0036*** (0.0007)	-0.0010* (0.0006)	-0.0022*** (0.0003)	-0.0005 (0.0003)
bank capital _{<i>jt</i>} [text]	0.0022*** (0.0007)	0.0021*** (0.0007)	0.0005 (0.0006)	0.0012*** (0.0003)	0.0002 (0.0003)
past focus _{<i>jt</i>} [text]	0.0001 (0.0004)	-0.0000 (0.0004)	0.0008** (0.0003)	0.0002 (0.0002)	0.0004** (0.0002)
present focus _{<i>jt</i>} [text]	-0.0004 (0.0004)	-0.0002 (0.0004)	0.0008*** (0.0003)	-0.0003** (0.0002)	0.0004*** (0.0001)
future focus _{<i>jt</i>} [text]	0.0023*** (0.0004)	0.0024*** (0.0004)	-0.0007** (0.0003)	0.0011*** (0.0002)	-0.0004** (0.0002)
past focus _{<i>jt</i>} * financial stability _{<i>jt</i>} [text]		-0.0003 (0.0003)			
present focus _{<i>jt</i>} * financial stability _{<i>jt</i>} [text]		0.0003 (0.0004)			
future focus _{<i>jt</i>} * financial stability _{<i>jt</i>} [text]		0.0008** (0.0004)			
academic focus _{<i>jt</i>} [text]	-0.0023*** (0.0004)	-0.0023*** (0.0004)	-0.0006 (0.0004)	-0.0006*** (0.0002)	-0.0003 (0.0002)
debt-to-gdp ratio _{<i>t</i>}	-0.0244** (0.0111)	-0.0244** (0.0111)	0.1715*** (0.0101)	-0.0442*** (0.0048)	0.0829*** (0.0051)
loan-to-deposit ratio _{<i>t</i>}	-0.0136*** (0.0002)	-0.0136*** (0.0002)	0.0077*** (0.0002)	-0.0059*** (0.0001)	0.0039*** (0.0001)
Sample Period	1960-2020	1960-2020	1960-2020	1960-2020	1960-2020
Macro Controls	YES	YES	YES	YES	YES
Financial Controls	YES	YES	YES	YES	YES
Housing Controls	YES	YES	YES	YES	YES
District FEs	YES	YES	YES	YES	YES
Standard Errors	NW	NW	NW	NW	NW
Adj. R-squared	0.3032	0.3032	0.1120	0.2665	0.2040
N	263,613	263,613	263,613	263,613	263,613

Notes: The dependent variable in each regression is a measure of annual asset returns. The returns are taken from Jordà et al. (2016) and include total equity, total bond, risky assets, and safe assets, as specified in the “return type” row of the table. All controls that include [text] indicate that they are text features measured using zero shot classification. For instance “inflation [text]” is the classification score for whether a sequence describes inflation. The features “past focus,” “present focus,” and “future focus” classify the tense of a sequence to determine whether the speaker was discussing the past, present, or future. “Academic focus” classifies whether a speaker was discussing an academic debate or the academic literature. Macro controls include the output gap and consumer price inflation. Financial controls include the log of total loans to the non-financial sector, the short term interest rate, and an indicator for whether a financial crisis occurred in a given year. The financial controls, housing controls, debt-to-gdp ratio, and loan-to-deposit ratio are taken from the *Macroeconomic History Database*, introduced by Jordà et al. (2016). Note that *j* indexes paragraph and *t* indexes time. Standard errors are either Newey-West (NW) or clustered (CL) at the district level. * $p < .1$, ** $p < .05$, *** $p < .01$. Table 14 in the Appendix presents results with lagged macroeconomic and financial controls.

We find a positive association between discussion of inflation or employment, and the returns to equity, bonds, and safe assets. This holds even in the presence of macroeconomic and financial controls, and even when text factors and controls are lagged by one year. Financial stability discussion appears to have a negative association with equity and risky asset returns, but not bond and safe asset returns. In contrast, financial crisis discussion has a negative association for all asset return categories. The same holds for bank liquidity, but the estimated coefficient is insignificant for safe asset returns. Bank capital discussion has the opposite effect; a positive association with equity and risky asset returns with insignificant estimated coefficients for bonds and safe asset returns.

It is also interesting that speech tense appears to impact asset returns. Notably, past or present focus is positively associated with bond and safe asset returns. A present focus negatively affects risky asset returns. Equities and risky assets appear to respond positively to a future (forward-looking) focus; bonds and safe assets negatively. Finally, an academic focus is also negatively associated with equity and risky asset returns.

4.3 Structural Breaks in Text Features

The subsample results in Tables 4-6 suggest changes over time in some of the associations documented above. In this subsection, we conduct formal structural break tests using the multiple break test of Bai (1997) and Bai and Perron (1998).¹⁵ The results are presented in Table 8.

The identified breaks coincide with important developments. The earliest identified breaks, in 1969Q2 (cosine similarity between employment and inflation) precedes the recession that took place from December 1969-November 1970. This period coincided with Fed tightening; hence it is not surprising that the cosine similarity between financial stability and inflation experiences a structural break in 1969Q4, the quarter that the Fed funds

¹⁵The tests are performed separately on each univariate text feature index series using the sequential test option in Eviews with 15% trimming, the HAC option, a degrees of freedom adjustment, and where the error distribution can differ across the break subsamples. The maximum number of allowable breaks is set to five.

Structural Breaks for Text Feature Series

Variable	Break dates
Dual Mandate	1996 Q3 and 2010 Q4
Monetary Policy	1996 Q3 and 2005 Q4
Non-Dual Mandate	1996 Q3 and 2010 Q4
Cosine Similarity Employment and Inflation	1969 Q2
Financial Stability	1995 Q3
Financial Crisis	1970 Q2
Cosine Similarity Monetary Policy and Financial Stability	None
Cosine Similarity Banking Regulation and Financial Stability	1980 Q1, 1992 Q1
Cosine Similarity Financial Stability and Inflation	1969 Q4 and 2010 Q4
Academic focus	1996 Q3 and 2010 Q1
Speech Tense: Past Focus	2012 Q3
Speech Tense: Present Focus	1971 Q3 and 2012 Q3
Speech Tense: Future Focus	1971 Q3 and 2008 Q3

Table 8: All text feature series are quarterly averages of the standardized classification scores from BERT for each paragraph in our text corpus. The tests for structural change are estimated over the full sample (1960-2020) with 15% trimming. The break dates listed represent the first date of the new regime.

rate peaked at 9%, followed by a break in the financial crisis index in 1970Q2.

There are two structural breaks in the cosine similarity between banking regulation and financial stability, in 1980Q1 and 1992Q1. The first of these quarters saw inflation peak at 14.8% and in March the Depository Institutions Deregulation and Monetary Control Act of 1980, giving the Fed more oversight of non-member banks, was passed. In addition, 1992Q1 is the first quarter the FDIC Improvement Act (commonly known as FDICIA) was in effect, having been signed into law in December 1991. This law brought prompt corrective action (PCA) into supervision and was a response to the wave of commercial bank failures that had occurred over the previous decade, the savings and loan crisis, and the resulting strain

on the FDIC’s coffers.¹⁶

As a result of the Great Moderation and a prolonged period of low inflation, there was little discussion of the price stability side of the dual mandate throughout the 1980s. That changed in the mid-1990s, and the financial stability feature has a single break, in 1995Q3, coinciding with the first quarter where the “dual mandate emerges as common parlance”. Prior to that quarter, the phrase “dual objective” was used. In addition, four features have the first of two breaks in 1996Q3: dual mandate, non-dual mandate, monetary policy, and academic focus. During this time (late 1995 and 1996), the Fed was having extensive discussions about changing Fed policy to specify an explicit inflation target, with Vice-Chair Alan Blinder arguing in opposition. In addition to the first structural break in 1996Q3, the academic focus has a second structural break in 2010Q1, the quarter that Ben Bernanke was renewed to a second term as Fed chairman. It is also interesting that the second structural break for both the dual mandate and non-dual mandate series, as well as the cosine similarity between financial stability and inflation, is in 2010Q4. This was the first quarter following the passage of the Dodd-Frank Act, which codified financial stability as a statutory concern, as well as the European Stability Mechanism and the Fed’s second round of quantitative easing (QE2). In addition, according to the Federal Reserve’s own history, the September 21, 2010 meeting was the first time a reference to the employment side of the mandate appeared explicitly in the FOMC statement.¹⁷

The structural breaks identified in the speech tense features also link to some very important episodes. The first break date in the present and future focus features, 1971Q3, is the quarter that President Nixon and a number of his advisers, including then-Fed Chairman Arthur Burns and future-Fed Chairman Paul Volcker, crafted the policy that ended dollar convertibility to gold implemented wage and price controls, marking “the beginning

¹⁶While FDICIA applied to the bank supervision practices of the Federal Deposit Insurance Corporation (FDIC), it had implications for the Federal Reserve’s oversight as well, as the Fed has supervisory responsibility for all bank holding companies in the US.

¹⁷See <https://www.federalreservehistory.org/essays/humphrey-hawkins-act> for documentation and further discussion related to the dates referred to in this paragraph

of the end of the Bretton Woods international monetary system and temporarily halt[ing] inflation”.¹⁸ The second break in the future focus feature, 2008Q3, was a critical period as the global financial crisis was unfolding. The last month of this quarter saw the government sponsored enterprises, Fannie Mae and Freddie Mac, go into conservatorship, the collapse of Lehman Brothers and Washington Mutual, and the near-failure of insurer AIG, as well as the September 29th rejection by the House of Representatives of the Emergency Economic Stabilization Act (it passed in October).

The date of the single break in the past focus and the second break in the present focus features, 2012Q3, includes the FOMC meeting where the Fed’s third round of quantitative easing was announced, as well as a prolonging of the expected period that the fed funds rate would remain at the zero lower bound. It might be related to the emergence of non-conventional monetary policy.

4.4 Taylor Rule Regressions

Based on the previous results, we next attempt to shed light on the time variation in the historical conduct of monetary policy through the lens of the Taylor framework. Therefore, we build on the seminal work by Taylor (1993, 1999) and conduct extended Taylor rule regressions to estimate policy rules based on both macroeconomic time series data and semantic variables. Existing work has documented low frequency time variation in the Fed’s monetary policy rule (Clarida et al., 2000; Orphanides, 2003; Boivin, 2006; Hamilton et al., 2011; Bianchi et al., 2022; Bauer et al., 2022). Our exercises aim to uncover key drivers based on the Federal Reserve’s own speeches that can rationalize changes in its monetary policy rule. Specifically, we investigate whether non-dual mandate related concerns, as articulated in speeches of Federal Reserve officials, attribute to changes in the policy rule.

¹⁸<https://www.federalreservehistory.org/essays/gold-convertibility-ends>

4.4.1 Specification of the Monetary Policy Rule

In a recent paper, Carvalho et al. (2021) argue in favor of using ordinary least squares (OLS) to estimate the Taylor rule and document that the endogeneity bias in OLS estimates is small. We also use OLS and first replicate the estimation results in Table 2 of Carvalho et al. for the period before 2008 using the same data. We then extend the analysis to the period from 2008 to 2020 with more recent data.

We use the same interest rate rule as in Clarida et al. (2000), which allows for interest rate smoothing. However, we also estimate versions of an extended Taylor rule, which includes text features from Federal Reserve speeches.¹⁹

$$r_t = \theta_0 + \theta_{1,1}r_{t-1} + \theta_{1,2}r_{t-2} + \theta_2 E[\pi_{t+1}] + \theta_3 E[x_{t+1}] + \theta_{4,i}\tau_t^i + e_t. \quad (3)$$

The dependent variable r_t is the interest rate. For the subsample prior to 2008 we use the federal funds rate for the interest rate and for the subsample from 2008 to 2020, when monetary policy was constrained by the zero lower bound, we either use the federal funds rate or the end-of-quarter shadow rate from Wu and Xia (2016). As explanatory variables we use lagged values of the interest rate, as well as lagged values of the Greenbook forecast for inflation, $E[\pi_{t+1}]$, and for the output gap, $E[x_{t+1}]$. We use τ_t^i for variables with text features and the superscripts indicate the specific features used: $i = nd$ refers to non dual mandate, $i = fs$ financial stability, and $i = mf$ refers to financial stability should be achieved using monetary policy. The choice of including a policy inertia term can be justified by both empirical evidence and theoretical studies (Coibion and Gorodnichenko, 2012). Importantly, the inclusion of interest rate smoothing increases the explanatory power due to an excessive volatility in interest rates predicted by the standard Taylor rule.

¹⁹See Peek et al. (2016) for a study of the ternary mandate using an extended Taylor rule and dictionary-based methods to measure financial stability concerns from FOMC meeting transcripts. In a related paper Istrefi et al. (2021) show that a more negative tone in Federal Reserve speeches on financial stability topics is associated with a more accommodative policy stance.

4.4.2 Extended Taylor Rule Regressions

Following Carvalho et al., we define $\rho \equiv \theta_{1,1} + \theta_{1,2}$, $\beta \equiv \theta_2/(1 - \rho)$, $\gamma \equiv \theta_3/(1 - \rho)$ and $\pi^* \equiv (\theta_0 - (1 - \rho)rr^*)/((1 - \rho)(1 - \beta))$, where rr^* is the equilibrium real interest rate and π^* is the inflation target.

Columns (1)-(3) of Table 9 are a replication of the results in Table 2 of Carvalho et al., which entails the estimation of equation (3) without the inclusion of text features for the subsamples: Volcker-Greenspan, (1960Q1-1979Q2), Greenspan-Bernanke (1987Q3-2007Q4), and Post-Volcker (1986Q1-2007Q4). We report the estimates for the coefficients β on the inflation gap and γ on the output gap, as well as ρ for the interest rate smoothing and π^* for the inflation target. The coefficient for the inflation gap is highly significant and above one for all three subsamples during the period from 1960 to 2007, meaning that the “Taylor principle” of raising the nominal interest rate by more than one-for-one is satisfied, which is an important condition for the existence of a stable inflation rate in macro models. The coefficient for the output gap is also highly significant and close to one.

Column (4) repeats the same exercise for the Post-Crisis subsample (2008Q1-2020Q4) using first the federal funds rate as dependent variable. Due to the zero lower bound episode there is little variation in the dependent variable and we replace it in column (5) with the Wu and Xia shadow federal funds rate that also captures the effects of quantitative easing. In fact, we find that the estimated coefficient on the output gap and the smoothing parameter in column (5) are more in line with the Greenspan-Bernanke subsample, while the estimated coefficient on the inflation gap turns insignificant. The finding of an insignificant coefficient for the inflation gap for recent years has been established in earlier studies (see, e.g., Bauer et al. (2022)) and may be attributed to the low variation in the Greenbook inflation forecast.

Next, we estimate equation (3), including the text features in the explanatory variable τ_t^i . Column (6) reports the estimated coefficients for the non dual mandate text feature variable τ_t^{nd} in the Post-Crisis subsample, which is defined analogously as $\delta^{nd} \equiv \theta^{nd}/(1 - \rho)$. We find that more non dual mandate related talk is negatively associated with the shadow

federal funds rate, that is non dual mandate talk enters the expanded policy rule in an accommodative fashion for the Post-Crisis subsample. In unreported regressions for the three Pre-Crisis subsamples, the coefficient estimates for δ^{nd} are insignificant.²⁰ This indicates that the text feature is more important for the Post-Crisis subsample, when the policy discussion was shaped by the GFC experience.

To explore the role played by financial stability discussions, we estimate equation (3) for the Post-Crisis subsample, including the variables τ_t^{fs} and τ_t^{mf} . The results are reported in columns (7) and (8). Interestingly, the financial stability text feature is associated with more accomodative policy decisions, while the monetary financial text feature, which captures whether monetary policy should be used to achieve financial stability, is associated with less accomodative policy decisions. This suggests that policy makers shift to a more accommodative monetary policy to address financial stability problems. At the same time, policy makers are mindful about the use of monetary policy to lean against the wind when financial risks are building up. When comparing the results in columns (6) and (8) we find that the sum of the estimated coefficients for the financial stability and monetary financial text features almost coincides with the estimated coefficient for the non dual mandate text feature. Moreover, using the lagged text features in column (9) suggests that the shift to a more accomodative monetary policy is associated with contemporaneous financial stability problems, while the leaning against the wind rationale is of a longer-term nature.

4.5 Monetary Policy Regimes

Taken together, our semantic variables can shed light on the time variation in the historical conduct of monetary policy and help us to identify financial stability-related concerns as an important determinant. To corroborate our findings, we link our text features with the asset valuation–monetary policy regimes shown in Bianchi et al. (2022). Specifically, we run the

²⁰These regression results are available from the authors on request.

regression with the specification

$$p_t = \theta_0 + \theta_i \tau_t^i + e_t \quad (4)$$

on the sample 1961Q1–2016Q1 of Bianchi et al. (2022), where p_t is the probability of the Fed being in the hawkish monetary policy regime, and τ_t^i contains relevant text features generated from the model. We make use of the text temporal focus, and include a new definition of current financial crisis concern, as the product of financial crisis related content score and a present temporal focus.

Table 10 presents the empirical results. We find that the hawkish monetary policy regime is significantly associated with more discussion about using monetary policy to achieve the financial stability goal, and non dual mandate related text. On the other hand, a higher fraction of discussions on current financial crisis, or discussions about monetary policy and inflation-employment themes, is negatively correlated with the hawkish monetary policy stance. Overall, the results are consistent with our extended Taylor rule regression results, and the monetary policy regimes found in the literature.

5 Conclusion

This paper considers the central bank mandates and how much of the monetary policy stances are driven by the mandates, gleaned through Federal Reserve officials’ own words as articulated in their speeches. We assemble the most comprehensive collection of speeches to-date and apply state-of-the-art natural language processing methods to extract the most significant concerns from each speech. By measuring the similarity between passages from different speeches, we are able to identify semantic similarities between different speeches given by speakers from Federal Reserve districts at various points in time, and of Presidents of the same district over time.

After providing an overview of the text features and content contained in the speeches, we

then use a regression framework to verify and interpret the results uncovered by our natural language processing (NLP) model. We find the states of the economy and discussion topics that are most closely associated with concerns about financial stability. The results indicate strong associations between dual mandate-related concerns and discussion of financial stability, as well as variation over the financial cycle in the use of speech tense and academic focus. The identified semantic variables significantly predict returns of broad asset classes, even after controlling for macroeconomic and financial controls, confirming the information content of officials' speeches.

We further demonstrate that the semantic variables summarized from the Federal Reserve speeches can interpret the time varying reaction function of monetary policy. We show the key text features based on the Federal Reserve's own speeches that can rationalize changes in its monetary policy rule in the historical sample. These findings are supported by the structural breaks in the textual features identified by the Bai and Perron (1998) multiple structural break test. The identified structural breaks coincide with notable changes either in the Fed communication strategy, or in the monetary policy regimes. Furthermore, we relate the regime switches between hawkish and dovish monetary policy stances by the Federal Reserve to the semantic variables generated by the NLP model, on dual-mandate and non-dual-mandates topics.

Returning to the question that motivated our research, we find that during tightening periods there is more likely to be discussion about the use of monetary policy to achieve financial stability, as well as non-dual mandate topics; during easing periods the discussion is more likely to focus on monetary policy and the dual mandate topics of inflation and employment. Among the non-dual mandate topics we consider, discussion of financial crises has a negative effect on all asset returns, while discussion of financial stability only significantly affects equity and risky asset returns. We also find that tense matters for asset returns, with a focus on either the past or present being associated with positive bond and safe asset returns and a focus on the future being better for equity and risky asset returns.

References

- Apel, M. and M. Blix Grimaldi**, “How Informative are Central Bank Minutes?,” *Review of Economics*, 2014, 65.
- Apel, Mikael, Marianna Blix Grimaldi, and Isaiah Hull**, “How Much Information Do Monetary Policy Committees Disclose? Evidence from the FOMC’s Minutes and Transcripts,” *Working Paper No. 381, Sveriges Riksbank*, 2019.
- Bahdanau, Dzmitry, Kyunghyun Cho, and Yoshua Bengio**, “Neural Machine Translation by Jointly Learning to Align and Translate,” in Yoshua Bengio and Yann LeCun, eds., *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- Bai, Jushan**, “Estimating multiple breaks one at a time,” *Econometric theory*, 1997, 13 (3), 315–352.
- **and Pierre Perron**, “Estimating and testing linear models with multiple structural changes,” *Econometrica*, 1998, pp. 47–78.
- Bauer, Michael D., Carolin E. Pflueger, and Adi Sunderam**, “Perceptions about Monetary Policy,” *Mimeo*, 2022.
- Bernanke, Ben S. and Mark Gertler**, “Should Central Banks Respond to Movements in Asset Prices?,” *The American Economic Review*, 2001, 91 (2), 253–257.
- **and –**, “Should Central Banks Respond to Movements in Asset Prices?,” *American Economic Review*, 2001, 91 (2), 253–257.
- Bianchi, Francesco**, “Regime Switches, Agents’ Beliefs, and Post-World War II U.S. Macroeconomic Dynamics,” *The Review of Economic Studies*, 10 2012, 80 (2), 463–490.
- **, Martin Lettau, and Sydney C Ludvigson**, “Monetary policy and asset valuation,” *The Journal of Finance*, 2022, 77 (2), 967–1017.
- Boivin, Jean**, “Has U.S. Monetary Policy Changed? Evidence from Drifting Coefficients and Real-Time Data,” *Journal of Money, Credit and Banking*, 2006, 28 (5), 1149–1173.
- Bowman, Samuel R., Gabor Angeli, Christopher Potts, and Christopher D. Manning**, “A large annotated corpus for learning natural language inference,” in “Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing” Association for Computational Linguistics Lisbon, Portugal September 2015, pp. 632–642.
- Byrne, David, Robert Goodhead, Michael McMahon, and Conor Parle**, “The Central Bank Crystal Ball: Temporal information in monetary policy communication,” *Mimeo*, 2022.
- Caballero, Ricardo and Alp Simsek**, “Prudential Monetary Policy,” *CEPR Discussion Paper No. 13832*, 2020.

- Carvalho, Carlos, Fernanda Nechio, and Tiago Tristão**, “Taylor rule estimation by OLS,” *Journal of Monetary Economics*, 2021, *124*, 140–154.
- Cer, Daniel et al.**, “Universal Sentence Encoder for English,” in “Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations” Association for Computational Linguistics Brussels, Belgium November 2018, pp. 169–174.
- Cho, Kyunghyun et al.**, “Learning Phrase Representations using RNN Encoder–Decoder for Statistical Machine Translation,” in “Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)” Association for Computational Linguistics Doha, Qatar October 2014, pp. 1724–1734.
- Clarida, Richard, Jordi Gali, and Mark Gertler**, “Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory,” *The Quarterly Journal of Economics*, 2000, *115* (1), 147–180.
- Coibion, Olivier and Yuriy Gorodnichenko**, “Why Are Target Interest Rate Changes so Persistent?,” *American Economic Journal: Macroeconomics*, 2012, *4* (4), 126–162.
- Conneau, Alexis and Guillaume Lample**, “Cross-lingual Language Model Pretraining,” in Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett, eds., *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, 2019, pp. 7057–7067.
- , **Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes**, “Supervised Learning of Universal Sentence Representations from Natural Language Inference Data,” in “Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing” Association for Computational Linguistics Copenhagen, Denmark September 2017, pp. 670–680.
- Correa, R., K. Garud, J.M. Londono, and N. Mislang**, “Sentiment in Central Banks’ Financial Stability Reports,” *Review of Finance*, 2021, *75* (1), 85–120.
- Dai, Andrew M. and Quoc V. Le**, “Semi-supervised Sequence Learning,” in Corinna Cortes, Neil D. Lawrence, Daniel D. Lee, Masashi Sugiyama, and Roman Garnett, eds., *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, 2015, pp. 3079–3087.
- Dell’Ariccia, Giovanni, Luc Laeven, and Gustavo A. Suarez**, “Bank Leverage and Monetary Policy’s Risk-Taking Channel: Evidence from the United States: Bank Leverage and Monetary Policy’s Risk-Taking Channel,” *The Journal of Finance*, April 2017, *72* (2), 613–654.
- Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova**, “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding,” in Jill

- Burstein, Christy Doran, and Thamar Solorio, eds., *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, Association for Computational Linguistics 2019, pp. 4171–4186.
- Federal Reserve Bank of St. Louis**, “FRED,” <https://fred.stlouisfed.org/>. Accessed: 2022-05.
- Gentzkow, M., B. Kelly, and M. Taddy**, “Text as Data,” *Journal of Economic Literature*, 2019, 57.
- Hamilton, James D., Seth Pruitt, and Scott C. Borger**, “Estimating the Market-Perceived Monetary Policy Rule,” *American Economic Journal: Macroeconomics*, 2011, 3 (3), 1–28.
- Hill, Felix, Kyunghyun Cho, and Anna Korhonen**, “Learning Distributed Representations of Sentences from Unlabelled Data,” in “Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies” Association for Computational Linguistics San Diego, California June 2016, pp. 1367–1377.
- Hochreiter, Sepp and Jürgen Schmidhuber**, “Long Short-Term Memory,” *Neural Computation*, November 1997, 9 (8), 1735–1780.
- Howard, Jeremy and Sebastian Ruder**, “Universal Language Model Fine-tuning for Text Classification,” in “Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)” Association for Computational Linguistics Melbourne, Australia July 2018, pp. 328–339.
- Howitt, Peter**, “What have central bankers learned from modern macroeconomic theory?,” *Journal of Macroeconomics*, 2012, 34 (1), 11–22.
- Hull, Isaiah**, *Machine Learning for Economics and Finance in TensorFlow 2: Deep Learning Models for Research and Industry*, Berkeley, CA: Apress, 2021.
- Istrefi, Klodiana, Florens Odendahl, and Giulia Sestri**, “FED Communication on Financial Stability Concerns and Monetary Policy Decisions: Revelations from Speeches,” *Banco de España Documentos de Trabajo. N.º 2110*, 2021.
- Jiménez, Gabriel, Steven Ongena, José-Luis Peydró, and Jesús Saurina**, “Hazardous Times for Monetary Policy: What Do Twenty-Three Million Bank Loans Say About the Effects of Monetary Policy on Credit Risk-Taking?,” *Econometrica*, 2014, 82 (2), 463–505.
- Kiros, Ryan, Yukun Zhu, Ruslan Salakhutdinov, Richard S. Zemel, Antonio Torralba, Raquel Urtasun, and Sanja Fidler**, “Skip-Thought Vectors,” in “Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2” NIPS’15 MIT Press Cambridge, MA, USA 2015, p. 3294–3302.

- Liu, Yinhan, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov**, “RoBERTa: A Robustly Optimized BERT Pretraining Approach,” *CoRR*, 2019, *abs/1907.11692*.
- Lo, Kyle, Lucy Lu Wang, Mark Neumann, Rodney Kinney, and Daniel Weld**, “S2ORC: The Semantic Scholar Open Research Corpus,” in “Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics” Association for Computational Linguistics Online July 2020, pp. 4969–4983.
- Loughran, T. and B. McDonald**, “When Is a Liability Not a Liability? Textual Analysis, Dictionaries, and 10Ks,” *Journal of Finance*, 2011, *LXVI*, 35–36.
- Luong, Thang, Hieu Pham, and Christopher D. Manning**, “Effective Approaches to Attention-based Neural Machine Translation,” in “Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing” Association for Computational Linguistics Lisbon, Portugal September 2015, pp. 1412–1421.
- Mankiw, N. Gregory**, “The Macroeconomist as Scientist and Engineer,” *Journal of Economic Perspectives*, December 2006, *20* (4), 29–46.
- McCann, Bryan, James Bradbury, Caiming Xiong, and Richard Socher**, “Learned in Translation: Contextualized Word Vectors,” in I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, eds., *Advances in Neural Information Processing Systems*, Vol. 30 Curran Associates, Inc. 2017.
- Meltzer, Allan H.**, *A History of the Federal Reserve*, Vol. 2 of Book 1, 1951–1969, University of Chicago Press, 2009.
- Nikolsko-Rzhevskyy, Alex, David H. Papellb, and Ruxandra Prodan**, “Deviations from rules-based policy and their effects,” *Journal of Economic Dynamics and Control*, December 2014, *49*, 4–17.
- Orphanides, Athanasios**, “Monetary Policy Rules Based on Real-Time Data,” *American Economic Review*, September 2001, *91* (4), 964–985.
- , “Historical monetary policy analysis and the Taylor rule,” *Journal of Monetary Economics*, 2003, *50* (5), 983–1022.
- Peek, Joe, Eric S. Rosengren, and Geoffrey M. B. Tootell**, “Does Fed policy reveal a ternary mandate?,” Working Papers 16-11, Federal Reserve Bank of Boston September 2016.
- Peters, Matthew E., Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer**, “Deep Contextualized Word Representations,” in Marilyn A. Walker, Heng Ji, and Amanda Stent, eds., *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers)*, Association for Computational Linguistics 2018, pp. 2227–2237.

- Pushp, Pushpankar Kumar and Muktabh Mayank Srivastava**, “Train Once, Test Anywhere: Zero-Shot Learning for Text Classification,” *CoRR*, 2017, *abs/1712.05972*.
- Rajpurkar, Pranav, Jian Zhang, Konstantin Lopyrev, and Percy Liang**, “SQuAD: 100,000+ Questions for Machine Comprehension of Text,” in “Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing” Association for Computational Linguistics Austin, Texas November 2016, pp. 2383–2392.
- , **Robin Jia, and Percy Liang**, “Know What You Don’t Know: Unanswerable Questions for SQuAD,” in Iryna Gurevych and Yusuke Miyao, eds., *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 2: Short Papers*, Association for Computational Linguistics 2018, pp. 784–789.
- Reimers, Nils and Iryna Gurevych**, “Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks,” in Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan, eds., *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, Association for Computational Linguistics 2019, pp. 3980–3990.
- Sachs, Jeffrey D.**, “International Policy Coordination: The Case of the Developing Country Debt Crisis,” *Working Paper No. 2287*, June 1987.
- Salimans, Alec Radford Karthik Narasimhan Tim and Ilya Sutskever**, “Improving Language Understanding with Unsupervised Learning,” *Technical Report, Open AI*, 2018.
- Schroff, Florian, Dmitry Kalenichenko, and James Philbin**, “FaceNet: A Unified Embedding for Face Recognition and Clustering,” *CoRR*, 2015, *abs/1503.03832*.
- Schularick, Moritz, Lucas ter Steege, and Felix Ward**, “Leaning against the Wind and Crisis Risk,” *American Economic Review: Insights*, 2021, 3 (2), 199–214.
- Shapiro, Adam H. and Daniel J Wilson**, “Taking the Fed at its Word: A New Approach to Estimating Central Bank Objectives Using Text Analysis,” *The Review of Economic Studies*, 2022, 92 (5), 2768–2805.
- Sims, Christopher A. and Tao Zha**, “Were There Regime Switches in U.S. Monetary Policy?,” *American Economic Review*, March 2006, 96 (1), 54–81.
- Sutskever, Ilya, Oriol Vinyals, and Quoc V. Le**, “Sequence to Sequence Learning with Neural Networks,” in “Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2” NIPS’14 MIT Press Cambridge, MA, USA 2014, p. 3104–3112.
- Taylor, John B.**, “Discretion versus policy rules in practice,” *Carnegie-Rochester Conference Series on Public Policy*, 1993, 39, 195–214.
- , “A Historical Analysis of Monetary Policy Rules,” January 1999.

- van der Maaten, Laurens and Geoffrey Hinton**, “Visualizing Data using t-SNE,” *Journal of Machine Learning Research*, 2008, 9 (86), 2579–2605.
- van Dieijen, Myrthe and Robin L. Lumsdaine**, “What Say They About Their Mandate? A Textual Assessment of Federal Reserve Speeches,” *Mimeo*, 2019.
- Vaswani, Ashish et al.**, “Attention is All you Need,” in Isabelle Guyon et al., eds., *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, 2017, pp. 5998–6008.
- Vincent, Pascal, Hugo Larochelle, Isabelle Lajoie, Yoshua Bengio, and Pierre-Antoine Manzagol**, “Stacked Denoising Autoencoders: Learning Useful Representations in a Deep Network with a Local Denoising Criterion,” *J. Mach. Learn. Res.*, December 2010, 11, 3371–3408.
- Vollmer, Uwe**, “Monetary policy or macroprudential policies: What can tame the cycles?,” *Journal of Economic Surveys (forthcoming)*, 11 2021, pp. 1–29.
- Wang, Alex et al.**, “GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding,” in “7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019” OpenReview.net 2019.
- Wang, Kexin, Nils Reimers, and Iryna Gurevych**, “TSDAE: Using Transformer-based Sequential Denoising Auto-Encoder for Unsupervised Sentence Embedding Learning,” *CoRR*, 2021, *abs/2104.06979*.
- Williams, Adina, Nikita Nangia, and Samuel Bowman**, “A Broad-Coverage Challenge Corpus for Sentence Understanding through Inference,” in “Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)” Association for Computational Linguistics New Orleans, Louisiana June 2018, pp. 1112–1122.
- Wu, Jing Cynthia and Fan Dora Xia**, “Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound,” *Journal of Money, Credit, and Banking*, 2016, 48 (2-3), 253–291.
- Wu, Yonghui et al.**, “Google’s Neural Machine Translation System: Bridging the Gap between Human and Machine Translation,” *CoRR*, 2016, *abs/1609.08144*.
- Yang, Zhilin, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le**, “XLNet: Generalized Autoregressive Pretraining for Language Understanding,” in H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, eds., *Advances in Neural Information Processing Systems*, Vol. 32 Curran Associates, Inc. 2019.
- Yin, Wenpeng, Jamaal Hay, and Dan Roth**, “Benchmarking Zero-shot Text Classification: Datasets, Evaluation and Entailment Approach,” in Kentaro Inui, Jing Jiang,

Vincent Ng, and Xiaojun Wan, eds., *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, Association for Computational Linguistics 2019, pp. 3912–3921.

Yogatama, Dani, Chris Dyer, Wang Ling, and Phil Blunsom, “Generative and Discriminative Text Classification with Recurrent Neural Networks,” *CoRR*, 2017, *abs/1703.01898*.

Zhang, Jingqing, Piyawat Lertvittayakumjorn, and Yike Guo, “Integrating Semantic Knowledge to Tackle Zero-shot Text Classification,” in Jill Burstein, Christy Doran, and Thamar Solorio, eds., *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, Association for Computational Linguistics 2019, pp. 1031–1040.

Zhu, Yukun et al., “Aligning Books and Movies: Towards Story-Like Visual Explanations by Watching Movies and Reading Books,” in “2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015” IEEE Computer Society 2015, pp. 19–27.

Òscar Jordà, Moritz Schularick, and Alan M. Taylor, *Macrofinancial History and the New Business Cycle Facts*, University of Chicago Press, May 2016.

A Appendix

A.1 Regression Tables

Table 9: Taylor Rule Regressions

	(1) Volek- Greenspan 1986Q1-2005Q4	(2) Greenspan- Bernanke 1987Q3-2007Q4	(3) Post- Volek 1986Q1-2007Q4	(4) Extended Sample, 2008Q1-2020Q4 Federal Funds Rate	(5) Extended Sample 2008Q1-2020Q4 Shadow Rate	(6) Extended Sample 2008Q1-2020Q4 Shadow Rate	(7) Extended Sample 2008Q1-2020Q4 Shadow Rate	(8) Extended Sample 2008Q1-2020Q4 Shadow Rate	(9) Extended Sample 2008Q1-2020Q4 Shadow Rate
β	1.32*** (0.20)	1.19*** (0.32)	1.24*** (0.28)	-0.44* (0.26)	-1.09 (0.76)	-1.15 (0.89)	-0.49 (0.49)	-0.77 (0.68)	-1.35 (0.96)
γ	0.99*** (0.10)	0.91*** (0.21)	1.06*** (0.15)	0.28*** (0.08)	0.68*** (0.25)	0.48* (0.28)	0.43** (0.19)	0.49** (0.21)	0.66* (0.38)
ρ	0.64*** (0.10)	0.82*** (0.05)	0.74*** (0.07)	0.63*** (0.11)	0.86*** (0.05)	0.89*** (0.04)	0.83*** (0.05)	0.84*** (0.05)	0.87*** (0.05)
π^*	1.73 (1.86)	0.81 (5.09)	1.08 (3.35)	2.18*** (0.40)	2.40*** (0.71)	2.19*** (0.80)	3.03*** (0.69)	3.00*** (0.84)	2.61*** (0.82)
non dual mandate _t [text]									
financial stability _t [text]									
monetary financial _t [text]									
financial stability _{t-1} [text]									
monetary financial _{t-1} [text]									
Adj. R-squared	0.91	0.96	0.91	0.84	0.93	0.94	0.94	0.94	0.94
N	80	82	88	52	52	52	52	52	52
RMSE	0.68	0.43	0.67	0.32	0.41	0.40	0.39	0.39	0.40

Notes: We estimate version of the specification in equation (3). The dependent variable is the federal funds rate in columns (1)-(4) and the Wu and Xia (2016) end-of-quarter shadow federal funds rate in columns (5)-(9). All controls that include [text] indicate that they are text features measured using zero shot classification. The text feature “non dual mandate [text]” is the classification score for whether a sequence is non dual mandate related, the text feature “financial stability [text]” is the classification score for whether a sequence is non dual mandate related, and the text feature “monetary financial [text]” is the classification score for whether a sequence supports the view that financial stability should be achieved using monetary policy. As additional controls we use lagged values of the interest rate, as well as lagged values of the Greenbook forecast for inflation and for the output gap. Robust standard errors are reported in parenthesis. * $p < .1$, ** $p < .05$, *** $p < .01$.

Table 10: Monetary Policy Regimes Regressions

	(1)	(2)	(3)	(4)	(5)	(6)
	hawkish	hawkish	hawkish	hawkish	hawkish	hawkish
financial crisis current	-0.058** (0.029)	-0.070** (0.028)	-0.069** (0.028)	-0.058* (0.034)	-0.070** (0.030)	-0.069** (0.033)
monetary financial	0.170*** (0.041)	0.167*** (0.045)	0.221*** (0.041)	0.170*** (0.063)	0.167** (0.071)	0.221*** (0.053)
employment inflation		-0.092* (0.046)			-0.092 (0.075)	
monetary policy			-0.199*** (0.049)			-0.199*** (0.076)
non dual mandate	0.104** (0.043)			0.104 (0.070)		
financial stability			-0.020 (0.034)			-0.020 (0.054)
No of obs	221.000	221.000	221.000	221.000	221.000	221.000
Adj. R2	0.091	0.082	0.134			
NW(4) Std. Err.	NO	NO	NO	YES	YES	YES

Notes: We estimate a version of the specification in equation (4). The dependent variable is the probability of the Fed being in a hawkish monetary policy regime. All controls are text features measured using zero shot classification. The text feature “financial crisis current” is the classification score for whether a sequence is discussing the financial crisis with a current temporal focus, the text feature “monetary financial” is the classification score for whether a sequence supports the view that financial stability should be achieved using monetary policy, the text feature “non dual mandate” is the classification score for whether a sequence is non dual mandate related, and other text features are the classification scores for whether a sequence is related to a particular discussion, including “employment/inflation”, “monetary policy” and “financial stability”. Standard errors are reported in parenthesis. * $p < .1$, ** $p < .05$, *** $p < .01$.

Table 11: Federal Reserve Speech Focus: Financial Stability (Lagged Macro and Finance Controls)

	(1)	(2)	(3)	(4)	(5)
inflation _{jt} [text]	0.0427*** (0.0029)	0.0427*** (0.0101)	0.0434*** (0.0029)	-0.0353*** (0.0058)	0.0628*** (0.0033)
employment _{jt} [text]	0.1204*** (0.0021)	0.1204*** (0.0065)	0.1206*** (0.0021)	0.1580*** (0.0050)	0.1111*** (0.0023)
financial crisis _{jt} [text]	0.0888*** (0.0029)	0.0888*** (0.0150)	0.0882*** (0.0029)	0.0669*** (0.0066)	0.0952*** (0.0032)
bank liquidity _{jt} [text]	0.1427*** (0.0040)	0.1427*** (0.0126)	0.1408*** (0.0040)	0.1775*** (0.0096)	0.1348*** (0.0044)
bank capital _{jt} [text]	0.2997*** (0.0043)	0.2997*** (0.0154)	0.3008*** (0.0043)	0.3458*** (0.0104)	0.2902*** (0.0047)
past focus _{jt} [text]	-0.1177*** (0.0020)	-0.1177*** (0.0080)	-0.1162*** (0.0020)	-0.1059*** (0.0043)	-0.1198*** (0.0022)
present focus _{jt} [text]	0.0738*** (0.0017)	0.0738*** (0.0042)	0.0744*** (0.0017)	0.0516*** (0.0035)	0.0799*** (0.0019)
future focus _{jt} [text]	0.0900*** (0.0021)	0.0900*** (0.0071)	0.0899*** (0.0021)	0.0632*** (0.0046)	0.0946*** (0.0023)
academic focus _{jt} [text]	0.1615*** (0.0021)	0.1615*** (0.0075)	0.1603*** (0.0021)	0.1564*** (0.0053)	0.1607*** (0.0023)
debt-to-gdp ratio _{t-1}	0.2862*** (0.0283)	0.2862*** (0.0654)		-0.1479 (0.2667)	0.2521*** (0.0367)
loan-to-deposit ratio _{t-1}	0.0038*** (0.0006)	0.0038*** (0.0016)		0.0079** (0.0035)	0.0025*** (0.0008)
Sample Period	1960-2020	1960-2020	1960-2020	1960-1983	1984-2020
Macro Controls	YES	YES	NO	YES	YES
Financial Controls	YES	YES	NO	YES	YES
Housing Controls	YES	YES	NO	YES	YES
Month x Year FEs	NO	NO	YES	NO	NO
District FEs	YES	YES	YES	YES	YES
Standard Errors	NW	CL	NW	NW	NW
Adj. R-squared	0.5245	0.5245	0.5273	0.5088	0.5287
N	263,613	263,613	263,613	49,262	214,351

Notes: The dependent variable in all regressions is the zero shot classification score for the term “financial stability.” A higher score indicates that the text in a given sequence of words – such as a paragraph or a sentence – describes a topic related to financial stability. All controls that include [text] indicate that they are text features measured using zero shot classification. For instance “inflation [text]” is the classification score for whether a sequence describes inflation. The features “past focus,” “present focus,” and “future focus” classify the tense of a sequence to determine whether the speaker was discussing the past, present, or future. “Academic focus” classifies whether a speaker was discussing an academic debate or the academic literature. Macro controls include the output gap and consumer price inflation. Financial controls include the log of total loans to the non-financial sector, the short term interest rate, and an indicator for whether a financial crisis occurred in a given year. The financial controls, housing controls, debt-to-gdp ratio, and loan-to-deposit ratio are taken from the *Macroeconomy Database*, introduced by Jordà et al. (2016). All non-text controls are lagged, as indicated by the time subscripts. Standard errors are either Newey-West (NW) or clustered (CL). * $p < .1$, ** $p < .05$, *** $p < .01$.

Table 12: Federal Reserve Speech Advocacy for the Use of Monetary Policy to Achieve Financial Stability (Lagged Macro and Finance Controls)

	(1)	(2)	(3)	(4)	(5)
inflation _{jt} [text]	0.2087*** (0.0024)	0.2087*** (0.0132)	0.1941*** (0.0024)	0.2053*** (0.0055)	0.2066*** (0.0026)
employment _{jt} [text]	-0.2042*** (0.0027)	-0.2042*** (0.0126)	-0.1991*** (0.0026)	-0.1851*** (0.0060)	-0.2106*** (0.0029)
financial crisis _{jt} [text]	0.1146*** (0.0031)	0.1146*** (0.0094)	0.1154*** (0.0030)	0.1483*** (0.0073)	0.1122*** (0.0034)
bank liquidity _{jt} [text]	-0.1472*** (0.0045)	-0.1472*** (0.0163)	-0.1386*** (0.0044)	-0.1751*** (0.0113)	-0.1350*** (0.0049)
bank capital _{jt} [text]	0.0487*** (0.0044)	0.0487*** (0.0175)	0.0478*** (0.0043)	0.0783*** (0.0114)	0.0389*** (0.0047)
past focus _{jt} [text]	0.0365*** (0.0026)	0.0365*** (0.0087)	0.0317*** (0.0025)	0.0121** (0.0055)	0.0396*** (0.0029)
present focus _{jt} [text]	-0.0015 (0.0023)	-0.0015 (0.0095)	-0.0009 (0.0022)	0.0336*** (0.0048)	-0.0136*** (0.0025)
future focus _{jt} [text]	0.0623*** (0.0026)	0.0623*** (0.0085)	0.0606*** (0.0025)	0.0211*** (0.0058)	0.0716*** (0.0028)
academic focus _{jt} [text]	-0.0658*** (0.0027)	-0.0658*** (0.0123)	-0.0609*** (0.0026)	-0.0940*** (0.0069)	-0.0575*** (0.0029)
debt-to-gdp ratio _{t-1}	-0.1819*** (0.0547)	-0.1819 (0.3508)		1.8531*** (0.4910)	-0.0470 (0.0727)
loan-to-deposit ratio _{t-1}	-0.0174*** (0.0011)	-0.0174** (0.0079)		-0.0519*** (0.0068)	-0.0231*** (0.0015)
Sample Period	1960-2020	1960-2020	1960-2020	1960-1983	1984-2020
Macro Controls	YES	YES	NO	YES	YES
Financial Controls	YES	YES	NO	YES	YES
Housing Controls	YES	YES	NO	YES	YES
Month x Year FEs	NO	NO	YES	NO	NO
District FEs	YES	YES	YES	YES	YES
Standard Errors	NW	CL	NW	NW	NW
Adj. R-squared	0.0970	0.0970	0.1499	0.1354	0.0993
N	263,613	263,613	263,613	49,262	214,351

Notes: The dependent variable in all regressions is the cosine similarity between the contextualized embedding for a given sequence of words and the contextualized embedding that corresponds to the sequence “monetary policy should be used to achieve financial stability.” A higher score indicates that a sequence is more likely to be advocating for the use of monetary policy to achieve financial stability. All controls that include [text] indicate that they are text features measured using zero shot classification. For instance “inflation [text]” is the classification score for whether a sequence describes inflation. The features “past focus,” “present focus,” and “future focus” classify the tense of a sequence to determine whether the speaker was discussing the past, present, or future. “Academic focus” classifies whether a speaker was discussing an academic debate or the academic literature. Macro controls include the output gap and consumer price inflation. Financial controls include the log of total loans to the non-financial sector, the short term interest rate, and an indicator for whether a financial crisis occurred in a given year. The financial controls, housing controls, debt-to-gdp ratio, and loan-to-deposit ratio are taken from the *Macrohistory Database*, introduced by Jordà et al. (2016). All non-text controls are lagged, as indicated by the time subscripts. Standard errors are either Newey-West (NW) or clustered (CL) at the district level. * $p < .1$, ** $p < .05$, *** $p < .01$.

Table 13: Federal Reserve Speech Advocacy for the Use of Banking Regulation to Achieve Financial Stability (Lagged Macro and Finance Controls)

	(1)	(2)	(3)	(4)	(5)
inflation _{<i>t</i>} [text]	-0.3011*** (0.0023)	-0.3011*** (0.0267)	-0.2891*** (0.0023)	-0.1905*** (0.0045)	-0.3215*** (0.0027)
employment _{<i>t</i>} [text]	-0.2022*** (0.0023)	-0.2022*** (0.0129)	-0.1884*** (0.0023)	-0.1563*** (0.0047)	-0.2107*** (0.0026)
financial crisis _{<i>t</i>} [text]	0.2225*** (0.0030)	0.2225*** (0.0190)	0.2130*** (0.0030)	0.1112*** (0.0060)	0.2451*** (0.0034)
bank liquidity _{<i>t</i>} [text]	-0.1550*** (0.0053)	-0.1550*** (0.0222)	-0.1469*** (0.0051)	-0.1661*** (0.0115)	-0.1466*** (0.0058)
bank capital _{<i>t</i>} [text]	0.3772*** (0.0053)	0.3772*** (0.0218)	0.3577*** (0.0051)	0.2942*** (0.0115)	0.3859*** (0.0057)
past focus _{<i>t</i>} [text]	-0.0016 (0.0024)	-0.0016 (0.0062)	-0.0038 (0.0023)	-0.0165*** (0.0047)	-0.0023 (0.0027)
present focus _{<i>t</i>} [text]	0.0747*** (0.0021)	0.0747*** (0.0050)	0.0678*** (0.0020)	0.0815*** (0.0040)	0.0711*** (0.0024)
future focus _{<i>t</i>} [text]	-0.0341*** (0.0024)	-0.0341*** (0.0050)	-0.0291*** (0.0023)	-0.0286*** (0.0048)	-0.0335*** (0.0027)
academic focus _{<i>t</i>} [text]	-0.0252*** (0.0026)	-0.0252 (0.0169)	-0.0229*** (0.0025)	-0.0182*** (0.0056)	-0.0239*** (0.0028)
debt-to-gdp ratio _{<i>t-1</i>}	0.0766 (0.0522)	0.0766 (0.2167)	0.0766 (0.2167)	2.1566*** (0.3869)	0.2281*** (0.0684)
loan-to-deposit ratio _{<i>t-1</i>}	0.0012 (0.0011)	-0.0012 (0.0063)	-0.0012 (0.0063)	-0.0114** (0.0051)	-0.0089*** (0.0015)
Sample Period	1960-2020	1960-2020	1960-2020	1960-1983	1984-2020
Macro Controls	YES	YES	NO	YES	YES
Financial Controls	YES	YES	NO	YES	YES
Housing Controls	YES	YES	NO	YES	YES
Month x Year FEs	NO	NO	YES	NO	NO
District FEs	YES	YES	YES	YES	YES
Standard Errors	NW	CL	NW	NW	NW
Adj. R-squared	0.1886	0.1886	0.2243	0.1372	0.2021
N	263,613	263,613	263,613	49,262	214,351

Notes: The dependent variable in all regressions is the cosine similarity between the contextualized embedding for a given sequence of words and the contextualized embedding that corresponds to the sequence “monetary policy should be used to achieve financial stability.” A higher score indicates that a sequence is more likely to be advocating for the use of monetary policy to achieve financial stability. All controls that include [text] indicate that they are text features measured using zero shot classification. For instance “inflation [text]” is the classification score for whether a sequence describes inflation. The features “past focus,” “present focus,” and “future focus” classify the tense of a sequence to determine whether the speaker was discussing the past, present, or future. “Academic focus” classifies whether a speaker was discussing an academic debate or the academic literature. Macro controls include the output gap and consumer price inflation. Financial controls include the log of total loans to the non-financial sector, the short term interest rate, and an indicator for whether a financial crisis occurred in a given year. The financial controls, housing controls, debt-to-gdp ratio, and loan-to-deposit ratio are taken from the *Macrohistory Database*, introduced by Jordà et al. (2016). All non-text controls are lagged, as indicated by the time subscripts. Standard errors are either Newey-West (NW) or clustered (CL) at the district level. * $p < .1$, ** $p < .05$, *** $p < .01$.

Table 14: Federal Reserve Speech Impact on Asset Returns (Lagged Macro and Financial Controls)

	(1)	(2)	(3)	(4)	(5)
inflation _{<i>t</i>} [text]	0.0008* (0.0005)	0.0013*** (0.0005)	-0.0009*** (0.0004)	0.0004** (0.0002)	-0.0004** (0.0002)
employment _{<i>t</i>} [text]	-0.0017*** (0.0004)	0.0001 (0.0004)	-0.0005 (0.0003)	0.0003* (0.0002)	-0.0000 (0.0002)
financial stability _{<i>t</i>} [text]	0.0008* (0.0004)	0.0013*** (0.0004)	-0.0009*** (0.0003)	0.0004** (0.0002)	-0.0004** (0.0002)
financial crisis _{<i>t</i>} [text]	0.0005 (0.0005)	-0.0022*** (0.0005)	-0.0014*** (0.0004)	-0.0011*** (0.0002)	-0.0012*** (0.0002)
bank liquidity _{<i>t</i>} [text]	-0.0002 (0.0007)	-0.0005 (0.0007)	-0.0009* (0.0005)	-0.0007*** (0.0003)	-0.0008*** (0.0003)
bank capital _{<i>t</i>} [text]	-0.0012* (0.0007)	-0.0010 (0.0007)	0.0016*** (0.0005)	-0.0003 (0.0003)	0.0010*** (0.0003)
past focus _{<i>t</i>} [text]	0.0002 (0.0004)	-0.0002 (0.0004)	-0.0000 (0.0003)	0.0002 (0.0002)	0.0000 (0.0002)
present focus _{<i>t</i>} [text]	0.0010*** (0.0004)	0.0009** (0.0004)	0.0003 (0.0003)	0.0003* (0.0001)	0.0001 (0.0001)
future focus _{<i>t</i>} [text]	-0.0007* (0.0004)	-0.0005 (0.0004)	0.0002 (0.0003)	-0.0000 (0.0002)	0.0003 (0.0002)
past focus _{<i>t</i>} * financial stability _{<i>t</i>} [text]	-0.0008*** (0.0003)	-0.0008*** (0.0003)	-0.0008*** (0.0003)	-0.0008*** (0.0003)	-0.0008*** (0.0003)
present focus _{<i>t</i>} * financial stability _{<i>t</i>} [text]	0.0006 (0.0004)	0.0006 (0.0004)	0.0006 (0.0004)	0.0006 (0.0004)	0.0006 (0.0004)
future focus _{<i>t</i>} * financial stability _{<i>t</i>} [text]	-0.0001 (0.0004)	-0.0001 (0.0004)	-0.0001 (0.0004)	-0.0001 (0.0004)	-0.0001 (0.0004)
academic focus _{<i>t</i>} [text]	-0.0010** (0.0004)	-0.0003 (0.0004)	0.0002 (0.0003)	-0.0003* (0.0002)	0.0000 (0.0002)
debt-to-gdp ratio _{<i>t-1</i>}	0.4355*** (0.0112)	0.3342*** (0.0098)	0.0885*** (0.0079)	0.2294*** (0.0046)	-0.0258*** (0.0040)
loan-to-deposit ratio _{<i>t-1</i>}	0.0065*** (0.0002)	0.0060*** (0.0002)	0.0029*** (0.0002)	0.0056*** (0.0001)	0.0005*** (0.0001)
Sample Period	1960-2020	1960-2020	1960-2020	1960-2020	1960-2020
Macro Controls	YES	YES	YES	YES	YES
Financial Controls	YES	YES	YES	YES	YES
Housing Controls	YES	YES	YES	YES	YES
District FEs	YES	YES	YES	YES	YES
Standard Errors	NW	NW	NW	NW	NW
Return Type	Equity	Equity	Bond	Risky	Safe
Adj. R-squared	0.2988	0.3345	0.2555	0.4187	0.2996
N	263,613	263,613	263,613	263,613	263,613

Notes: The dependent variable in each regression is a measure of annual asset returns. The returns are taken from Jordà et al. (2016) and include total equity, total bond, risky assets, and safe assets, as specified in the “return type” row of the table. All controls that include [text] indicate that they are text features measured using zero shot classification. For instance “inflation [text]” is the classification score for whether a sequence describes inflation. The features “past focus,” “present focus,” and “future focus” classify the tense of a sequence to determine whether the speaker was discussing the past, present, or future. “Academic focus” classifies whether a speaker was discussing an academic debate or the academic literature. Macro controls include the output gap and consumer price inflation. Financial controls include the log of total loans to the non-financial sector, the short term interest rate, and an indicator for whether a financial crisis occurred in a given year. The financial controls, housing controls, debt-to-gdp ratio, and loan-to-deposit ratio are taken from the *Macroeconomy Database*, introduced by Jordà et al. (2016). All non-text controls are lagged, as indicated by the time subscripts. Standard errors are either Newey-West (NW) or clustered (CL) at the district level. * $p < .1$, ** $p < .05$, *** $p < .01$.

A.2 Sequence-to-Sequence Modeling

Sequence-to-sequence (S2S) modeling poses challenges that are not present in other NLP tasks. Among other problems, dense neural networks require fixed-length inputs and outputs and, thus, cannot be used for S2S tasks, such as machine translation, which requires variation in input and output length. This is true even for a given translated sentence, which may be best expressed using a different number of words in two languages. Consequently, models that are suitable for text classification purposes, such as dense neural networks, are not usable for S2S tasks.

Early breakthroughs in S2S modeling centered around the use of a variant of recurrent neural networks (RNN) called a long-short term memory (LSTM) model (Hochreiter and Schmidhuber, 1997).²¹ LSTMs explicitly treat input data, such as words in a sentence, as a sequence, rather than as a set of features.²² This allows for a more parsimonious parameterization than a dense neural network would permit. LSTMs are also able to handle variable length input sequences, which provides an advantage over dense neural networks for S2S modeling.

The initial innovation in S2S modeling with LSTMs involved the use of an encoder-decoder architecture (Sutskever et al., 2014; Cho et al., 2014). The encoder maps a sequence of symbols to a latent vector, which can be viewed as a compressed representation of the input text. The decoder then maps the latent vector to a sequence of symbol predictions. Since the output sequence must also be permitted to have a variable length in many applications, the model is trained to output an end-of-sentence ([EOS]) token, which terminates the sequence of predictions.

LSTM-based models with an encoder-decoder architecture provided the initial means of performing high-quality machine translation and also generated many spillover benefits for

²¹See Apel et al. (2019) for an application of LSTM models to natural language processing tasks in economics.

²²LSTMs also allow for long-term dependence between words in a sequence and correct a version of the vanishing gradient problem. See Hull (2021) for an overview of LSTMs in the context of economics.

related NLP tasks. However, the introduction of the attention mechanism (Bahdanau et al., 2015) and later the Transformer model (Vaswani et al., 2017) fundamentally changed how sequences are modeled in S2S contexts and in NLP more broadly. This provided a foundation for the set of NLP tools we use in this paper to measure the content of central bank communications.

A.3 The Transformer Model

In this subsection, we provide a detailed overview of the Transformer model and its advantages over earlier S2S modeling techniques.

Time Complexity. One advantage of the Transformer model is that it reduces the computational complexity of certain components of the training process relative to LSTM-based S2S models. In particular, the attention mechanism has a time complexity of $\mathcal{O}(T^2D)$; whereas recurrent operations, such as LSTM cells, have a time complexity of $\mathcal{O}(TD^2)$. This suggests that Transformer models will tend to have a training time advantage when embeddings are large.

Another important advantage of Transformer models is that they require $\mathcal{O}(1)$ sequential operations; whereas LSTM-based S2S models require $\mathcal{O}(T)$. This implies a substantial training time advantage for Transformer models, since they can parallelize operations that must be performed in sequence for LSTM-based S2S models.

Attention in Transformer Models. Transformer models apply three distinct forms of scaled dot product attention: 1) encoder-decoder attention; 2) encoder self-attention; and 3) decoder self-attention. Encoder-decoder attention uses the query vectors, \mathcal{Q} , from the previous decoder layer and key and value vectors, \mathcal{K} and \mathcal{V} , from the current encoder layer.

Self-attention, in contrast, is applied to words in the same sequence and in the same encoder or decoder layer. We make a distinction between encoder and decoder self-attention because decoder self-attention only uses the sequence of words up to and including the

word being predicted; whereas encoder self-attention uses the entire sequence. Figure 2 in Section 2 illustrates the self-attention mechanism from the RoBERTa model applied to word sequences in the Federal Reserve speeches in our corpus.²³

The Attention Mechanism Bahdanau et al. (2015) argued that the latent vector in LSTM-based encoder-decoder models created a bottleneck that made it difficult to improve model architecture and training. As a solution to this problem, they proposed using the attention mechanism, which eliminated the need to encode the entire input sequence in a single latent vector. For a given symbol, such as a word, the attention mechanism determines which symbols are related to it without explicitly considering the temporal ordering of the sequence. This allows for symbols that are not close together to be closely related. Luong et al. (2015) demonstrated how this could be used effectively on machine translation tasks.

In practice, attention proved to be an invaluable construct for NLP. We will briefly describe the attention mechanism below in an LSTM encoder-decoder context, focusing on the scaled dot product variant, which was later used in the Transformer model. For concreteness, consider the passage given in the quote below, taken from a speech given by Gary Stern, then-President of the Minneapolis Federal Reserve Bank, in January 2009. We will first convert the sequence of words to a sequence of embedding vectors, as shown in Equation (5).

“Commercial banks have long been thought of, and indeed have functioned as, the backup source of liquidity for many other financial institutions and markets. Banks continue to play this role, but it has become more challenging today to do so because some lenders find themselves capital-constrained as a result of recent losses and/or sizable, unanticipated additions to their balance sheets of formerly off-balance-sheet instruments.”

²³(The RoBERTa model is a “robustly optimized” version of the BERT model, which is described in detail below.)

In Equation (5), M is the embedding dimension and T is the sequence length. The mapping between words and embeddings is constructed outside of the model, typically through a separate training process.²⁴ In a minimal LSTM model with a single cell, the sequence of embeddings is processed in order with each step yielding a hidden state, which is then combined with the next input embedding in the sequence. Equation (6) provides the sequence of hidden states, which have the same dimension as the embedding vector in this architecture.

$$\left\{ e_1, e_2, \dots, e_t, \dots, e_T \right\} = \left\{ \begin{array}{c} \begin{bmatrix} e_{11} \\ e_{12} \\ \vdots \\ e_{1M} \end{bmatrix} \\ \begin{bmatrix} e_{21} \\ e_{22} \\ \vdots \\ e_{2M} \end{bmatrix} \\ \dots \\ \begin{bmatrix} e_{t1} \\ e_{t2} \\ \vdots \\ e_{tM} \end{bmatrix} \\ \dots \\ \begin{bmatrix} e_{T1} \\ e_{T2} \\ \vdots \\ e_{TM} \end{bmatrix} \end{array} \right\} \quad (5)$$

$$\left\{ h_1, h_2, \dots, h_t, \dots, h_T \right\} = \left\{ \begin{array}{c} \begin{bmatrix} h_{11} \\ h_{12} \\ \vdots \\ h_{1M} \end{bmatrix} \\ \begin{bmatrix} h_{21} \\ h_{22} \\ \vdots \\ h_{2M} \end{bmatrix} \\ \dots \\ \begin{bmatrix} h_{t1} \\ h_{t2} \\ \vdots \\ h_{tM} \end{bmatrix} \\ \dots \\ \begin{bmatrix} h_{T1} \\ h_{T2} \\ \vdots \\ h_{TM} \end{bmatrix} \end{array} \right\} \quad (6)$$

The relationship between the contemporaneous hidden state, h_t , the previous hidden state, h_{t-1} , and the input embedding, e_t , is given by Equation (7). Note that \mathcal{W}^E and \mathcal{W}^H are shape-preserving linear transformations of e_t and h_{t-1} and \mathcal{G} is an elementwise nonlinear activation function.

$$h_t = \mathcal{G} (\mathcal{W}^E e_t + \mathcal{W}^H h_{t-1}) \quad (7)$$

In a standard encoder-decoder architecture, the terminal state, which we denote h_T , is a fixed-length vector that encodes a summary of the entire sequence. Attention modifies

²⁴For some models, such as BERT – which we use in most NLP exercises – we will use embeddings based on sub-word units, such as the WordPiece embeddings (Wu et al., 2016). This allows for the use of full words, individual characters, and multi-character strings. The word “growing,” for example, can be split into the word “grow” and the multi-character string “ing.”

the standard LSTM construction by retaining the hidden states and scoring them. It then applies the softmax function to the hidden states, and then multiplies the softmaxed states by the original (untransformed) states. The procedure introduces three additional sets of trainable weights: \mathcal{W}^Q , \mathcal{W}^K , and \mathcal{W}^V .

For each hidden state h_t , there are associated query (q_t), key (k_t), and value (v_t) vectors, which are computed as $\mathcal{W}^i h_t$, where $i \in \{Q, K, V\}$. Stacking those row vectors into matrices \mathcal{Q} , \mathcal{K} , and \mathcal{V} , we can define scaled dot product attention for h_t in Equation (8), where $D = \dim(k_t)$.

$$\mathcal{Z} = \sigma \left(\frac{\mathcal{Q}\mathcal{K}^T}{\sqrt{D}} \right) \mathcal{V} \quad (8)$$

Note that σ is a rowwise softmax function, defined in Equation (9), where X_t is row t in $\mathcal{Q}\mathcal{K}^T/\sqrt{D}$.

$$\sigma(X_{td}) = \exp(X_{td}) / \sum_{d \in D} \exp(X_{td}) \quad (9)$$

The elementwise product of the row vectors \mathcal{Q}_t and \mathcal{K}_s measures the extent to which they are related to or *attend to* each other. Dividing by \sqrt{D} improves computational performance, but is not strictly necessary. Applying the rowwise softmax function σ amplifies the strength of strong associations and also normalizes the sum of the attention weights to be equal to one.

The weights are then multiplied by \mathcal{V} , which is a matrix of linear transformations of the hidden vectors. Each row of the vector \mathcal{Z}_t is a weighted sum of the embedding vectors, where each weight depends on the extent to which a given hidden vector attends to another. Since each hidden vector is most closely associated with a specific embedding in the sequence (i.e., h_t is closest to e_t), attention provides us with a contextualized embedding for each input word. That is, rather than using a fixed embedding, we incorporate the context of other words that are most closely related to it in a sentence.

To be consistent with the subsection that follows on the Transformer model, we have provided a description of how to compute attention for all hidden vectors. However, for

most LSTM-based encoder-decoder models, we will exclusively use \mathcal{Z}_T , which is the row vector associated with the final hidden state, h_T . \mathcal{Z}_T is sometimes called the *context* vector and is concatenated with h_T and passed to the decoder. Similar to the earlier variant of S2S models, such as Bahdanau et al. (2015) and Luong et al. (2015), the decoder takes the output of the encoder as an input and then sequentially predicts symbols until it terminates with an ([END]) token.

Multi-headed Attention. One innovation of the Transformer model is to make use of multi-headed attention, which is enabled by the removal of sequential elements from the model. This amounts to an h -way partition of the query, key, and value matrices, such that $\mathcal{W}^j = \{\mathcal{W}_1^j, \dots, \mathcal{W}_h^j\}$ for $j \in \{Q, K, V\}$, where $\mathcal{W}_i^j \in \mathbb{R}^{D \times D/h}$. For the original transformer model, $M = 512$, $D = 64$, which yields a model with $h = 8$ attention heads. The model then attends to each subspace separately and in parallel, yielding a computational time that is similar to that of a single-headed model.

Positional Encodings. Rather than using sequential elements like recurrence or convolution, the Transformer modifies input embeddings by encoding positional information. For each input embedding in a sequence, positional encodings are generated using Equation (10).

$$p(t, m) = \begin{cases} \sin\left(\frac{t}{10000^{2m/M}}\right) & \text{if } m \text{ even} \\ \cos\left(\frac{t}{10000^{2m/M}}\right) & \text{if } m \text{ odd} \end{cases} \quad (10)$$

The positional encodings are then added to the input embeddings to create positional embeddings, which contain both information about word features and position within the embedding and sequence. The positional embedding is given in Equation (11) and is the input to the model.

$$\left\{ \tilde{e}_1, \dots, \tilde{e}_T \right\} = \left\{ \begin{array}{c} \left[\begin{array}{c} e_{11} + p(1, 1) \\ e_{12} + p(1, 2) \\ \vdots \\ e_{1M} + p(1, M) \end{array} \right] \dots \left[\begin{array}{c} e_{T1} + p(T, 1) \\ e_{T2} + p(T, 2) \\ \vdots \\ e_{TM} + p(T, M) \end{array} \right] \end{array} \right\} \quad (11)$$

A.4 Text Features

Table 15: Mean Text Feature Values by Federal Reserve District

District	Past Focus	Present Focus	Future Focus	Financial Stability	Academic Focus
ATL	-0.04	-0.0	0.04	-0.05	-0.07
BOS	0.09	-0.03	-0.04	-0.03	-0.01
CHI	-0.05	0.0	-0.02	-0.09	-0.1
CLE	0.02	-0.0	-0.07	-0.04	-0.05
DAL	0.11	-0.14	-0.05	-0.11	-0.0
FRB	-0.0	0.03	-0.0	0.07	0.08
KC	0.1	0.03	0.06	0.06	0.02
MIN	0.02	-0.08	-0.06	-0.04	-0.06
NY	0.16	0.17	0.15	0.19	0.19
PHL	0.01	0.01	0.01	0.01	-0.04
RIC	0.09	-0.07	0.21	0.37	0.18
SF	0.09	-0.05	-0.05	-0.14	-0.02
STL	0.06	-0.1	-0.14	-0.05	-0.05

The table above provides the mean values of selected text features for each Federal Reserve district bank over the entire sample period. All text features are measured at the paragraph level. The first column lists the Federal Reserve district. The next three columns provide measures of tense. Note that tense usage is not mutually exclusive and not all passages have a clear focus on a single tense. As such, it is possible for a given district bank to have positive or negative scores for all three tenses. The remaining two columns provide mean values for features that 1) indicate whether a paragraph is about financial stability; and 2) indicate whether a paragraph references academic work or an academic discussion.

Table 16: Mean Text Feature Values by Speaker (1/2)

Speaker	Past Focus	Present Focus	Future Focus	Financial Stability	Academic Focus
alan greenspan	0.13	0.01	-0.08	0.05	0.21
alfred broadus	0.06	-0.08	0.21	0.32	0.09
alfred hayes	0.35	0.1	0.2	0.21	0.29
alice rivlin	0.09	0.09	0.22	-0.01	0.08
anthony santomero	0.01	0.02	0.12	0.1	0.02
ben bernanke	-0.07	-0.07	0.03	0.15	0.16
bob mcteer	0.2	-0.28	-0.14	-0.26	-0.11
braddock hickman	-0.02	-0.12	-0.05	0.13	0.03
bruce maclaury	0.04	-0.06	-0.07	-0.14	-0.11
cathy minehan	0.15	-0.14	0.28	0.37	0.18
charles evans	-0.01	-0.01	0.12	0.05	0.05
charles plosser	-0.04	0.03	0.14	0.16	-0.03
daniel tarullo	0.06	0.03	0.11	0.24	0.14
darryl francis	0.14	-0.05	-0.22	-0.09	-0.03
david eastburn	0.15	-0.07	-0.04	-0.26	-0.09
delos johns	0.12	-0.08	-0.11	-0.01	0.08
dennis lockhart	0.12	0.05	0.08	-0.04	-0.04
donald kohn	0.1	0.01	0.15	0.23	0.24
edward boehne	0.08	-0.02	0.11	0.27	0.11
edward gramlich	-0.06	-0.04	-0.13	0.08	0.18
edward kelley	0.05	0.08	0.2	0.29	0.26
eliot swan	0.2	0.14	0.16	0.06	0.05
elizabeth duke	-0.12	0.01	0.02	0.11	-0.1
eric rosengren	0.08	-0.01	-0.09	-0.1	-0.04
esther george	0.12	0.0	0.08	0.15	0.01
frederic mishkin	-0.1	-0.12	-0.13	-0.02	0.15
frederick deming	-0.06	-0.11	-0.24	-0.07	-0.12
gary stern	0.13	0.06	0.23	0.22	0.17
gerald corrigan	0.03	0.07	0.07	0.11	0.06
harry shuford	-0.18	0.01	-0.27	0.06	-0.15
hugh galusha	0.04	-0.11	-0.16	-0.18	-0.1
jack guynn	0.11	-0.12	0.08	0.1	0.02
james bullard	-0.01	-0.1	0.05	-0.03	0.01
janet yellen	-0.13	-0.02	0.04	0.04	-0.04
jeffrey lacker	0.1	-0.06	0.21	0.38	0.21
jeremy stein	-0.09	-0.05	0.05	0.12	0.1

Notes: The table above provides the mean values of selected text features for each speaker over the entire sample period. All text features are measured at the paragraph level. The first column lists the speaker's name. The next three columns provide measures of tense. Note that tense usage is not mutually exclusive and not all passages have a clear focus on a single tense. As such, it is possible for a given speaker to have positive or negative scores for all three tenses. The remaining two columns provide mean values for features that 1) indicate whether a paragraph is about financial stability; and 2) indicate whether a paragraph references academic work or an academic discussion.

Table 17: Mean Text Feature Values by Speaker (2/2)

Speaker	Past Focus	Present Focus	Future Focus	Financial Stability	Academic Focus
jerome powell	-0.13	0.14	0.13	0.2	0.08
jerry jordan	0.04	-0.04	-0.17	-0.11	-0.08
john balles	0.08	-0.1	-0.03	-0.11	-0.01
john williams	-0.03	-0.03	0.03	-0.03	-0.0
karen horn	-0.39	0.07	-0.04	0.27	-0.1
karl bopp	-0.01	0.01	-0.21	-0.1	-0.08
kevin warsh	0.21	0.01	0.08	0.11	0.24
lael brainard	-0.17	0.03	-0.04	0.27	0.01
laurence meyer	0.03	0.0	0.01	0.17	0.21
lawrence roos	0.14	-0.08	-0.06	-0.26	-0.08
lee hoskins	0.18	-0.02	-0.02	0.02	0.03
loretta mester	-0.18	0.04	0.3	0.31	-0.06
mark olson	-0.08	0.11	-0.03	0.21	0.22
mark willes	0.02	-0.02	-0.07	-0.21	-0.19
michael moskow	-0.05	0.01	-0.05	-0.11	-0.12
monroe kimbrel	0.03	-0.0	-0.09	-0.25	-0.11
narayana kocherlakota	0.0	-0.17	0.1	0.19	0.01
paul volcker	0.01	0.07	-0.06	-0.13	-0.16
randall kroszner	-0.08	0.05	0.01	0.13	0.18
richard fisher	0.09	-0.1	-0.02	-0.08	0.02
robert forrestal	-0.12	0.0	0.05	-0.02	-0.08
robert parry	0.14	-0.0	-0.11	-0.25	-0.06
roger ferguson	0.03	0.05	0.05	0.22	0.27
sandra pianalto	-0.1	0.08	0.01	-0.05	-0.09
sarah raskin	0.09	0.1	0.05	0.02	0.03
stanley fischer	-0.06	0.02	0.02	0.1	0.13
susan bies	-0.02	0.18	0.07	0.22	0.3
susan phillips	-0.0	0.25	0.14	0.25	0.26
theodore roberts	0.06	-0.15	-0.15	0.04	-0.03
thomas hoenig	0.1	0.04	0.06	0.04	0.02
thomas melzer	-0.0	-0.07	-0.16	-0.09	-0.1
timothy geithner	0.15	0.2	0.14	0.17	0.23
wilbur fulton	0.3	-0.2	-0.23	0.08	0.09
william dudley	0.21	0.15	0.2	0.21	0.2
william mcdonough	0.06	0.2	0.1	0.19	0.13
william poole	0.01	-0.2	-0.08	0.09	-0.04
willis winn	0.0	-0.08	-0.1	-0.07	-0.06

Notes: The table above provides the mean values of selected text features for each speaker over the entire sample period. All text features are measured at the paragraph level. The first column lists the speaker's name. The next three columns provide measures of tense. Note that tense usage is not mutually exclusive and not all passages have a clear focus on a single tense. As such, it is possible for a given speaker to have positive or negative scores for all three tenses. The remaining two columns provide mean values for features that 1) indicate whether a paragraph is about financial stability; and 2) indicate whether a paragraph references academic work or an academic discussion.

Recent Working Papers:

For a complete list of Working Papers published by Sveriges Riksbank, see www.riksbank.se

The Macroeconomic Effects of Trade Tariffs: Revisiting the Lerner Symmetry Result <i>by Jesper Lindé and Andrea Pescatori</i>	2019:363
Biased Forecasts to Affect Voting Decisions? The Brexit Case <i>by Davide Cipullo and André Reslow</i>	2019:364
The Interaction Between Fiscal and Monetary Policies: Evidence from Sweden <i>by Sebastian Ankargren and Hovick Shahnazarian</i>	2019:365
Designing a Simple Loss Function for Central Banks: Does a Dual Mandate Make Sense? <i>by Davide Debortoli, Jinill Kim and Jesper Lindé</i>	2019:366
Gains from Wage Flexibility and the Zero Lower Bound <i>by Roberto M. Billi and Jordi Galí</i>	2019:367
Fixed Wage Contracts and Monetary Non-Neutrality <i>by Maria Björklund, Mikael Carlsson and Oskar Nordström Skans</i>	2019:368
The Consequences of Uncertainty: Climate Sensitivity and Economic Sensitivity to the Climate <i>by John Hassler, Per Krusell and Conny Olovsson</i>	2019:369
Does Inflation Targeting Reduce the Dispersion of Price Setters' Inflation Expectations? <i>by Charlotte Paulie</i>	2019:370
Subsampling Sequential Monte Carlo for Static Bayesian Models <i>by David Gunawan, Khue-Dung Dang, Matias Quiroz, Robert Kohn and Minh-Ngoc Tran</i>	2019:371
Hamiltonian Monte Carlo with Energy Conserving Subsampling <i>by Khue-Dung Dang, Matias Quiroz, Robert Kohn, Minh-Ngoc Tran and Mattias Villani</i>	2019:372
Institutional Investors and Corporate Investment <i>by Cristina Cella</i>	2019:373
The Impact of Local Taxes and Public Services on Property Values <i>by Anna Grodecka and Isaiah Hull</i>	2019:374
Directed technical change as a response to natural-resource scarcity <i>by John Hassler, Per Krusell and Conny Olovsson</i>	2019:375
A Tale of Two Countries: Cash Demand in Canada and Sweden <i>by Walter Engert, Ben Fung and Björn Segendorf</i>	2019:376
Tax and spending shocks in the open economy: are the deficits twins? <i>by Mathias Klein and Ludger Linnemann</i>	2019:377
Mind the gap! Stylized dynamic facts and structural models <i>by Fabio Canova and Filippo Ferroni</i>	2019:378
Financial Buffers, Unemployment Duration and Replacement Labor Income <i>by Mats Levander</i>	2019:379
Inefficient Use of Competitors' Forecasts? <i>by André Reslow</i>	2019:380
How Much Information Do Monetary Policy Committees Disclose? Evidence from the FOMC's Minutes and Transcripts <i>by Mikael Apel, Marianna Blix Grimaldi and Isaiah Hull</i>	2019:381
Risk endogeneity at the lender/investor-of-last-resort <i>by Diego Caballero, André Lucas, Bernd Schwaab and Xin Zhang</i>	2019:382
Heterogeneity in Households' Expectations of Housing Prices – Evidence from Micro Data <i>by Erik Hjalmarsson and Pär Österholm</i>	2019:383
Big Broad Banks: How Does Cross-Selling A Affect Lending? <i>by Yingjie Qi</i>	2020:384
Unemployment Fluctuations and Nominal GDP Targeting <i>by Roberto Billi</i>	2020:385
FAQ: How do I extract the output gap? <i>by Fabio Canova</i>	2020:386

Drivers of consumer prices and exchange rates in small open economies <i>by Vesna Corbo and Paola Di Casola</i>	2020:387
TFP news, stock market booms and the business cycle: Revisiting the evidence with VEC models <i>by Paola Di Casola and Spyridon Sichliridis</i>	2020:388
The costs of macroprudential deleveraging in a liquidity trap <i>by Jiaqian Chen, Daria Finocchiaro, Jesper Lindé and Karl Walentin</i>	2020:389
The Role of Money in Monetary Policy at the Lower Bound <i>by Roberto M. Billi, Ulf Söderström and Carl E. Walsh</i>	2020:390
MAJA: A two-region DSGE model for Sweden and its main trading partners <i>by Vesna Corbo and Ingvar Strid</i>	2020:391
The interaction between macroprudential and monetary policies: The cases of Norway and Sweden <i>by Jin Cao, Valeriya Dinger, Anna Grodecka-Messi, Ragnar Juelsrud and Xin Zhang</i>	2020:392
Withering Cash: Is Sweden ahead of the curve or just special? <i>by Hanna Armelius, Carl Andreas Claussen and André Reslow</i>	2020:393
Labor shortages and wage growth <i>by Erik Frohm</i>	2020:394
Macro Uncertainty and Unemployment Risk <i>by Joonseok Oh and Anna Rogantini Picco</i>	2020:395
Monetary Policy Surprises, Central Bank Information Shocks, and Economic Activity in a Small Open Economy <i>by Stefan Laséen</i>	2020:396
Econometric issues with Laubach and Williams' estimates of the natural rate of interest <i>by Daniel Buncic</i>	2020:397
Quantum Technology for Economists <i>by Isaiah Hull, Or Sattath, Eleni Diamanti and Göran Wendin</i>	2020:398
Modeling extreme events: time-varying extreme tail shape <i>by Bernd Schwaab, Xin Zhang and André Lucas</i>	2020:399
The Effects of Government Spending in the Eurozone <i>by Ricardo Duque Gabriel, Mathias Klein and Ana Sofia Pessoa</i>	2020:400
Narrative Fragmentation and the Business Cycle <i>by Christoph Bertsch, Isaiah Hull and Xin Zhang</i>	2021:401
The Liquidity of the Government Bond Market – What Impact Does Quantitative Easing Have? Evidence from Sweden <i>by Marianna Blix Grimaldi, Alberto Crosta and Dong Zhang</i>	2021:402
Five Facts about the Distributional Income Effects of Monetary Policy <i>by Niklas Amberg, Thomas Jansson, Mathias Klein and Anna Rogantini Picco</i>	2021:403
When domestic and foreign QE overlap: evidence from Sweden <i>by Paola Di Casola and Pär Stockhammar</i>	2021:404
Dynamic Macroeconomic Implications of Immigration <i>by Conny Olovsson, Karl Walentin, and Andreas Westermark</i>	2021:405
Revisiting the Properties of Money <i>by Isaiah Hull and Or Sattath</i>	2021:406
The cost of disinflation in a small open economy vis-à-vis a closed economy <i>by Oleksandr Faryna, Magnus Jonsson and Nadiia Shapovalenko</i>	2021:407
On the Performance of Cryptocurrency Funds <i>by Daniele Bianchi and Mykola Babiak</i>	2021:408
The low-carbon transition, climate commitments and firm credit risk <i>by Sante Carbone, Margherita Giuzio, Sujit Kapadia, Johannes Sebastian Krämer, Ken Nyholm and Katia Vozian</i>	2022:409
Seemingly Irresponsible but Welfare Improving Fiscal Policy at the Lower Bound <i>by Roberto M. Billi and Carl E. Walsh</i>	2022:410
Pension Reform and Wealth Inequality: Evidence from Denmark <i>by Torben M. Andersen, Joydeep Bhattacharya, Anna Grodecka-Messi and Katja Mann</i>	2022:411

Inflation Targeting or Fiscal Activism? <i>by Roberto M. Billi</i>	2022:412
Trading volume and liquidity provision in cryptocurrency markets <i>by Daniele Bianchi, Mykola Babiak and Alexander Dickerson</i>	2022:413
DISPERSION OVER THE BUSINESS CYCLE: PASSTHROUGH, PRODUCTIVITY, AND DEMAND <i>by Mikael Carlsson, Alex Clymo and Knut-Eric Joslin</i>	2022:414
Electoral Cycles in Macroeconomic Forecasts <i>by Davide Cipullo and André Reslow</i>	2022:415
The Curious Incidence of Monetary Policy Across the Income Distribution <i>by Tobias Broer, John Kramer and Kurt Mitman</i>	2022:416



Sveriges Riksbank
Visiting address: Brunkebergs torg 11
Mail address: se-103 37 Stockholm

Website: www.riksbank.se
Telephone: +46 8 787 00 00, Fax: +46 8 21 05 31
E-mail: registratorn@riksbank.se