

How You Say It Matters: Text Analysis of FOMC Statements Using Natural Language Processing

By Taeyoung Doh, Sungil Kim, and Shu-Kuei Yang

The Federal Reserve has increasingly used public statements to shape expectations about future policy actions. This practice has become more prevalent since the Great Recession, when the nominal short-term interest rate reached its effective lower bound. To provide further policy easing, the Federal Open Market Committee (FOMC) turned toward explicit forward guidance about the future path of the policy rate as well as the amount and composition of large-scale asset purchases in their post-meeting statements. Although these statements sometimes included quantitative information, such as a specific threshold for the unemployment rate that would make an increase in the federal funds rate appropriate, they also included more nuanced, qualitative descriptions of economic conditions.

Measuring the effects of these qualitative communications is not straightforward. Previous research has found that public communications by policymakers can affect financial market conditions (Eberly, Stock, and Wright 2019). However, most of these studies identify the effects of central bank communications through changes in bond prices during short windows around policy announcements. Although changes in quantitative information such as the policy rate and the pace of

Taeyoung Doh is a senior economist at the Federal Reserve Bank of Kansas City. Sungil Kim is a research associate at the bank. Shu-Kuei Yang is a software engineer at the bank. This article is on the bank's website at www.KansasCityFed.org

asset purchases can be more directly translated into bond prices, it is often difficult to determine the effect of changes in qualitative descriptions of economic conditions. For example, when the FOMC communicates a “subdued outlook for inflation” in its post-meeting statement, it is not immediately clear how market participants might adjust their expectations of future monetary policy actions.

In this article, we use a natural language processing tool to provide a new measure of how changes in qualitative descriptions of the economy in post-meeting statements affect bond prices. These changes may or may not be communicated jointly with changes in quantitative information. Our measure yields two key findings: First, qualitative descriptions of economic conditions and the balance of risk can have as much of an effect on bond prices as quantitative information about the target policy rate. In some cases, information about the factors that played into the Committee’s assessment of economic conditions makes a substantial difference in our measure of the overall tone of the statement—that is, whether it connotes policy easing or tightening relative to the previous statement—and that tone can affect financial market conditions even if no policy action is taken. Second, our new measure of the effects of central bank communications is generally correlated with alternative measures in prior research based solely on bond price data. Our measure is particularly well correlated with medium-term policy expectations.

Section I describes how we assess the tone in FOMC statements using information that staff of the Board of Governors of the Federal Reserve System prepares for the FOMC before each meeting. Section II explains how we quantify the overall tone of the post-meeting statements using a natural language processing tool. Section III shows that qualitative information contained in FOMC statements has significant effects on financial market conditions from March 2004 to December 2014.

I. Assessing the Tone of FOMC Statements

Because FOMC statements can signal the future path of interest rates or plans for large-scale asset purchases, financial market participants watch these statements closely and react to any unexpected information contained in them. This unexpected information is not necessarily quantitative, such as the size of a policy rate change or the

value of intended asset purchases; the overall “tone” of the statement, which does not have a numeric value by itself, can influence market participants as well. For example, market participants might interpret the tone of a statement as optimistic or pessimistic about the future irrespective of the quantitative information it contains and behave accordingly. As a result, identifying the tone of new qualitative information in a monetary policy announcement is critical to anticipating how markets will react to it.

Since March 2004, staff at the Board of Governors of the Federal Reserve System, in consultation with the Chair, has prepared drafts of typically three (but occasionally four) alternative versions of the Committee’s post-meeting statement along with the rationale for each alternative.¹ These alternative statements, referred to as Alt. A, Alt. B, Alt. C, and Alt. D, serve as a basis for the discussion of policy options at FOMC meetings and help the Committee formulate its policy stance. In addition to proposing possible policy actions, the alternative statements describe the economic and financial conditions that might motivate them. Although the official FOMC statement describing the actual policy decision is released at the conclusion of each FOMC meeting, the draft alternatives are released to the public only after five years, along with the meeting transcripts. Documents released with the alternative statements contain additional information about their rationale and possible implications for financial market conditions.²

Each alternative statement has a somewhat different tone and is written to capture a range of possible Committee views on the proper stance of policy. In general, Board staff writes Alt. B as representative of the likely consensus of the Committee, Alt. A as suggestive of an easier policy stance than Alt. B, and Alt. C (or Alt. D, which suggests a tighter policy stance than Alt. C when it is prepared) as suggestive of a tighter policy stance than Alt. B.

The official FOMC statement is not necessarily identical to any of the three draft alternatives. Committee members may interpret the incoming data and appropriate policy stance differently, and they discuss and negotiate the wording in the released statement before and during FOMC meetings. As a result, the official statement may not be exactly the same as any of the three alternative statements.

These semantic differences make alternative statements almost ideal for identifying the tone of the official post-meeting FOMC statement. Because alternative statements intentionally signal “more” or “less”

monetary policy accommodation to financial markets, we can compare the semantic similarity of the official statement with alternative statements to identify its tone. In addition, the detailed discussion of the rationale for the policy positions in alternative statements can help us easily map the semantic differences across statements to differences in the degree of policy accommodation signaled by each statement.

II. Using a Natural Language Processing Model to Quantify the Tone of FOMC Statements

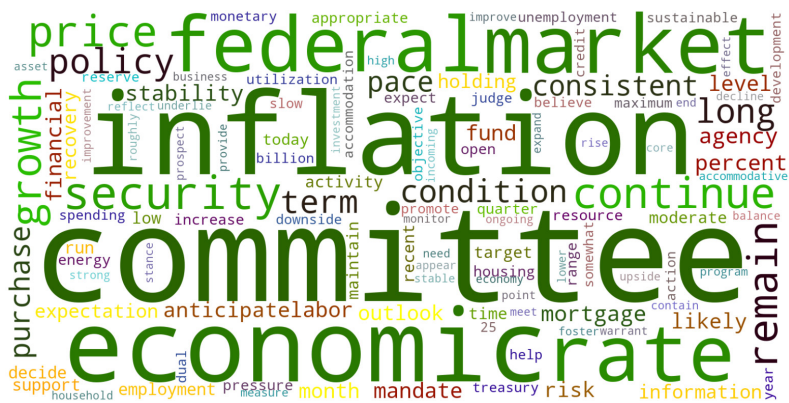
Identifying the tone of FOMC statements requires a model that can capture both quantitative and qualitative information on the policy stance. Although many statistical models can easily classify the policy implications of quantitative information such as a change in the policy rate, classifying qualitative descriptions is not straightforward.

A class of models that uses natural language processing (NLP) can help overcome this challenge by highlighting certain patterns in the distribution of words in a given text or speech to learn the sentiment of those words. The patterns the NLP model highlights can then be used to determine the overall tone of a text. One common strategy is to assess the overall tone of a document by counting the frequency with which words classified as having a positive or negative tone appear in the document. For example, many researchers have classified the tone of documents using a list of words with negative implications in financial contexts developed by Loughran and McDonald (2011).

However, this classification strategy alone may not be useful for FOMC statements. Figures 1 and 2 show word clouds of the most commonly used words in statements that announced policy tightening (Alt. C and Alt. D) versus statements that announced policy easing (Alt. A), respectively, from March 2004 to December 2014.³ Together, the word clouds suggest that FOMC statements may not have sufficient word variation, making it difficult to construct a dictionary of words classified as having specific tones. For example, “inflation” is the second most frequently used word in all the alternative statements, though the contextual meaning of inflation likely differs across these statements.

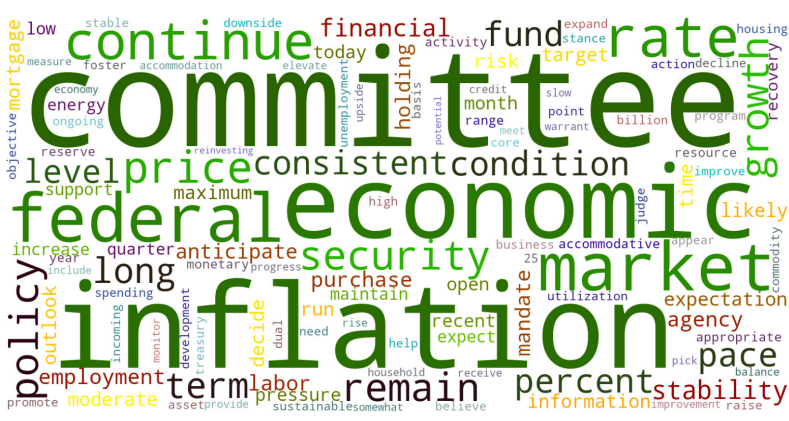
Instead, classifying the tone of a word used in an FOMC statement may require a method that can account for the context in which the word is used. Such a method would be consistent with the

Figure 1
Cloud of Frequently Used Words in Alternative Statements that Suggest Policy Easing



Notes: More frequently used words are represented by larger text size. We count the frequency of words only in Alt. A.
Source: Authors' calculations.

Figure 2
Cloud of Frequently Used Words in Alternative Statements that Suggest Policy Tightening



Notes: More frequently used words are represented by larger text size. We count the frequency of words only in Alt. C and Alt. D.
Source: Authors' calculations.

“distributional representation hypothesis” in linguistics, which suggests that the meaning of a particular word is clarified by the words surrounding it. For example, “broke” in “he broke the law” represents a different, negative sentiment from “broke” in “he broke the world record,” which has a positive sentiment. Simply examining the frequency with which the word “broke” appears in a document would not be sufficient to classify the document’s tone.

Fortunately, recent developments in NLP models allow us to build rich features that can detect the contextual meaning of words above and beyond simple frequency patterns. One such model is the Universal Sentence Encoder (USE) developed by Google researchers, in which a computer algorithm generates a numerical representation of a text document (Cer and others 2018). Specifically, the USE converts any given input document into a numerical vector. Because the USE algorithm is trained to model the meaning of sequences of words, rather than just individual words, its numeric representations are “context-aware”—that is, even the same word can be represented by different numeric values depending on the context in which it is used. Hence, the distance between the numerical representations of two sentences can capture semantic differences between them even if they contain many overlapping words. For example, consider the following three sentences:

- S1) How old are you?
- S2) What is your age?
- S3) How are you?

S1) and S3) have similar words but ask very different questions. In contrast, S1) and S2) have no overlapping words but ask essentially the same question. A model trained to compare the meaning of two sentences based only on the frequency with which similar words appear would erroneously suggest S1) is most similar to S3). However, when the USE is tasked with scoring the similarity of the three sentences, it correctly identifies that S1) is more similar to S2) than S3).

The ability of the USE to discriminate the contextual semantic meaning of words makes it particularly useful for identifying the overall tone of a FOMC statement. Although “inflation” is one of the most frequently used words in all of the alternative statements, its

sentiment—whether it connotes an increase or decrease in the outlook for inflation—changes depending on the words around it or its order in the sentence. As an example, Table 1 compares a sentence from the FOMC statement released after the November 2005 meeting with a counterfactual sentence in which the order of two clauses describing the development in the inflation outlook is swapped. The released version discounts the rise in energy and other costs and ends by emphasizing the stability of core inflation and longer-term inflation expectations. However, the counterfactual version reverses the order, beginning with the stability of core inflation and then tempering that stability with the rise in energy costs. By changing the order of the clauses, the counterfactual paragraph emphasizes the rise in energy and other costs over the stability of core inflation and longer-term inflation expectations. Although humans can recognize these semantic differences qualitatively, they may not agree on the magnitude of the semantic difference across different descriptions of economic conditions. The USE provides an automatic way to quantify such a semantic difference.⁴

As a result, we follow the method detailed in Doh, Song, and Yang (2020) and identify the tone of a post-meeting FOMC statement by comparing its numeric representation from the USE algorithm to the numeric representations of the alternative statements. We classify the tone of the post-meeting statement as less accommodative if it is semantically more similar to Alt. C (or Alt. D if available) than Alt. A. As a first step, we calculate the “similarity score” between the numeric representations of two texts generated by the USE algorithm. The similarity score lies between 0 and 1, where a value of 1 indicates that the two documents are identical. Because we are interested in new information from one FOMC statement relative to another, the *dissimilarity* between documents is more informative for our purpose. Thus, we subtract the similarity score from 1 to yield the “semantic distance” between the documents, which measures how dissimilar they are.

As a second step, we use this measure of semantic distance to, in turn, construct a measure of tone that is bounded between -1 and 1 . If the released statement is indistinguishable from the alternative statement suggesting policy easing based on the numeric representation generated by the USE algorithm, our tone measure takes the value of -1 ; if the released statement is indistinguishable from the alternative statement suggesting policy tightening, our measure takes the value of 1 . In

Table 1

Inflation Language in the November 2005 FOMC Statement

Released version	Counterfactual version
The cumulative rise in energy and other costs has the potential to add to inflation pressures; however, core inflation has been relatively low in recent months and longer-term inflation expectations remain contained.	Core inflation has been relatively low in recent months and longer-term inflation expectations remain contained; however, the cumulative rise in energy and other costs has the potential to add to inflation pressures.

Source: Board of Governors of the Federal Reserve System.

addition, we measure the semantic distance between each post-meeting FOMC statement and the previous post-meeting FOMC statement to capture new information in the more recent statement. We then follow Ke, Kelly, and Xiu (2019) and multiply this measure of new information by our measure of tone to generate a measure of the change in the policy stance from the previous FOMC meeting.

Because monetary policy actions are transmitted through financial markets to the real economy, policymakers are interested in assessing the effect of central bank communications on financial markets. Economic theory suggests financial markets respond only to unexpected information—also known as the “surprise component” of a statement. As a result, we need a measure of the surprise component of the tone of a statement to isolate the change in bond prices that occurred in response to the released statement. Unfortunately, we do not have a measure of the market’s expectation of the tone of the upcoming FOMC statement comparable to asset price data. However, we can reasonably assume that the market expectation of the statement’s tone will lie somewhere between the tone of the less accommodative and more accommodative statements. Thus, some weighted average of the tones in these two statements will provide the market expectation of the tone in the upcoming post-meeting statement. We can then subtract this measure from the estimated tone of the official statement to get a measure of the surprise component of the statement.

Although these weights are unknown, they can be estimated by examining how bond yields change during a narrow event window around a policy announcement (10 minutes before and after the announcement). When the tone of a released FOMC statement is unexpectedly less accommodative, bond yields increase; when the tone is unexpectedly more accommodative, bond yields decline. If FOMC

statements with similar tones induce different market responses, we attribute the difference to a difference in market expectations. Although this method looks similar to the method in Swanson (2020) and Nakamura and Steinsson (2018) of identifying policy surprises based purely on the high-frequency responses of bond markets around policy announcements, we can measure the magnitude of surprises not just by the size of bond market responses but also by information in FOMC statements. For example, we can study what part of the statement can make it more or less accommodative in tone, which drives the market's reaction to FOMC announcements.

III. Empirical Analysis of FOMC Statements (March 2004–December 2014)

Using the USE algorithm and following the methods detailed in Doh, Song, and Yang (2020), we find that the overall tone of a statement has as least as much of an effect on financial market conditions as quantitative information such as the size of the rate cut. This analysis is based on a sample of publicly available alternative versions of FOMC statements for 87 scheduled FOMC meetings from March 2004 through December 2014. When multiple versions of alternative statements that suggest policy tightening or policy easing exist, we use the most extreme version (for example, Alt. D instead of Alt. C) to identify the tone of the released statement.

To illustrate our results in more detail, we consider two examples in which the NLP model provides additional insights on the policy stance. Our first example is from September 2007, when alternative FOMC statements differed not only in the size of the rate cut suggested but also in their characterizations of the balance of risk. Table 2 shows a summary of the policy decisions, stated rationale, and risk assessment of all four alternative statements. Based only on the decision for the target policy rate, the alternative statements appear to get progressively less accommodative as they move from A to D: Alt. A and B suggest a 50 basis point cut in the federal funds rate target, Alt. C suggests a 25 basis point cut, and Alt. D suggests no change to the federal funds rate target at all. However, the statements also differ meaningfully in later paragraphs that describe the outlook and the balance of risk: first, Alt. B and Alt. C provide the same description of the outlook, while Alt.

Table 2
Alternative Language for the September 2007 FOMC Statement

Component	Alternative A	Alternative B	Alternative C	Alternative D
Policy decision	Lower the federal funds rate target by 50 basis points	Lower the federal funds rate target by 50 basis points	Lower the federal funds rate target by 25 basis points	Keep the federal funds rate target
Rationale	Tighter credit conditions and the intensification of the housing correction appear likely to exert appreciable restraint on economic growth. Moreover, the potential for significant spillovers from credit market disruptions to business and household spending poses a risk to the outlook.	Economic growth was moderate during the first half of the year, but the tightening of credit conditions has the potential to intensify the housing correction and to restrain economic growth more generally.	Economic growth was moderate during the first half of the year, but the tightening of credit conditions has the potential to intensify the housing correction and to restrain economic growth more generally.	Economic growth was moderate during the first half of the year. Financial market conditions have deteriorated in recent weeks, leading to tighter credit and an intensification of the housing correction. These developments have the potential to restrain growth in economic activity. Nonetheless, the economy seems likely to continue to expand in a moderate pace over coming quarters, supported by solid growth outside the housing sector and a robust global economy.
Assessment of risk	Even after today's action, the Committee judges that the downside risks to economic growth outweigh the upside risks to inflation.	The Committee will continue to closely follow timely indicators of economic prospects and will act as needed to foster price stability and sustainable economic growth.	Even after today's action, the Committee judges that the downside risks to economic growth outweigh the upside risks to inflation.	In the current circumstances, the Committee judges that the downside risks to economic growth are now roughly balanced by the upside risks to inflation.

Source: Board of Governors of the Federal Reserve System.

A (Alt. D) provides a more pessimistic (optimistic) outlook; second, Alt. A and Alt. C acknowledge that “the downside risks to economic growth outweigh the upside risks to inflation,” while Alt. B and Alt. D sound neutral regarding the risk assessment. Thus, even though Alt. C suggests a smaller rate cut than Alt. B, it may not necessarily be less accommodative in tone overall.

To assess how this qualitative risk assessment affects the tone of the released FOMC statement, we compare the similarity scores among alternative statements based on their USE representations. Importantly, the released FOMC statement adopted the rate cut and risk assessment language from Alt. B, so we compare the similarity scores of the released statement with the other alternatives. Table 3 shows that the released statement was semantically more similar to Alt. C than Alt. A, even though the released statement and Alt. A cut the rate by 50 basis points, while Alt. C cut the rate by 25 basis points. This result suggests the qualitative description of the outlook matters as much as quantitative decisions on the rate cut in determining the overall tone of the statement.

In addition, Table 3 shows that the stated rationale for the rate decision generates semantic differences across alternative statements even when the rate decision itself is the same. Although Alt. A and Alt. B both cut the rate by 50 basis points, Alt. A points to significant spillovers from credit markets to business and household spending and emphasizes downside risks. In doing so, Alt. A describes an outlook consistent with additional easing in the future relative to Alt. B, generating significant semantic differences between the two alternative statements. In this sense, how the Committee signals future policy actions matters as much as its current policy action.⁵

Our second example is from October 2013, when alternative statements differed in their interpretations of the strength in incoming data. As Table 4 shows, Alt. A downplays the signal from strong incoming data by describing the effect of the temporary government shutdown on the interpretation of the data; other versions do not have that sentence. The paragraphs describing incoming data are otherwise similar across alternative statements. Accordingly, Table 5 shows that when we apply similarity scoring based only on the frequency of words used in the text, the sentiment of the alternative statements does not appear to differ much. However, their USE representations are quite different,

Table 3

Similarity Scoring of Alternative Statements for the September 2007 FOMC Meeting

Similarity (A, C)	Similarity (A, FOMC)	Similarity (C, D)	Similarity (C, FOMC)
0.990	0.968	0.897	0.983

Note: FOMC denotes the post-meeting statement, and A, C, and D denote alternative statements.
Source: Doh, Song, and Yang (2020).

Table 4

Alternative Language Describing Outlook for the October 2013 FOMC Meeting

Alternative A	FOMC	Alternative C
<p>The effects of the temporary shutdown of the federal government have made the evolution of economic conditions during the intermeeting period somewhat more difficult to assess. However, information received since the Federal Open Market Committee met in September generally suggests that economic activity has been expanding at a modest pace. Indicators of labor market conditions have shown some further improvement but the unemployment rate remains elevated. Available data suggest that household spending and business fixed investment advanced, but that the recovery in the housing sector has slowed in response to higher mortgage rate. Fiscal policy is restraining economic growth. Apart from fluctuations due to changes in energy prices, inflation has been running below the Committee's longer-run objective, even though longer-term inflation expectations have remained stable.</p>	<p>Information received since the Federal Open Market Committee met in September generally suggests that economic activity has continued to expand at a moderate pace. Indicators of labor market conditions have shown some further improvement, but the unemployment rate remains elevated. Available data suggest that household spending and business fixed investment advanced, while recovery in the housing sector slowed somewhat in recent months. Fiscal policy is restraining economic growth. Apart from fluctuations due to changes in energy prices, inflation has been running below the Committee's longer-run objective, but longer-term inflation expectations have remained stable.</p>	<p>Information received since the Federal Open Market Committee met in September generally suggests that economic activity has continued to expand at a moderate pace. Indicators of labor market conditions have some further improvement; in particular, the unemployment rate, though still elevated, has continued to decline. Household spending and business fixed investment advanced, and the housing sector has continued to strengthen, even though mortgage rates have risen on balance in recent months and fiscal policy is restraining economic growth. Apart from fluctuations due to changes in energy prices, inflation has been running somewhat below the Committee's longer-run objective, but longer-term inflation expectations have remained stable.</p>

Note: FOMC denotes the post-meeting statement, and alternatives A and C denote alternative statements.
Source: Board of Governors of the Federal Reserve System.

Table 5
Similarity Scoring of Alternative Statements for the October 2013 FOMC Meeting

Algorithm	Similarity (A, FOMC)	Similarity (C, FOMC)
USE	0.895	0.990
Word frequency	0.975	0.972

Note: FOMC denotes the post-meeting statement, and A and C denote alternative statements.
 Source: Doh, Song, and Yang (2020).

Table 6
Correlation with Policy Surprise Measures from Swanson (2020)

Policy instrument	Correlation of surprise in Doh, Song, and Yang (2020) and Swanson (2020)
Federal funds rate	0.20
Forward guidance	0.52
Asset purchases	-0.12

Notes: Asset purchase shocks in Swanson (2020) are normalized to associate a negative number with policy tightening comparable to the increase in the interest rate, so the sign is flipped even though the underlying relation should be similar.

Sources: Swanson (2020) and Doh, Song, and Yang (2020).

as the sentence on the effect of the temporary government shutdown changes the contextual meaning of the first paragraph substantially.⁶ This example shows once again that changing qualitative descriptions of the outlook can materially change the tone of the FOMC statement.

Because monetary policy influences the real economy through financial markets, it is important to know how the tone we identify using text analysis is transmitted to asset prices, including stock returns and bond returns. To answer this question, we compare our measure of the tone of FOMC statements with measures from previous research based on high-frequency asset price data only. In these measures, an upward jump in Treasury yields during a narrow window around an announcement is associated with a surprisingly less accommodative tone, and a downward movement in Treasury yields is associated with a surprisingly more accommodative tone. Our measure of tone is highly correlated with these other measures, suggesting the tone we identify using alternative statements matches the tone financial markets perceive (Doh, Song, and Yang 2020).

In addition to incorporating qualitative information, our measure has another advantage over measures in previous research. When bond yields with different maturities respond differently to FOMC

announcements, we can use this cross-sectional difference to decompose policy surprises along multiple dimensions. For example, we can assess how announcements related to asset purchases influence longer-term (10-year and longer) bond yields and how forward guidance about the future path of interest rates influences near- to medium-term (two- to five-year) bond yields. We compare our text-based policy surprise measure with multiple dimensions of monetary policy surprises in Swanson (2020) constructed in this way. Table 6 shows that our measure is highly correlated with forward guidance on the future path of the interest rate, and modestly correlated with large-scale asset purchases. Given that our measure is sensitive to alternative descriptions of the outlook, this finding suggests that market participants associate changes in the outlook mostly with medium-term policy expectations.

Conclusion

Central bank communications about current and future policy actions have increasingly received attention as a policy tool. Many studies evaluate the effects of quantitative information from policy statements based on the response of bond prices during narrow event windows around the release of FOMC statements. However, this approach may not fully capture the potential effect of qualitative descriptions of the economy or the rationale for a policy decision contained in these statements.

We overcome the limitation of event studies based on asset price responses by using a natural language processing tool to directly quantify the policy stance from texts. Specifically, we identify the tone of a post-meeting statement by quantifying how close it is semantically to alternative versions of statements, whose more or less accommodative tones can be determined based on the rationale given for each alternative in FOMC documents. Our analysis of post-meeting FOMC statements illustrates an important role for qualitative statement language. In particular, we find that information about the Committee's assessment of risk is as important as a quantitative decision on the target policy rate. In addition, we find that including information about how the Committee interprets incoming data in a statement can substantially influence a statement's tone. Our results suggest that *how* the Committee describes its rationale for the policy action may matter as much as the policy action itself.

Endnotes

¹Before 2004, the Board staff prepared alternative draft language for some parts of the statement (such as “maintaining [increasing/reducing] the federal funds rate”) but did not provide alternative versions of the full statement. See Board of Governors of the Federal Reserve System (2004).

²The Board declassifies documents after five years and publicly releases them each January. Thus, as of November 2020, the latest available alternative statements were those prepared for the December 2014 meeting and included in a document called “Tealbook,” which had a “Part A” and a “Part B.” Tealbook Part A was formerly called “Greenbook,” while Part B was called “Bluebook.” Occasionally, Board staff provides a fourth policy alternative, Alt. D. In these cases, Alt. D is intended to be less accommodative than Alt. C. When available, we use Alt. D instead of Alt. C to identify the tone of the post-meeting statement.

³Lucca and Trebbi (2009) classify the tone of words used in the post-meeting FOMC statement by the systematic co-frequency of those words with “hawkish” (associated with policy tightening) or “dovish” (associated with policy easing) in documents found by the Google search engine. However, the contextual meaning of words depends on word ordering and various qualifiers not easily detected by this method.

⁴The similarity scoring is 1 based on word frequency but 0.991 in the USE. Although the difference does not seem to be material even in the USE, the similarity score is generally high and less variable in the USE. For example, the average similarity score between the released post-meeting statement and Alt. A is 0.977, and the standard deviation is 0.022. Hence, the difference between 1 and 0.991 can still make a meaningful difference once we construct our measure of the tone based on the counterfactual paraphrase.

⁵Gürkaynak, Sack, and Swanson (2005) highlight this point using high-frequency bond price changes as instruments for the effect of changing language in FOMC statements.

⁶The alternative statements have other slight differences—for instance, Alt. A adds a more dovish tone to the threshold of forward guidance on the duration of the effective lower bound. However, when we calculate the semantic differences paragraph by paragraph, we find that the first paragraph is most critical in generating the semantic difference between alternative statements.

References

- Board of Governors of the Federal Reserve System. 2004. "Monetary Policy Alternatives." *Bluebook*, March 11.
- Cer, Daniel, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St. John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, Yun-Hsuan Sung, Brian Strope, and Ray Kurzweil. 2018. "Universal Sentence Encoder." Association for Computational Linguistics, *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, Brussels, November. Available at <https://doi.org/10.18653/v1/D18-2029>
- Doh, Taeyoung, Dongho Song, and Shu-Kuei Yang. 2020. "Deciphering Federal Reserve Communication via Text Analysis of Alternative FOMC Statements." Federal Reserve Bank of Kansas City, Research Working Paper no. 20-14, October. Available at <https://doi.org/10.18651/RWP2020-14>
- Eberly, Janice C., James H. Stock, and Jonathan H. Wright. 2019. "The Federal Reserve's Current Framework for Monetary Policy: A Review and Assessment." National Bureau of Economic Research, working paper no. 26002, June. Available at <https://doi.org/10.3386/w26002>
- Gürkaynak, Refet S., Brian Sack, and Eric T. Swanson. 2005. "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statement." *International Journal of Central Banking*, vol. 1, no. 1, pp. 55–93. Available at <https://doi.org/10.2139/ssrn.633281>
- Ke, Zheng Tracy, Bryan T. Kelly, and Dacheng Xiu. 2019. "Predicting Returns with Text Data." National Bureau of Economic Research, working paper no. 26186, November. Available at <https://doi.org/10.3386/w26186>
- Loughran, Tim, and Bill McDonald. 2011. "When Is a Liability Not a Liability? Textual Analysis, Dictionaries, and 10-Ks." *Journal of Finance*, vol. 66, no. 1, pp. 35–65. Available at <https://doi.org/10.1111/j.1540-6261.2010.01625.x>
- Lucca, David O., and Francesco Trebbi. 2009. "Measuring Central Bank Communication: An Automated Approach with Application to FOMC Statements." National Bureau of Economic Research, working paper no. 15367, September. Available at <https://doi.org/10.3386/w15367>
- Nakamura, Emi, and Jón Steinsson. 2018. "High-Frequency Identification of Monetary Non-Neutrality: The Information Effect." *Quarterly Journal of Economics*, vol. 133, no. 3, pp. 1283–1330. Available at <https://doi.org/10.1093/qje/qjy004>
- Swanson, Eric T. 2020. "Measuring the Effects of Federal Reserve Forward Guidance and Asset Purchases on Financial Markets." *Journal of Monetary Economics*, forthcoming. Available at <https://doi.org/10.1016/j.jmoneco.2020.09.003>