
**WORKING PAPER
N. 160
MARCH 2022**

**Monetary policy and financial markets:
evidence from Twitter traffic**

By Donato Masciandaro, Davide Romelli and Gaia Rubera

This Paper can be downloaded without charge from The Social Science
Research Network Electronic Paper Collection:
<http://ssrn.com/abstract=3823185>

Monetary policy and financial markets: evidence from Twitter traffic*

Donato Masciandaro

Bocconi University

Davide Romelli[†]

Trinity College Dublin

Gaia Rubera

Bocconi University

March 25, 2022

Abstract

The monetary policy announcements of three major central banks: the European Central Bank, the US Federal Reserve and the Bank of England trigger significant discussions about monetary policy on Twitter. Using machine learning techniques we identify Twitter messages related to monetary policy around the release of policy decisions and build an hourly measure of similarity between tweets on monetary policy and the transcripts of announcements. We show that large changes in the similarity between tweets and central bank communication around the time of the announcement correspond to spikes in stock market volatility and jumps in sovereign yields. These findings suggest that social media discussions on central bank communication are a good proxy for monetary policy surprises.

JEL Classifications: E44, E52, E58, G14, G15, G41.

Keywords: monetary policy, central bank communication, financial markets, social media, Twitter, US Federal Reserve, European Central Bank, Bank of England

* We thank Adriana Manna, Jacopo Maria Madurno and Romano Vincenzo Tarsia for excellent research assistance, and Marianna Blix Grimaldi, Michael Ehrmann, Laurent Ferrara, Etienne Farvaque and seminar participants at the Computational and Financial Econometrics 2020 Virtual Conference, Western Economic Association International 2021 Virtual conference, University of Lille, Banque de France Real-Time Data Analysis, Methods and Applications Virtual Conference, 3rd Monash-Warwick-Zurich Text-as-Data Workshop for helpful discussions, comments, and suggestions. Davide Romelli gratefully acknowledge financial support from the 2020 Trinity College Dublin Arts and Social Sciences Benefactions Fund. Any errors or omissions are solely those of the authors.

[†]Corresponding author: Department of Economics, Trinity College Dublin, Arts Building, Dublin 2, Ireland. email: romellid@tcd.ie.

“The actions of central banks are no longer cloaked in silence, and perhaps never will be again. Whereas in the past silence was seen as a guarantee of independence, today this is achieved by giving an explicit account of one’s actions.”

—PAOLO BAFFI, GOVERNOR OF THE BANK OF ITALY, 31 MAY 1979

1 Introduction

Central banks’ actions have never been more closely monitored than nowadays, when social media enables monetary policy announcements to be communicated widely to the public. While central bank press releases are known to be closely monitored by financial market participants, there is less research into understanding whether the wider public’s reactions to central bank communication is aligned with financial markets’ reactions. In this paper, we present a novel approach to examining the link between social media traffic and asset prices volatility and returns following the release of monetary policy decisions. Using a high-frequency event study approach, we investigate whether changes in the similarity between Twitter messages and policy decisions in the two hours surrounding an announcement are systematically related to asset price reactions.

To conduct our analysis, we collect all tweets related to monetary policy in the four days around the monetary policy announcements of three of the most important central banks of the world, i.e. the European Central Bank (ECB), the U.S. Federal Reserve Bank (Fed) and the Bank of England (BoE), between January 2011 and February 2020. As the initial sample of close of half a million Twitter messages also contains tweets that are not related to monetary policy decisions, we use a Machine Learning algorithm to identify all tweets that directly refer to these events. This strategy allows us to identify a sample of 228,348 tweets that discuss topic related to the monetary policy announcement of the three central banks in our sample.

Using a natural language processing algorithm, we then compute a hourly measure of the similarity between the text of the central bank announcement and Twitter traffic on the monetary policy stance. We use this measure of similarity to infer the alignment of Twitter users with the monetary policy announcements before and after monetary policy events. Next, we use the time variation in text similarity to estimate the change in Twitter

users' perception around a policy announcement.

We show that changes in similarity matter and have an important effect on financial markets volatility and returns. In particular, announcements characterised by higher changes in the measure of similarity before and after the announcement are associated with higher stock market volatility and absolute returns, particularly following ECB press conferences. We also find that changes in tweets similarity are only linked to higher stock market volatility following announcements made by the Federal Reserve Bank, but no effect is found for the Bank of England.

Our results also show a link between changes in similarity and sovereign yields. We find that changes in similarity following ECB press conferences are associated to larger changes in sovereign yields' realised variance and absolute returns for the four major Euro area countries included in our sample, i.e. France, Germany, Italy and Spain. The effect is stronger for longer-term sovereign bonds, suggesting that the market surprise captured by the changes in tweets similarity is more likely to be reflected in longer maturity assets. Similar results are obtained when looking at the sovereign yields volatility and absolute returns following press releases made by the Bank of England, while little evidence is found for US Treasury yields. Overall, the results suggest that social media reactions can be a good proxy for monetary policy surprises.

This work is related to a growing literature that studies the effects of central bank communication on financial markets. For instance, [Gürkaynak et al. \(2005\)](#) show that the Fed's monetary policy actions have important but differing effects on asset prices, with statements having a much greater impact on longer-term Treasury yields. Similar results have been found by [Brand et al. \(2006\)](#) for the ECB case. More recently, [Jarociński and Karadi \(2020\)](#) show the differential impact of information about monetary policy and the central bank's assessment of the economic outlook on interest rates and stock prices. These works, however, obtain information on monetary policy surprises by extracting factors from changes in the yields of risk-free rates at different maturities or by looking at the co-movement of interest rates and stock prices around policy announcements. To the best of our knowledge, this is the first paper to propose an indirect measure of monetary policy surprise captured by social media reactions to monetary policy announcements and show its link to asset price volatility and returns.

At the same time, research related to textual similarity of central bank communication has only looked at the similarity between subsequent statements, normally taking place every one or two months, and not between tweets and statements, as is the focus of this paper. For example, [Acosta and Meade \(2015\)](#) document how the similarity of FOMC post-meeting statements has increased over time. [Amaya and Filbien \(2015\)](#) find similar results for the case of the ECB. [Ehrmann and Talmi \(2020\)](#) use the variation in the drafting process at the Bank of Canada and document an increase in market volatility when substantial changes in press releases occur after a period of similar statements.

The link between Twitter messages on monetary policy and market reactions has been studied recently in relation to US President Trump’s tweets on US monetary policy ([Camous and Matveev, 2019](#); [Bianchi et al., 2019](#)). [Lüdering and Tillmann \(2020\)](#) analyse Twitter messages during the Taper tantrum period and find that shocks to the share of discussions related to the “tantrum”, “QE” and “data” are associated with significant changes in asset prices. Recently, Twitter has also been used to study the communication policies of European central banks ([Korhonen and Newby, 2019](#)) or to build real-time measures of consumers’ inflation expectations ([Angelico et al., 2021](#)).

However, one important contribution of our paper with respect to the existing research is the use of Twitter data to study asset prices reactions around monetary policy announcements. To the best of our knowledge, the only paper in economics which uses our methodological approach is [Giavazzi et al. \(2020\)](#), which compute measures of textual similarity between the tweets of German voters and the ones of the main German parties. As such, we provide a new method to understand the extent to which changes in the similarity between social media discussions and policy decisions around monetary policy announcements are associated to asset price volatility and returns, which might affect the way monetary policy decisions are communicated to the markets and via social media.

The remainder of this paper is structured as follows. A review of the related literature is provided in Section 2. Section 3 introduces our database of central bank communication events, discusses the Twitter data, presents the methodology used to construct the measure of similarity and describes the intraday data on equity and sovereign bonds. Section 4 presents the empirical findings, while Section 5 concludes.

2 Related literature

Our paper is related to two strands of literature. The first deals with the importance and effects of central bank communication. In the 1970s and 1980s, central banks were shrouded in monetary mystique and secrecy (Goodfriend, 1986). However, the development of modern monetary policy theory naturally produced a shift in communication from secrecy towards transparency (Eijffinger and Masciandaro, 2014) and central bank communication has gained momentum (Blinder et al., 2008). As a matter of fact, all central banks in advanced economies have taken major steps to incorporate communication strategies into their decision-making processes (Ehrmann and Fratzscher, 2005).

The increased importance of communication for policy makers is mirrored in the rapid development of the academic literature on this topic. This literature sheds light on the impact of central bank communication on macroeconomic variables, such as exchange rates (Jan Jansen and De Haan, 2004; Fratzscher, 2008; Conrad and Lamla, 2010; Gürkaynak et al., 2021), interest rates (Gürkaynak, 2005; Gürkaynak et al., 2005; Lucca and Trebbi, 2009; Hayo and Neuenkirch, 2011; Lamla and Sturm, 2013; Neuenkirch, 2013; Altavilla et al., 2014; Lucca and Moench, 2015; Altavilla et al., 2019; Hansen et al., 2019), asset prices (Hayo et al., 2010; Rosa, 2011; Cieslak and Schrimpf, 2019; Ehrmann and Talmi, 2020; Gorodnichenko et al., 2021) and real variables (Hansen and McMahon, 2016), as well as future monetary policy decisions (Bennani et al., 2020). Apart from focusing on the impact of monetary policy announcements, this literature has also stressed the importance of focusing communication on other aspects, such as inflation (Čihák et al., 2012) or financial stability (Born et al., 2014; Correa et al., 2021).

Another important aspect of central bank communication strategy is its consistency. Jansen and de Haan (2013) analyse whether the ECB uses consistent language in its communication and find an overall consistency, even though its communication seems flexible enough to adapt to changing circumstances. Acosta and Meade (2015) study the similarity of FOMC post-meeting statements and show that they have become more similar over time, especially since the global financial crisis. Nevertheless, FOMC statements have also become more complex since the onset of unconventional monetary policy, as shown by Hernández-Murillo et al. (2014). More recently, language processing algorithms have been used to identify differences between subsequent FOMC statements (Doh et al., 2020).

Together with consistency, the role of language is also crucial (see [Gerlach, 2004](#); [Hansen and McMahon, 2016](#); [Kawamura et al., 2016](#), among others). In this context, computational linguistic tools have been used to analyse monetary policy communication ([Bailey and Schonhardt-Bailey, 2008](#); [Lucca and Trebbi, 2009](#); [Schonhardt-Bailey, 2013](#); [Hansen and McMahon, 2016](#); [Hansen et al., 2018](#); [Schmeling and Wagner, 2019](#); [Hubert and Fabien, 2017](#); [Bailliu et al., 2021](#)).

The second burgeoning literature to which this paper contributes is the one on social media interactions and monetary policy. A part of this literature has focused on the use of social media as a further communication tool by central banks. [Korhonen and Newby \(2019\)](#) examine the extent to which European central banks maintain an institutional Twitter account and analyse their tweeting activity. They find that central banks' Twitter activity has no relation to citizens' online participation and that communication on financial stability has increased more in comparison to the one on monetary policy. Looking at the United States, [Conti-Brown and Feinstein \(2020\)](#) undertake the first systematic analysis of the Fed's participation on Twitter and find that the Fed is more engaged on Twitter than other independent agencies. [Gorodnichenko et al. \(2021\)](#) analyse the Federal Reserve System communication on Facebook and Twitter and its effectiveness. In the case of the Fed, Twitter appears to be more popular and gains greater public engagement. They show that market participants do update their inflation expectations based on information contained in the Fed's social media posts. However, they find no evidence of stock market reactions to the Fed communication on social media.

Other work has focused on tweets about monetary policy made by Twitter users, other than central banks themselves. For example, [Azar and Lo \(2016\)](#) create a new dataset of tweets that cite the Federal Reserve to understand how investors on social media behave around FOMC meeting dates. Their results suggest that tweets do contain information which can be used to predict returns and to build portfolios that outperform the benchmark market portfolio. [Meinusch and Tillmann \(2015\)](#) and [Lüdering and Tillmann \(2020\)](#) analyse the Federal Reserve's *taper tantrum* period during April and October 2013, and capture information on the debate among market professional during this period. Their results show that both the revisions of expectations of market participants and shocks to selected topics discussed in the tweets lead to significant changes on U.S. bond yields, ex-

change rates and stock prices. Similarly, [Stiefel and Vivès \(2019\)](#) study the extent to which changes in the belief about an intervention of the ECB during the summer 2012 explain the sudden reduction in government bonds spreads for distressed countries in the euro area. Finally, [Ehrmann and Wabitsch \(2022\)](#) analyse tweets about the ECB to understand the extent to which its communication is received by non-experts and how it affects their views. Their results suggest that Twitter also serves as a platform for controversial discussions about monetary policy.

To the best of our knowledge, our paper is the first to investigate the effect of social media interaction on asset prices in the hours surrounding the monetary policy announcements.

3 Tweets on Monetary Policy as Market Sentiment Metrics: Methodology and Data

In this section, we provide information on the steps followed for the construction of the database on monetary policy announcements and Twitter messages, which we then use to compute the similarity between tweets and central bank communication.

3.1 Monetary policy communication

We first create a database of time-stamped communication on monetary policy decisions by three central banks: the European Central Bank, the U.S. Federal Reserve Bank and the Bank of England. Our sample period runs from January 2011 through February 2020.¹

As our empirical analysis focuses on the variation in tweets similarity around monetary policy announcements, we collect information on scheduled policy announcements from central bank websites. This choice is motivated by the fact that scheduled events might be able to attract social media traffic both before and after announcements, while this is not possible for some of the policy news announced during unscheduled events.

Our database includes: 1) 89 monetary policy decisions (MPDs) made by the ECB at

¹The decision to start our analysis in 2011 is motivated by the limited availability of tweets before 2011, while we choose to stop our analysis in February 2020 to exclude the extraordinary measures taken by central banks since the starting of the Covid-19 pandemic.

13:45 Frankfurt time and 89 transcripts of the press conference which begins at 14:30 and ends at 15:30; 2) 71 MPDs statements released immediately after the policy decision taken by the Fed during FOMC meetings and 3) 94 MPDs announcements released by the BoE at 12:00 London time.²

3.2 Tweets on monetary policy

This section briefly describes the procedure followed for the extraction of the full sample of tweets between January 2011 and February 2020.

3.2.1 Identification of the tweets of interest

We use the “Get Old Tweets” module in Python, to collect all Twitter messages related to monetary policy and published between 48 hours prior to a scheduled monetary policy announcement to 48 hours after it. In order to identify the tweets of interest, we first conduct an extensive analysis of the keywords and hashtags used by Twitter users to discuss monetary policy decisions for a sample of 10 monetary policy announcements. Next, we automatise the selection of tweets by collecting all Twitter messages that: (a) mentioned the official Twitter account of the central bank, e.g. @bankofengland; (b) contained a hashtag followed by the central bank’s acronym, e.g. #ecb; or (c) contained a hashtag followed by the surname of the chair of the central bank, e.g. #yellen.³ Table 1 presents an overview of the keywords used to extract tweets.⁴ The overall number of tweets collected during this first round of the selection process is 467,777.

Table 1: Overview of the keywords used for Twitter messages extraction

Central bank	Keywords			
European Central Bank	@ecb	#ecb	#trichet	#draghi #lagarde
Federal Reserve	@federalreserve	#fed	#bernanke	#yellen, #powell
Bank of England	@bankofengland	#boe, #bankofengland	#carney	

²Since the FOMC meeting of March 19-20, 2013, FOMC statements are released at 14:00 New York time. Before this date, Fed press statements were released at either 12:30 or 14:15. The exact timing of each press release has been taken into consideration for the extraction on the associated tweets. The timing of these events has been double-checked with the database provided in [Cieslak and Schrimpf \(2019\)](#).

³Given the high number of tweets potentially associated with the surname of the former Governor of the Bank of England, Mervyn King (#King), we decided to exclude this hashtag from the search.

⁴The hashtag #interestrates might have been an ideal candidate for capturing discussions related to monetary policy. However, we excluded it from the set of keywords as it would have been automatically associated to all three banks and might have only created noise.

3.2.2 Selection of relevant tweets


As not all the tweets collected might be considered relevant, we trained a Machine Learning algorithm on a manually labelled training set to isolate relevant tweets. To do so, we first selected a random sample of 3,000 tweets and we asked two research assistants to independently classify tweets as relevant, i.e. related to monetary policy announcements, or irrelevant. Details on the guiding principle for the selection of the relevant tweets are presented in Appendix A. At the end of the classification process, we considered as relevant the following set of tweets: 1) considered relevant by both the research assistants and the authors of the paper; or 2) classified as relevant by one of the two research assistant, and validated by the authors. This screening process allowed us to identify 782 relevant, i.e. 26% of the sub-sample, and 2,218 irrelevant tweets.

After the manual labelling of the sub-sample of tweets, we used the classification to train an algorithm able to identify all relevant tweets within the entire corpus of 467,777 messages. By doing so, we were able to identify 228,348 tweets which focused their discussion on the monetary policy decisions of the European Central Bank, the Federal Reserve Bank and the Bank of England in each of their scheduled announcements between January 2011 and February 2020. Figure 1 shows the distribution of Twitter messages in the two days around monetary policy announcements. Not surprisingly, the number of tweets spikes in the 2 hours surrounding a press release and this provides us with a first confirmation that the supervised classifier that we have used to train the algorithm achieves a clear identification of relevant tweets.

Looking at the unrelated tweets, we also find that some of the tweets that contained one of the acronyms indicated in Table 1, e.g. #ecb, were in a language different from English and we removed them from our sample. Eliminating these tweets had an additional advantage, as some of the pre-processing steps needed for the next steps of the analysis relied on pre-existing dictionaries that were developed only for the English language.

We also gather some anecdotal evidence on other irrelevant Twitter messages. Firstly, there seems to be a shared attitude of hostility towards central banks, especially in the US. More specifically, many tweets in our sample complained about what they call “Fed interventionism”. In their words “#Government interference with #InterestRates distorts accuracy of vital information, increasing mistakes of market participants”. Also, we found

Figure 1: Twitter traffic and Central bank communication



Note: The figure shows the distribution of the overall number of Twitter messages related to monetary policy created in the hours (± 48 hours) surrounding a central bank scheduled communication. The solid line refers to the tweets related to the ECB, the dashed line to those mentioning the Fed, while the dashed-dotted marks refer to the BoE.

many tweets revealing the presence of a sub-culture grounded in the far-right tradition of “libertarian conservatism” among US Twitter users. Such a community complains not only for the interventionism of the central bank, but also for the very existence of central banks as institutions, especially the FED (dubbed “the creature from Jekyll Island”). Though we are reporting these information here, we did not label any of these tweets as relevant.

3.3 Twitter similarity metrics

After the collection of all relevant monetary policy announcements and tweets, we compute an hourly similarity index between the Twitter messages on monetary policy and the transcripts of announcements by transforming the two corpus of text into vectors with doc2vec, a deep learning technique. Details on the pre-processing and technique used are reported in Appendix B. Here, we briefly summarize the method.

For each hour surrounding a monetary policy announcement, we create documents on: i) central bank transcripts and ii) tweets. Transcript contain the text of the monetary policy decision released at a specific date and time by one of the three central banks in our sample, while the tweet documents aggregate the text of all the tweets related to monetary policy published in a given hour around an announcement.


Given these documents, we use doc2vec (Le and Mikolov, 2014), an unsupervised deep learning algorithm that learns how to represent each document with a unique vector. We then measure similarity between documents as the cosine of the angle between the two corresponding vectors, i.e. the normalized inner product of the two vectors, for a certain central bank press release p and a certain group of tweets on monetary policy m at hour h of day t :

$$\cos \theta_{mp_{t,h}} = \frac{\overrightarrow{m_{t,h}} \cdot \overrightarrow{p_{t,h}}}{\|\overrightarrow{m_{t,h}}\| \|\overrightarrow{p_{t,h}}\|}. \quad (1)$$

In order to control whether the computed similarity is indeed a valid measure of the evolution of Twitter users' discussion on monetary policy announcements, we test the validity of our similarity measure by looking at six representative announcements (two for each central bank) in our sample. These events are: 1) the launch of the OMT programme by the ECB on September 6, 2012; 2) the adoption of the expanded asset purchase programme on January 22, 2015 by the ECB; 3) the Fed announcement on the program to taper its bond-buying program on December 18, 2013; 4) the Fed decision to raise US interest rates for the first time since 2006 on December 16, 2015; 5) the BoE announcement to cut rates for the first time since 2009 on August 4, 2016; and 6) the BoE warning of the possibility of earlier and larger rate hikes for the UK on February 8, 2018. The evolution of the similarity measure around these events is reported in Figure 2.

In line with anecdotal evidence, the two announcements made by the ECB had been largely anticipated by the market participants and the general public. The low volatility of the similarity measure in the hours surrounding these events support this evidence. In addition, we can notice that the similarity measure in the hour of the announcement ($h=0$) is higher for the ECB than for the events reported for both the Fed and the BoE. Indeed, the 4 selected announcements for these latter central banks had been less anticipated by the markets and this brought to an higher variation of the similarity measure around monetary policy events.

Figure 2: Similarity measure validation: key events



Note: The figure shows the evolution of the similarity measure in the 8 hours around six selected monetary policy decisions. The shaded area marks the time span that goes from the hour prior to an announcement until the hour after it.

3.4 High-frequency data and asset price volatility and returns

Our empirical analysis aims at understanding the link between changes in similarity and asset prices' volatility and returns using high frequency data. To do so, we extract high-frequency one-minute data on stock market indices and sovereign yields from Refinitiv. The data availability and the coverage of maturities for government bonds differs by country. For the euro area, we have data on stock market indices for France, Germany, Italy and Spain as well as the EURO STOXX50 index and the EURO STOXX Banks Index (SX7E), which is the stock price index for the biggest banks in the euro area. We also obtained sovereign yields with maturities ranging from 1 to 30 years for these four major Euro area countries. For the US, we have the stock price indices for the Dow Jones, Nasdaq and S&P 500 and Treasury yields with maturities of 2, 5, 10 and 30 years, while for the UK we have high-frequency data for the FTSE 100 stock market index and Gilts yields with maturities

of 1, 2, 5, 10, 15 and 30 years.

Since our goal is to assess the sensitivity of asset prices to central bank communication, we first compute the realised variance of stock returns and sovereign yields around monetary policy announcements. Let τ denote the time of a communication event, and $\tau^- = \tau - h^-$ and $\tau^+ = \tau + h^+$ the time before and after the event. If we divide the interval $h^+ + h^-$ into N sub-intervals of length $\Delta = \frac{h^+ + h^-}{N}$, then the Realized Variance (RV) of asset prices around event τ is computed as:

$$\text{RV}_\tau(\tau^-, \tau^+, N) = \sum_{i=0}^N r_{\tau+i\Delta}^2 \quad (2)$$

where $r_{\tau+i\Delta}^2 = (p_{\tau+i\Delta} - p_{\tau+(i-1)\Delta})^2$ and p is the log of an asset price. In our baseline estimations we construct the realized variance by summing up the squared value of the one-minute returns over an event window: from 15 min before to 15 min after a monetary policy decision (for example between 13:30 and 14:00 for the European Central Bank press releases).⁵

An alternative way to assess the sensitivity of asset prices to monetary policy announcements is to look at their returns. Given that our measure of changes in similarity does not capture information on tweets sentiment, we focus our attention on the absolute value of returns. In particular, we compute the absolute value of returns following [Altavilla et al. \(2019\)](#) who measure returns as the percentage variation in the median price between the 15-25 minutes following a press release and the 10-20 minutes prior to it. Following their approach, returns associated to ECB press conferences are computed using the median price in the 14:15-14:25 interval as the pre-conference window and the median price in the 15:40-15:50 interval as the post-conference window.

⁵For the ECB press conference, we compute the realised volatility between 14:15, i.e. 15 minutes before the conference, and 15:45, i.e. 15 minutes after the end of the event.

4 Monetary Policy, tweets and asset price volatility and returns

This section identifies how changes in the similarity of tweets around monetary policy announcements affect asset price volatility and returns.

4.1 Asset price volatility

We start by presenting the high-frequency identification strategy that exploits the link between changes in similarity and asset price volatility. The estimation takes the following form:

$$RV_{\tau,MPD} = \alpha_y + \beta_1 \Delta \text{Similarity}_{h,MPD} + \epsilon_{\tau,MPD}; \quad (3)$$

where $RV_{\tau,MPD}$ is the realised variance of returns around the time τ of a communication event MPD , i.e. -15 and +15 minutes around an announcement. The main explanatory variable is the change in the similarity measure between the hour post and prior to a monetary policy decision, computed as $\Delta \text{Similarity}_{h,MPD} = \text{Similarity}_{h,MPD} - \text{Similarity}_{h-1,MPD}$. In addition, we add year fixed effects, α_y , to absorb common time-variation in asset prices reactions to monetary policy announcements within a year. To avoid assigning excessive weight to monetary policy events that attracted limited social media traffic, we use a weighted least squares approach, weighting each event by the number of tweets in the hours surrounding a monetary policy announcement.

4.1.1 Stock market volatility

Table 2 shows the realised variance of the stock market indices for the four biggest economies of the Euro area, i.e. France, Germany, Italy and Spain, and two Euro area stock indices for blue chip companies (STOXX50E) and banks (SX7E) in a 30-minute window around each announcement.

As discussed in [Altavilla et al. \(2019\)](#), the ECB policy decisions are announced in two separate steps. At 13:45 Central European Time (CET) a brief press release summarizes the policy decision without providing any explanation and rationale for the decision. Then, at 14:30 CET the ECB President reads the introductory statement, which explains the

Table 2: Changes in similarity on Euro area stock market indices volatility

Panel A: Press release window						
	(1)	(2)	(3)	(4)	(5)	(6)
	CAC 40	DAX	FTSE MIB	IBEX	STOXX50E	SX7E
Δ Similarity	0.534 (2.311)	0.280 (1.851)	2.132 (1.822)	1.156 (3.182)	1.194 (2.196)	2.473 (3.003)
Observations	89	89	89	89	89	89
R-squared	0.267	0.257	0.369	0.199	0.312	0.404

Panel B: Press conference window						
	(1)	(2)	(3)	(4)	(5)	(6)
	CAC 40	DAX	FTSE MIB	IBEX	STOXX50E	SX7E
Δ Similarity	4.757*** (1.269)	4.060*** (1.043)	5.692*** (1.709)	5.414*** (1.651)	5.533*** (1.429)	9.671*** (2.920)
Observations	89	89	89	89	89	89
R-squared	0.462	0.467	0.362	0.314	0.478	0.344

Note: The dependent variable is the realised variance of the stock market indices of major Euro area countries: CAC 40 for France, DAX for Germany, FTSE MIB for Italy and IBEX for Spain, as well as the EURO STOXX50 (STOXX50E) and EURO STOXX Banks (SX7E) indices for European blue chip companies and banks in the 30-minute window around each event. Δ Similarity is the change in the similarity index between Twitter traffic and the policy announcement in the hour post as compared to the hour prior to the release of a policy decision. Year fixed-effects dummies are included, but not reported. Robust standard errors in parentheses. ***, **, and * denote statistical significance at the 1%, 5% and 10% levels, respectively.

rationale behind the decision. Usually, the introductory statement is read out in about 15 minutes and the conference continues with a follow up question-and-answer session of the ECB President with journalists that lasts for about 45 minutes. Until December 2014, press releases only provided information related to policy rates decisions, disregarding therefore announcements on non-standard measures. Between January 2015 and January 2016 press releases mentioned the adoption of further measures, but did not provide its details, which were announced during the press conference. Finally, starting from March 2016, the content of the decisions on non-standard policy measures has been summarized in the press release, but all the details were provided during the introductory statement to the press conference. This staggered procedure, motivates our decision to provide two estimates for all results related to the ECB: one for the press release window (panel A) and one for the press conference window (Panel B).

The results reported in Panel A of Table 2 suggest that changes in similarity around press releases are not associated with variations in volatility in European stock market indices. Consistent with the idea that more information are provided during the ECB

Table 3: Changes in similarity on US and UK stock market indices volatility

	United States			United Kingdom
	(1)	(2)	(3)	(4)
	Dow Jones	Nasdaq	S&P 500	FTSE 100
Δ Similarity	0.755* (0.392)	0.715** (0.351)	0.745** (0.370)	0.064 (0.048)
Observations	71	71	71	94
R-squared	0.571	0.611	0.572	0.251

Note: The dependent variable is the realised variance of US and UK stock market indices in the 30-minute window surrounding a monetary policy announcement. Δ Similarity is the change in the similarity index between Twitter traffic and the policy announcement in the hour post as compared to the hour prior to the release of a policy decision. Year fixed-effects dummies are included in all specifications. Robust standard errors in parentheses. ***, **, and * denote statistical significance at the 1%, 5% and 10% levels, respectively.

press conference, the results presented in Panel B show a positive and strongly statistically significant coefficient for the changes in similarity across all estimations. This suggests that a large change in the similarity of tweets and policy measures before and after the press conference is associated with higher stock market volatility. These results support the idea that press conferences and in particular, Q&A sessions might facilitate market participants' information processing and be associated with higher trading activity (Hayo et al., 2020).


Table 3 reports estimates for the US and UK stock market indices. The results are consistent for the Federal Reserve press releases, as stock market volatility spikes when the change in similarity around the event is larger. Announcements made by the Bank of England, on the other hand, do not seem to affect the realised variance of the FTSE 100 Index.

4.1.2 Sovereign yields volatility

Next, we estimate the impact of changes in similarity on sovereign yields' realised variance. Figures 3 and 4 summarise the results for French, German, Italian and Spanish sovereign yields at different maturities.⁶ These figures show some interesting patterns. Similar to the results shown in Table 2, Figure 3 shows that changes in the measure of similarity do not seem to affect sovereign bonds volatility in the 30-minute window surrounding ECB


⁶See Appendix Tables C.1-C.4 for information on the estimations obtained for each country at different maturities.

Figure 3: Changes in similarity and European sovereign yields volatility during ECB press releases



Note: The figure show the coefficient of $\Delta\text{Similarity}$ in Eq. (3). The dependent variable is the volatility of European sovereign yields during ECB press release window [13:30–14:00]. Year fixed effects are included. 90% confidence intervals are presented.

Figure 4: Changes in similarity and European sovereign yields volatility during ECB press conferences



Note: The figure show the coefficient of $\Delta\text{Similarity}$ in Eq. (3). The dependent variable is the volatility of European sovereign yields during ECB press conference window [14:15–15:45]. Year fixed effects are included. 90% confidence intervals are presented.

press releases. The only exception in this case is represented by the realised variance of the German sovereign yields with a 20-year maturity, where larger changes in similarity are associated with higher volatility. On the other end, the regression results in Figure 4 highlight a positive and statistically significant effect of changes in similarity on sovereign yields volatility at longer maturities, i.e. from 5 to 30 years. Importantly, the coefficient of interest increases in magnitude for sovereign bonds characterized by longer maturities. This evidence suggests that the market surprise captured by the change in the measure of similarity is reflected more in the volatility of longer maturity bonds.

Table 4: US and UK sovereign yield volatility

Panel A: United States						
	(1)	(2)	(3)	(4)		
	2-year	5-year	10-year	30-year		
Δ Similarity	-0.000 (0.001)	0.009** (0.004)	0.022 (0.014)	0.027 (0.060)		
Observations	71	71	62	71		
R-squared	0.130	0.361	0.228	0.239		

Panel B: United Kingdom						
	(1)	(2)	(3)	(4)	(5)	(6)
	1-year	2-year	5-year	10-year	15-year	30-year
Δ Similarity	-0.005 (0.013)	0.009 (0.006)	0.055** (0.022)	0.201*** (0.074)	0.366** (0.165)	0.908 (0.609)
Observations	94	94	94	94	94	94
R-squared	0.408	0.155	0.372	0.348	0.282	0.455

Note: The dependent variable is the realised variance of United States and United Kingdom sovereign yields at different maturities in the 30-minute window surrounding a monetary policy announcement. Δ Similarity is the change in the similarity index between Twitter traffic and the policy announcement in the hour post as compared to the hour prior to the release of a policy decision. Year fixed-effects dummies are included in all specifications. Robust standard errors in parentheses. ***, **, and * denote statistical significance at the 1%, 5% and 10% levels, respectively.

The results for sovereign yield volatility of US and UK government bonds are reported in Panel A and Panel B of Table 4, respectively. Consistent with the results presented in Figure 4, the magnitude of the coefficients increases at longer maturities. However, we find that changes in the measure of tweets similarity surrounding the Fed monetary policy announcements have a limited effect on Treasuries volatility. While the result, as the coefficient is only statistically significant for 5-year bonds.

Overall, the estimates presented in this section suggest that changes in the similarity

of tweets related to monetary policy around announcements are associated with higher sovereign yields volatility, particularly for the the ECB press conferences and the BoE press releases with larger effects on longer term maturities.

4.2 Asset price returns

The results presented so far focused on asset price volatility. In this section, we explore the link between changes in the measure of similarity and asset price returns. The estimation takes the following form:

$$|r|_{\tau,MPD} = \alpha_y + \beta_1 \Delta \text{Similarity}_{h,MPD} + \epsilon_{\tau,MPD}; \quad (4)$$

where $|r|_{\tau,MPD}$ is the absolute value of returns obtained by computing the absolute percentage variation in the median price between the 15-25 minutes following a press release and the 10-20 minutes prior to it.⁷ Similar to the estimations presented in section 4.1, we also add year fixed effects, α_y , and use a weighted least squares approach, weighting each event by the number of tweets in the hours surrounding a monetary policy announcement.

4.2.1 Stock market returns

Table 5 reports the estimations on the impact of changes in tweets similarity using high frequency data on European stock market indices. While the results presented in Table 2 suggested that large changes in similarity were associated to higher stock market volatility during ECB press conferences, the results presented in Table 5 do not show a statistically significant relationship between changes in similarity and stock market returns. In particular, we only find a positive effect on the absolute value of returns for the Spanish stock market index and the Stoxx index for Banks around the press releases window.

Similarly, the results presented in Table 6 show the absence of any link between changes in tweets similarity and stock markets returns for both the United States and the United Kingdom, around the monetary policy announcements made by the Federal Reserve Bank

⁷As both stock market indices and sovereign yields can experience positive or negative returns following monetary policy announcement. We focus our analysis on the absolute value of returns as our measure of tweets similarity does not capture information on the direction of monetary policy decisions or changes in Twitter users sentiment, but only how close tweets are related to monetary policy announcements.

Table 5: Changes in similarity on Euro area stock market indices returns

Panel A: Press release window						
	(1)	(2)	(3)	(4)	(5)	(6)
	CAC 40	DAX	FTSE MIB	IBEX	STOXX50E	SX7E
Δ Similarity	2.088 (1.612)	1.169 (1.270)	3.259 (2.099)	3.235* (1.806)	2.455 (1.739)	6.208** (2.779)
Observations	89	89	89	89	89	89
R-squared	0.328	0.334	0.447	0.420	0.370	0.564

Panel B: Press conference window						
	(1)	(2)	(3)	(4)	(5)	(6)
	CAC 40	DAX	FTSE MIB	IBEX	STOXX50E	SX7E
Δ Similarity	1.107 (1.032)	1.216 (1.058)	1.024 (1.020)	0.290 (0.810)	1.142 (1.056)	-0.033 (1.079)
Observations	89	89	89	89	89	89
R-squared	0.129	0.106	0.163	0.163	0.106	0.230

Note: The dependent variable is the absolute value of returns of the stock market indices of major Euro area countries, i.e. CAC 40 for France, DAX for Germany, FTSE MIB for Italy and IBEX for Spain, as well as the EURO STOXX50 and EURO STOXX Banks indices for the Euro area using high-frequency one-minute data. Returns are computed as the percentage variation in the median price between the 15-25 minutes following a press release and the 10-20 minutes prior to it. Δ Similarity is the change in the similarity index between Twitter traffic and the policy announcement in the hour post as compared to the hour prior to the release of a policy decision. Year fixed-effects dummies are included, but not reported. Robust standard errors in parentheses. ***, **, and * denote statistical significance at the 1%, 5% and 10% levels, respectively.

Table 6: US and UK stock market indices returns


	United States			United Kingdom
	(1)	(2)	(3)	(4)
	Dow Jones	Nasdaq	S&P 500	FTSE 100
Δ Similarity	-0.203 (0.907)	0.240 (0.751)	-0.046 (0.839)	-0.335 (0.478)
Observations	71	71	71	94
R-squared	0.180	0.208	0.174	0.296

Note: The dependent variable is the absolute value of returns of US and UK stock market indices using high-frequency one-minute data. Returns are computed as the percentage variation in the median price between the 15-25 minutes following a press release and the 10-20 minutes prior to it. Δ Similarity is the change in the similarity index between Twitter traffic and the policy announcement in the hour post as compared to the hour prior to the release of a policy decision. Year fixed-effects dummies are included in all specifications. Robust standard errors in parentheses. ***, **, and * denote statistical significance at the 1%, 5% and 10% levels, respectively.

and the Bank of England, respectively. Overall, these results suggest that larger changes in similarity are not associated to large returns between the pre and post-event windows.

4.2.2 Sovereign bonds returns

Figure 5: Changes in similarity on European sovereign bonds returns during ECB press releases




Note: The figure summarizes the effect of changes in similarity on the absolute change of European sovereign bond returns during the ECB press release window. The absolute value of yield changes is computed as the absolute value of the percentage variation in the median price between the 15-25 minutes following a press release and the 10-20 minutes prior to it, i.e. using the median price in the 14:00-14:10 and 13:25-14:35 intervals, respectively. Year fixed effects are included. 90% confidence intervals are presented.

The benchmark estimates presented in Figures 3-4 showed how larger changes in tweets similarity around ECB press conferences were associated to higher sovereign yield volatility. We now turn to sovereign bond returns.⁸ Figure 5 show that changes in tweets similarity are not associated to sovereign bond returns around ECB press releases. However, when we focus our attention to the press conference window (Figure 6), we observe a positive relationship between changes in similarity and the absolute value of sovereign bond returns. This effect becomes statistically significant for sovereign bonds characterized by longer maturities.

Finally, Table 7 presents the estimates for the absolute change in sovereign yields for the United States and the United Kingdom. The results presented in Panel A confirm previous findings, suggesting that changes in similarity are not associated to changes in sovereign yields, following monetary policy announcements made by the Federal Reserve Bank. For the Bank of England, the coefficients are positive and statistically significant

⁸See Appendix Tables C.5-C.8 for information on the estimations obtained for each country at different maturities.

Figure 6: Changes in similarity on European sovereign bond returns during ECB press conferences



Note: The figure summarizes the effect of changes in similarity on the absolute value of European sovereign bond returns during the ECB press conference window. The absolute value of returns is computed as the absolute value of the percentage variation in the median price between the 15-25 minutes following a press conference and the 10-20 minutes prior to it, i.e. in the 15:40-15:50 and 14:15-14:25 intervals, respectively. Year fixed effects are included. 90% confidence intervals are presented.

for all maturities except the 30-years one. Consistent with our previous results, we find that higher market surprises are associated with large absolute price change in sovereign bonds with larger effects on securities with maturities up to 20 years.

Table 7: Changes in similarity on US and UK sovereign bond returns

Panel A: United States				
	(1)	(2)	(3)	(4)
	US 2Y	US 5Y	US 10Y	US 2Y 30y
Δ Similarity	-0.029 (0.018)	0.042 (0.059)	0.046 (0.121)	-0.229 (0.257)
Observations	67	71	71	71
R-squared	0.165	0.180	0.166	0.119

Panel B: United Kingdom						
	(1)	(2)	(3)	(4)	(5)	(6)
	UK 1Y	UK 2Y	UK 5Y	UK 10Y	UK 15Y	UK 30Y
Δ Similarity	0.039* (0.023)	0.145** (0.060)	0.488** (0.189)	1.000*** (0.362)	1.082*** (0.410)	0.788 (0.641)
Observations	94	94	94	94	94	94
R-squared	0.514	0.443	0.429	0.426	0.419	0.451

Note: The dependent variable is the absolute change in United States and United Kingdom sovereign bonds returns at different maturities using high frequency one-minute data. Returns are computed as the percentage variation in the median price between the 15-25 minutes following a press release and the 10-20 minutes prior to it. Δ Similarity is the change in the similarity index between Twitter traffic and the policy announcement in the hour post as compared to the hour prior to the release of a policy decision. Year fixed-effects dummies have been included, but not report. Robust standard errors in parentheses. ***, **, and * denote statistical significance at the 1%, 5% and 10% levels, respectively.

5 Conclusion

In this paper, we propose a novel approach to identify the role of monetary policy communication by examining reactions on social media to central bank announcements and how these are related to financial market volatility and returns. In particular, we employ machine learning techniques to compute measures of textual similarity between press releases and tweets about the monetary policy decisions of three major central banks: the European Central Bank, the Federal Reserve Bank and the Bank of England.

Our results point to a significant role of changes in market expectations in the hours surrounding monetary policy announcements. We find that large changes in the similarity between tweets and monetary policy decision before and after announcements are associated with higher bond and stock market volatility and returns, suggesting that our Twitter-based measure is a good proxy for monetary policy surprises.

The novel data and empirical strategy in this paper also sheds new light on the importance of changes in social media discussions and market reactions following monetary policy announcements, something which could potentially impact the channels of central bank communication.

References

- Acosta, M. and E. E. Meade (2015). Hanging on every word: Semantic analysis of the fomc’s post-meeting statement. Technical report, Board of Governors of the Federal Reserve System (US).
- Altavilla, C., L. Brugnolini, R. S. Gürkaynak, R. Motto, and G. Ragusa (2019). Measuring euro area monetary policy. *Journal of Monetary Economics* 108, 162–179.
- Altavilla, C., D. Giannone, and M. Lenza (2014). The financial and macroeconomic effects of omt announcements.
- Amaya, D. and J.-Y. Filbien (2015). The similarity of ECB’s communication. *Finance Research Letters* 13, 234–242.
- Angelico, C., J. Marcucci, M. Miccoli, and F. Quarta (2021, February). Can we measure inflation expectations using Twitter? Temi di discussione (Economic working papers) 1318, Bank of Italy, Economic Research and International Relations Area.
- Azar, P. D. and A. W. Lo (2016). The wisdom of twitter crowds: Predicting stock market reactions to fomc meetings via twitter feeds. *The Journal of Portfolio Management* 42(5), 123–134.
- Bailey, A. and C. Schonhardt-Bailey (2008). Does deliberation matter in fomc monetary policymaking? the volcker revolution of 1979. *Political Analysis*, 404–427.
- Bailliu, J., X. Han, B. Sadaba, M. Kruger, et al. (2021). Chinese monetary policy and text analytics: Connecting words and deeds. Technical report, Bank of Canada.
- Bennani, H., N. Fanta, P. Gertler, and R. Horvath (2020). Does central bank communication signal future monetary policy in a (post)-crisis era? the case of the ecb. *Journal of International Money and Finance* 104, 102167.
- Bhatia, S., J. H. Lau, and T. Baldwin (2016). Automatic labelling of topics with neural embeddings. *arXiv preprint arXiv:1612.05340*.
- Bianchi, F., T. Kind, and H. Kung (2019). Threats to central bank independence: High-frequency identification with twitter. Technical report, National Bureau of Economic Research.
- Blinder, A. S., M. Ehrmann, M. Fratzscher, J. De Haan, and D.-J. Jansen (2008). Central bank communication and monetary policy: A survey of theory and evidence. *Journal of economic literature* 46(4), 910–45.
- Born, B., M. Ehrmann, and M. Fratzscher (2014). Central bank communication on financial stability. *The Economic Journal* 124(577), 701–734.
- Brand, C., D. Buncic, and J. Turunen (2006). The impact of ECB monetary policy decisions and communication on the yield curve.
- Camous, A. and D. Matveev (2019). Furor over the fed: Presidential tweets and central bank independence. Technical report, Bank of Canada.

- Cieslak, A. and A. Schrimpf (2019). Non-monetary news in central bank communication. *Journal of International Economics* 118, 293–315.
- Čihák, M., D.-J. Jansen, and A. Bulíř (2012). Clarity of central bank communication about inflation.
- Conrad, C. and M. J. Lamla (2010). The high-frequency response of the eur-usd exchange rate to ECB communication. *Journal of Money, Credit and Banking* 42(7), 1391–1417.
- Conti-Brown, P. and B. D. Feinstein (2020). Twitter and the federal reserve. *Brookings Center on Regulation & Markets Working Paper*.
- Correa, R., K. Garud, J. M. Londono, and N. Mislav (2021). Sentiment in central banks’ financial stability reports. *Review of Finance* 25(1), 85–120.
- Doh, T., D. Song, S.-K. X. Yang, et al. (2020). Deciphering federal reserve communication via text analysis of alternative fomc statements. Technical report.
- Ehrmann, M. and M. Fratzscher (2005). Central bank communication: different strategies, same effectiveness. *ECB Working*.
- Ehrmann, M. and J. Talmi (2020). Starting from a blank page? semantic similarity in central bank communication and market volatility. *Journal of Monetary Economics* 111, 48–62.
- Ehrmann, M. and A. Wabitsch (2022). Central bank communication with non-experts â a road to nowhere? *Journal of Monetary Economics*.
- Eijffinger, S. and D. Masciandaro (2014). *Modern Monetary Policy and Central Bank Governance*. Edward Elgar Publishing.
- Fratzscher, M. (2008). Communication and exchange rate policy. *Journal of Macroeconomics* 30(4), 1651–1672.
- Gerlach, S. (2004). Interest rate setting by the ECB: words and deeds.
- Giavazzi, F., F. Ighaut, G. Lemoli, and G. Rubera (2020). Terrorist attacks, cultural incidents and the vote for radical parties: Analyzing text from twitter. Technical report, National Bureau of Economic Research.
- Goodfriend, M. (1986). Monetary mystique: Secrecy and central banking. *Journal of Monetary Economics* 17(1), 63–92.
- Gorodnichenko, Y., T. Pham, and O. Talavera (2021). The voice of monetary policy. Technical report, National Bureau of Economic Research.
- Gorodnichenko, Y., T. Pham, O. Talavera, et al. (2021). ‘liked’, ‘shared’, ‘commented’: Central bank communication on facebook and twitter. Technical report.
- Gürkaynak, R. S. (2005). Using federal funds futures contracts for monetary policy analysis.

- Gürkaynak, R. S., A. H. Kara, B. Kısacıkoglu, and S. S. Lee (2021). Monetary policy surprises and exchange rate behavior. *Journal of International Economics* 130, 103443.
- Gürkaynak, R. S., B. P. Sack, and E. T. Swanson (2005). Do actions speak louder than words? the response of asset prices to monetary policy actions and statements. *International Journal of Central Banking* 1, 55–93.
- Hansen, S. and M. McMahon (2016). Shocking language: Understanding the macroeconomic effects of central bank communication. *Journal of International Economics* 99, S114–S133.
- Hansen, S., M. McMahon, and A. Prat (2018). Transparency and deliberation within the fomc: a computational linguistics approach. *The Quarterly Journal of Economics* 133(2), 801–870.
- Hansen, S., M. McMahon, and M. Tong (2019). The long-run information effect of central bank communication. *Journal of Monetary Economics* 108, 185–202.
- Hayo, B., K. Henseler, M. S. Rapp, and J. Zahner (2020). Complexity of ecb communication and financial market trading.
- Hayo, B., A. M. Kutan, and M. Neuenkirch (2010). The impact of us central bank communication on european and pacific equity markets. *Economics Letters* 108(2), 172–174.
- Hayo, B. and M. Neuenkirch (2011). Canadian interest rate setting: The information content of canadian and us central bank communication. *Southern Economic Journal* 78(1), 131–148.
- Hernández-Murillo, R., H. Shell, et al. (2014). The rising complexity of the fomc statement. *Economic Synopses* (23).
- Hubert, P. and L. Fabien (2017). Central bank sentiment and policy expectations.
- Jan Jansen, D. and J. De Haan (2004). Look who’s talking: ECB communication during the first years of emu. *De Nederlandsche Bank, manuscript*.
- Jansen, D.-J. and J. de Haan (2013). An assessment of the consistency of ECB communication. *Central Bank Communication, Decision Making, and Governance: Issues, Challenges, and Case Studies*, 183.
- Jarociński, M. and P. Karadi (2020). Deconstructing monetary policy surprises—the role of information shocks. *American Economic Journal: Macroeconomics* 12(2), 1–43.
- Kawamura, K., Y. Kobashi, M. Shizume, and K. Ueda (2016). Strategic central bank communication: discourse and game-theoretic analyses of the bank of japan’s monthly report.
- Korhonen, I. and E. Newby (2019). Mastering central bank communication challenges via twitter. Technical report, BoF Economics Review.

- Lamla, M. J. and J.-E. Sturm (2013). Interest rate expectations in the media and central bank communication. *Central Bank Communication, Decision Making, and Governance: Issues, Challenges, and Case Studies*, 101–111.
- Le, Q. and T. Mikolov (2014). Distributed representations of sentences and documents. In *International conference on machine learning*, pp. 1188–1196. PMLR.
- Lucca, D. O. and E. Moench (2015). The pre-fomc announcement drift. *The Journal of Finance* 70(1), 329–371.
- Lucca, D. O. and F. Trebbi (2009). Measuring central bank communication: an automated approach with application to fomc statements. Technical report, National Bureau of Economic Research.
- Lüdering, J. and P. Tillmann (2020). Monetary policy on twitter and asset prices: Evidence from computational text analysis. *The North American Journal of Economics and Finance* 51, 100875.
- Meinusch, A. and P. Tillmann (2015). Quantitative easing and tapering uncertainty: Evidence from twitter. Technical report, Joint Discussion Paper Series in Economics.
- Mikolov, T., K. Chen, G. Corrado, and J. Dean (2013). Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Neuenkirch, M. (2013). Monetary policy transmission in vector autoregressions: A new approach using central bank communication. *Journal of Banking & Finance* 37(11), 4278–4285.
- Rosa, C. (2011). Words that shake traders: The stock market’s reaction to central bank communication in real time. *Journal of Empirical Finance* 18(5), 915–934.
- Schmeling, M. and C. Wagner (2019). Does central bank tone move asset prices? *Available at SSRN 2629978*.
- Schonhardt-Bailey, C. (2013). *Deliberating American monetary policy: A textual analysis*. MIT Press.
- Stiefel, M. and R. Vivès (2019). ‘whatever it takes’ to change belief: Evidence from twitter.

Appendices

A Criteria for tweets manual classification

This section provides a summary of the guiding principle used for the selection of the relevant tweets:

- **Conservativeness:** We focused on those tweets whose object was either monetary policy or relate to financial sector supervision (or both). However, we adopted a conservative approach by selecting only those tweets that we were absolutely confident to be pertinent. For a tweet to be pertinent, it had to contain a description or a judgement over the course of action (performed or expected) of one of the central banks in our sample, stated as explicitly as possible.
- **Machine thinking:** Given that our work is supposed to be replicated in an automated way on a large scale of tweets, we performed our assessment accordingly. Here follow the main implications:
 1. We decided to include only those tweets which were self-explanatory. As far as we know, the machine is going to operate on a tweet by tweet basis (sort of row by row) reading exclusively the text body. Therefore, we decided to:
 - (a) exclude those tweets which were part of a larger Twitter thread and were difficult to understand in isolation. Such types of tweets were answers to a previous tweet or to a chain of previous tweets and their meaning was clear only considering the whole larger context of the thread.
 - (b) exclude those tweets which included an image or an URL address and which required the image or the URL to be used in order to be properly understood.
 2. Those tweets whose language style was excessively difficult to grasp due to metaphorical phrasing, abbreviations, use of slang language etc. have not been included.
- **Accuracy and netiquette:** We deliberately excluded those tweets which were excessively generic or that used an offensive and non-acceptable language.
- **Advertising, updates and market trends:** Many tweets in our sample had advertising goals. Also, there were tweets whose objective was to update traders about the latest news and market trends. We excluded both categories from our analysis.
- **Papers, conferences and other policy documents:** We found some tweets which were referring to policy documents, papers and conferences. We typically did not include them.

B Text Processing Details

As discussed in the main text, we compute similarity between the tweets related to monetary policy and the text of the monetary policy announcements by transforming the two corpus of text into vectors using doc2vec, a deep learning model described below. As we are interested in how the change in similarity affects asset prices after press releases, we gathered the tweets published in the interval between 48 hours before and 48 hours after an announcement. The tweets were then split into 1 hour segments around the monetary policy communication events.

We then measure similarity as the cosine similarity between each one-hour corpus of Twitter messages and the nearest monetary policy decision. Before proceeding with doc2vec, we pre-process all our corpus of text. In the following paragraphs we provide details on these steps.

B.1 Text Preprocessing

With pre-processing we reduce the number of words, and hence the computational time necessary to run the doc2vec model, without losing relevant information. We follow standard procedures in text pre-processing with different libraries in Python. We pre-processed the text in both the central bank press releases and Twitter messages by lowercasing all words. For tweets, we also removed all URLs and mentions to other Twitter users. For central bank transcripts, we removed standard introductions to speakers. We then broke streams of text into single words called “token”. Thereafter, we eliminated stop words, i.e. words that occur frequently in our corpus but have little meaning, and punctuations. We do this using “word.tokenize” from the Python module NLTK. We also removed all tokens that consisted only of non-alphanumeric characters. We also remove all tokens that consist of non-alphanumeric characters only, and remove emoticons, links, @, and # symbols.

Next, we lemmatized the words using WordNetLemmatizer from the Python module NLTK. Lemmatization entails reducing words to a common root form, called a “lemma”, to limit the presence of synonyms. Then we performed stemming, which implies conflating the various forms of a word into a common representation known as the stem. For instance, as a result of this process, the words “ate” and “eating” are both reduced to the common stem “eat”. Stemming and lemmatization rely on pre-existing dictionaries for the English language, which explains why we eliminated non-English tweets from our corpus. We relied on Porter Stemmer in the Python module NLTK for our stemming. Finally, we introduced collocation – the combination of two words that have higher probabilities of co-occurring together than separately. For instance, the tokens “federal” and “reserve” have higher chances of co-occurring as the bigram “federal reserve” than separately. In this case, collocations transform the two separate tokens into just one: “federal_reserve”. We used BigramCollocationFinder in NLTK. We then use the pre-processed corpus of text to train the doc2vec model.

B.2 Vector representation: doc2vec

After pre-processing our tweets and transcripts we have two types of “documents”: the transcripts of the monetary policy decisions and full set of one-hour tweets created in the 48 hours around central bank communications. The monetary policy decision document is the text of press release or the transcript of a press conference. A tweet document is

the text of all the tweets posted on a certain day-hour window around the central bank communication of reference.

Following [Giavazzi et al. \(2020\)](#), our approach consists of using neural networks to compute vector representations of words, including their context, through embedding. To perform this task, [Mikolov et al. \(2013\)](#) propose using word2vec, which learns word embeddings and aims to predict the occurrence of a word given the surrounding words (context). In this model, every word is mapped to a unique vector, which is represented by a column in weight matrix W . The algorithm constructs a vocabulary from the input corpus and then learns word representations by training a neural network language model. The model is trained using stochastic gradient descent with back propagation. When the model converges, it represents words as word embeddings – meaningful, real-value vectors of configurable dimensions (usually 150-500 dimensions). The neural network learns a word’s embedding based on its contexts in different sentences. As a result, the words that occur in similar contexts are mapped onto close vectors.

As an extension of word2vec, [Le and Mikolov \(2014\)](#) introduced doc2vec to learn embeddings of sentences and documents (or sentence embeddings), not just words. By treating each document as a word token, the word2vec methodology is used to learn document embeddings ([Bhatia et al., 2016](#)). As with word2vec, training occurs through back propagation. Each iteration of the algorithm is called an “epoch”, and its purpose is to increase the quality of the output vectors. This type of document embedding allows for texts to be represented as dense, fixed-length feature vectors that take their semantic and syntactic structure into account. We used a freely available implementation of the doc2vec algorithm included in the GENSIM Python module and asked for 300-dimensional vectors.

C Appendix Tables

Table C.1: Changes in similarity on French sovereign yields volatility

Panel A: Press release window						
	(1)	(2)	(3)	(4)	(5)	(6)
	1-year	2-year	5-year	10-year	20-year	30-year
Δ Similarity	0.014 (0.017)	0.119 (0.092)	-0.065 (0.128)	0.480 (0.322)	0.342 (0.528)	0.677 (0.960)
Observations	88	89	89	89	89	89
R-squared	0.418	0.427	0.316	0.399	0.433	0.591

Panel B: Press conference window						
	(1)	(2)	(3)	(4)	(5)	(6)
	1-year	2-year	5-year	10-year	20-year	30-year
Δ Similarity	0.043 (0.056)	0.074 (0.119)	0.307 (0.309)	0.869*** (0.234)	1.305* (0.674)	4.740** (2.183)
Observations	88	89	89	89	89	89
R-squared	0.126	0.485	0.336	0.465	0.359	0.469

Note: The dependent variable is the realised variance of French sovereign yields at different maturities in the 30-minute window surrounding a monetary policy announcement. Δ Similarity is the change in the similarity index between Twitter traffic and the policy announcement in the hour post as compared to the hour prior to the release of a policy decision. Year fixed-effects dummies are included in all specifications. Robust standard errors in parentheses. ***, **, and * denote statistical significance at the 1%, 5% and 10% levels, respectively.

Table C.2: Changes in similarity on German sovereign yields volatility

Panel A: Press release window						
	(1)	(2)	(3)	(4)	(5)	(6)
	1-year	2-year	5-year	10-year	20-year	30-year
Δ Similarity	-0.010 (0.007)	0.008 (0.009)	-0.007 (0.055)	0.151 (0.096)	1.190** (0.498)	-2.213 (1.348)
Observations	88	89	89	89	89	89
R-squared	0.329	0.289	0.348	0.669	0.713	0.360

Panel B: Press conference window						
	(1)	(2)	(3)	(4)	(5)	(6)
	1-year	2-year	5-year	10-year	20-year	30-year
Δ Similarity	0.009 (0.011)	0.023*** (0.008)	0.041* (0.024)	0.230** (0.100)	1.270* (0.713)	4.072*** (1.284)
Observations	89	89	89	89	89	89
R-squared	0.182	0.475	0.382	0.374	0.652	0.561

Note: The dependent variable is the realised variance of German sovereign yields at different maturities in the 30-minute window surrounding a monetary policy announcement. Panel A presents the results obtained focusing on the press release window [13:30–14:00], while the press conference estimates [14:15–15:45] are presented in Panel B. Δ Similarity is the change in the similarity index between Twitter traffic and the policy announcement in the hour post as compared to the hour prior to the release of a policy decision. Year fixed-effects dummies are included in all specifications. Robust standard errors in parentheses. ***, **, and * denote statistical significance at the 1%, 5% and 10% levels, respectively.

Table C.3: Changes in similarity on Italian sovereign yields volatility

Panel A: Press release window						
	(1)	(2)	(3)	(4)	(5)	(6)
	1-year	2-year	5-year	10-year	20-year	30-year
Δ Similarity	-0.007 (0.021)	0.003 (0.033)	-0.069 (0.077)	-0.491 (0.436)	-0.353 (0.327)	-1.403 (0.967)
Observations	88	89	89	89	82	89
R-squared	0.452	0.551	0.210	0.187	0.300	0.279

Panel B: Press conference window						
	(1)	(2)	(3)	(4)	(5)	(6)
	1-year	2-year	5-year	10-year	20-year	30-year
Δ Similarity	-0.291 (0.176)	0.023 (0.025)	0.107** (0.047)	0.946*** (0.329)	2.065*** (0.611)	3.711*** (1.063)
Observations	88	89	89	89	82	89
R-squared	0.404	0.504	0.526	0.595	0.605	0.652

Note: The dependent variable is the realised variance of Italian sovereign yields at different maturities in the 30-minute window surrounding a monetary policy announcement. Panel A presents the results obtained focusing on the press release window [13:30–14:00], while the press conference estimates [14:15–15:45] are presented in Panel B. Δ Similarity is the change in the similarity index between Twitter traffic and the policy announcement in the hour post as compared to the hour prior to the release of a policy decision. Year fixed-effects dummies are included in all specifications. Robust standard errors in parentheses. ***, **, and * denote statistical significance at the 1%, 5% and 10% levels, respectively.

Table C.4: Changes in similarity on Spanish sovereign yields volatility

Panel A: Press release window					
	(1)	(2)	(3)	(4)	(5)
	2-year	5-year	10-year	20-year	30-year
Δ Similarity	-0.029 (0.025)	-0.047 (0.152)	-0.119 (0.226)	-0.777 (0.571)	-0.911 (0.599)
Observations	89	89	89	89	89
R-squared	0.296	0.339	0.310	0.420	0.256

Panel B: Press conference window					
	(1)	(2)	(3)	(4)	(5)
	2-year	5-year	10-year	20-year	30-year
Δ Similarity	-0.008 (0.021)	0.104*** (0.031)	0.150*** (0.057)	0.987*** (0.369)	2.097*** (0.788)
Observations	89	89	89	89	89
R-squared	0.374	0.596	0.350	0.414	0.588

Note: The dependent variable is the realised variance of Spanish sovereign yields at different maturities in the 30-minute window surrounding a monetary policy announcement. Panel A presents the results obtained focusing on the press release window [13:30–14:00], while the press conference estimates [14:15–15:45] are presented in Panel B. Δ Similarity is the change in the similarity index between Twitter traffic and the policy announcement in the hour post as compared to the hour prior to the release of a policy decision. Year fixed-effects dummies are included in all specifications. Robust standard errors in parentheses. ***, **, and * denote statistical significance at the 1%, 5% and 10% levels, respectively.

Table C.5: Changes in similarity on French government bond returns

Panel A: Press release window						
	(1)	(2)	(3)	(4)	(5)	(6)
	1-year	2-year	5-year	10-year	20-year	30-year
Δ Similarity	-0.049 (0.051)	-0.095 (0.105)	0.098 (0.325)	0.976 (0.722)	1.686 (1.053)	1.911 (1.574)
Observations	88	89	89	89	89	89
R-squared	0.264	0.240	0.302	0.591	0.665	0.707

Panel B: Press conference window						
	(1)	(2)	(3)	(4)	(5)	(6)
	1-year	2-year	5-year	10-year	20-year	30-year
Δ Similarity	0.002 (0.012)	0.042** (0.020)	0.256** (0.107)	0.527** (0.207)	0.724** (0.276)	3.372** (1.417)
Observations	67	89	89	87	89	89
R-squared	0.407	0.252	0.355	0.433	0.499	0.482

Note: The dependent variable is the absolute value of French sovereign bond returns at different maturities using high frequency one-minute data. Δ Similarity is the change in the similarity index between Twitter traffic and the policy announcement in the hour post as compared to the hour prior to the release of a policy decision. Year fixed-effects dummies are included in all specifications. Robust standard errors in parentheses. ***, **, and * denote statistical significance at the 1%, 5% and 10% levels, respectively.

Table C.6: Changes in similarity on German bond returns

Panel A: Press release window						
	(1)	(2)	(3)	(4)	(5)	(6)
	1-year	2-year	5-year	10-year	20-year	30-year
Δ Similarity	-0.050 (0.058)	-0.011 (0.113)	-0.243 (0.305)	0.704 (0.580)	0.546 (0.789)	-0.355 (0.483)
Observations	88	89	89	89	89	89
R-squared	0.124	0.301	0.308	0.626	0.662	0.120

Panel B: Press conference window						
	(1)	(2)	(3)	(4)	(5)	(6)
	1-year	2-year	5-year	10-year	20-year	30-year
Δ Similarity	0.010 (0.014)	0.031 (0.019)	0.041* (0.024)	0.159 (0.163)	2.074** (0.896)	6.063** (3.016)
Observations	79	87	89	89	89	88
R-squared	0.347	0.096	0.382	0.259	0.477	0.160

Note: The dependent variable is the absolute value of German sovereign bond returns at different maturities using high frequency one-minute data. Δ Similarity is the change in the similarity index between Twitter traffic and the policy announcement in the hour post as compared to the hour prior to the release of a policy decision. Year fixed-effects dummies are included in all specifications. Robust standard errors in parentheses. ***, **, and * denote statistical significance at the 1%, 5% and 10% levels, respectively.

Table C.7: Changes in similarity on Italian bond returns

Panel A: Press release window						
	(1)	(2)	(3)	(4)	(5)	(6)
	1-year	2-year	5-year	10-year	20-year	30-year
Δ Similarity	0.029 (0.031)	-0.265 (0.195)	-0.767 (0.477)	-1.706 (1.129)	-1.837 (1.575)	-2.507 (2.031)
Observations	88	89	89	89	82	89
R-squared	0.481	0.215	0.151	0.207	0.247	0.325

Panel B: Press conference window						
	(1)	(2)	(3)	(4)	(5)	(6)
	1-year	2-year	5-year	10-year	20-year	30-year
Δ Similarity	-0.069 (0.051)	0.134 (0.097)	0.657** (0.299)	1.735** (0.728)	4.529** (1.978)	6.685** (3.093)
Observations	73	89	89	89	82	88
R-squared	0.274	0.402	0.413	0.382	0.506	0.499

Note: The dependent variable is the absolute value of Italian sovereign bond returns at different maturities using high frequency one-minute data. Δ Similarity is the change in the similarity index between Twitter traffic and the policy announcement in the hour post as compared to the hour prior to the release of a policy decision. Year fixed-effects dummies are included in all specifications. Robust standard errors in parentheses. ***, **, and * denote statistical significance at the 1%, 5% and 10% levels, respectively.

Table C.8: Changes in similarity on Spanish bond returns

Panel A: Press release window					
	(1)	(2)	(3)	(4)	(5)
	2-year	5-year	10-year	20-year	30-year
Δ Similarity	0.013 (0.155)	-0.364 (0.382)	-0.318 (0.721)	-0.864 (1.087)	-2.815 (1.838)
Observations	89	89	89	89	89
R-squared	0.165	0.258	0.495	0.447	0.327

Panel B: Press conference window					
	(1)	(2)	(3)	(4)	(5)
	2-year	5-year	10-year	20-year	30-year
Δ Similarity	0.059* (0.034)	0.209 (0.136)	0.557** (0.232)	2.133*** (0.736)	4.687** (1.947)
Observations	89	87	88	89	87
R-squared	0.404	0.381	0.390	0.459	0.543

Note: The dependent variable is the absolute value of Spanish sovereign bond returns at different maturities using high frequency one-minute data. Δ Similarity is the change in the similarity index between Twitter traffic and the policy announcement in the hour post as compared to the hour prior to the release of a policy decision. Year fixed-effects dummies are included in all specifications. Robust standard errors in parentheses. ***, **, and * denote statistical significance at the 1%, 5% and 10% levels, respectively.