



Nr. 2020/05

Above, but close to two percent.
Evidence on the ECB's inflation target
using text mining

Johannes Zahner
MACIE, Philipps-Universität Marburg

Marburg Centre for Institutional Economics • Coordination: Prof. Dr. Elisabeth Schulte
c/o Research Group Institutional Economics • Barfuessertor 2 • D-35037 Marburg

Phone: +49 (0) 6421-28-23196 • Fax: +49 (0) 6421-28-24858 •
www.uni-marburg.de/fb02/MACIE • macie@wiwi.uni-marburg.de

Above, but close to two percent. Evidence on the ECB's inflation target using text mining*

Johannes Zahner[†]

October 1, 2020

Abstract

Due to its official mandate, the European Central Bank (ECB) is assumed to maximize an implied objective function that leads it to pursue inflation with a subordinate focus on supporting the general economic policy of the European Union. This objective is – by its very nature – difficult to quantify. My paper tries to decipher information regarding the ECB's objective through the use of text mining on all public speeches between 2002 and 2020. The estimation of a sentiment index through a 'bag-of-words'-approach yields the following results. First, the findings of my analysis suggest a concave objective regarding the inflation rate. The implied inflation target is best summarized as an inflation rate of '*above, but close to 2%*'. Deviations from this target lead to a reduction in the sentiment of the institutions' communication. Second, my findings suggest a convex objective towards output growth and a linear objective towards the unemployment rate, with a preference for higher GDP growth and employment independently of the current level. Furthermore, the hierarchical order in the the European Central Bank (ECB)'s mandate does not always appear to be consistent with my findings. Deviations from its primary objective, the inflation rate, appear to be of no greater concern than deviations in its subordinate objective. Third, in periods of heightened uncertainty, there is an additional decrease in the sentiment of speech. Last, over the last two decades, speeches have become more pessimistic, even when controlling for macroeconomic conditions.

Keywords: Sentiment Analysis, ECB, Monetary Policy, Public Perception.

JEL Codes: E53, E58, E61

*I would like to thank Elisabeth Schulte and Marc Steffen Rapp for sparking my interest in this topic. Furthermore, I am grateful to Bernd Hayo and Matthias Neuenkirch for invaluable input. The paper also benefitted from comments by Peter Tillmann, Carsten Hefeker and other participants of the MAGKS Doctoral Colloquium 2020.

[†]Corresponding author. Marburg Centre for Institutional Economics (MACIE), Philipps-University Marburg. Email: johannes.zahner@wiwi.uni-marburg.de. Address: Barfuessertor 2 D-35037 Marburg, Germany.

"We communicate what we are trying to achieve, and how we go about achieving it. In practical terms, this means that communication revolves around providing a narrative about the economy and the outlook for price stability relative to our objective." Praet (2014)

1 Introduction

The relevance of narratives and perceptions in explaining macroeconomic phenomena is gaining importance in the economic discourse (Shiller (2017)). Lorenzo Bini Smaghi, Member of the Executive Board of the European Central Bank (ECB), summarizes this importance with respect to monetary policy by stating that *"policy-makers (including central banks) have to take [...] perceptions seriously. In economics, perceptions shape reality"* (Smaghi (2007)). Since the effectiveness of monetary policy depends on the belief of the general public in the inflation target (Diron and Mojon (2005)), the perception is directly proportional to the ECB's success in maintaining its target – price stability. As a result, the ECB is thought to nurture a perception around the primary objective of maintaining price stability in the euro area since its founding. Thus, the central bank invests enormous efforts to shape a narrative around this institution and thereby the perception of the ECB itself. For example, after an internal evaluation of its monetary policy strategy in May 2003 the central bank clarified its inflation objective and introduced the component *"close to [2 % inflation]"* to its existing definition of price stability. This clarification was intended to eliminate the ECB's previous ambiguity regarding optional inflation and further anchor inflation expectations.

However, some ambiguity with respect to the central bank's objective function remains (Miles et al. (2017)). Commonly assumed to contain a some relationship towards inflation and economic activity, specific quantitative measurements of the function's shape are unknown. This paper follows Shapiro and Wilson (2019)'s work in approximating a central bank's objective function using semantic information from the ECB's communication. In other words, by the use of text mining, my paper aims to reveal the narrative of the ECB's objective function that is being nurtured by its members.

In recent years, with the restriction of the effective lower bound, communication has become an even more relevant monetary policy tool (Lucca and Trebbi (2009)). The ECB has various communication tools to shape the public perception that target a broader audience. Alongside its blog, podcast, and 'ECB Listens'-events, the primary means of ECB communication to a wider audience are still public speeches. Consequently, over the last two decades, members of the ECB Executive Board gave more than 2,000 speeches, and thereby provide an abundance of qualitative information regarding the narrative they attempt to nutritionalise. This paper quantifies this qualitative data through the use of text mining. Quantifying the communication allows for the approximation of otherwise intangible objectives such as the inflation target.

To approximate such latent variables, I rely on a 'bag-of-words' approach to quantify the respective narrative for each speech. Taking advantage of the heterogeneity in economic conditions – i.e. inflation and economic activity – allows me to estimate the communicated objective function of the ECB. Ultimately, I will be able to identify the respective bliss point towards the variables in this function. This means that speeches becomes more pessimistic (optimistic), above and below the maximum (minimum) – implying an articulated target. My findings suggest that the ECB communicates an inflation target that is *"above, but close to 2%"*. I find further evidence that the subordinate objective of promoting the general economic policy of the European Union (EU) is best aligned with a relationship where the central bank favors better economic conditions,

unconditionally of the present level. Furthermore, my findings suggest that speeches response equally strong to changes in economic activity as to changes in the inflation. Moreover, speeches become more pessimistic over time, even when controlling for macroeconomic and financial variables.

The contribution of this work is to be found in the areas of text mining, monetary policy, and public perception. Using modern text mining techniques, this paper contributes to a better understanding of the perception of monetary policy in the euro area. It is, to my best knowledge, the first work on the monetary policy objective function, using text analysis on speeches the euro area.

This paper is structured in the following order. The second section provides a brief literature overview of text mining applications in the central bank communication literature. Next, the empirical approach to estimate the latent inflation target is introduced. The text-mining approach used to derive the left-hand side variable, as well as the underlying datasets for the right-hand side variables, are then introduced and discussed in depth in the fourth section. The results and their implications are addressed in the fifth section before the final section concludes this paper.

2 Literature review

The relevance of communication as a monetary policy instrument has increased substantially in central banking. Blinder et al. (2008) identify four rationalizations that require communication by a central bank, including the degradation of asymmetric information between the general public and the central bankers. The content of this asymmetric information could originate from a variety of sources. It may regard the outcome of previous policy votes (Meade (2005)), the content of committee members' deliberations (Hansen, McMahon, and Prat (2017)), or the central bankers' risk balance assessment (Hanson and Stein (2015)). As a result, there is a myriad of applications where tone variations in monetary policy communication are used to evaluate market responses.¹ Reactions to changes in sentiment were found by asset price markets (Schmeling and Wagner (2019)), sovereign yield spreads (Falagiarda and Reitz (2015)), and short-term interest futures (Rosa and Verga (2008)). Furthermore, improvements in the predictability of future monetary policy decisions, using sentiment analysis (Apel, Grimaldi, and Hull (2019); Baranowski, Bennani, and Doryń (2020)), have been found.

The number of applications using of text mining in analyzing central bank communication has increased substantially thanks to the rise in computational power and the abundance of available text. Inherent in those papers – as in mine – is the assumption that market participants believe a central bank obeys a coherent systemic approach that allows for the inference with respect to its parameters. However, while earlier text mining literature assumed central bankers would communicate coherently and observe market reactions to changes in communication, my analysis concentrates on changes in communication in response to variation in economic conditions. To be precise, my paper concentrates on the institutionalized objective function underlying the ECB and thus the communicated objectives towards inflation and economic activity that it implies. Since the perception per se – i.e. how the audience perceives the communication – is difficult to measure, my analysis focuses instead on the central bank's narrative through its communication – i.e. how it wants to be perceived.

¹The terms 'tone' and 'sentiment' are used substitutable throughout this paper. In section 4.1, sentiment will be defined as a measure of optimism in one's language.

The closest papers in this respect are by Shapiro and Wilson (2019) and Paloviita et al. (2020). Shapiro and Wilson (2019) provide the econometric model applied in this paper. While their work focuses on transcript of Federal Open Market Committee (FOMC) meetings – and therefore does not emphasize the narrative –, my analysis uses public information from the Executive Board members of its European counterpart. Their results differ from mine predominantly on the dimension of magnitude in the inflation target. The FOMC had no explicit inflation target before 2012. Nevertheless, it has long been assumed to be close to ECB’s. However, Shapiro and Wilson (2019) find an implicit inflation objective of 1.5% for the FOMC, substantially lower than the expected 2%. Their estimated inflation target prevails at 1.5%, even when considering public speeches and a dual mandate – the most analogous regression to mine. Contrary, my findings suggest that the ECB – which had an explicit inflation objective since its beginning – maintains a communicated inflation target above 2%.

Paloviita et al. (2020) replicate the results for the euro area using the ECB’s introductory statements to its monthly press conference. In section 4.1, I show that the use of these statements, in contrast to speeches, has the disadvantage of increasing similarity over time, which might explain the discrepancy in result. While Paloviita et al. (2020) find an implicit 1.7% inflation target, my results indicate a target that is significantly higher than 2%.

3 Econometric model

This paper aims to empirically identify the latent inflation target that is communicated by the ECB by the use of sentiment analysis. I follow the econometric approach for identification presented by Shapiro and Wilson (2019). In short: the method approximates a central bank’s loss by a semantic index. Assuming a traditional New-Keynesian loss function, this, in turn, allows estimating the central banks’ perceived inflation target – using information on current inflation.

Following the textbook literature (e.g. Galí (2015); Walsh (2017)), a central bank is modelled as minimizing a loss function. Assuming that its members share one function – or at least support the same narrative –, the utility loss L_t at time t can be described as

$$L_t = \hat{\pi}_t^2 + \phi \hat{y}_t \quad (1)$$

where $\hat{\pi}_t$ (\hat{y}_t) represent deviations in inflation (economic activity) and ϕ the relative weight on economic activity. Note that both, the inflation gap and the gap in economic activity, are latent variables, i.e. can not be directly observed. They are defined as the difference of the observable current inflation π_t (current economic activity y_t) and the target inflation π^* (target economic activity y^*) – the variable at interest in this study. Consequently, equation (1) can be written with $\hat{\pi}_t = \pi_t - \pi^*$ and $\hat{y}_t = y_t - y^*$.

The inclusion of an inflation and an economic activity term in the monetary policy loss function stems from the official mandate ECB set out in the Treaty on the Functioning of the European Union. The Treaty specifies two objectives that the ECB must adhere to in hierarchical order. Article 123 §1 sets out the primary objective of maintaining price stability in the medium term. A subordinate objective of promoting the general economic policies of the European Union is laid out in Article 3. This is usually interpreted as an objective regarding output. However, as one of the European Union’s explicit targets relates to employment levels, the

unemployment rate is sometimes used as an appropriate indicator (e.g. Molodtsova and Papell (2012)). In section 5.4, I present findings with respect to both hypotheses, a measure for output and unemployment as economic activity terms, demonstrating irrelevance of this choice to the overall results.

One essential presumption in the following analysis is that public statements of the central bank contain information that can be approximated by the tone of those discourses. In the communication literature on monetary policy, this principle has been applied numerous times. For instance, Bennani and Neuenkirch (2017) demonstrate that central bankers' sentiment in public communication is affected by current economic conditions, and Bennani et al. (2019) provide evidence that this sentiment could be used as a predictor for future interest rate decisions. My paper continues with this practice of suggesting that speeches are a communication device used to inform the public about a central bank's objective and, thus, the words chosen in a speech can approximate that objective. Assuming that the loss function from equation (1) can be approximated by a sentiment index $S_{i,t}$, the sentiment for speaker i at time t can be written as follows:

$$\begin{aligned} S_{i,t} &= \gamma_i L_t \\ &\approx \beta_{0,i} + \beta_1 \hat{\pi}_t^2 + \beta_2 \hat{y}_t + \varepsilon_t. \end{aligned} \quad (2)$$

The introduction of a speaker-specific constant allows for personal and unique communication nuances between the speaker that are unrelated to the objective. It may also represent different perceptions on the shared narrative of the central bank. Thereby the constant ensures that the Executive Board's composition does not affect the estimated loss function. The necessity of this speaker specific term is further elaborated in section 4.1, where the heterogeneity of the Executive Boards' communication is addressed. Disassembling the latent variables into its components, equation (2) can be reformulated as:

$$\begin{aligned} S_{i,t} &= \beta_{0,i} + \beta_1 (\pi_t - \pi^*)^2 + \beta_2 (y_t - y_t^*) + \varepsilon_t \\ &= \Phi_i + \beta_1 \pi_t^2 + \Omega \pi_t + \beta_2 y_t + \varepsilon_t \end{aligned} \quad (3)$$

with $\Phi_i = \beta_{0,i} + \beta_1 \pi^{*2} + \beta_2 y^*$ and $\Omega = -2\beta_1 \pi^*$. Thus, the communicated inflation target can be effectively calculated as:

$$\pi^* = -\frac{1}{2} \frac{\Omega}{\beta_1} \quad (4)$$

Assuming that the perceived quadratic loss function can be approximated by the sentiment index, equation (4) enables me to calculate the global minimum or maximum in the loss function with respect to the inflation rate. In other words, if the members of the ECB's Executive Board communicate an objective, which is, say, concave to the inflation rate, the sentiment in its speeches should be the highest – the most positive – if the current inflation is equal to π^* . At this point, the deviation from the inflation target is effectively zero, which in turn minimizes equation (1).

4 Data

Following from the previous section, a sentiment indicator ($S_{i,t}$) as well as data on the objectives – inflation (π_t) and economic activity (y_t) – is required. Hence this section's focus is twofold. The estimation of the

Figure 1: Similarity of ECB speeches.



Note: The plot above illustrates three similarity indices estimated by Amaya and Filbien (2015)’s measurement. Similarity is estimated as the occurrence of the same bigrams in two consecutive months (left plot) and within a month in different speeches (right plot) against time. Consequently, a high similarity (near one) is a sign for homogeneous texts, and the similarity is low (near zero) for heterogeneous text. The appendix provides a brief overview as to what bigrams are and a summary of Amaya and Filbien (2015)’s index. The author’s calculations for all three indices are available upon request.

dependent variable, i.e. the sentiment index, is discussed first. The second part discusses on the right-hand side macroeconomic variables.

4.1 Sentiment Index

Studies quantifying central bank communication differ mainly in two aspects: the underlying text corpus and the method of quantifying the information in the text.

With respect to the underlying text corpus, most research on the ECB focuses on the organization’s introductory statements to the monthly press conference following its monetary policy decision. While this permits for coherent observations – with constant time, location and communicator – it has two shortcomings: few observations and increasing similarity. Using exclusively press conference statements disregards much of the potentially relevant information the central bank provides to market participants. To be explicit: There are ten times the number of speeches compared to press conferences throughout my observation period. This larger amount of observation renders the narrative quantification simpler. In addition, speeches remain heterogeneous in their language throughout the observation period as opposed to the press conference texts. Amaya and Filbien (2015) find the similarity of of the ECB’s press statements increased about fivefold between 1998 and 2014. A graphical illustration of both observations is provided in figure 1, where each red square presents one ECB’s introductory statement and each green triangle a speech. It is apparent that the public speeches between 2002 and 2020 do not exhibit such non-stationarity behavior in the similarity index. Although not as abundant as studies on press conferences, there exists literature that investigates the ECB’s

communication beyond press statements. Significant contributions were made through the analysis of media coverage (Hayo, Kutan, and Neuenkirch (2008); Bennani et al. (2019)), tweets (Masciandaro, Romelli, and Rubera (2020)) and speeches (Gertler and Horvath (2018); Bennani and Neuenkirch (2017)). My work is closely related to the latter two studies, however, I use a novel dataset that enables me to evaluate all of the central bank's Executive Board public speeches. The speech corpus used in this analysis is larger in size, covers a longer time period and is uniform in terms of the organization represented. Thus differences in tone nuances can be attributed to the individual, rather than the speakers organisation.

For the calculation of the sentiment index, I rely on the speech dataset provided by the ECB itself (European Central Bank (2019)). The speech corpus contains (all) public speeches made by members of the Executive Board between January 1998 and January 2020. The dataset provides information on 2150 speeches, including the name of the speaker, the title, the subtitle, the date, and, of course, the text of the speech itself. Due to the missing observation in the dataset, I am forced to discard speeches held before January 2002. In addition, restrictions on word-, sentence- and speech-level are implemented.

To begin with, I convert all words to lower case and remove numbers and special signs. On the sentence level, I follow the convention of Shapiro and Wilson (2019) to include only sentences concerning economic content. Therefore, each sentence must contain at least one economic word or phrase as defined in the Oxford Dictionary of Economics (Black, Hashimzade, and Myles (2017)). This restriction intends to remove the non-economic parts of the speech – like obligatory thank-you remarks in the beginning. Although the dictionary is quite exhaustive, with more than 3,500 words and phrases, almost one-third of the sentences in the corpus do not pass this restriction. On the speech level, my restriction follows Bennani and Neuenkirch (2017). Thereby, each speech must contain at least 25 of the 'economic sentences' defined above. Imposing these restrictions yields a corpus of 1,915 speeches and roughly 170,000 sentences from 21 board members over a time horizon of 18 years.

There are different ways to reduce the dimensionality of such a dataset – i.e. to quantify text. The interested reader is referred to more comprehensive literature on textual analysis such as Gentzkow, Kelly, and Taddy (2019) and Bholat et al. (2015). The three most common methods are based on manually rating text by hand, enabling a computer to rate text according to a predefined dictionary and unsupervised machine learning techniques. This paper relies on the dictionary approach – often called a 'bag-of-words' – to calculate a sentiment index, using a predefined lexicon to count positive and negative terms. In recent years, the 'bag-of-words' approach on central bank communication text has been applied to the ECB (Tobback, Nardelli, and Martens (2017)), the Bundesbank (Tillmann and Walter (2018)), the FOMC (Apel, Grimaldi, and Hull (2019)) and the Riksbank (Apel and Grimaldi (2014)), among others. Regarding the lexicon, Loughran and McDonald (2011) provide evidence that the frequently applied Harvard dictionary causes misleading results when applied in the financial sector. Instead, they developed their lexicon, identifying around 2,700 words that are classified as either positive or negative. Although their dictionary is heavily biased towards negative terms – the updated dictionary contains 2,355 negative terms and only 354 positive terms – I will later show that the average sentiment estimated with this lexicon is not significantly different from zero.²

Following the literature on central bank communication (Bholat et al. (2015)) the sentiment index $S_{i,t}$ is estimated by the relative fraction of positive minus negative terms. After removing positive and negative words that are preceded by "not" – as recommended by Loughran and McDonald (2016) – the sentiment for

²An example text for illustration purposes is provided in appendix A.1

Table 1: Statistic summary of Sentiment Index

	Number of sentences	Negative terms	Positive terms	Sentiment score $S_{i,t}$
Mean	91.14	47.78	36.66	-0.06
Std. dev.	(51.76)	(38.69)	(23.73)	(0.36)

each speech is estimated to be the average sentiment of all sentences j in that speech:

$$S_{t,i} = \frac{1}{n} \sum_{j=1}^n \frac{\#positive_j - \#negative_j}{\#positive_j + \#negative_j} \quad (5)$$

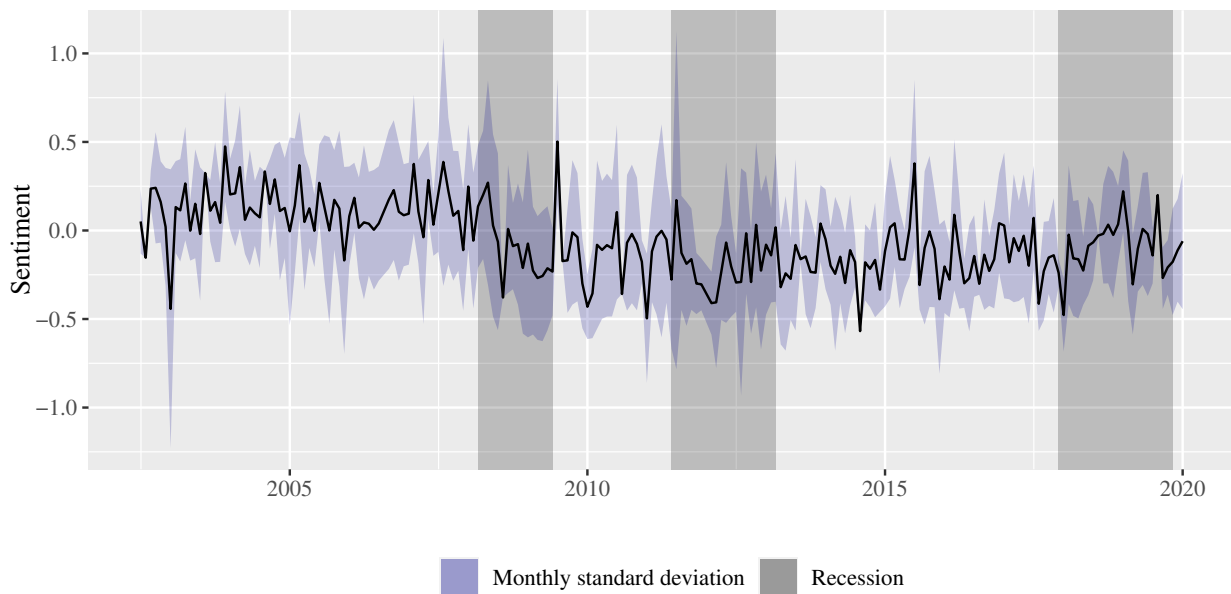
Using Loughran and McDonald (2011)'s dictionary, the average speech contains approximately 48 negative terms and 37 positives. The resulting mean sentiment score is therefore -0.06 . However, with a standard deviation of 0.36, the average sentiment is not significantly different from zero. The index's statistical summary can be found in table 1.

In order to assess time-dependent developments, the average monthly sentiment and its standard deviation are graphically illustrated in figure 2. Two noteworthy observations arise. First, there seems to be no adjustment in the sentiment during economic downturns (grey bars). One might expect the sentiment of those speeches to reflect the economic condition during the crisis, however, there seems to be no evidence in this graphical representation. Second, there appears to be a shift from predominantly positive communication in the 2000s to largely negative communication in the 2010s. Given that monetary policy in the 2010s was dominated by the presence of the effective lower bound and inflation rates that are considerably below the explicit target, this observation provides first anecdotal evidence that the sentiment index measures the narrative of the institutions' objective and therefore contains useful information. This interpretation should be regarded with caution, as given the small number of observations per month this trend is not statistically significant.

Although the index appears to be stationary, a unit root test (Dickey and Fuller (1979)) on the sentiment scores cannot be rejected. To control for such non-stationarity, I include a delayed sentiment term into the regression. Since speeches do not occur on a regular basis (e.g. once every month), there are two possibilities to include such a term. In the baseline regression, the sentiment term from the previous speech is included, averaging the sentiment score when more than one speech was held on a single day. However, suspecting that a central bankers sentiment in a particular speech might be affected by the 'general' mood of her peers during speech writing, I regress the lagged sentiment variable with the average sentiment of the previous month in one robustness check.

There are three potential shortcomings in the text corpus. Table 3 in the appendix provides an extended summary statistic for each Executive Board member which highlights the first two shortcomings discussed here. First, the underlying text corpus varies with regard to the number of speeches per speaker. Some Executive Board members (in particular Jean-Claude Trichet) deliver far more speeches than others and are therefore more likely to influence overall perception. Second, the speaker's average sentiment varies widely within the Executive Board as illustrated in figure 3. This may suggest that there are members in the ECB's Executive Board who are generally more positive (or negative) in their communication. However, note that the statistical rule of mean reversal appears to adhere to the sentiment scores, since the average score of the three most frequent speakers (Jean-Claude Trichet, Benoit Coeure, and Mario Draghi) can be found at the

Figure 2: Monthly sentiment score.

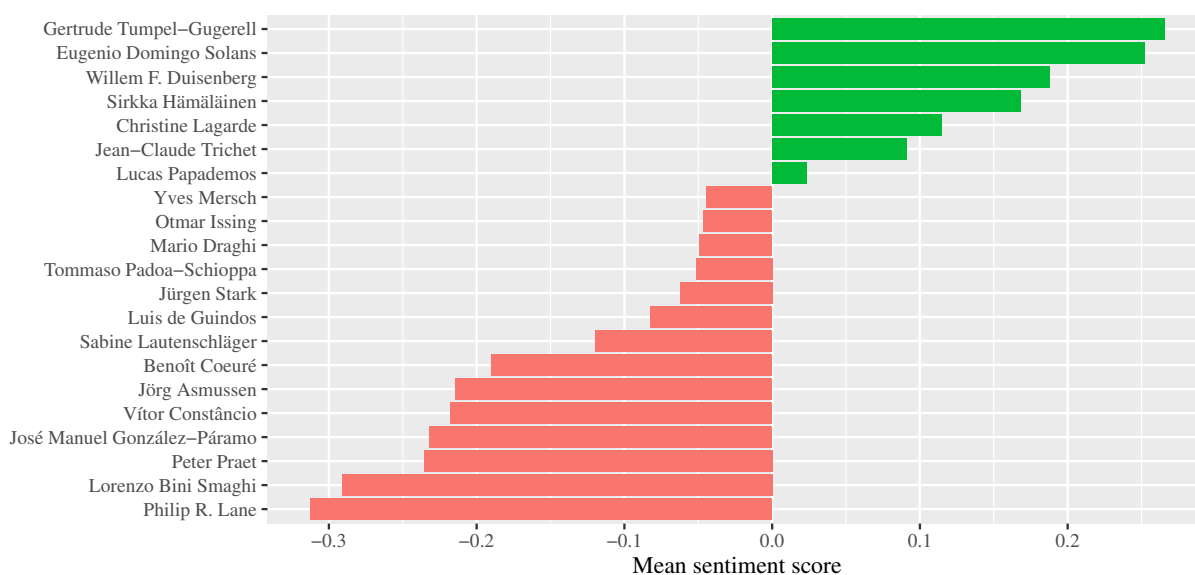


Note: The plot above shows the monthly average sentiment index over time. In addition the standard deviation for each month is illustrated by the shading around the curve. Recessions are illustrated according to the organization of Economic Development (OECD) based recession indicator for the euro area. The author's calculations are available upon request.

center of the distribution. Consequently, this indicates that variations in sentiment between speakers may be due to limited observations, and not to speaker's heterogeneity. Nevertheless, with those limitations in mind, the deviations are considerable. The most extreme (Philip Lane with -0.31 and Gertrude Tumpel-Gugerell with 0.27) are almost 0.6 index points apart. In absolute terms, this would imply a difference of 57 positive expressions between those two in an otherwise average speech. Note that the difference is substantial with an average of only 37 positive terms per speech (see table 1). Therefore, in order to control for such diversity in communication, I include a speaker-specific constant in the regression model (see equation (2) in the previous section) and thereby measure only the deviation from their mean.

Finally, it is not clear whether the information provided is of relevance. In particular, it is important to question whether the speaker provides information on the objective function of the institution. Hansen, McMahon, and Prat (2017) argue that the setting of a public speech might prevent a speaker from revealing true preferences. However, Jong and Esch (2014) argue that a public event may force the speaker to express her opinion on the official role she represents. That is particularly important for a central banker trying to shape the narrative of her central bank. In this way, the sentiment index may represent the objective function as perceived by the employee. Implicit in this argument is the assumption that the members of the ECB Executive Board have a coherent narrative – and therefore an inflation objective – in mind when communicating, and only differ in word choice due to individual semantic preferences and their perception of the shared narrative. The result would be the heterogeneity of speeches for which I control.

Figure 3: Speaker Sentiment



Note: The plot above is an illustration of the table 3 in the appendix. The author’s calculations are available upon request.

4.2 Inflation, Output and Unemployment

This second subsection focuses on the variables used to represent the objectives – i.e. inflation, output, and unemployment – and the need for real-time data in the analysis. In combination with the day-specific sentiment scores, this enables me to replicate accurately the information set available to the general public when listening to a speech.

Orphanides (2001) provides evidence that monetary policy regressions with revised data sets can lead to inaccurate findings. Since the ECB has acknowledged significant revisions of macroeconomic variables (ECB (2010)), it is important to use unrevised data – real-time data – for the question at hand. I utilize real-time information provided by the ECB through its Real Time Database. The database collects revisions of macroeconomic variables in the Monthly Bulletin. By treating each revision as new information to the general public, I am able to create a real time dataset for the macroeconomic variables of interest. This approach has been used in other work with an interest in the perception of monetary policy (e.g. Gross and Zahner (2020)).

The Harmonized Index of Consumer Prices (HICP), as stated in the official objective,³ expressed in annual growth rates, is used for the underlying analysis to measure the inflation rate. Economic activity is expressed either in terms of output, through the euro-wide annual gross domestic product (GDP) growth rate and the output-gap – calculated using a Hodrick-Prescott (HP) filter –, or in terms of the seasonally adjusted unemployment rate. One assumption, at least implicitly, is that this macro-economic dataset is exogenous to the speech corpus. This assumption is likely to be true, since macroeconomic information collection is done independently from speech writing. However, the prospect of a revision of a macro-economic variable in the

³<https://www.ecb.europa.eu/mopo/strategy/pricestab/html/index.en.html> (accessed 2020-05-01)

context of the forthcoming speech can not be completely excluded.

The regression analysis includes a number of control variables. As mentioned before, the speaker's specific sentiment as well as lagged sentiment is controlled for. In addition, one regression incorporates a time variable, capturing the development in the general sentiment as highlighted in the previous section. Moreover, the presence of a stability objective, independent of inflation and economic activity, as a tertiary objective reemerged following the central banks' response during the financial crisis (Peek, Rosengren, and Tootell (2016); Kaefer (2014)). Using text-mining techniques, Wischnewsky, Jansen, and Neuenkirch (2019) find evidence that deteriorating financial stability affect the monetary policy of FOMC. Consequently, to control for such a potential tertiary objective, three additional variables are included in the analysis. First, the EURO STOXX 50 Volatility (VSTOXX) is added in linear and square terms to measure the implied volatility for the euro area stock market. Second, the european uncertainty index of Baker, Bloom, and Davis (2016) is included as a benchmark in logarithmic terms for political instability. Finally, the 3-month standard deviation of the EURO STOXX 50 is added to capture medium-term volatility in the asset market.

5 Results

This section presents the regression results. The ECB's speeches are used to quantify the central banks' narrative, allowing for the calculation of its implied inflation target. The regressions are estimated versions of equation (3), using around 1.900 speeches from 21 speakers between January 2002 and January 2020. All results can be found in table 2, with the implied inflation target estimates as calculated in equation (4) and a 95% confidence interval at the bottom of the table.⁴

As discussed in the previous section, there are two potential narratives under consideration with different representations for the ECB's subordinate objective. This section starts with a discussion on the findings for the output growth rate as representation. The regression results can be found in table 2 in the first three columns. Then the findings for a central bank with an unemployment objective are discussed. The results are in column four to six. Afterwards, the result of the the ECB targeting both output and unemployment, as presented in the seventh column, are discussed. Robustness checks with respect to (1) the exponentiation of the economic activity term, (2) an alternative economic activity terms, (3) the lagged sentiment indicator, and (4) the resorts of the central bankers conclude this section. My findings appear to be robust.

5.1 Inflation and output

To begin with, a naive model is estimated, including only the inflation rate and output growth as inputs, neglecting the speakers heterogeneity and other controls like the tertiary financial stability objective. The results are shown in column one. There are several noteworthy observations by this naive model. First, the inflation coefficient is statistically significant and economically relevant. The coefficient suggests that an increase in the inflation by 1% leads to an increase of the sentiment index by 0.18. The positive sign of the inflation coefficient, combined with the negative sign of the squared inflation coefficient indicates a concave

⁴The inflation target and the confidence interval were estimated using a non-parametric case resembling bootstrap method, sampling the regression 5000 times with replacement. The implied inflation target represents the average of the estimated implied inflation targets. The 95% confidence interval of these estimates is presented in this paper. Note that extrem outliers were removed ($\sim 0.1\%$).

Table 2: Regression Results

	Sentiment Index $S_{i,t}$						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
π_t	0.118** (0.050)	0.071** (0.031)	0.079*** (0.025)	0.041 (0.034)	0.038 (0.026)	0.024 (0.021)	0.052* (0.030)
π_t^2	-0.019 (0.012)	-0.014 (0.009)	-0.015** (0.007)	-0.007 (0.012)	-0.007 (0.008)	-0.004 (0.006)	-0.011 (0.008)
Δy_t	0.049*** (0.010)	0.041*** (0.012)	0.042*** (0.010)				0.030** (0.012)
u_t				-0.043*** (0.014)	-0.028*** (0.011)	-0.036*** (0.007)	-0.024*** (0.009)
S_{t-1}	0.117*** (0.022)	0.079*** (0.022)	0.059*** (0.021)	0.100*** (0.020)	0.077*** (0.021)	0.055*** (0.020)	0.052** (0.021)
Year			-0.012* (0.007)			-0.016*** (0.004)	-0.016*** (0.005)
EPU_{Europe}			-0.089*** (0.022)			-0.070*** (0.016)	-0.070*** (0.016)
Speaker Dummy	No	Yes	Yes	No	Yes	Yes	Yes
Inflation target	3.4	2.9	3.2	2.8	2.2	2.1	2.2
95% CI	[3.3, 3.5]	[2.7, 3.1]	[3.0, 3.5]	[2.2, 3.4]	[1.9, 2.4]	[1.7, 2.5]	[2.0, 2.5]
Observations	1,749	1,749	1,749	1,892	1,892	1,892	1,749
Adjusted R ²	0.057	0.212	0.225	0.062	0.208	0.222	0.228

Note: Robust standard errors (clustered by speaker) are recorded in parentheses. The variable y_t denotes log GDP. Significance levels are given in the following form: * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

relationship towards inflation. This translates to around 16 more positive words throughout an ordinary speech, although the negative quadratic term may partly offset this effect. Importantly, the concavity in the inflation relationship is consistent across all specifications in this paper. The existence of a such concave inflation target implies the existence of a global maximum for that objective. Based on the results in the naive regression model, the communicated inflation target of the ECB is more than 3%. This is well above the official target of 'close, but below 2%'.

Second, the output growth coefficients is significant and suggests a positive correlation between production and sentiment, i.e. one additional percentage point of GDP growth increases an average speech by approximately four positive terms. Therefore, when comparing only the linear coefficients, the hierarchical order of the ECB's objective appears to be manifested in it's speeches sentiment. However, this might be driven by the heterogenous communication between speakers, as we will see below. Finally, for the present speech, the tone of the previous speech – delayed sentiment index – is a highly statistically significant and relevant approximate, even if the speakers differ. The Durbin-Watson test suggests that with the inclusion of


this delayed index, no autocorrelation remains in the residual. Taking into consideration that much of the variation in the sentiment of a particular speech arises from the semantic differences between the speakers, the following regression will correct for this.

In the second column, the regression result controlling for speaker-specific nuances is presented. The main findings are the following. First, the speakers generally enter the regression results with the anticipated signs. The most pessimistic speaker is Lorenzo Smaghi (-0.135***), whereas Gertrude Tumpel-Gugerell (0.419***) is the most optimistic speaker after controlling for economic conditions. The inclusion has significant implication on the delayed sentiment index, which halves in magnitude. Second, controlling for speaker-specific sentiment has a significant impact on the magnitude of the implied inflation target, which decreases to 2.9%. This suggests that the previous finding may genuinely have been affected by the heterogeneity of the members in the ECB's Executive Board. Using bootstrapping on the perceived inflation target, the 95% confidence interval can be estimated at [2.7%, 3.1%]. This interval clearly rejects the null hypothesis that the target is articulated to be zero. It also rejects the interpretation that the implied inflation target does not significantly exceed 2%. Third, the output coefficient remains positive, statistically significant and only decreases only to a small degree in magnitude. As a result, the ECB appears to be only slightly more responsive to deviations in inflation than production, the difference is not high. A back-on-the-envelope calculation shows that an increase in GDP by 1% has almost as great an effect on the sentiment index than an increase in inflation from 1% to 2%. Fourth, the regression's explanatory power increases more than fivefold, suggesting the necessity to incorporate speaker-specific information in the regression model. Note that to economize on space the only measure for the goodness-of-fit provided here is the adjusted R^2 but the Akaike information criterion confirms these findings.

Finally, I include the additional measures of financial stability and the time-dependent dummy. The results are qualitatively and quantitatively consistent with the previous ones. The inflation rate becomes highly significant – as does the squared inflation term – and the implied inflation target increases slightly to 3.2%. A graphic illustration of the objective, as estimated in this regression, is provided in figure 5. The left bin-scatter plot illustrates the perceived inflation target. The concavity of the communicated objective itself (red line) is visually highlighted. The estimated relationship towards output is illustrated on the right side of the above mentioned bin-scatter plot. The positive correlation between sentiment and GDP growth rates is evident. The plot provides further evidence on the stronger reaction of sentiment to changes in production than inflation. GDP growth remains a highly significant variable.

In addition a potential tertiary objective towards financial stability, as identified for the FOMC (Wischnewsky, Jansen, and Neuenkirch (2019)), is added to the regression. However, neither the VSTOXX, nor the standard deviation of the euro area stock market have a significant impact. Only Baker, Bloom, and Davis (2016)'s uncertainty index is statistically significant. A 1% increase in the index decreases the respective speech by about six positive words, indicating that the discourse is becoming extra pessimistic in times of heightened uncertainty. Furthermore, there appears to be a negative trend in sentiment over time. This finding is significant and economically relevant, with an average of almost two fewer optimistic terms used per speech every year. The decline in sentiment over time is consistent with the anecdotal evidence presented in figure 2 and may be indicative of the central bankers' increasingly challenging role of maintaining inflation in the euro area while being restricted by the effective lower bound.

Figure 4: Bin-Scatter Plot of Output Regression (3)



Note: The illustration of a 'bin-scatter-plot' represents a practical alternative to the more conventional scatter plot. The data points are grouped into bins, and each bin is averaged. In addition, the size of each point is proportional to the number of data points within the respective bin. The objective (red continuous line) is illustrated as of regression results (3). The author's calculations are available upon request.


In my findings, the adjusted R^2 is substantially higher than that by Shapiro and Wilson (2019) on their public communication dataset of the FOMC. As mentioned previously, the analysis of Shapiro and Wilson (2019) focuses on the FOMC's non-public communication. There, up to a quarter of the variation can be explained by their regression ($R^2 \sim 0.24$). However, the model's explanatory power decreases to $R^2 \sim 0.06$ when using speech data. Therefore, with regard to explanatory power, my findings ($R^2 \sim 0.22$) on public speeches by the ECB seem to be in line with their results on the private correspondence of the FOMC.

Overall, the results of the regressions suggest that a narrative is communicated with a concave relationship for the inflation rate, while the objective regarding the output rate is well approximated with a linear increasing function. This finding implies the presence of an inflation rate, which minimizes the loss function of the ECB. According to my results, the implied inflation target is around 3%. Above and below this price increase, the institution's sentiment in speeches becomes more negative. In addition, speaker-specific effects have a huge impact on the sentiment score and financial uncertainty tend to influence the speeches. In the next step, I will substitute the GDP growth rate with the unemployment rate in order to test the consistency of the findings.

5.2 Inflation and unemployment

As discussed in section 4.2, there remains ambiguity regarding the subordinate objective of ECB. In the previous subsection, the findings when interpreting the aim of promoting the general economic policies in the euro area as an explicit production growth target were presented. As mentioned in section 2, the unemployment rate is an alternative metric for this subordinate objective. Qualitatively and quantitatively, my results on the inflation target appear to be independent of the choice in the economic activity objective.

Figure 5: Bin-Scatter Plot of Unemployment Regression (6)



Note: The illustration of a 'bin-scatter-plot' represents a practical alternative to the more conventional scatter plot. The data points are grouped into bins, and each bin is averaged. In addition, the size of each observation is proportional to the number of data points within the respective bin. The objective (red continuous line) The objective (red continuous line) is illustrated as of regression results (3). The author's calculations are available upon request.

However, using of employment data reduces the statistical power of the inflation coefficients. The results can be found in table 2 in columns four to six.

Embedding the unemployment rate in the regression instead of the production growth rate yields several findings. First, the implied inflation target does not diverge from previous findings and may be best described as '*above, but close to 2%*'. Again, it is noteworthy that the implied inflation target decreases, once I control for the respective speaker. Although they deviate in quantitative terms, the findings of the preceding subsection can be qualitatively confirmed. The implied inflation targets are consistently between 2% and 3%. The inflation objective – including speaker-specific effects, financial stability controls and time effects – is graphically illustrated in the left bin-scatter plot in figure 5.

Second, the indicated relationship towards unemployment is negative and highly statistically significant, indicating that the ECB communicates a preference towards lower unemployment. Analogous to my previous findings, the reaction of the sentiment to changes in unemployment tends to be economically relevant. An increase in the unemployment rate by 1% decreases an average speech by up to three positive terms. The relationship of the unemployment rate is illustrated on the right-hand side in figure 5. Subsequently, the integration of financial stability and a time index yield findings that parallel those in the output-case. The depreciation of sentiment over the observation horizon is now highly statistically significant and in terms of magnitude very similar to that was previously found.

5.3 Inflation, output, and unemployment

In the last regression, GDP growth is embedded together with the unemployment rate in the regression. The results are presented in the last column in table 2. Interestingly, both economic activity coefficients – GDP growth and unemployment rate – remain highly statistically significant, but decrease slightly in magnitude. Contrary, the inflation rate yields only slight significant result anymore and further matched by the economic activity coefficients in terms of economic relevance. The back-of-the-envelope calculation (ignoring significant levels) shows that GDP growth would have to increase by only 0.5% (or the unemployment rate decrease by 1%) in order to have the same impact on the communicated loss as an increase in the inflation rate from 1% to 2%. This might indicate that central bankers are reacting to changes in GDP and unemployment rather than inflation, or at least construct a narrative that is more concerned with these variables. Nevertheless, the implied inflation target remains above, albeit close to 2%. Also its confidence interval [2.0, 2.5] still rejects the 'below 2% inflation' target.


5.4 Robustness checks

This final subsection covers a set of robustness checks conducted for my analysis. To economize on space, the regression results are shown in appendix A.4. The first two robustness check concern the presence of a non-linear economic activity term (columns one and two) and the choice of output measurement (column three). This section concludes with the presentation of an alternative delayed sentiment index (columns four and five) and the inclusion of task-specific central banker controls (last two columns).

In equation (1), following standard textbooks, \hat{y}_t enters the objective function with a linear term. However, Walsh (2017) presents both a linear version (by Barro and Gordon (1983) and used by Shapiro and Wilson (2019)) and a non-linear version. Following Shapiro and Wilson (2019), I use the linear model in my baseline regression. However, in this subsection I validate my results using a squared economic activity in addition to the linear one, demonstrating that the results qualitatively hold independently of this assumption. Including this non-linear term has the interesting side effect of enabling me to calculate the optimal level of output and employment, i.e. the implicit bliss point for the economic activity term as well. Including the non-linear term yields the following results: On the one hand, there is only a slight variation in the implicit inflation target, being 2.9%. On the other hand, the non-linear output coefficient is statistically significant (as is the linear coefficient), indicating a narrative with a convex production objective for ECB. The implicit bliss point is a growth rate of -0.6%. With few exceptions, this growth rate is below the lowest reported GDP growth rate in the euro area. In other words, the Executive Board at the ECB communicates more pessimistically until production deteriorates by 0.6% year-on-year, i.e. the central bankers seem to favor a positive output-gap, almost independently of the current economic situation. Note that the implication is therefore very similar to the linear case. This objective is illustrated in figure 6. When the squared unemployment rate is included, the coefficients do not yield the same statistical significance anymore. Since u_t^2 and u_t are highly correlated ($\sim 99\%$), this insignificance is likely to be driven by multicollinearity. Hence, for the unemployment rate, the linear specification seems preferable. Finally, the remaining variables join the model with similar sign, magnitude and significance level as they did in the baseline regression.

The second robustness reviews the validity of GDP growth as representation for the economic activity

Figure 6: Bin-Scatter Plot of Robustness Check Regression (1)



Note: The illustration of a 'bin-scatter-plot' represents a practical alternative to the more conventional scatter plot. The data points are grouped into bins, and each bin is averaged. In addition, the size of each point is proportional to the number of data points within the respective bin. The objective (red continuous line) is illustrated as of the robustness regression results (1). The author's calculations are available upon request.

objective. The use of the output-gap is an alternative to the previously applied growth rate. The output-gap is estimated using a HP-filter and replaces the growth rate in this regression. While inflation becomes a highly statistically significant variable, the output-gap is not significant. In addition, the response of the sentiment index to uncertainty as measured by the Baker, Bloom, and Davis (2016)'s index increases. The inflation target in this regression can be calculated as 2.8% inflation.

The third robustness check concerns the delayed sentiment variable. The sentiment of the previous speech was included in the baseline regression. However, in order to capture the general mood when giving the speech and test the robustness of this choice, I substitute this term with the average sentiment of the previous month. The findings remain the same. The result indicates – similar to the previously used lagged variable – a significant positive impact. Both, output and unemployment, remain significant variables in this regression, while inflation remains significant in the output specification.

The final robustness check addresses the speaker specific effects. One might argue that the resorts (departments) held by a central banker influences the sentiment of this banker's speech. In other words, a central banker speaking of, say, "Risk Management" may be forced to choose a more pessimistic language than her peer speaking on the "New ECB Premises Project". Therefore, I reconstruct the resorts held by each central banker at the time of her speech and include the corresponding dummy in the regression. Although most resorts do not yield a significant impact on sentiment, the hypothesis of irrelevance is rejected by a joint significant test. Nevertheless, the results indicate that the inclusion of speaker-specific dummies is preferable in terms of explanatory power.

6 Conclusion

With its price stability mandate in the euro area, the ECB is tasked to pursue inflation with a subordinate goal to stabilize economic activity. The general public's perception of this objective function is important for the effectiveness of the central bank's monetary policy, which in turn incentivizes the ECB to construct an appropriate narrative around this objective function. Consequently, my paper employs a text mining approach suggested by Shapiro and Wilson (2019) to approximate the objective function and its determinants. Taking all public speeches of the ECB's Executive Board between 2002 and 2020, this paper builds upon a growing body of academic work concentrating on quantifying qualitative text data.

My results suggest that the ECB nurtures a narrative with an inflation target beyond its stated objective. The following key results are obtained by estimating a sentiment index through a 'bag-of-words'-approach and using this index to approximate the relationship of the ECB to its objectives. First, the findings of my analysis suggest a concave objective regarding the inflation rate. The implied optimum for inflation – the ECB's inflation target – is best described as *'above, but close to 2%'*. Deviations from this value lead to a reduction in the sentiment of the institution's communication. This finding contrasts the results obtained by text analysis on press conference statements by Paloviita et al. (2020), who found a communicated target of $\sim 1.6\%$. However, as pointed out in this paper, the non-stationarity in similarity in the press-conference statements, in comparison to the speeches, makes me confident in the presented results. Furthermore, my results appear to be robust across a myriad of model-specifications and determinants. Second, including output growth or the unemployment rate as a subordinate aim of promoting the European Union's general economic policies yields qualitatively similar results. My findings suggest a potential concave relationship between the central bankers sentiment and output, as well as a linear relationship towards the unemployment rate. Independent of the specification, my findings suggest a communicated preference for higher GDP growth and lower levels of unemployment, regardless of the current economic condition. The hierarchical order in the the ECB's mandate does not always appear to be consistent with my findings. Deviations from its primary objective inflation appear to be of no greater concern than deviations in its subordinate objective. Third, over the last two decades, the tone of speeches has deteriorated, even when controlling for economic condition. This might be indicative for the increasing challenging role of the ECB after approaching the effective lower bound. Fourth, during times of crisis and heightened uncertainty, central bankers appear extra pessimistic. Finally, the explanatory power of public communication of the ECB is surprisingly high, especially when compared to previous studies on the FOMC.

This paper provides a first analysis of the ECB's narrative on its objective function using text mining on the institutions' speeches. Future research may further examine the tremendous amount of information provided by the speeches and the heterogeneity of the speakers. One possibility would be to extend the analysis beyond its sentiment by incorporating meta-information provided by the text of the speech. Evidence of the impacts of complexity and length of press conference statements have previously been observed (Smales and Apergis (2017); Jansen (2010)). Further potential might be in the heterogeneity of communication between central bankers or even national central banks (Hayo and Neuenkirch (2013); Tillmann and Walter (2018)). This could provide an interesting comparison of the difference in narratives across countries. In addition, it would be one way of testing whether there is evidence for the hypothesis that a central banker communicates the inflation target of a central bank, or vice versa.

References

- Amaya, Diego and Jean-Yves Filbien (2015). “The Similarity of ECB’s Communication”. In: *Finance Research Letters* 13, pp. 234–242.
- Apel, Mikael and Marianna Blix Grimaldi (2014). “How informative are central bank minutes?” In: *Review of Economics* 65.1, pp. 53–76.
- Apel, Mikael, Marianna Blix Grimaldi, and Isaiah Hull (2019). “How Much Information Do Monetary Policy Committees Disclose?: Evidence from the FOMC’s Minutes and Transcripts”. In: *Sveriges Riksbank Working Paper Series*.
- Baker, Scott R, Nicholas Bloom, and Steven J Davis (2016). “Measuring Economic Policy Uncertainty”. In: *The quarterly journal of economics* 131.4, pp. 1593–1636.
- Baranowski, Paweł, Hamza Bennani, and Wirginia Doryń (2020). “Do ECB Introductory Statements help to predict Monetary Policy: Evidence from Tone Analysis”. In: *NBP Working Paper* 323.
- Barro, Robert J. and David B. Gordon (1983). “Rules, discretion and reputation in a model of monetary policy”. In: *Journal of Monetary Economics* 12.1, pp. 101–121.
- Bennani, Hamza, Nicolas Fanta, Pavel Gertler, and Roman Horvath (05/2019). “Does Central Bank Communication Signal Future Monetary Policy? The Case of the ECB”. In: *Working Papers IES* 2019/12.
- Bennani, Hamza and Matthias Neuenkirch (2017). “The (Home) Bias of European Central Bankers: New Evidence based on Speeches”. In: *Applied Economics* 49.11, pp. 1114–1131.
- Bholat, David M., Stephen Hansen, Pedro M. Santos, and Cheryl Schonhardt-Bailey (2015). “Text Mining for Central Banks”. In: *SSRN Electronic Journal*.
- Black, John, Nigar Hashimzade, and Gareth Myles (2017). *A Dictionary of Economics*. Oxford University Press.
- Blinder, Alan S., Michael Ehrmann, Marcel Fratzscher, Jakob de Haan, and David-Jan Jansen (04/2008). “Central Bank Communication and Monetary Policy: A Survey of Theory and Evidence”. In: *ECB Working Paper Series* 170.
- Dickey, David A. and Wayne A. Fuller (1979). “Distribution of the Estimators for Autoregressive Time Series with a Unit Root”. In: *Journal of the American Statistical Association* 74.366a, pp. 427–431.
- Diron, Marie and Benoit Mojon (2005). “Forecasting the Central Bank’s Inflation Objective is a Good Rule of Thumb”. In: *ECB Working Paper Series* 564.
- ECB (2010). *ECB Monthly Bulletin September*. ECB. URL: <https://www.ecb.europa.eu/pub/pdf/mobu/mb201009en.pdf>.
- European Central Bank (10/2019). *Speeches dataset*. URL: <https://www.ecb.europa.eu/press/key/html/downloads.en.html>.
- Falagiarda, Matteo and Stefan Reitz (2015). “Announcements of ECB Unconventional Programs: Implications for the Sovereign Spreads of Stressed Euro Area Countries”. In: *Journal of International Money and Finance* 53, pp. 276–295.

- Gali, Jordi (06/09/2015). *Monetary Policy, Inflation, and the Business Cycle*. Princeton University Press. 296 pp.
- Gentzkow, Matthew, Bryan Kelly, and Matt Taddy (2019). “Text as Data”. In: *Journal of Economic Literature* 57.3, pp. 535–574.
- Gertler, Pavel and Roman Horvath (2018). “Central Bank Communication and Financial Markets: New high-frequency Evidence”. In: *Journal of Financial Stability* 36, pp. 336–345.
- Gross, Jonas and Johannes Zahner (2020). “What’s on the ECB’s Mind? Monetary Policy before and after the Global Financial Crisis”. In: *MAGKS Joint Discussion Paper Series in Economics*.
- Hansen, Stephen, Michael McMahon, and Andrea Prat (2017). “Transparency and Deliberation Within the FOMC: A Computational Linguistics Approach”. In: *The Quarterly Journal of Economics* 133.2, pp. 801–870.
- Hanson, Samuel G. and Jeremy C. Stein (2015). “Monetary policy and long-term real rates”. In: *Journal of Financial Economics* 115.3, pp. 429–448.
- Hayo, Bernd, Ali M Kutan, and Matthias Neuenkirch (2008). *Communicating with many Tongues: FOMC speeches and US Financial Market Reaction*. Tech. rep. Joint discussion paper series in economics.
- Hayo, Bernd and Matthias Neuenkirch (2013). “Do Federal Reserve Presidents communicate with a Regional bias?” In: *Journal of Macroeconomics* 35, pp. 62–72.
- Jansen, David-Jan (2010). “Does the Clarity of Central Bank Communication affect Volatility in Financial Markets? Evidence from Humphrey-Hawkins Testimonies”. In: *Contemporary Economic Policy* 29.4, pp. 494–509.
- Jong, Eddie de and F.A.W.J. van Esch (2014). “Culture matters: French,ÄGerman conflicts on European Central Bank Independence and Crisis Resolution”. In: *Financial Cultures and Crisis Dynamics*.
- Kaefer, Benjamin (2014). “The Taylor Rule and Financial Stability – A literature review with Application for the Eurozone”. In: *Review of Economics* 65.2.
- Loughran, Tim and Bill McDonald (2011). “When Is a Liability Not a Liability? Textual Analysis, Dictionaries, and 10-Ks”. In: *The Journal of Finance* 66.1, pp. 35–65.
- (2016). “Textual Analysis in Accounting and Finance: A Survey”. In: *Journal of Accounting Research* 54.4, pp. 1187–1230.
- Lucca, David O. and Francesco Trebbi (09/2009). “Measuring Central Bank Communication: An Automated Approach with Application to FOMC Statements”. In: *NBER Working Papers* 15367.
- Masciandaro, Donato, Davide Romelli, and Gaia Rubera (2020). “Tweeting on Monetary Policy and Market Sentiments: The Central Bank Surprise Index”. In: *SSRN Electronic Journal*.
- Meade, Ellen E. (2005). “The FOMC: Preferences, Voting, and Consensus”. In: *Federal Reserve Bank of St. Louis Review* 87.2.
- Miles, David, Ugo Panizza, Ricardo Reis, and Angel Ubide (2017). *And yet it moves: Inflation and the great Recession*. International Center for Monetary and Banking Studies/Centre for Economic.

- Molodtsova, Tanya and David Papell (2012). “Taylor Rule Exchange Rate Forecasting During the Financial Crisis”. In: *NBER Working Paper* 18330.
- Orphanides, Athanasios (2001). “Monetary policy rules based on real-time data”. In: *American Economic Review* 91.4, pp. 964–985.
- Paloviita, Maritta, Markus Haavio, Pirkka Jalasjoki, Juha Kilponen, Ilona Vänni, et al. (2020). “Reading between the lines-Using text analysis to estimate the loss function of the ECB”. In: Peek, Joe, Eric S Rosengren, and Geoffrey Tootell (2016). “Does Fed policy reveal a Ternary Mandate?” In: *Federal Reserve Bank of Boston Working Papers* 16-11.
- Praet, Peter (2014). “Current issues and challenges for central bank communication”. In: *Speech at the conference on ”The ECB and Its Watchers XV”*. Vol. 12.
- Rosa, Carlo and Giovanni Verga (2008). “The Impact of Central Bank Announcements on Asset Prices in Real Time”. In: *International Journal of Central Banking* 4.2, pp. 175–217.
- Schmeling, Maik and Christian Wagner (2019). “Does Central Bank Tone Move Asset Prices?” In: *CEPR Discussion Paper* 13490.
- Shapiro, Adam H. and Daniel J. Wilson (2019). “Taking the Fed at its Word: A New Approach to Estimating Central Bank Objectives Using Text Analysis”. In: *Federal Reserve Bank of San Francisco, Working Paper Series*, pp. 01–74.
- Shiller, Robert J (2017). “Narrative economics”. In: *American Economic Review* 107.4, pp. 967–1004.
- Smaghi, Lorenzo Bini (2007). “Globalisation and public perceptions”. In: *Meeting of the GEPA-GPA-GSPA*. Vol. 12.
- Smales, L.A. and N. Apergis (2017). “Does More complex Language in FOMC Decisions impact Financial Markets?” In: *Journal of International Financial Markets, Institutions and Money* 51, pp. 171–189.
- Tillmann, Peter and Andreas Walter (2018). “ECB vs Bundesbank: Diverging Tones and Policy Effectiveness”. In: *MAGKS Joint Discussion Paper Series in Economics* 201820.
- Tobback, Ellen, Stefano Nardelli, and David Martens (2017). “Between Hawks and Doves: Measuring Central Bank Communication”. In: *ECB Working Paper Series*.
- Walsh, Carl E. (05/12/2017). *Monetary Theory and Policy*. MIT Press Ltd. 688 pp.
- Wischnewsky, Arina, David-Jan Jansen, and Matthias Neuenkirch (2019). “Financial Stability and the Fed: Evidence from Congressional Hearings”. In:

A Appendix

A.1 Sample text

The two exemplary texts here provide an illustration of the 'bag-of-words' approach and Loughran and McDonald (2011)'s dictionary. **Positive** terms are highlighted in bold and blue, *negative* terms in italic and red color, and negations are underlined.

"[...] There is no *doubt* in my mind that the ultimate goal of economic policy, which is to improve the standard of living and quality of life of all members of the community, can only be **achieved** by an **effective** interplay between policies with a view to **achieving** simultaneously **efficiency**, **stability** and equity. [...]"

- Eugenio Solans (2003)⁵

"[...] *Unfortunately*, many euro area countries entered the financial *crisis* and the economic *downturn* with *unnecessarily weak* fiscal balances, having *missed* the **opportunity** presented by past years' revenue windfalls to consolidate their budgets. [...]"

- José Manuel González-Páramo (2009)⁶

These exemplary text contains 1 (5) negative and 4 (1) positive negative terms. However, note that the negative term 'doubt' is preceded by a negation and therefore does not affect the sentiment index. The respective sentiment index for this sentences would therefore be: $S_{Solans,2003} = \frac{5-0}{5} = 1$ and $S_{Gonzalez-Paramo,2009} = \frac{1-5}{1+5} = -2/3$.

A.2 Similarity index

The method of estimating the similarity between documents is borrowed from Amaya and Filbien (2015). Amaya and Filbien compare two preceding ECB press conferences by counting the occurrence of bigrams. Bigrams are all combinations in one sentence of two succeeding words. For example, the sentence "*Inflation expectations are higher than expected*" has the bigrams "inflation-expectations", "expectations-are", "are-higher" and so on. By extracting all possible bigrams in the speech, they measure the similarity between two press conferences that take place in month t as the number of bigrams that occur in both press conferences, divided by all bigrams that occur in the two conferences:

$$sim_{t+1} = \frac{bigrams_t \cap bigrams_{t+1}}{bigrams_t \cup bigrams_{t+1}} \quad (6)$$

As this paper compares speeches at a much higher frequency in this paper, I have adjusted the similarity index by Amaya and Filbien (2015) by aggregating the bigrams from all individual i speeches within one

⁵<https://www.ecb.europa.eu/press/key/date/2003/html/sp030307.en.html> (accessed 2020-05-01)

⁶<https://www.ecb.europa.eu/press/key/date/2009/html/sp090206.en.html> (accessed 2020-05-01)

month and then comparing them with the the bigrams of the following month:

$$sim_{t+1} = \frac{\sum_i bigrams_{i,t} \cap \sum_i bigrams_{i,t+1}}{\sum_i bigrams_{i,t} \cup \sum_i bigrams_{i,t+1}} \quad (7)$$

To detect whether speeches within a month become more similar with each other, I calculated a similarity index within a month (sim_t^{wm}) as follows:

$$sim_t^{wm} = \frac{bigrams_{i,t} \cap \sum_i bigrams_{j,t}}{bigrams_{i,t} \cup \sum_i bigrams_{j,t}} \quad (8)$$

with $i \neq j$. The results are presented in the paper in section 4.1 and can be provided upon request.

A.3 Speaker summary

Table 3: Summary statistics per speaker

	Speakers	Speeches	Mean # of sentences	Mean sentiment
1	Jean-Claude Trichet	318	89	0.090
2	Mario Draghi	183	81	-0.050
3	Benoit Cure	179	100	-0.190
4	Yves Mersch	149	69	-0.040
5	Gertrude Tumpel-Gugerell	148	74	0.270
6	Vitor Constancio	121	121	-0.220
7	Jose Manuel Gonzalez-Paramo	117	96	-0.230
8	Peter Praet	117	94	-0.240
9	Lorenzo Bini Smaghi	107	122	-0.290
10	Lucas Papademos	87	112	0.020
11	Juergen Stark	77	93	-0.060
12	Sabine Lautenschlaeger	76	70	-0.120
13	Otmar Issing	47	106	-0.050
14	Juerg Asmussen	45	60	-0.210
15	Luis de Guindos	40	71	-0.080
16	Willem F. Duisenberg	25	59	0.190
17	Eugenio Domingo Solans	20	93	0.250
18	Tommaso Padoa-Schioppa	16	104	-0.050
19	Philip R. Lane	8	184	-0.310
20	Christine Lagarde	6	49	0.110
21	Sirkka Haemaelaeinen	6	102	0.170
		1,892	91	-0.059

A.4 Robustness Check Regressions

Table 4: Regression Results Robustness Checks

	Sentiment Index $S_{i,t}$						
	<i>Squared Economic Activity</i>		<i>Output-Gap</i>	<i>Monthly ARI</i>		<i>Resorts</i>	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
π	0.055*** (0.020)	0.025 (0.021)	0.053** (0.026)	0.072*** (0.021)	0.020 (0.020)	0.077*** (0.026)	0.019 (0.019)
π^2	-0.009 (0.006)	-0.004 (0.006)	-0.007 (0.008)	-0.013** (0.006)	-0.003 (0.006)	-0.014* (0.008)	-0.002 (0.006)
Δy_t	0.014*** (0.005)			0.039*** (0.007)		0.038*** (0.010)	
Δy_t^2	0.011*** (0.003)						
u_t		0.001 (0.128)			-0.033*** (0.011)		-0.039*** (0.011)
u_t^2		-0.002 (0.007)					
Δy_t			-0.006 (0.011)				
S_{t-1}	0.056*** (0.020)	0.055*** (0.020)	0.073 (0.052)			0.056** (0.022)	0.049** (0.020)
$S_{t-1,month}$				0.120*** (0.041)	0.115*** (0.020)		
year	-0.012** (0.006)	-0.016*** (0.004)	-0.007 (0.007)	-0.010** (0.004)	-0.014* (0.007)	-0.011 (0.009)	-0.017** (0.007)
EPU_{Europe}	-0.076*** (0.022)	-0.075*** (0.019)	-0.112*** (0.022)	-0.083*** (0.016)	-0.067*** (0.023)	-0.080*** (0.022)	-0.070*** (0.019)
Inflation target	2.91	3.18	2.83	3.7	4.11	2.84	3.8
Observations	1,892	1,892	1,892	1,749	1,892	1,749	1,892
Adjusted R ²	0.221	0.222	0.213	0.225	0.223	0.233	0.228

Note:

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