

'Whatever it Takes' to Change Belief: Evidence from Twitter

Michael Stiefel
Rémi Vivès

WP 2019 - Nr 07

‘Whatever it takes’ to change belief: Evidence from Twitter*

Michael Stiefel[†]

Department of Economics, University of Zurich

Rémi Vivès[‡]

Aix-Marseille Univ., CNRS, EHESS, Centrale Marseille, AMSE

March 2019

Abstract

The sovereign debt literature emphasizes the possibility of avoiding a self-fulfilling default crisis if markets anticipate the central bank to act as lender of last resort. This paper investigates the extent to which changes in belief about an intervention of the European Central Bank (ECB) explain the sudden reduction of government bond spreads for the distressed countries in summer 2012. We study Twitter data and extract belief using machine learning techniques. We find evidence of strong increases in the perceived likelihood of ECB intervention and show that those increases explain subsequent decreases in the bond spreads of the distressed countries.

JEL classification: E44, E58, D83, F34

Keywords: Self-fulfilling default crisis, unconventional monetary policy, Twitter data.

*We are grateful to Frédéric Dufourt, Mathias Hoffmann, Charles Lai Tong and Alain Venditti for helpful discussions and suggestions as well as their constant encouragement and support. We also would like to thank, without implicating, Marcel Aloy, Jess Benhabib, Hamza Bennani, Hagop Boghazdeklian, Renaud Bourlès, Sebastian Dörr, Gilles Dufrénot, Martin Mandler, Egor Maslov, Elias Moor, Sebastien Laurent, Ralph Ossa, Céline Poilly, Frank-Alexander Raabe, Morgan Raux, Lorenzo Rotunno, Grigorios Spanos, Peter Tillmann and Tobias Wekhof for useful comments and suggestions. This paper also benefited from seminar participants at the 3rd ICFBP EMU, April 2018, University of Konstanz, May 2018, 14th CEUS Workshop WHU, May 2018, SoFiE Summer School in Brussels, June 2018, 20th Workshop for Young Economists ZEW, July 2018, 2nd CARMA conference València, July 2018, EDGE Jamboree Munich, September 2018, 43rd SAEe Madrid, December 2018. This work was supported by the French National Research Agency Grant ANR-17-EURE-0020.

[†]address: Zürichbergstrasse 14, CH-8032 Zürich, Switzerland, e-mail address: michael.stiefel@econ.uzh.ch

[‡]corresponding author, address: 5-9 Boulevard Bourdet, CS 50498 - 13205 Marseille Cedex 1, France, e-mail address: remi.vives@univ-amu.fr

1. Introduction

In summer 2012, when the eurozone was on the verge of breaking up, the European Central Bank (ECB) gradually announced the Outright Monetary Transactions (OMT) program¹ which gives the central bank the possibility to buy an unlimited amount of short-term government debt in secondary markets under certain conditions. Even though such purchases were never made, sovereign yields of the distressed countries (Greece, Ireland, Italy, Portugal, Spain) fell during this time period.

Pioneered by Calvo (1988), the sovereign debt literature emphasizes the emergence of multiple equilibria when economic fundamentals worsen, namely a fundamental and a self-fulfilling default equilibrium. More precisely, the former equilibrium (i.e. the “good” equilibrium) prices the true level of economic fundamentals whereas the latter equilibrium (i.e. the “bad” equilibrium) generates a self-fulfilling debt crisis triggered by pessimistic investors (Cole and Kehoe 1996, 2000). Applying this framework to the European sovereign debt crisis, Corsetti and Dedola (2016) and Roch and Uhlig (2018) show that this self-fulfilling default crisis can be avoided if markets anticipate that the central bank will act as the lender of last resort.

In this paper, we analyze the extent to which changes in belief about an intervention of the ECB explain the sudden reduction of government bond spreads for the distressed countries in the eurozone, as suggested by the literature on self-fulfilling default crises. To study this change, we follow a direct approach by extracting belief from Twitter data. At first, we document that there were large increases in the volume of tweets around important dates of central bank communication, showing that Twitter was used to both communicate and interpret the ECB’s actions. Then, we create a belief index of the perceived likelihood of central bank intervention using techniques from natural language processing. This analysis reveals that the belief index jumps at two important days of ECB communication: the day of ECB president Mario Draghi’s ‘Whatever it takes’ speech and the day of the OMT program announcement. These large increases in our belief index coincide with large decreases in the sovereign spreads of the distressed countries on the same and the following day. We also find that, to a smaller degree, our belief index is sensitive to other events, such as information leaks and rumors. Using a pooled panel estimator, we show further that a one-standard deviation increase in the lagged belief index is associated with a six basis point reduction in the spreads of the crisis countries. To corroborate our findings, we also compare changes

¹The OMT program was officially announced after the meeting of the ECB Governing Council on September 6th 2012. However, the ECB communication had already changed in the previous two months, so that the literature includes earlier speeches as part of the OMT announcement (Altavilla et al. 2016; Ambler, Rumler, et al. 2017; Falagiarda and Reitz 2015; Krishnamurthy et al. 2017; Van Der Heijden et al. 2018).

in beliefs of individual users with several tweets around the dates of ECB communication and also detect strong increases in the belief index at our two identified key dates. Finally, the following robustness checks demonstrate that our results are robust: using spreads of sovereign Credit Default Swaps (CDS), forming alternative versions of our belief index and controlling for the users' level of information using the number of followers similarly to Gholampour and Van Wincoop (2017).

We make three contributions. First, this paper demonstrates that we can learn from social media data how the public receives news announcements and how, in turn, these announcements influence belief formation and confidence building. Second, we show that capturing this belief formation can improve upon typical event studies. Event studies can be problematic if there is anticipation before or a delayed reaction afterwards that is not inside the event window. By creating an index over the full time horizon, our procedure can both capture rumors and information leaks outside the event window as well as distinguish the importance of different announcements. Third, even though we are not formally putting such a model to the data, our results are consistent with the theoretical prediction that a central bank, which credibly commits to an intervention as lender of last resort, can eliminate self-fulfilling equilibria. The 'Whatever it takes' episode is widely recognized as a turning point in the sovereign debt crisis in the eurozone, a common narrative in the popular press ("5 years ago, Draghi saved the euro in one sentence" (Les Echos 2017)) as well as in the economic literature (e.g. Corsetti (2015) in his Schumpeter Lecture). We motivate the channel through which this speech has affected bonds, namely a change in belief leading to the perception of the ECB as a central bank which is willing to intervene to reduce government bond spreads. Over this three-month horizon, our analysis associates a 180 basis points reduction in the 10-year bond spreads of distressed countries due to our identified change in belief.

Related literature

This paper is related to three branches of the economic literature. First, studies have analyzed whether sovereign risk is priced according to "fundamentals", or whether sentiments and market coordination play a role as well. Second, recent work has investigated the effect of central bank communication on financial markets, and more specifically the announcement of central bank programs on government bond spreads. Finally, a new branch of literature has started to use social media data such as Twitter to analyze financial fluctuations, and also to model the expectation formation about monetary policy.

Several studies have documented that one cannot explain the large increases in government bond spreads leading to the eurozone crisis using fundamental factors alone, such as debt and GDP dynamics; see for instance De Grauwe and Ji (2012) and Di Cesare et al.

(2013). As an explanation, De Grauwe and Ji (2012) highlight miscoordination among market participants, while Di Cesare et al. (2013) point to a perceived break up risk of the eurozone as a potential channel. Bocola and Dovis (2016) provide a quantitative decomposition of the self-fulfilling and fundamental parts of Italy’s sovereign risk. Using the model of Cole and Kehoe (2000), they indirectly infer beliefs from observed changes in the maturity structure of government bonds. They find that 12 percent of the Italian spread is explained by rollover risk.

Given the importance of central bank actions in the aftermath of the financial crisis, many studies have investigated the effects of unconventional monetary policy measures. For the eurozone, event studies have found that just the announcement of central bank policies leads to sizeable effects on government bond yields (see among others Falagiarda and Reitz (2015) and Szczerbowicz (2015)). Fendel and Neugebauer (2018) document that the main announcement effects occurs with a delay of one day. The closest paper to ours is Altavilla et al. (2016) who study the financial and macroeconomic effects of the OMT program announcements. Focusing on the financial effects in an event study, they show that the OMT announcements triggered a reduction of about 200 basis points in the 2-year government bond yield of Italy and Spain. Furthermore, using a multi-country VAR model and constructing a counterfactual scenario without bond buying program announcements, they show that the announcements had significant effects on Italian and Spanish growth rates.

Our approach differs from the above paper in several dimensions. First, we shed light on the *channel* through which the OMT announcement affected spreads, namely belief. Second, as has been criticized by D’Amico (2016) in the discussion of Altavilla et al. (2016), estimating the financial effect with event dummies does not take into account the expectation formation process in between and after the days of ECB announcements. In this paper, we address precisely this concern by modeling belief through the entire event period which allows us to study the possible effects of anticipation, rumors and information leaks.

Another branch of the literature on central bank communication has used tools by computational linguistics to infer different dimensions of central bank communication (Hansen and McMahon 2016; Hansen, McMahon, and Prat 2014). Like these papers, we apply machine learning methods to classify text. We focus, however, on extracting information from responses to the central bank communication, not by applying them to the policymakers directly, and by only investigating a unique and pre-specified dimension of the text instead of modeling different topics.

A burgeoning literature studies the impact of Twitter sentiment on financial fluctuations. Pioneered by Bollen et al. (2011) and Zhang et al. (2011), this literature shows that the

general public mood of Twitter users can predict stock market indices.

Gholampour (2017) develops a financial dictionary to proxy the daily sentiment and disagreements of investors to predict financial fluctuations. In the same spirit, Gholampour and Van Wincoop (2017) highlight that Twitter is an important source of information to predict the euro-dollar exchange rate and show that informed traders share their information on the microblogging social networking platform. Concerning monetary policy, Azar and Lo (2016) perform a sentiment analysis of tweets referring to the Federal Reserve. They show that Twitter sentiment has a large impact on asset prices.

Meinusch and Tillmann (2017) were the first to infer beliefs about monetary policy from Twitter. They investigate the extent to which long-term bond yields and the exchange rate are sensitive to changes in belief about the Federal Reserve's exit from quantitative easing. In order to proxy those beliefs, the authors label and aggregate tweets from April to October 2013, thereby distinguishing the users' opinions on whether the Federal Reserve will taper soon or late. Using a VAR-X model, they identify a belief shock. Their results show that changes in belief have strong and persistent effects on bond yields and exchanges rates. While our paper is similar in spirit, our extracted belief is not about the timing of a central bank action, but rather about the type of the ECB and its willingness to intervene at all. Moreover, we are interested in the differential effect of this belief on the distressed countries compared to other eurozone countries.

The remainder of the paper is organized as follows. In section 2, we detail the succession of the ECB key events over summer 2012. Our Twitter and financial data are described in section 3. In section 4, we present our belief index that is used for the empirical analysis in section 5. We discuss the results and robustness checks in section 6. We look at a different dimension of our Twitter data and compare belief before and after the key dates for individual users in section 7. We conclude the paper in section 8.

2. Setting

In this section, we first describe the background of our study, the severity of the sovereign debt crises during summer 2012 and the debate about ECB interventions. Then, we highlight in detail the key ECB actions in this timespan which culminated in the official announcement of the OMT program. Finally, we explain why this background is well suited to analyze our research question.

Our horizon of study, July to September 2012, captures the moment when the sovereign debt crisis in the eurozone was hitting Italy and Spain. This was a critical time for both countries since the financing costs had seen dramatic increases in a short amount of time:

the spreads to Germany amounted to less than 100 basis points in 2010, while in summer 2012 they reached close to 600 basis points. In addition, this was also a decisive moment for the eurozone as a whole: the debt crises in Greece, Ireland and Portugal had led to new institutions like the European Stability Mechanism (ESM), but also sparked tensions in other member countries. There was severe resistance to the so-called rescue packages, especially in Germany. A lurking bailout among the larger economies in Spain or Italy would have outsized the already agreed upon emergency funding schemes, and sparked further conflict among the member countries.

The ECB was under increasing pressure to intervene because of the severity of the crisis, but it had remained rather passive until then, mostly focusing on bank liquidity measures.²³ Although the FED and the Bank of England had already purchased large amounts of government debt as part of their unconventional monetary policy, the ECB had not started a large scale bond buying program as either quantitative easing or as a lender of last resort. Article 123 of the Treaty on the Functioning of the European Union forbids the ECB to directly purchase debt from its member countries. Article 125 of the same treaty - the “no-bailout-clause” - states that no member state is accountable for the debt of other member countries. For this reason, an ECB intervention even in secondary markets for government bonds caused legal concerns.

We now explain the three key communication actions by the ECB that were undertaken in this time horizon, and that are nowadays regarded as a fundamental change of ECB policy. Typically, the three actions are jointly regarded as a gradual announcement of the OMT program (Altavilla et al. 2016). However, each event is fundamentally different and might therefore also have affected market expectations in different ways.

On July 26th, talking to financial market participants at the *Global Investor Conference* in London, Mario Draghi made the following remarks:

“Within our mandate, the ECB is ready to do whatever it takes to preserve the euro. And believe me, it will be enough.” (ECB 2012a)

We want to emphasize that Draghi did not choose the official statements around Governing

²See Falagiarda and Reitz (2015) and Szczerbowicz (2015) for detailed accounts of the ECB’s actions.

³In fact, there were two other ECB programs involving purchases of government bonds on secondary markets, the Securities Market Program (SMP) in 2010, which was replaced by the OMT program, and the expanded Asset Purchase Program (APP) in 2015. Both programs are quite different to the OMT. According to the ECB, both SMP and APP target both public and private securities with the objective of ensuring the monetary policy transmission and price stability. In contrast, the OMT program is specifically designed to reduce yields of distressed countries. Furthermore, the SMP is different because it featured de-facto limits on the purchased amounts of debt and the ECB had seniority on the bonds it purchased (Bruegel 2012). The APP is a quantitative easing program in which government bonds are purchased but not specifically targeting a specific country, and not allowing for an unlimited amount.

Council meetings for this remark, but an external event with financial market participants. He also directly addresses them and their belief. Neither in this quote nor in the full speech is a direct reference to a new policy program, so the impact of the speech crucially hinges on the market participants' interpretation.

On August 2nd, at the regular meeting of the ECB Governing Council and the subsequent press conference, Draghi went a step further to link his statement to a possible action by the ECB but remained very vague about a specific program. Specifically, Draghi said,

“The Governing Council (...) may undertake outright open market operations of a size adequate to reach its objective.” (ECB 2012b)

Questioned by a journalist whether his ‘whatever it takes’ speech was about bond buying by the ECB, Draghi responded:

“Have you read the speech? Had you read it, you would have seen that there is no reference whatsoever to a bond buying programme.” (ECB 2012b)

As the journalists interrogate him further about whether his remarks were then misinterpreted by markets which seemed to expect the ECB to become active, Draghi then responded:

“I like these remarks very much. And they were not misinterpreted. Markets simply took their actions based on their expectations following these remarks. That is what happened. And these expectations are what they are.” (ECB 2012b)

Those quotes illustrate again that within the first two main communication events, the ECB did not commit to, but only hinted at, a specific program such that the consequences of Draghi's words effectively depend on the market participants' interpretation. To therefore truly capture those announcement effects, it is necessary to measure the market participants' response to those statements in contrast to event studies which just give a dummy variable for such an announcement day.

Finally, on September 6th, the last main ECB action in this time horizon, the Governing Council officially announced the OMT program. This program gives the ECB possibility to buy an unlimited amount of short-term bonds on secondary sovereign debt markets under certain conditions. Until now, this program has never been activated. Even though the ECB had formally proposed a program, this did not stop discussions about its legality and whether the ECB would actually commit to it. In fact, prominent politicians and lawyers had appealed to the German Constitutional Court, arguing that this program was beyond the ECB's mandate, and those appeals were declared invalid only in 2016. In this regard, the impact of the program still hinged on the expectations of the market participants after the announcement.

We argue that this period around the announcement of the OMT program is an ideal setting to study our research question. While the ECB was very active in *communicating* its intentions, actual purchases within this program were never made. Thus, the change in spreads can be attributed to changes in belief and not to large purchases by the central bank. At that time, Italy and Spain were not part of any rescue package by the ESM or EFSM, i.e. they were fully dependent on private lenders to finance their expenses. Additionally, other confounding factors are minimal within this time horizon. As the work by Altavilla et al. (2016) has shown, controlling for other economic news does not change the effect of the OMT announcements in an event study.

3. Data

In this section, we start by justifying the use of Twitter and explaining how our dataset has been constructed. Then, we detail the financial data used in this study.

3.1. *Twitter*

Twitter data presents several interesting features for our analysis. Firstly, Twitter is a large source of opinionated data of individual users. As a microblogging social networking platform, Twitter allowed 200 million monthly users in 2012 to express their opinions on different topics through short public messages. A tweet must be concise (with a limit of 140 characters in 2012) which makes it possible to extract a simple opinion from it. Furthermore, Twitter is a large source of high frequency data. This allows users to quickly react to news and events, and in fact, more than 50 percent of tweets in 2012 came from mobile devices. Finally, a third interesting feature of Twitter data is that users include policymakers, financial journalists, and also traders. Thus, relevant information from all different disciplines is shared on this platform.

To construct our dataset, we use web scraping techniques that allow us to extract data directly from the Twitter website based on date and keywords. We collect tweets from July 2nd to October 1st 2012. Each tweet contains at least two of the following keywords: “Draghi”, “ECB”, “bailout”. For each tweet, we gather information about the text content, the number of retweets and favorites, but also the user name, user id and tweet id. This method allow us to gather 42,685 unique English tweets from 11,506 accounts after filtering by language and day of the week. We keep only English tweets because English is the main language both in financial markets and on Twitter. Furthermore, we only look at tweets posted on weekdays, since the volume of tweets on weekends is quite low and the daily indices

are then fully consistent with the financial data. We add the number of retweets to each tweet in order to control for the number of retweets. For instance, a tweet that has been retweeted 10 times counts for 11 tweets. This mechanically gives a higher weight to tweets which have been retweeted, which can be regarded as a sign of importance. Our data set is now composed of 49,522 English tweets, as most tweets have not been retweeted.

The daily number of tweets ranges from 13 to 9,601 with a mean of 750 tweets and a large daily variance (see Table A2 for quantiles). Moreover, 90 percent of users have tweeted 8 times or less and 52 percent have only tweeted once. A recurrent question when gathering data based on combinations of keywords is about the relevance of the extracted information. In Figure 1, we plot the daily number of tweets. We can observe three main peaks from the series. The first peak, on July 26th, corresponds to the “Whatever it takes” speech by Mario Draghi. The second peak corresponds to the meeting of the Governing Council of the ECB on August 2nd. Finally, the third peak, on September 6th, relates to the OMT announcement conference. The fact that we can recover the key events of our time horizon from the daily number of tweets shows that Twitter was used to spread and interpret the news from those events. An interesting feature in Figure 1 is that we can observe other peaks that are not directly related to the days of ECB communication. Based on the content of the tweets, we can identify that they are also linked to rumors and information leaks. For instance, before the official announcement of the OMT program, Mario Draghi announced to cancel his participation in the Jackson Hole conference at the end of August which led to rumors that the ECB was “up to something big” (see tweet example below). Another example is that on September 3rd, only three days, before the announcement, Mario Draghi spoke to members of the European Parliament in Brussels behind closed doors, but on Twitter and in newspapers there are rumors that Draghi said that buying short-term debt did not breach the EU treaty (El País 2012).

3.2. Financial and Macroeconomic Data

We construct our series of government bond spreads using data from Bloomberg. We retrieve government bond yields for 11 European countries (Austria, Belgium, Finland, France, Germany, Greece, Ireland, Italy, Netherlands, Portugal and Spain) for maturities of 2, 5 and 10 years over our time horizon of interest, 2nd July until October 1st. Due to data unavailability, we only consider the 10-year maturity spreads for Greece and the 2-year and 5-year maturity spreads for Ireland. For all maturities, the bond spreads are computed relative to Germany.

Furthermore, we consider the European counterpart of the VIX index: the V2TX index.




Fig. 1. Timeline of the Daily Number of Tweets

Note: the figure shows the number of tweets on the left axis. The three key ECB communication events on July 26th, August 2nd and September 6th are highlighted with a vertical bar.

It is an uncertainty index based on the EURO STOXX 50 realtime option prices. We also use the Citi Economic Surprise Index (CESI) for the eurozone. The latter captures macroeconomic surprises by comparing consensus expectations and the state of the economy. This index is often used to control for macroeconomic fundamentals at a daily frequency.

Finally, for robustness checks, we also use the time series on sovereign Credit Default Swaps from Datastream at 2-,5- and 10-year maturity for the same set of countries except Greece, where the corresponding series are not available.

4. A Belief Index on Central Bank Intervention

In this section, we first explain how we construct the belief index. We then show descriptive statistics for this belief index and how it relates to the financial data.

4.1. Construction

To extract a measure for the perceived credibility of central bank intervention, we assign a label to each tweet and then compute a daily aggregate. A tweet is assigned the label “1” if an intervention by the ECB is considered to be likely. Similarly, if an intervention by the ECB is not considered to be likely, the tweet is assigned the label “-1”. We give a neutral label, “0”, to those tweets which do not express an opinion about central bank intervention.

To give an example, consider the following tweets:⁴

Draghi’s not kidding. My take on his comments and expectations of ECB bond buying

ECB ‘willing to buy bonds of weaker EU nations’ says Draghi — It’s a start

With the cancellation of Draghi trip to Jackson Hole, ECB is up to something big

Draghi reportedly told EU Parliament ECB can buy 3 year bonds and bond purchases are not state financing

The first two examples indicate a clear opinion that the ECB is willing to intervene and receive the label “1”. The last two examples are also labeled “1” and show the rumors and information leaks on Twitter.

Now, consider the following examples that express an opinion that the ECB will not intervene and therefore receive the label “-1”:

⁴For this exposition, we remove links/hashtags from the tweets to ease readability.

Chatter that other ECB policymakers don't agree with President Draghi's statements yesterday and bond buying is unlikely to be restarted.

Boy, did Draghi blow it today. I was wrong. I thought he and Bernanke were on the same page. Now ECB has lost credibility.

“DRAGHI SAYS ECB MAY UNDERTAKE OUTRIGHT OPEN MARKET OPERATIONS” ... ecb is simply not allowed to do that!

BofA: The ECB will never be able to enforce the centerpiece of its news bailout plan

Further examples including neutral tweets are provided in the appendix in Table A1.

We randomly split the tweets into two different sets. The first set, consisting of 20 percent of the tweets, is labeled manually. Based on this manually labeled data set, we now employ a double cross-validation procedure to select a machine learning model that can evaluate the remaining unlabeled tweets. This double cross-validation procedure consists of two layers. In the first layer for model assessment, we randomly split this manually labeled set further into a training set (90 percent) and a test set (10 percent). In the second layer for model selection, we train a machine learning classification model to this training set, as explained further below.⁵

The machine learning in our context faces the challenge of learning from textual data. A popular method in this context is the “*n-gram*” approach. At first, preprocessing steps clean the text from links and hashtags.⁶ Then, a so-called count vectorizer creates a dictionary in which all words (“tokens”) are contained. A tweet can then be regarded as a collection of items in this dictionary. The *n-gram* method allows to group together *n* consecutive tokens (in the order in which they appear in the tweet) as an *n-tupel* so that the final dictionary contains unique words and combinations of those words up to the number *n*. In many instances, *n-tupel* allow to catch more meaning. For example, the words in “The ECB will buy government bonds soon” and “The ECB will not buy government bonds soon” are the same with the exception of the single word “not”.⁷ To measure the importance of an item

⁵See Figure A1 in the appendix for an illustration of the procedure.

⁶Stemming or lemmatizing the text do not improve the results. Therefore, we decide to continue our analysis without applying these methods. Stemming is a process that allows to reduce words to their word stems while lemmatization is a process that groups different inflected forms of a word to one single unit.

⁷There is a trade-off between allowing for higher *n-grams* - that is allowing for more combinations of tokens to capture more meaning - and overfitting since one then allows for many features specific to a single tweet. This trade-off is solved by cross-validating the model on a hold-out set to determine which *n-gram* model performs best, as explained below.

in this dictionary, we use the tf-idf (term frequency - inverse document frequency) statistic. This statistic assigns a higher value, if this item appears more often in the tweet, but less so, if this item also appears more often in all the other tweets. The final dataset is then a large matrix in which each row is a tweet, and each column corresponds to a dictionary item. The matrix entries correspond to the tf-idf score of each item in the tweet. Since not every word appears in each tweet, this is a sparse matrix, which makes our machine learning approach computationally feasible.

More formally, the problem is to assign one of three categorical targets (the labels “-1”, “0”, “1”) to each tweet based on the explanatory variables which are given here by the tf-idf scores for each dictionary element. For this multiclass classification problem with a large sparse matrix, a popular classifier is Support Vector Machines (SVM).⁸ When this supervised learning model is fit to our manually labeled data set, it essentially fits a hyperplane to separate the different tweets in a high-dimensional space. This model is then used to predict the labels for the remaining 80 percent of the tweets in the data set that has not been manually labeled.

In the model selection part, we need to find the right parameters for the SVM-classifier (also called to “hypertune” the parameters), as well as to determine which n-gram model to choose. We proceed with grid search cross-validation. The training set is again randomly split into 5 different folds, each of which contains 20 percent of the data set.⁹ The classifier is then trained on four folds with different parameters and is tested on the remaining fold. This is repeated five times until each fold has been used for testing once. The classifier then uses the parameters which, on average, perform best in this cross-validation task. This cross-validation procedure was also used to determine that a trigram model, allowing for dictionary items of up to three tokens, performs best. Finally, we now apply the selected classifier on the test set which, of course, has not been used during the training, for model assessment. This allows us to compute an accuracy score for the machine learning which is approximately 93 percent.¹⁰

4.2. Descriptive Statistics

Given the set of labeled tweets, we now proceed to compute daily statistics. There are 40 percent of tweets considering ECB intervention to be likely, 50 percent of neutral tweets and 10 percent of tweets considering that an intervention of the ECB is unlikely. Since there

⁸We also tried other kinds of machine learning classifiers such as Logistic Regression, Random Forest, Naive Bayes, k-Nearest Neighbors. However, their performances in terms of predictive accuracy, recall and precision, were worse than what were obtained by the SVM classifier.

⁹See Figure A1.

¹⁰The classifier also performed well in terms of precision and recall for all three classes.

is a large variance in the daily amount of tweets, we do not consider the mean of labels per day to be a relevant statistic, because it would give a relatively higher weight to tweets on days with a small volume of tweets. Instead, we suggest the following two statistics:

$$\Delta\text{Belief}_t = \sum_i \text{Tweet}_{i,t} \cdot \mathbf{1}' - \sum_i \text{Tweet}_{i,t} \cdot \mathbf{-1}' \quad (1)$$

$$\text{Belief}_t = \sum_{j=1}^t \Delta\text{Belief}_j \quad (2)$$

That is, we compute the daily sum of the labels per day (ignoring the neutral tweets) and interpret this as changes in belief.

As we see in the timeline of the number of tweets in Figure 1, users seem to predominantly respond to new events and information. Clearly, when there is no new information and therefore a low volume of tweets, this does not mean that the previous events and changes in belief are no longer important. With this in mind, we obtain the final belief index in levels, Belief_t as the cumulative sum of the previous changes. This interpretation of daily tweets as changes is also implicit in the literature about stock market predictions using Twitter Sentiment. For example, Gholampour (2017) associates stock market changes with Twitter Sentiment which then suggests a relationship between the stock market in levels and the accumulated Twitter sentiment. However, in the robustness part, we also consider different ways of constructing a belief index, such as the mean, and show that our results are qualitatively unaffected. Summary statistics and quantiles are reported in Table A3 in the appendix.

As one can see in Figure 2, our belief index has an upward trend with two strong peaks, one at the day of the ‘Whatever it takes’ speech and another large spike at the day of the official OMT-announcement. Interestingly, in spite of a large volume of tweets on August 2nd when the Governing Council of the ECB met, we do not find a change in belief on this day. As we outlined in the setting, the ECB communication on August 2nd was ambiguous. On the one hand, the ECB did not announce any new specific program which might look like a step back after the ‘Whatever it takes’ speech in the previous week. On the other hand, Mario Draghi signaled that the ECB might engage in a bond buying program without outlining formal details. From the Twitter data, we can clearly observe two types of reactions. First, there were tweets expressing disappointment which mirrors the behavior of the 10-year maturity spreads in Figure 2. Second, there were tweets welcoming Draghi’s statement about




Fig. 2. Bond Spreads and Belief index

Note: this figure shows the government bonds spreads of Spain and Italy in basis points relative to Germany (left hand side) and the belief index created from Twitter (right hand side).

outright open market operations and therefore regarding ECB intervention to still be likely to occur. The latter reaction slightly dominates in the aggregate in our analysis. We note that this behavior of our belief index is different from the typical event studies which regard the 2nd August as one part of the OMT-announcement, see for instance Altavilla et al. (2016).

Finally, we also observe that our belief index is sensitive to rumors and information leaks. For instance, the change in belief on August 16th is associated with tweets about a statement by German Chancellor Angela Merkel supporting the ECB’s approach for reducing borrowing costs of indebted countries. As shown in the example tweets, Draghi’s cancellation of the Jackson Hole meeting and his appearance in front of the European Parliament in Brussels led to speculations that the ECB was preparing a program. Unlike an event study analysis, our approach allows to capture the effects of those events.

We also run a principal component analysis in order to see by how much our belief index is correlated with the first principal component of the spreads. The results show that the correlation between our belief index and the first principal component is 0.91, and this first principal component explains 70 percent of the variation in spreads. Hence, even though we just look at one dimension which might explain the variation in government bond spreads and other factors like uncertainty and macroeconomic surprises might play a role as well, our belief index seems to capture the relevant parts of the actual movement in the data.

In the next section, we dig deeper into how changes in the belief index are related to changes of the government bond spreads, and how it might have differently affected crisis and non-crisis countries.

5. Empirical Strategy

We want to understand whether changes in our belief index can explain the changes in the government bond spreads of the distressed countries. We estimate the following pooled panel regression:

$$\Delta s_{it} = \alpha_0 + \beta_1 \Delta \text{Belief}_{t-1} + \beta_2 \Delta \text{Belief}_{t-1} \times \text{Crises}_i + X_t + D_t + u_{it}, \quad (3)$$

in which Δs_{it} is the change in the spread of country i on day t , computed relative to Germany. ΔBelief_t is the standardized change in our belief index. We further interact ΔBelief_t with Crisis_i , an indicator variable for the countries that faced a sovereign debt crisis (Greece, Italy, Ireland, Portugal and Spain). This interaction term is our coefficient of interest and shows the differential effect of changes in belief on the crisis countries relative to the non-crisis

countries (Austria, Belgium, Finland, France and the Netherlands). To account for possible endogeneity, we mainly focus on the lag of the changes in belief, $\Delta\text{Belief}_{t-1}$. Looking at the effect of the belief index in lags is not unusual. Most event studies allow for a two-day window and Fendel and Neugebauer (2018) document that the announcement effects of ECB unconventional policies seem to occur with a lag in general.

X_t is a set of control variables. We control for other common factors that could equally explain changes in government bond spreads. We are using the European uncertainty index $V2TX_t$ based on the EURO STOXX 50. Furthermore, to control for changes to macroeconomic fundamentals at a daily frequency, we are using the CESI macroeconomic surprise index CESI_t . Ideally, we would like to control with a surprise index for each country in our data set, unfortunately, such a news index does not exist for every country.

Finally, D_t is an event dummy variables that controls for the three ECB announcement dates: July 26th, August 2nd and September 6th. We use two specifications based on the event dummy: a one-day event window and a two-day event window, which allows effects to also occur on the subsequent day. All variables are standardized to simplify the interpretation.

6. Results and Robustness

6.1. Results

In Table 1, we present the results from estimating equation (3) using data for 10-year maturity bonds without adding any other regressors. We expect a positive change in the belief index to be associated with a reduction in the spreads of the crisis countries. In column (1), we focus on the effect of changes in belief on the same day. We see that a one-standard deviation change in the belief index has a sizable and significant negative effect of -6.7 basis points on the 10 years bond spreads of the crisis countries. We also run this regression with a lagged change in belief index in columns (2). The coefficient stays significant with only a slightly smaller magnitude of approximately -6 basis points. From this finding, we conclude that the result is not driven by reverse causality of changes in spreads feeding into our belief index. We also see that the coefficient of the change in the belief index on the non-crisis countries is negligible and not significant. In column (3), we remove this regressor in order to add time fixed effects without having a multicollinearity problem. We cluster the standard errors by the crisis countries and time.¹¹ The magnitude and the significance

¹¹We believe that those are the dimensions among which standard errors might be correlated. It turns out that clustering rather reduces the standard errors. The results are robust, however, to not including them or including them individually.

of the results remain unchanged. In Table A4 in the appendix, we show that the results also go through for bonds of maturities of 5 or 2 years. However, the latter are not directly comparable due to data unavailability for single countries, Greece is only contained in the table for 10-year maturity, and Ireland only for 2 and 5-year maturity bonds. In column (4) we include both the contemporaneous and lagged change in belief, a similar specification to a two-day event window in an event study. Both coefficients on the interaction terms indicate an economically meaningful effect on the bond spreads of the crisis countries. On average, a standard deviation change in the belief index reduces their bond spreads by 5.7 basis points on the same day and by 4.7 basis points on the following day.

We now add further control variables. We focus on the lag in the change in belief index which is conservative since this effect is smaller than the contemporaneous effect and it rules out reverse causality concerns. Results are shown in Table 2. In column (1) and (2), we add the European uncertainty index V2TX and we also control for macroeconomic surprises using the CESI index. Those factors enter with a positive sign but their respective coefficients are not significant. However, our coefficient of interest is still highly significant and remains of the same size. Those results show that the change in the sovereign bond spreads of the European crisis countries were not driven by a change in uncertainty or macroeconomic fundamentals. In further regressions not reported here, we find that the result also holds when using the Scotti macroeconomic surprise index for the eurozone (Scotti 2016), or when further interacting the common uncertainty and macroeconomic surprise factors with the crisis countries.

Finally, in column (3) and (4), we estimate equation (3) adding an event dummy variable to control for the ECB announcements dates: July 26th, August 2nd and September 6th. We include one-day and a two-day event dummies in column (3) and (4) respectively. These variables enter with expected large and significant negative signs. The effect of a change in the belief index on the crisis countries when using a one-day dummy variable is unchanged, but the latter is reduced by half when one increases the window of the dummy variable up to two days. However, the inclusion of these variables still leaves the coefficient at a sizable magnitude of a 3.2 basis point reduction, and is still significant at a one percent level. The result that the coefficient of a change in the belief index on the crisis countries remains significant shows that our approach captures more information than a simple dummy approach analysis around the three ECB announcement dates of summer 2012. This fact indicates that there is further relevant information in our belief index likely due to rumors, information leaks or by better accounting for varying impacts of different announcement dates.

Overall, our index Belief_t increased to approximately 15,000 at the end our three-month

Table 1: Regression Results for 10-year Government Bond Spreads

	(1)	(2)	(3)	(4)
ΔBelief_t	-0.812 (1.002)			-0.910 (1.006)
$\Delta\text{Belief}_t \times \text{Crisis}_i$	-6.741*** (1.470)			-5.668*** (1.487)
$\Delta\text{Belief}_{t-1}$		-0.195 (1.014)		-0.203 (1.006)
$\Delta\text{Belief}_{t-1} \times \text{Crisis}_i$		-6.028*** (1.489)	-6.028*** (1.008)	-4.747*** (1.487)
Time Fixed Effects	No	No	Yes	No
Clustered Standard Errors	No	No	Crisis + Time	No
Observations	585	585	585	585
R ²	0.073	0.050	0.303	0.104

*p<0.1; **p<0.05; ***p<0.01

Note: the dependant variable is the government bond spread with 10-year maturity for all euro area members with available data. ΔBelief_t is the standardized change in the belief index. We use daily data from July 2nd to October 1st. Crisis countries are Greece, Italy, Portugal and Spain. Other countries in the sample are Austria, Belgium, Finland, France and the Netherlands.

horizon. With a standard deviation of 830, we have roughly 18 standard deviation increases in this time. Multiplying this with the interaction coefficient for the crisis countries on the same and the following day (columns (4) in Table 1), this back-of-the-envelope calculation implies that, on average, the 10-year government bond spread of crisis countries was reduced by 180 basis points due to this change in belief. This effect is mainly driven by the ‘Whatever it takes’ speech (4 standard deviation increase on the same day and 2 standard deviation on the following day, accounting for an 63 basis points reduction on average) and on the day of the official OMT-announcement (7 standard deviation on the same day and 0.5 on the following day, accounting for a 78 basis points reduction on average). This result is in line with the finding by Altavilla et al. (2016) who report a 200 basis points reduction in the bond yields of Spain and Italy due to the OMT announcements.

6.2. Robustness

We propose three robustness checks. At first, we focus on an alternative measure of sovereign risk, namely spreads of Credit Default Swaps (CDS). Then, we propose alternative versions of our belief index. Finally, we refine our analysis controlling for the number of followers.

6.2.1. Credit Default Swaps

Credit Default Swaps are a financial derivative which allows lenders to insure against the risk of a default by the debtor. The buyer of the CDS has to pay a fee (“spread”) for this insurance, typically quoted in basis points. Naturally, this spread is higher when a default is considered to be more likely. Hence it will be similar to sovereign bond spreads. For our purpose CDS spreads are a good comparison for robustness, since we lack bond spreads for Ireland at 2-year and 5-year maturity but do have the CDS spread.

Table 3 shows the results for the CDS spreads with maturity of 10 years. A one-standard deviation increase in the belief index reduces the spreads of crises countries by between 3.6 and 4 basis points on the same day and by between 1.6 and 2.5 basis points on the following day, depending on the specification. The results are significant throughout at the 5 percent level. Interestingly, in the regressions with CDS, the stand-alone term ΔBelief_t is also significant but of a smaller magnitude of less than one basis point and only at the 10 percent level. Table A5 in the appendix shows similar results for CDS spreads referring to 2 and 5-year maturities.

Table 2: Regression Results for 10-year Government Bond Spreads with Controls and Event Dummies

	(1)	(2)	(3)	(4)
$\Delta\text{Belief}_{t-1}$	0.139 (0.751)	-0.272 (0.430)	-0.270 (0.452)	0.331 (0.753)
$\Delta\text{Belief}_{t-1} \times \text{Crisis}_i$	-6.028*** (0.254)	-6.028*** (0.172)	-5.935*** (0.184)	-3.244*** (0.312)
ΔV2TX_t	5.073 (3.817)			
ΔCESI_t		1.383 (1.774)		
EventDummy_t			-1.408 (4.419)	
$\text{EventDummy}_t \times \text{Crisis}_i$			-12.404*** (1.848)	
$\text{TwoDay} - \text{EventDummy}_t$				-4.093 (3.875)
$\text{TwoDay} - \text{EventDummy}_t \times \text{Crisis}_i$				-16.869*** (1.625)
Time Fixed Effects	No	No	No	No
Clustered Standard Errors	Crisis + Time	Crisis + Time	Crisis + Time	Crisis + Time
Observations	585	585	585	585
R^2	0.120	0.052	0.060	0.084

*p<0.1; **p<0.05; ***p<0.01

Note: the dependant variable is the government bond spread with 10-year maturity for all euro area members with available data. ΔBelief_t is the standardized change in the belief index. We use daily data from July 2nd to October 1st. Crisis countries are Greece, Italy, Portugal and Spain. Other countries in the sample are Austria, Belgium, Finland, France and the Netherlands. ΔV2TX_t is the standardized change in the EURO STOXX 50 Volatility Index and ΔCESI_t is the standardized change in the Citigroup Economic Surprise Index. EventDummy_t is an indicator variable with the value 1 on July 26th, August 2nd and September 6th. $\text{TwoDay} - \text{EventDummy}_t$ is also an indicator variable which additionally includes the following day.

Table 3: Regression Results for 10-year Credit Default Swaps

	(1)	(2)	(3)	(4)
ΔBelief_t	-0.755* (0.437)			-0.793* (0.443)
$\Delta\text{Belief}_t \times \text{Crisis}_i$	-3.980*** (0.675)			-3.608*** (0.689)
$\Delta\text{Belief}_{t-1}$		-0.058 (0.456)		-0.018 (0.443)
$\Delta\text{Belief}_{t-1} \times \text{Crisis}_i$		-2.459*** (0.705)	-2.459*** (0.913)	-1.644** (0.689)
Time Fixed Effects	No	No	Yes	No
Clustered Standard Errors	No	No	Crisis + Time	No
Observations	650	650	650	650
R^2	0.112	0.031	0.402	0.125

*p<0.1; **p<0.05; ***p<0.01

Note: the dependant variable is the CDS spread with 10-year maturity for all euro area members with available data. ΔBelief_t is the standardized change in the belief index. We use daily data from July 2nd to October 1st. Crisis countries are Ireland, Italy, Portugal and Spain. Other countries in the sample are Austria, Belgium, Finland, France, Germany and the Netherlands.

6.2.2. Alternative belief indices

We now propose two robustness checks regarding the construction of the belief index. At first, we propose a new *definition* of the index. Second, we propose two alternative *measures* of the belief index.

Up to this point, we were considering both beliefs “1” and “-1”. We now turn to an alternative version of our belief index in which we only take into consideration the number of belief “1”.

$$\Delta\text{Belief}_t^{\text{alt}} = \sum_i \text{Tweet}_{i,t} \text{ '1'}$$

This alternative index is motivated theoretically if one assumes that the economy is in the self-fulfilling equilibrium to begin with and can either stay or switch to the fundamental “good” equilibrium. In other words, we consider the expression of a belief “1” as a change in belief itself.

Table ?? reports the results using the alternative belief index that only takes into account tweets that are labelled “1”. A one-standard deviation change in the lagged alternative belief index has a significant negative effect of -6.9 basis points on the 10 years bond spreads of the crises countries. This result is in line with the initial belief index.

In column (2) and (3), we propose alternative measures of the belief index. Mean_t is the mean of all labels on day t . PositiveRatio_t is the number of tweets with label “1” divided by the number of tweets with labels “0” and “-1”. We can observe that the corresponding coefficients have a slightly lower magnitude compared to the previous measure, but they are still of economically meaningful size and highly significant.

$$\text{Mean}_t = \frac{\Delta\text{Belief}_t}{\sum_i \text{Tweet}_{i,t}} \quad (4)$$

$$\text{Pos.Ratio}_t = \frac{\sum_i \text{Tweet}_{i,t} \text{ '1'}}{\sum_i \text{Tweet}_{i,t} \text{ '0'} + \sum_i \text{Tweet}_i \text{ '-1'}} \quad (5)$$

6.2.3. Controlling for the number of followers

Up to now, we have considered all the tweets to be equal. This is possibly problematic since users might differ in their level of information, in their importance as an information provider or as a market participant. In general, when working with such mainly anonymous

data, this is an unsolved question. Also, even if we had complete information on the users, it is conceptually not clear how to weight, for example, a journalist who possibly influences many market participants, the market participants themselves, or the actual decision-makers. One possible attempt in the literature has been to control for the level of information by using the number of followers (Gholampour and Van Wincoop 2017). This sets a threshold at 500 followers. Then, when a user’s number of followers is lower (larger) than 500, the user is considered to be uninformed (informed). We go a step beyond this by adding an additional restriction: we focus on the users that are active on the topic and have tweeted more than once. Results of estimating (3) taking into account only users with more than one tweet and more than 500 followers is shown in Table 5. Again, results hold with the same order of magnitude and are highly significant. Interestingly, Gholampour and Van Wincoop (2017) find different results when controlling for the users’ level of information. Our result is therefore consistent with a general change in belief during summer 2012, whatever the level of information of the users.

7. Changes in Belief for individual users

In this final section, we look at a different dimension of our Twitter data and compare belief before and after the key dates for single users. First, we sort the tweets for each user and compute the mean of labeled tweets before and after the central bank key events, like the ‘Whatever it takes’ speech. Second, we compute a t-test for paired samples to detect whether the difference in means is indeed pointing in a certain direction.

Since most of the users in our sample just tweet once and we require labeled tweets before and after key events, the number of users here is reduced. However, the sample still consists of more than 1,400 users with tweets before and after the ‘Whatever it takes’ speech and more than 3,300 users for the ECB Governing Council meetings on August 2nd and September 6th. We compute the mean of labeled tweets for each user, since the number of tweets varies and only this allows comparison across users.

Table 6 reports the results of our paired t-tests. For all the users with tweets before and after a key date, we compute the mean of labeled tweets after and before and then difference. We then compute the average and the standard deviation across all users, which is reported in the third and fourth column of Table 6. Interestingly, the mean difference is positive for July 26th and September 6th, meaning that users perceived an ECB intervention to be more likely after those events. For August 2nd, however, the mean difference is negative, meaning that users perceived an ECB intervention to be less likely after this event. We can formally test whether those mean differences are indeed significantly greater than zero for

Table 5: Regression Results for 10-year Government Bond Spreads, Belief Index only computed based on accounts with more than 500 followers

	(1)	(2)	(3)	(4)
ΔBelief_t	-0.833 (1.000)			-0.935 (1.003)
$\Delta\text{Belief}_t \times \text{Crisis}_i$	-6.798*** (1.466)			-5.667*** (1.481)
$\Delta\text{Belief}_{t-1}$		-0.245 (1.011)		-0.257 (1.003)
$\Delta\text{Belief}_{t-1} \times \text{Crisis}_i$		-6.230*** (1.483)	-6.230*** (1.065)	-4.932*** (1.481)
Time Fixed Effects	No	No	Yes	No
Clustered Standard Errors	No	No	Crisis + Time	No
Observations	585	585	585	585
R ²	0.075	0.054	0.305	0.108

*p<0.1; **p<0.05; ***p<0.01

Note: the dependant variable is the government bond spread with 10-year maturity for all euro area members with available data. ΔBelief_t is the standardized change in the belief index, computed after all tweets from accounts with less than 500 followers have been deleted. We use daily data from July 2nd to October 1st. Crisis countries are Greece, Italy, Portugal and Spain. Other countries in the sample are Austria, Belgium, Finland, France and the Netherlands.

Table 6: Paired t-tests

Date	N	Diff	SD	H1	p
26.07.2012	1462	0.255	0.532	> 0	< 0.001
02.08.2012	3470	-0.071	0.644	< 0	< 0.001
06.09.2012	3330	0.212	0.627	> 0	< 0.001

Note: this table shows results for paired sample t-tests at the three key dates of ECB communication. We compute the mean of labelled tweets after and before the key date and then take the difference. We then test whether the mean difference across all users is statistically greater than zero (for July 26th and September 6th) or smaller than zero (for August 2nd).

July 26th and September 6th using a one-directional t-test for paired samples. The p-values are smaller than any common significance level, thus we conclude that there are indeed within-user changes in belief at those key dates.

These results are consistent with the earlier, aggregate evidence that belief mainly changed at two events, the day of the ‘Whatever it takes’ speech and the day of the OMT-announcement. The ambiguous meeting of August 2nd, on the other hand, caused a small negative change in belief about central bank intervention of individual users.

8. Conclusion

This paper pursues a new approach to study the financial effects of the OMT program announcements. We apply a textual analysis to Twitter data in order to extract belief about the perceived likelihood of central bank intervention. We show that a belief index based on tweets spikes at two important dates of ECB communication - the day of the ‘Whatever it takes’ speech and the day of the official announcement of the OMT program. Empirically, our created belief index can account for sizeable decreases of the bond spreads of distressed countries in the eurozone. We contribute to the literature in three ways: First, in line with other recent work (Meinusch and Tillmann 2017), we show that social media data reveals useful information for the expectation formation about monetary policy. This contribution might also be relevant in different fields because survey data at high frequency is typically unavailable. Second, we show that this methodology can improve on simple event studies because it can account for an expectation formation over the full horizon and not only on event days. Comparing our belief index to the work by Altavilla et al. (2016), our results look similar to a dummy approach at two points in time - the day of the ‘Whatever it

takes' speech and the day of the official announcement of the OMT program. However, our results are markedly different for the ambiguous communication after the Governing Council on 2nd August, where we do not find evidence of aggregate changes in belief, and on an individual user level, even evidence for a negative change in belief. Furthermore, our belief index captures rumors and information leaks in advance of event days. Given those findings, we regard our belief index as a “microfoundation” of event dummies. Third, although we do not formally test a sovereign debt model, our results indicate that a credible commitment by a central bank to act as lender of last resort can be used as a coordination device in a sovereign debt crisis.

References

- Altavilla, Carlo, Domenico Giannone, and Michele Lenza (2016). “The Financial and Macroeconomic Effects of the OMT Announcements”. *International Journal of Central Banking* 12.3, pp. 29–57.
- Ambler, Steve, Fabio Rumler, et al. (2017). “The effectiveness of unconventional monetary policy announcements in the euro area: an event and econometric study”. *Oesterreichische Nationalbank Working Paper 212*.
- Azar, Pablo D and Andrew W Lo (2016). “The Wisdom of Twitter Crowds: Predicting Stock Market Reactions to FOMC Meetings via Twitter Feeds”. *Journal of Portfolio Management* 42.5, p. 123.
- Bocola, Luigi and Alessandro Dovis (2016). “Self-Fulfilling Debt Crises: A Quantitative Analysis”. *NBER Working Paper No. 22694*.
- Bollen, Johan, Huina Mao, and Xiaojun Zeng (2011). “Twitter mood predicts the stock market”. *Journal of computational science* 2.1, pp. 1–8.
- Bruegel (2012). *The SMP is dead. Long live the OMT*. URL: <http://bruegel.org/2012/09/the-smp-is-dead-long-live-the-omt/> (visited on 08/13/2018).
- Calvo, Guillermo A (1988). “Servicing the public debt: The role of expectations”. *The American Economic Review*, pp. 647–661.
- Cole, Harold L and Timothy J Kehoe (1996). “A self-fulfilling model of Mexico’s 1994–1995 debt crisis”. *Journal of international Economics* 41.3-4, pp. 309–330.
- (2000). “Self-fulfilling debt crises”. *The Review of Economic Studies* 67.1, pp. 91–116.
- Corsetti, Giancarlo (2015). “The mystery of the printing press”. *Schumpeter Lecture at the EEA*.

- Corsetti, Giancarlo and Luca Dedola (2016). “The mystery of the printing press: Monetary policy and self-fulfilling debt crises”. *Journal of the European Economic Association* 14.6, pp. 1329–1371.
- D’Amico, Stefania (2016). “Discussion of ‘The Financial and Macroeconomic Effects of the OMT Announcements’”. *International Journal of Central Banking* 12.3, pp. 59–68.
- De Grauwe, Paul and Yuemei Ji (2012). “Mispricing of sovereign risk and macroeconomic stability in the Eurozone”. *Journal of Common Market Studies* 50.6, pp. 866–880.
- Di Cesare, Antonio, Giuseppe Grande, Michele Manna, and Marco Taboga (2013). “Recent estimates of sovereign risk premia for euro-area countries”. *Bank of Italy Occasional Paper No. 128*.
- ECB (2012a). *Introductory statement to the press conference (with Q&A)*. URL: <https://www.ecb.europa.eu/press/pressconf/2012/html/is120802.en.html> (visited on 03/19/2018).
- (2012b). *Introductory statement to the press conference (with Q&A)*. URL: <https://www.ecb.europa.eu/press/pressconf/2012/html/is120802.en.html> (visited on 03/19/2018).
- El País, Newspaper (2012). *Draghi says ECB buying short-term bonds does not breach EU Treaty*. URL: https://elpais.com/elpais/2012/09/03/inenglish/1346696018_263197.html (visited on 08/14/2018).
- Falagiarda, Matteo and Stefan Reitz (2015). “Announcements of ECB unconventional programs: Implications for the sovereign spreads of stressed euro area countries”. *Journal of International Money and Finance* 53, pp. 276–295.
- Fendel, Ralf and Frederik Neugebauer (2018). “Country-specific Euro area government bond yield reactions to ECB’s non-standard monetary policy announcements”. *WHU Working Paper 18/02*.
- Gholampour, Vahid (2017). “What Can We Learn from Stock Market Tweets?” *Working Paper*.
- Gholampour, Vahid and Eric Van Wincoop (2017). “What can we Learn from Euro-Dollar Tweets?” *NBER Working Paper No. 23293*.
- Hansen, Stephen and Michael McMahon (2016). “Shocking language: Understanding the macroeconomic effects of central bank communication”. *Journal of International Economics* 99, S114–S133.
- Hansen, Stephen, Michael McMahon, and Andrea Prat (2014). “Transparency and deliberation within the FOMC: a computational linguistics approach”. *The Quarterly Journal of Economics*.

- Krishnamurthy, Arvind, Stefan Nagel, and Annette Vissing-Jorgensen (2017). “ECB policies involving government bond purchases: Impact and channels”. *Review of Finance* 22.1, pp. 1–44.
- Les Echos, Newspaper (2017). *Il y a cinq ans, Draghi sauvait l’euro en une phrase*. URL: <https://www.lesechos.fr/amp/57/2104457.php> (visited on 08/14/2018).
- Meinusch, Annette and Peter Tillmann (2017). “Quantitative Easing and Tapering Uncertainty: Evidence from Twitter”. *International Journal of Central Banking* 13.4, pp. 227–258.
- Roch, Francisco and Harald Uhlig (2018). “The Dynamics of Sovereign Debt Crises and Bailouts”. *Journal of International Economics* 114.136, pp. 1–13.
- Scotti, Chiara (2016). “Surprise and uncertainty indexes: Real-time aggregation of real-activity macro-surprises”. *Journal of Monetary Economics* 82, pp. 1–19.
- Szczerbowicz, Urszula (2015). “The ECB unconventional monetary policies: have they lowered market borrowing costs for banks and governments?” *International Journal of Central Banking* 11.4, pp. 91–127.
- Van Der Heijden, Minke, Roel Beetsma, and Ward Romp (2018). “‘Whatever it takes’ and the role of Eurozone news”. *Applied Economics Letters* 25.16, pp. 1166–1169.
- Zhang, Xue, Hauke Fuehres, and Peter A Gloor (2011). “Predicting stock market indicators through twitter ”I hope it is not as bad as I fear””. *Procedia-Social and Behavioral Sciences* 26, pp. 55–62.

Appendix

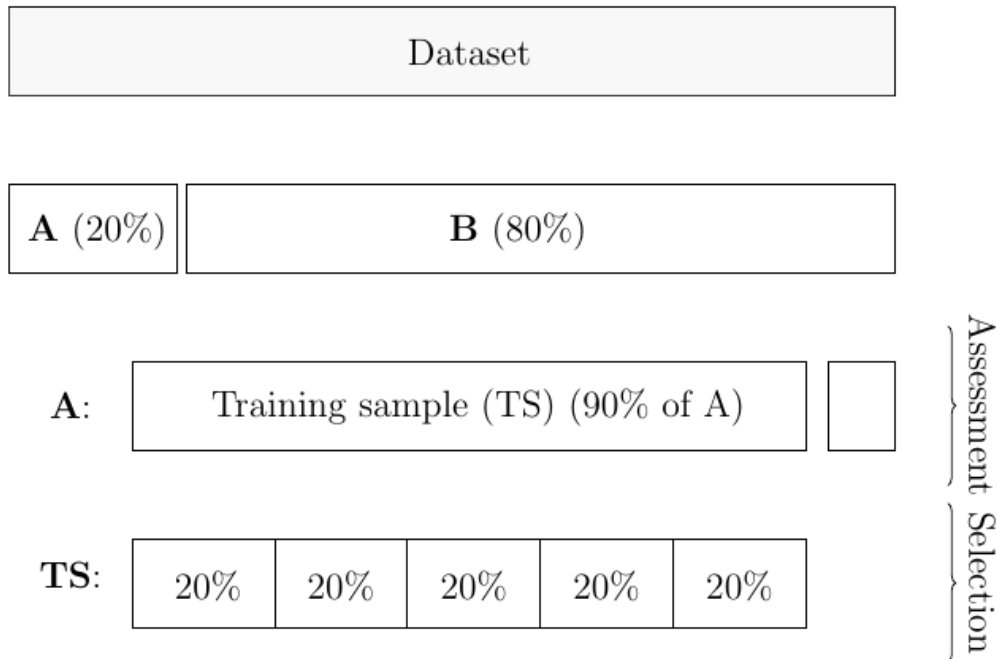
Table A1: Examples of tweets

Text	Label
Draghi's opens door to new ECB policy territory	1
Draghi skips Jackson Hole as ECB shapes bond plans - US News and World Report	1
Monti - - Notes Draghi commented on Time Lag Between Govt Action and fall in Spreads and Said ECB may Buy Bonds	1
Reports Of Talks Between Mario Draghi & Bundesbank Head Signal That ECB Measures R 4 Real	1
IMPORTANT..Draghi has opened poss 4 ECB bondbuys w/o EFSF/ESM actually being touched for sov reasons	1
Bloomberg: Draghi said to to give fellow ECB'ers 24 hours +/-s to digest rescue plan before Sept. 6 meeting	1
El Pais:No Taboos at Next ECB Meet:Notes comments from Draghi: they suggest relief measures for Italy & Spain could be on way citing sources	1
No Bazooka As ECB Backtracks: Draghi Won't Pursue Yield Caps, To Sterilize Bond Buys In SMP Continuation	-1
Debt crisis: ECB's Draghi Plan doused by rebellions in Germany and Greece	-1
Super Mario disappoints. Just wondering what Draghi meant by "ready to do whatever it takes to preserve the euro" #ECB talks Italian style..	-1
Global stocks tumble as investors reacted to disappointment that the ECB's Mario Draghi failed to match his words with prompt action	-1
So Draghi says that ECB buying EU member state sovereign bonds is *not* state aid?! Oh sure, whatever you say..	-1
Draghi Overpromised What the ECB Could Achieve	-1
ECB's Draghi gives no hint of bond buys, LTROs	-1
ECB Follows Words With More Words - the market fears Draghi has written a cheque he can't cash	-1
Draghi : "the EURO is irreversible"... again, same speech, no news... #ECB	-1
Sceptics abound as Mario Draghi's ECB bond 'bluff' electrifies global markets	-1
ECB Draghi is useless and his words mean nothing . Eu should break up . Eu does not know what to do at this point .	-1
Draghi says vote not unanimous – there was one dissent. Guess whooooo? #Germany #ECB	0 0
All Eyes on ECB's Draghi to Fight Crisis	0
I wonder if economists thought of what comes to everyone's mind when they say Mario Draghi's nickname, Super Mario #ECB #Europe #Eurozone	0
ECB President Draghi Speaks in 4 mins	0
Draghi: economic growth in euro area remains weak #ECB	0
ECB'S DRAGHI: There will be more transparency than before.	0

Table A2: Summary Statistics for daily Number of Tweets

N	Mean	Std.Dev.	0%	25%	50%	75%	100%
66	750	1653	13	93	155	823	9601

Fig. A1. Double cross-validation procedure



Note: this figure shows how our dataset is organized. The tweets in “A” are manually labeled while those in “B” will be predicted by a machine learning classifier. The last two layers describe our double cross-validation procedure. The set “A” is randomly split into a training set (90 percent) and a test set (10 percent). To select our model (to “hypertune” the parameters of the SVM classifier), we proceed with a grid search using a 5-fold cross-validation. The accuracy of the selected model is then assessed on the test set.

Table A3: Summary Statistics for ΔBelief_t

N	Mean	Std.Dev.	0%	25%	50%	75%	100%
66	231	829	-127	0	22	96	5695

Table A4: Regression Results for Government Bond Spreads at lower Maturities

	5-year-maturity			2-year-maturity		
	(1)	(2)	(3)	(4)	(5)	(6)
ΔBelief_t	-0.896 (0.763)			-0.455 (0.849)		
$\Delta\text{Belief}_t \times \text{Crisis}_i$	-6.830*** (1.119)			-4.946*** (1.245)		
$\Delta\text{Belief}_{t-1}$		-0.274 (0.802)			0.045 (0.863)	
$\Delta\text{Belief}_{t-1} \times \text{Crisis}_i$		-3.741*** (1.176)	-3.741*** (0.539)		-3.495*** (1.266)	-3.495*** (0.798)
Time Fixed Effects	No	No	Yes	No	No	Yes
Clustered Standard Errors	No	No	Crisis + Time	No	No	Crisis + Time
Observations	585	585	585	585	585	585
R ²	0.125	0.034	0.394	0.053	0.022	0.262

*p<0.1; **p<0.05; ***p<0.01

Note: the dependant variable is the government bond spread with 5-year and 2-year maturity respectively for all euro area members with available data. ΔBelief_t is the standardized change in the belief index. We use daily data from July 2nd to October 1st. Crisis countries are Ireland, Italy, Portugal and Spain. Other countries in the sample are Austria, Belgium, Finland, France and the Netherlands.

Table A5: Regression Results for CDS Spreads at lower Maturities

	5-year-maturity			2-year-maturity		
	(1)	(2)	(3)	(4)	(5)	(6)
ΔBelief_t	-0.552 (0.461)			-0.258 (0.674)		
$\Delta\text{Belief}_t \times \text{Crisis}_i$	-5.243*** (0.711)			-5.914*** (1.041)		
$\Delta\text{Belief}_{t-1}$		-0.096 (0.487)			0.112 (0.691)	
$\Delta\text{Belief}_{t-1} \times \text{Crisis}_i$		-3.072*** (0.752)	-3.072*** (1.057)		-4.079*** (1.067)	-4.079*** (1.255)
Time Fixed Effects	No	No	Yes	No	No	Yes
Clustered Standard Errors	No	No	Crisis + Time	No	No	Crisis + Time
Observations	650	650	650	650	650	650
R ²	0.143	0.042	0.417	0.081	0.033	0.280

*p<0.1; **p<0.05; ***p<0.01

Note: the dependant variable are the CDS spreads with 5-year and 2-year maturity respectively for all euro area members with available data. ΔBelief_t is the standardized change in the belief index. We use daily data from July 2nd to October 1st. Crisis countries are Ireland, Italy, Portugal and Spain. Other countries in the sample are Austria, Belgium, Finland, France, Germany and the Netherlands.