The ECB's Tracker

Nowcasting the Press Conferences of the ECB

Armando Marozzi*

(London School of Economics, European Central Bank)

*Emails: a.marozzi@lse.ac.uk and Armando.Marozzi@ecb.europa.eu.

This paper shall not be reported as representing the views of the ECB. The views expressed are those of the author and do not necessarily reflect those of the ECB.

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Abstract

This paper proposes an econometric framework for nowcasting the monetary policy stance and decisions of the European Central Bank (ECB) exploiting the flow of conventional and textual data that become available between two consecutive press conferences. Decompositions of the updated nowcasts into variables' marginal contribution are also provided to shed light on the main drivers of the ECB's reaction function at every point in time. In out-of-sample nowcasting experiments, the model provides an accurate tracking of the ECB monetary policy stance and decisions. The inclusion of textual variables contributes considerably to the gradual improvement of the model performance.

Keywords: Natural Language Processing, Dynamic Factor Model, Monetary Policy, Forecasting, Mixed frequencies

1. Introduction

Central banks usually announce changes to their monetary policy stance in press conferences. The interval between two consecutive press conferences, however, can be significantly long. Even if macroeconomic and/or financial conditions change abruptly, central bank watchers need to wait the nearest press conference to see confirmed or denied their expectations on central banks' decisions. Market participants and policy-makers can thus only rely on non-systematic strategies to interpret the future path of monetary policy.

This paper provides an econometric framework which enables interested parties to track systematically the real time evolution of the monetary policy stance and decisions of a central bank on the basis of the increasing amount of information that becomes available between two consecutive press conferences.

I focus on the European Central Bank (ECB). Since its inception, the "blind spot" between two consecutive ECB press conferences has increased. In 1999, the Governing Council used to take policy decisions twice a month with a press conference taking place only on the first meeting of the month. Afterwards, from November 2001 there was only one policy meeting per month followed by a press conference and, from January 2015, the frequency of monetary policy meetings changed again to occur every six weeks. At the same time, however, a wide range of conventional and non-conventional data are published between two consecutive Governing Council meetings at daily or monthly frequencies and contain valuable information on the expected path of monetary policy. Therefore, although the contemporaneous monetary policy stance and decisions of the incoming press conference are not available, they can be estimated exploiting higher-frequency variables that are released in a more timely manner.

First, I construct ECB field-specific dictionaries and apply them to subsets of the introductory statements of the press conferences to derive the indexes of ECB monetary policy stance, economic and inflation outlook. Second, I structure a textual dataset with around 300,000 documents into daily time series with macroeconomic information. Third, I model the total dataset containing around 140 variables as a Dynamic Factor Model (DFM) with flow and stock variables estimated at daily frequency. The DFM is augmented with an auxiliary equation that takes the specification of a multinominal logit with three possible monetary policy outcomes: ease, constant and hike. Last, I set up a "pseudo" Taylor rule to assess the performance of the DFM. Overall, the model produces three pieces of information: the nowcast of the ECB monetary policy stance, the forecast of the conditional probability that the ECB will actually take a monetary policy decision at time $t+1$ and the block of variables that drives the revision of each nowcast at every point in time.

The empirical results are the following. First, the model provides higher forecast accuracy than a benchmark model approximated by a "pseudo" Taylor rule; second, the inclusion of textual variables in the dataset contributes considerably to the improvement of the forecasting performance; third, the model provides an accurate tracking of the ECB monetary policy stance and decisions at key historical ECB announcements; last, the DFM proves to be useful in forecasting the EONIA rates from January 2008 to December 2009.

This paper refers to two strands of literature. The first one is on nowcasting. Since the release of Giannone et al. (2008)'s seminal paper, the nowcasting literature has significantly developed methodologically and empirically (see Bok et al., 2018 for a survey). In particular, this paper is directly related to Bańbura and Modugno (2014) as it draws on their Expectation Maximization (EM) algorithm to estimate the DFM and to Bańbura et al. (2013) for laying down the strategy to model mixed frequency flow and stock variables. Thorsrud (2020) and Cimadomo et al. (2020) come also close to this paper: the former employs textual variables for nowcasting Norwegian GDP, while the latter, among other things, proposes a mixed-frequency VAR to forecast the Fed Funds rate given the latest news on US economic conditions.

The second strand of literature studies forecasting interest rate decisions. Following Taylor (1993), it has become common to characterize central bank policy as an interest rate rule (the so-called Taylor-rule) that responds to inflation and the output gap or other combinations of macro variables (for a survey, see Wieland and Wolters, 2013). More recently, with the ascent of text-mining techniques, macroeconomists augmented the stylized Taylor rule with textual variables capturing central bank communication. In particular, many papers attempted to study whether ECB communication helps predict future monetary policy (Sturm and de Haan, 2011; Picault and Renault, 2017; Bennani and Neuenkirch, 2017; Bennani et al., 2020; Baranowski et al., 2021).

To the best of my knowledge, however, no study exploits the flow of information that becomes available between press conferences to continuously and systematically update the nowcasts of the monetary policy stance and decisions of a central bank. Therefore, the main contribution of this paper is the set-up of a unified econometric framework that is novel for describing central banks' reaction function.

The rest of the paper is organised as follows: Section 2 describes the datasets, the text-mining techniques and the estimated news topics. Section 3 details the models. Section 4 illustrates the results and Section 5 concludes the paper.

2. Data and Textual Methods

Section 2.1 derives three indexes from the introductory statements of the press conferences of the ECB: a monetary policy stance index, an economic outlook index and an inflation outlook index; Section 2.2 details the procedural steps to structure and quantify textual data; finally, Section 2.3 describes the total dataset used in the model.

2.1. Quantifying the Press Conferences of the ECB

The ECB's press conferences have a pivotal role in transmitting monetary policy information and are therefore best suited to extracting macro and policy signals. In particular, the aim is to derive indexes of monetary policy stance, economic and inflation outlook. While the first one is the variable to nowcast, the other two convey significant information for the path of monetary policy and, for this reason, will be part of the DFM.

To quantify these indexes, I proceed in three steps: first, I gather press conferences from January 2002 to December 2020 by webscraping the ECB's website. I then remove the Q&A part from every press conference and subset the texts into paragraphs. To further reduce the noise in the data, I apply some common pre-processing steps such as removing stopwords, punctuation, numbers and deleting general expressions with no economic content (e.g. greetings, welcome statements and the like).

Second, I apply Latent Dirichlet Allocation (LDA) model (Blei et al., 2003) to the corpus of press conferences to identify and cluster together paragraphs that belong to the same topic. This is possible since the fundamental idea of LDA is that documents are represented as a distribution of latent topics, where each topic is characterized by a distribution over words. More formally, let D denote the entire corpus composed by i documents and $N = \sum_{d=1}^{D} N_d$ total number of words, with K being the number of latent topics. A corpus D has a distribution of topics given by $\boldsymbol{\theta_d}$ and, in turn, each topic has a distribution of words denoted by $\boldsymbol{\varphi_k}$, with both $\boldsymbol{\theta_d}$ and $\boldsymbol{\varphi_k}$

assumed to have conjugate Dirichlet distributions with hyper parameters α and β. Note that the bold font indicates the vector version of the variables. Each document i in the corpus D is an iterated choice of topics $Z_{d,n}$ and words $W_{d,n}$ drawn from the multinomial distribution using $\boldsymbol{\theta_d}$ and $\boldsymbol{\varphi_k}$. This can be formally expressed with the joint distribution of all known and latent variables given the hyper parameters as follows:

$$P(\mathbf{W}_d, \mathbf{Z}_d, \boldsymbol{\theta}_d, \boldsymbol{\Phi}; \alpha, \beta) = \prod_{n=1}^{N} \underbrace{\underbrace{\underbrace{P(\mathbf{W}_{d,n} \mid \boldsymbol{\varphi}_k) P(\boldsymbol{\Phi}; \beta)}_{\text{word level}} \cdot P(\mathbf{Z}_{d,n} \mid \boldsymbol{\theta}_d)}_{\text{topic level}} \cdot P(\boldsymbol{\theta}_d; \alpha)}_{\text{corpus level}} \tag{1}$$

where $\boldsymbol{\Phi} = \{\boldsymbol{\varphi_k}\}_{k=1}^{K}$ is a $K \times N$ matrix. As equation (1) shows, there are three levels in the LDA model: a corpus level where every document is a mixture of latent topics, a topic level where every document has a probability to belong to a topic and a word level where every word has a probability to belong to a topic. The final step is to determine K, so far, assumed to be fixed. Perplexity cross-validation measure across Markov chain Monte Carlo (MCMC) iterations (Heinrich, 2009) and minimization of the average cosine distance of topics (Cao et al., 2009) suggest that 8 topics provide the best decomposition of the press conferences. The model is then estimated integrating $\boldsymbol{\theta_d}$ and $\boldsymbol{\varphi_k}$ out of equation (1) and using Gibbs sampling simulations as in Griffiths and Steyvers (2004). Further technical details are described in Appendix A. The estimated words' probabilities ($\boldsymbol{\varphi_k}$) allow me to identify the three topics of interest (monetary policy stance, economic outlook and inflation outlook). Then, once the topics are identified, I use the probability of each paragraph to belong to one of these three topics ($\boldsymbol{\theta_d}$) to combine together paragraphs with the same content in datasets that I denote PC^{MP} for introductory statements with monetary policy information, PC^{IO} for introductory statements with inflation content, and PC^{EO} for introductory statements with growth outlook indications[1].

The third step implies the quantification of these datasets. Drawing on Picault and Renault (2017)'s methodology, I create ECB-field specific dictionaries for each topic j. I first subset each dataset PC^j into combinations of words (hereafter n-gram) and manually classify each n-gram into the topic j (monetary policy, inflation and economic outlook) with tone i (hawkish, neutral,

[1] Since every paragraph p has a probability $\boldsymbol{\theta_d}^p$ to belong to any of the three topics, paragraph p is assigned to topic j only if the topic with the highest probability ($\boldsymbol{\theta_d}^{p,MAX}$) exceeds 35%.

dovish for monetary policy and positive, neutral, negative for inflation and economic outlook)[2]. I then compute the probability that every n-gram belongs to each of the corresponding category. Next, each n-gram is classified as positive, 1, or negative, -1, on the basis of which category has the highest probability and kept only if the probability is greater than 50%. In parallel, I create an ECB-specific dictionary for valence shifters (i.e., negators, amplifiers, de-amplifiers and adversative conjunctions). Once the polarized dictionary is constructed, I apply Rinker (2019)'s methodology to measure the tone of a document. The algorithm breaks each press conference into sentences and, in turn, each sentence into an ordered bag of words (w). The word i in each sentence j of paragraph k is then compared to the dictionaries of polarized words just described. These words, that is, the words in the jth sentence that are found in the dictionaries, form a polar cluster $\gamma_{i,j,k}$ which is a subset of a sentence where every polarized word $w^p_{i,j,k}$ in the cluster is preceded and succeeded by valence shifters that weight the impact of the reference word by a factor z set by the researcher. Amplifiers $w^a_{i,j,k}$ (de-amplifiers $w^d_{i,j,k}$) increase (decrease) the polarity by z in such a way that $w_{amp} = \sum[z \cdot (w_{neg} \cdot w^a_{i,j,k})]$ where $w_{neg} = (-1)^{2+\sum w^n_{i,j,k}}$. Amplifiers become de-amplifiers $w^d_{i,j,k}$ if there are an odd number of negators $w^n_{i,j,k}$ in the cluster. This is so because w_{neg} is positive for an even number of negators and negative otherwise; such a logic is based on the rule that two negatives equal a positive, three negatives a negative, and so on. As a result, negations can also change the sign of the polarized word. On the other hand, an adversative conjunction w_{advcon} *before* the polarized word up-weights the cluster by $1 + [z \cdot (w_{advcon})]$, whereas an adversative conjunction *after* the polarized word down-weights the cluster by $1 + [(w_{advcon} - 1) \cdot z]$. This resembles the belief that an adversative conjunction augments the weight of the next clause while reducing the weight attributed to the prior clause. Overall, the score for each sentence j is computed following the equation:

$$\psi_j = \frac{\gamma^s_{i,j,k}}{\sqrt{\sum w_{i,j,n}}} \tag{2}$$

where $\gamma^s_{i,j,k} = \sum[(1 + w_{amp} + w_{deamp}) \cdot w^p_{i,j,k} \cdot w_{neg}]$ is the sum of single polar clusters and $\sqrt{\sum w_{i,j,n}}$ is the square root of the total number of words in a sentence. To obtain the

[2] For instance, the j^{th} n-gram of PC^{IO} is assigned to one of the three categories for inflation: inflation outlook positive, inflation outlook neutral and inflation outlook negative. The same procedure holds for all the datasets.

mean of all sentences within a press conference I simply calculate the average sentiment score $PC_i^m = \frac{1}{n}\sum \psi_j$, where m indicates the three topics above.

This method departs from the frequently used Loughran and McDonald (2011) (LM) and the Apel and Blix Grimaldi (2012) (AG) dictionaries in at least two dimensions: on the one hand, it is crafted on the ECB communication on a specific topic, that is, it avoids misinterpreting words' tone because of a different context; on the other hand, it considers a mixture of single words as well as sequence of n words, thereby preventing from attributing the wrong sign to specific n-grams (e.g. "lower unemployment" vs "unemployment"). Besides, unlike Picault and Renault (2017), I calculate the average tone of a press conference controlling for valence shifters. Such a method provides a much more robust inference since it is able to handle linguistic nuances that approximate the natural language. Figure 1 displays the evolution of the economic outlook, inflation outlook and monetary policy stance indexes against relevant macroeconomic data.

Figure 1: Indexes Based on the Introductory Statements of the Press Conferences

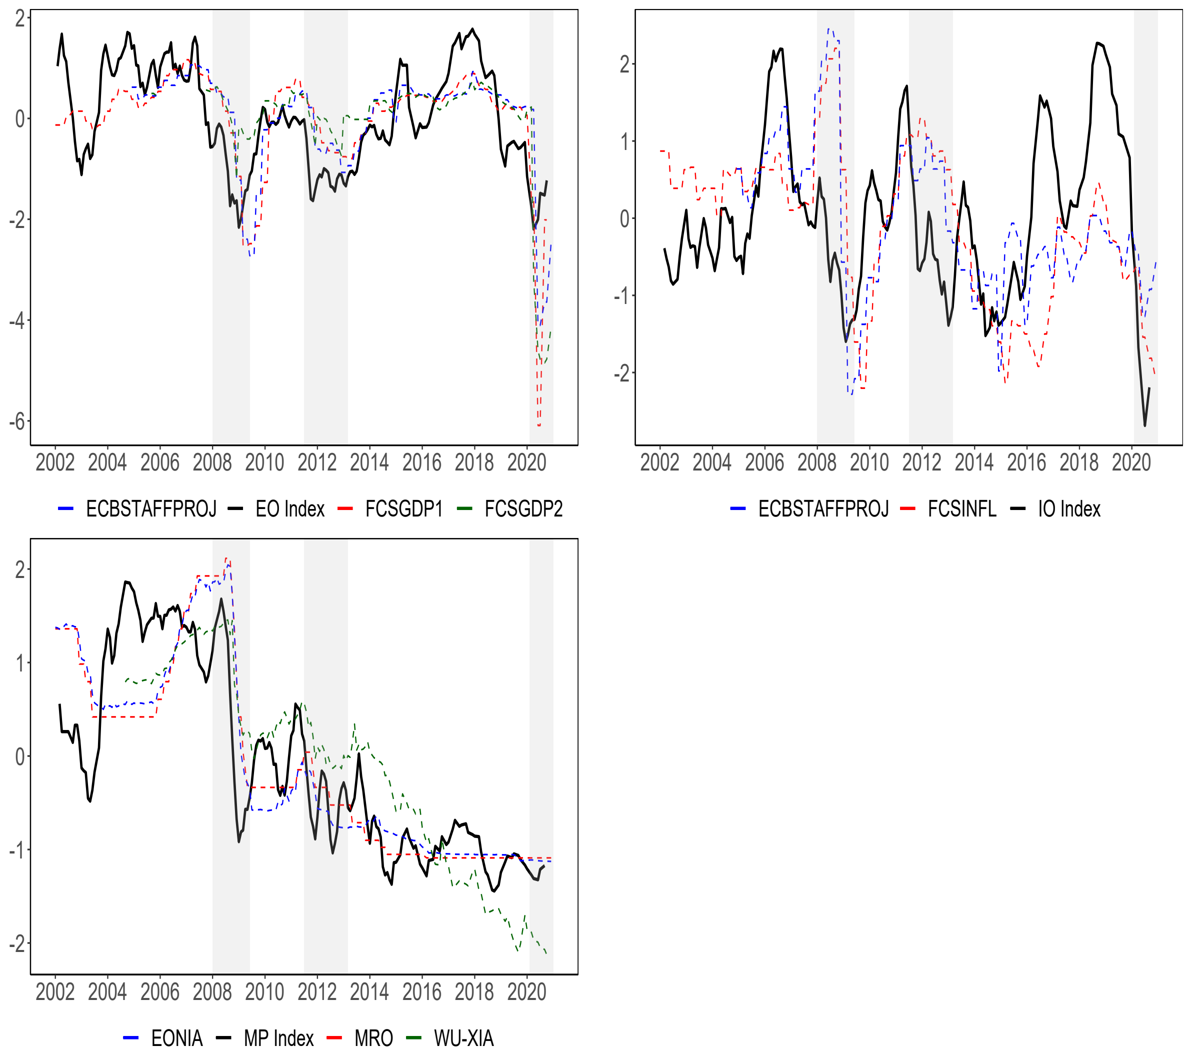

Note: The figure plots the time series of economic outlook, inflation outlook and monetary policy stance against key macro variables. The variables are standardized to facilitate the comparison. Since the indexes are somewhat jagged, I show the moving average.

In particular, the economic outlook index is strongly correlated with the ECB's GDP projections as well as the Bloomberg weighted-average of private (FCSGDP1) and official (FCSGDP2) GDP forecasts. The inflation outlook index, instead, underperformed the ECB (ECBSTAFF-PROJ) and private (FCSINFL) inflation forecasts in 2008 and 2013, while it outperformed them from 2016 to 2019. Lastly, the monetary policy stance index closely resembles the dynamics of the Euro overnight index average (EONIA), the marginal refinancing operations (MRO) and the Wu-Xia ECB shadow rate.

2.2. Words as Data: Augmenting the Nowcasting Exercise with Textual Data

In addition to conventional macro and financial data, I augment the model with information extracted from textual data at daily frequency. The dataset is taken from *Precise*'s database and contains around 300,000 documents from 19 September 2004 to 31 December 2020. This dataset includes newspapers' articles, online websites, magazines, TV news, etc. that explicitly mention the ECB in their content (see Appendix A for details).

To make this dataset applicable for time series analysis, I follow a similar procedure to the one employed in Section 2.1. After pre-processing the dataset, I decompose the textual corpus into news topics using LDA estimated with $K = 80$[3]. Out of 80 topics, I identify and label 60 of them which I reduce to 40 once filtering for monetary policy relevance[4]. The remaining 40 topics are then clustered into 7 macro topics on the basis of the similarity among topics. These 7 macro topics are: "Financial Crisis", "Eurexit", "European Banks", "Inflation Outlook","Economic Outlook","Monetary Policy Index" and "Fiscal Policy Index"[5]. Figure 2 shows the evolution of the topic probability of the first three topics. I refer to Appendix A for further details.

[3] Similar to the previous section, I estimate the optimal number of topics using perplexity cross-validation measure across Markov chain Monte Carlo (MCMC) iterations (Heinrich, 2009) and minimization of the average cosine distance of topics (Cao et al., 2009).

[4] I identify and label only 60 topics out of 80 since the remaining topics were not clearly identifiable.

[5] It should be noted that while the indexes for inflation and economic outlook in Section 2.1 represent the ECB internal assessment of those macro variables, the ones derived here reflect the tone of a much broader "audience".

Figure 2: Topic Probabilities for Daily Indexes

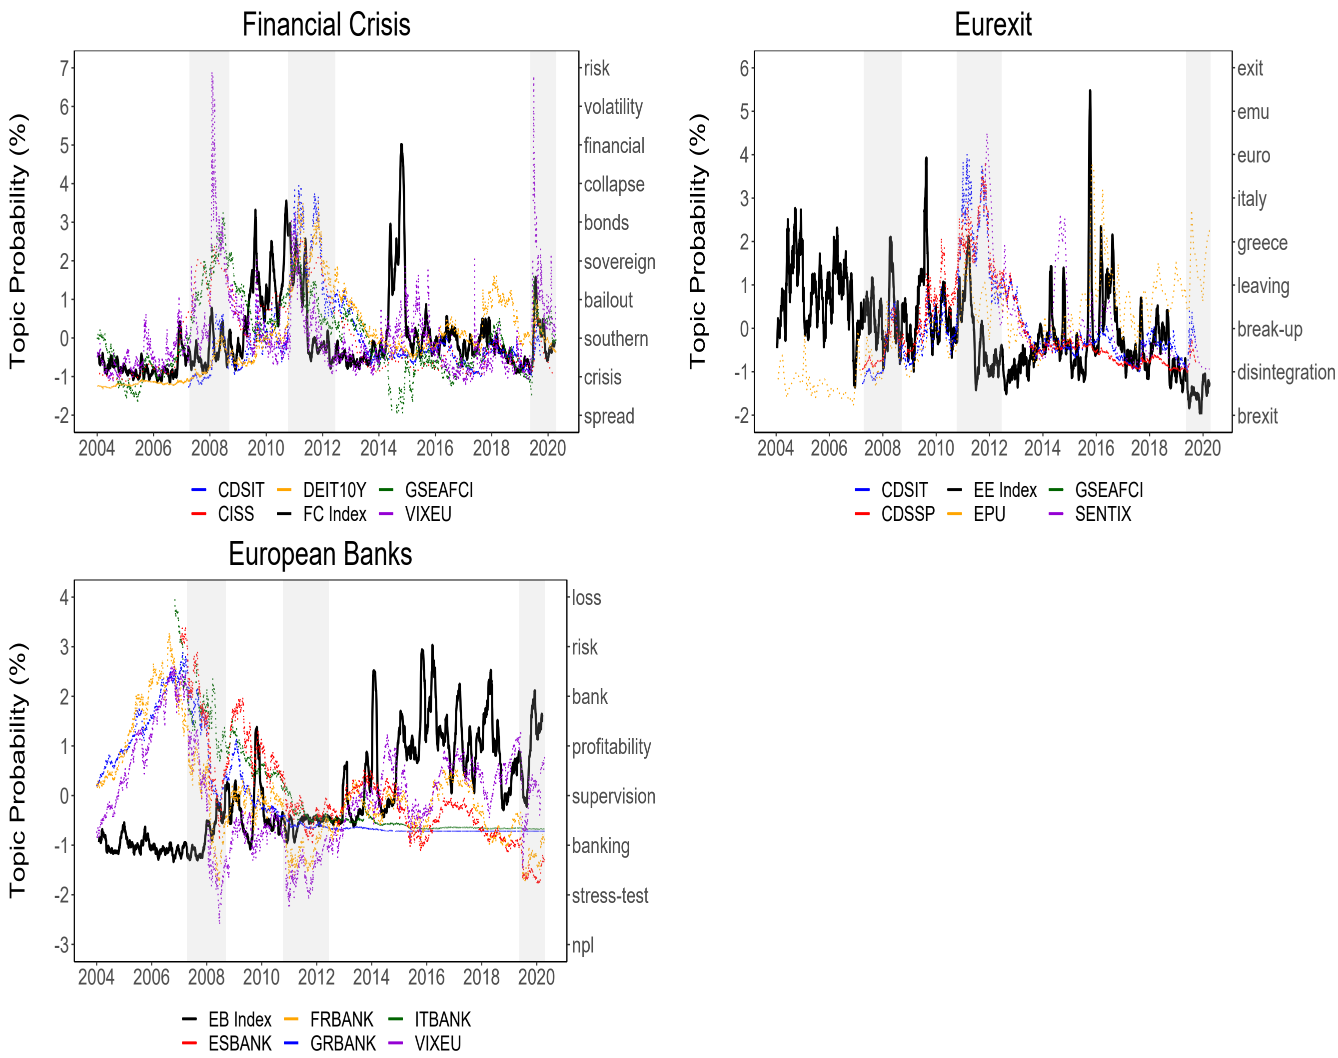

Note: The figure plots the evolution of the topic probabilities derived from the DTM for "Financial Crisis", "Eurexit" and "European banks". The variables are standardized to facilitate the comparison. The second y-axis shows the words with the highest probability to belong to a certain topic. Since the indexes are somewhat jagged, I display the moving average.

As the figure reports, the "Financial Crisis" and "Eurexit" series to a large extent correlate with a wide range of indicators from EU volatility Index (VIXEU) to the composite indicator of systemic stress (CISS), while the "European banks" index, especially from 2008 on, follows a similar evolution to a variety of bank-related indicators[6].

[6] In the figure's legend, "FRBANK" refers to FR CAC 40, "ITBANK" to MIB, "ESBANK" to IBEX 30 and "GRBANK" to FTSE/ATHEX Banks Index. Additionally, "GSEAFCI" indicates the Goldman Sachs index of financial conditions in the euro area, "EPU" stands for the Economic Policy Uncertainty Index for the euro area

While for the topics in Figure 2 the creation of field-specific dictionaries is not necessary as it would not add much information to the evolution of the topic probabilities, the remaining four topics require a quantification step to be made. To do so, I proceed in five steps that I describe in more detail in Appendix A: first, as in Section 2.1, I exploit LDA's topic probabilities to create content-specific datasets. Second, I apply word embedding (Stoltz and Taylor, 2019) to identify patterns of the co-occurrence of words within a "context window" centred around, respectively, "inflation", "economy", "monetary", and "fiscal", where these key words simply correspond to the topics of the field-specific dictionaries. Third, I compare the vector of words just obtained with the field-specific dictionaries in Section 2.1. If the word in the vector does not appear in the existing dictionary, I manually classify it to be either positive or negative and add it to the relative field-specific dictionary. For the fiscal policy index I use the same procedure but the comparison occurs with a dictionary of fiscal words developed in Marozzi (2021). Fourth, I apply equation (2) to score and aggregate each sentence for each content-specific dataset. Fifth, I filter out the bias that can come from the country of origin of the document by regressing the sentiment score on countries' dummies. Figure 3 presents the results.

and "SENTIX" is the Sentix Investor Confidence.

Figure 3: Quantified Indexes Derived from the Textual Dataset

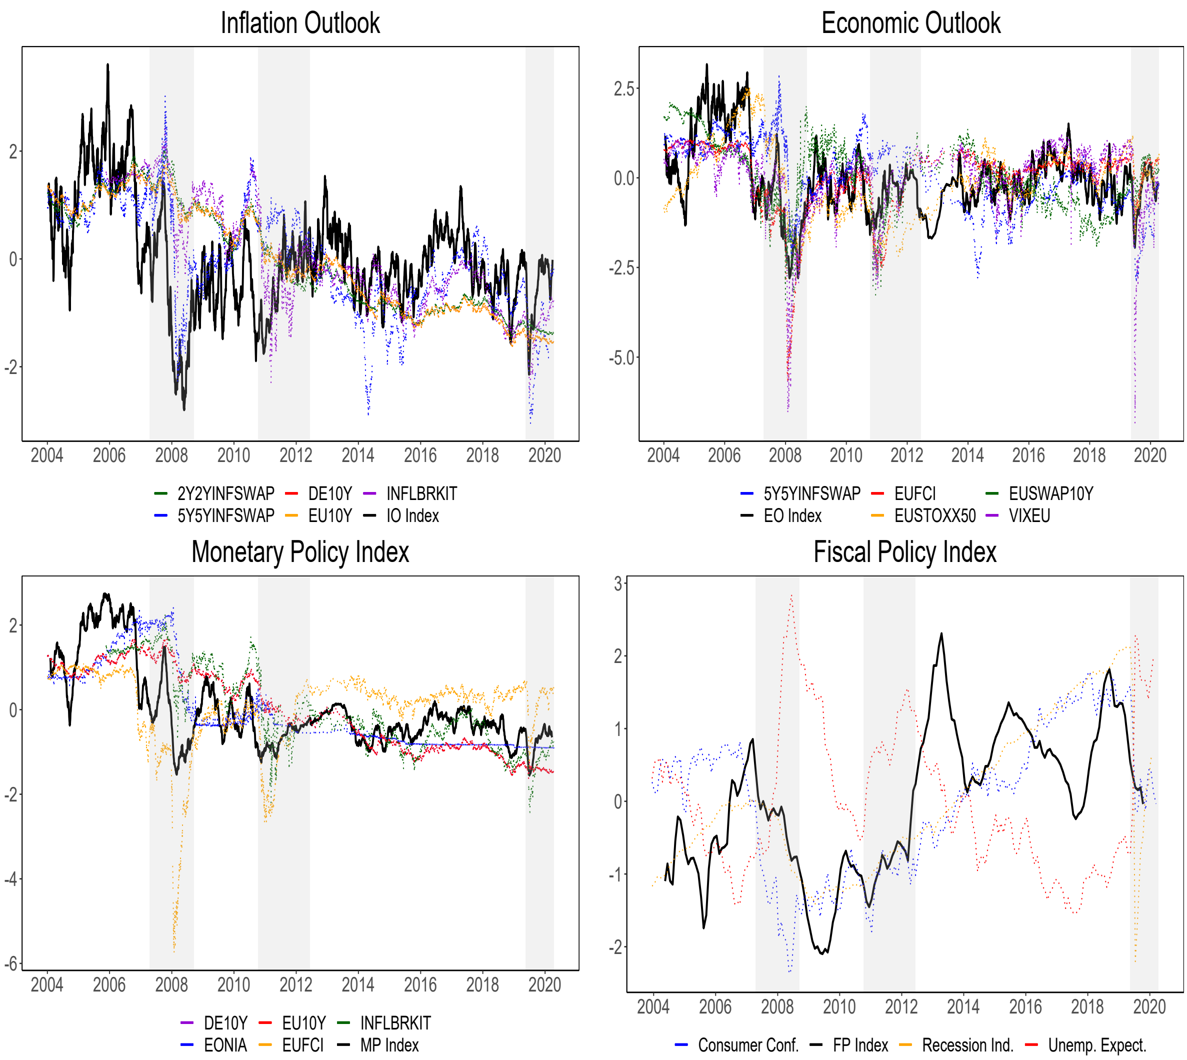

Note: The figure plots the quantified indexes of inflation outlook, economic outlook, monetary policy index and fiscal policy index derived from Precise's dataset. The variables are standardized to facilitate the comparison. Since the indexes are somewhat jagged, I display the moving average.

Notably, while the index for inflation outlook closely follows 5Y5Y inflation swap rates, the dynamics of the economic outlook index somewhat correlates with VIX EU and the European stock market index (EUSTOXX50). Moreover, the index for monetary policy resembles the evolution of a bunch of risk-free rates in the euro area. Instead, the index for fiscal policy appears to have its closest fit with the European Commission's Consumer Confidence Indicator[7].

[7] In the figure's legend, "INFLBRKIT" indicates the Italian 10Y inflation break-even rate, "EUFCI" stands for the OIS-LIBOR spread and "EUSWAP10Y" is Euro 10Y swaps.

2.3. Total Dataset

The total dataset consists of 140 macroeconomic indicators for the euro area at daily, monthly and irregular frequency, with the latter being denoted as the frequency of the introductory statements of the press conferences. Given the mixed frequency nature of the variables, I follow Mariano and Murasawa (2003) in writing the dataset at the highest frequency, i.e. daily, and therefore assuming that lower frequency variables are missing periodically[8]. Data are collected in December 2020 with the sample starting in January 2002 and transformed to induce stationarity. Since real-time vintages for every series are not available, the dataset is to be considered a "pseudo" real-time dataset.

To provide more detail on the structure of the dataset, the variables have been aggregated into eleven blocks with similar economic content and release day (Table 1)[9]. For a complete breakdown of the dataset I refer to Appendix B. Starting with the daily frequency, I group the variables into four blocks: *rates and spreads*, *financial*, *forecasts*, and *textual*. While *rates and spreads* contains government bond rates and spreads, risk-free rates, overnight index swaps (OIS), mortgage and break-even rates, the *financial* block includes stock indexes, banking and credit data and credit default swaps (CDS) for the main European countries. Moreover, the *forecasts* block comprises private and official forecasts for GDP, inflation and unemployment, whereas the *textual* block is composed of the variables derived in Section 2.2[10]. As for the monthly frequency, while *output* spans industrial production, unemployment and exchange rates variables, *price* contains various inflation indexes; *surveys* denotes Purchasing Managers Indexes (PMIs), EuroCOIN indicator and the European Commission's surveys; the *mixed* group indicates variables with the same release period ranging from sovereign CISS to indexes of financial stress and economic sentiment; ECB's loans to household, financial and non-financial institutions, ECB's holdings of securities, M3 and a range of ECB key rates are subsumed into the *monetary* group; the *US* block contains US variables ranging from PMIs to CPI. Finally, the last block for irregular data simply includes the textual variables extracted from press conferences in

[8] In so doing, I also avoid applying any transformation to the variables at irregular frequency, among which there is the variable to nowcast.

[9] Because the timing and order of data releases vary only slightly from month to month, I assume that the pattern of data availability is unchanged throughout the evaluation sample.

[10] To have forecasts in a unique release block, when forecasts were not available at a daily frequency, I filled missing observations with previous values.

Section 2.1.

Table 1: Total Dataset by Blocks

Block	Timing	Delay	Frequency	Number
Financial	Daily	No delay	Daily	16
Forecasts	Daily	No delay	Daily	6
Rates and Spreads	Daily	No delay	Daily	41
Textual Newspapers	Daily	No delay	Daily	7
Prices	Mid-month	One month	Monthly	9
Output	Mid-month	One month	Monthly	10
Surveys	End of month	No delay	Monthly	14
Mixed	End of month	One Month	Monthly	8
Monetary	End of month	One month	Monthly	19
US	Mixed	Mixed	Monthly	7
Textual PC	Press conferences	No Delay	Irregular	3

Note: The first column reports the block in which the released variable are included. The second column indicates the official dates of the publication. The third one reports the lag with which the data are released. The frequency of the data is reported in the fourth one, while in the last column is displayed the number of variables per group. Data have been collected from Haver and Bloomberg.

3. Methodology

Section 3.1 details the features of the Dynamic Factor Model (DFM) used to track in real time the developments of the ECB monetary policy stance and decisions and Section 3.2 explains the benchmark model employed to compare the performance of the DFM.

3.1. The Dynamic Factor Model

I build on Modugno (2013), Bańbura et al. (2013) and Bańbura and Modugno (2014) to develop a mixed-frequency Dynamic Factor Model with flow and stock variables. Let $Y_t^{k_z,n}$ be the collection of variables $n = \{f, s\}$ in Section 2.3 where f denotes a flow variable and s a stock variable at interval k of frequency $z = \{i, m, d\}$ where i stands for irregular, m for monthly and d for daily[11]. Let $F_t^{k_z,n}$ also denote the corresponding unobservable factor for $Y_t^{k_z,n}$. The measurement equation can be written as follows:

$$\begin{bmatrix} Y_t^{k_i,f} \\ Y_t^{k_m,f} \\ Y_t^{k_m,s} \\ Y_t^{d} \end{bmatrix} = \begin{bmatrix} \tilde{\Lambda}^{i,f} & 0 & 0 & 0 \\ 0 & \tilde{\Lambda}^{m,f} & 0 & 0 \\ 0 & 0 & \Lambda^{m,s} & 0 \\ 0 & 0 & 0 & \Lambda^{d} \end{bmatrix} \begin{bmatrix} \tilde{F}_t^{k_i,f} \\ \tilde{F}_t^{k_m,f} \\ F_t^{k_m,s} \\ F_t^{d} \end{bmatrix} + \begin{bmatrix} E_t^{k_i,f} \\ E_t^{k_m,f} \\ E_t^{k_m,s} \\ E_t^{d} \end{bmatrix}, \qquad (3)$$

where $\Lambda^{z,\cdot}$ are the factor loadings for each frequency z and $\tilde{\Lambda}^{z,f} = (\Lambda^{z,f}, 0)$ is needed to match $\tilde{F}_t^{k_z,f} = (F_t^{k_z,f}, \bar{F}_t^{k_z,f})$ that contains an additional auxiliary aggregator $\bar{F}_t^{k_z,f}$ for the flow variables; $E_{i,t}$ are idiosyncratic errors such that $E_{i,t} = i.i.d.N(0, \sum_E)$ where $\sum_E$ is diagonal. I leave technical details on aggregators in Appendix C. The transition equation with time-varying coefficients can then be written as:

$$\begin{bmatrix} I_{2r} & 0 & 0 & \mathcal{W}_t^{i,f} \\ 0 & I_{2r} & 0 & \mathcal{W}_t^{m,f} \\ 0 & 0 & I_r & \mathcal{W}_t^{m,s} \\ 0 & 0 & 0 & I_r \end{bmatrix} \begin{bmatrix} \tilde{F}_t^{k_i,f} \\ \tilde{F}_t^{k_m,f} \\ F_t^{k_m,s} \\ F_t^{d} \end{bmatrix} = \begin{bmatrix} \mathcal{I}_t^{i,f} & 0 & 0 & 0 \\ 0 & \mathcal{I}_t^{m,f} & 0 & 0 \\ 0 & 0 & \mathcal{I}_t^{m,s} & 0 \\ 0 & 0 & 0 & \Phi \end{bmatrix} \begin{bmatrix} \tilde{F}_{t-1}^{k_i,f} \\ \tilde{F}_{t-1}^{k_m,f} \\ F_{t-1}^{k_m,s} \\ F_{t-1} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 0 \\ U_t \end{bmatrix}, \qquad (4)$$

where $\mathcal{W}_t$ contain aggregation weights for every frequency higher than the daily one, $\mathcal{I}_t$ are selection matrices of time varying coefficients that take the value of zero the day after each release of data with frequency z and one elsewhere, Φ is a matrix of autoregressive coefficients for the daily factors that are assumed to follow a VAR(p) process and $U_{i,t} \sim i.i.dN(0, \sum_U)$ where $\sum_U$ is a diagonal matrix.

The model is then estimated using the methodology developed by Bańbura and Modugno (2014) that adapts from Dempster et al. (1977) the Expectation Maximization (EM) algorithm to handle arbitrary patterns of missing data. The algorithm consists of three steps: first, initial parameters are estimated via principal components on a sample of data where missing observations are handled using splines. Afterwards, the algorithm iterates until convergence to

[11] The interval k is defined as the average number of days between two releases.

a local maximum between the expectation step, where the missing data in the likelihood are filled in by the Kalman filter and smoother, and the maximization step where this likelihood with complete data is optimized (further details in Appendix C).

Equations (3) and (4) output the nowcasts of monetary policy stance. However, they are not informative on the expected probability that the ECB, conditional on the incoming data, will actually take a monetary policy decision. To add this important piece of information, I augment the model with a bridge equation whose specification is:

$$P(y_{t+h} = j|\hat{X}_{t+h}) = \Phi(\alpha + \beta' \hat{X}_{t+h}) \,. \tag{5}$$

Equation (5) is a multinomial logit model where y_{t+h} is a categorical variable that takes $j = 3$ values: y_{t+h} equals 1 if the ECB hikes at time $t+h$, 0 if there is no actual change and -1 if the ECB eases at time $t+h$[12]; α and β are, respectively, a constant and a vector of parameters; $\Phi(\cdot)$ denotes the cumulative distribution function of the logistic distribution; and $\hat{X}_{t+1}$ is a vector containing a set of predictors for which the DFM provided the forecasts at time $t+1$[13]. This bridge equation thus yields updated nowcasts of the conditional probability attached to each outcome whenever new information becomes available.

This model finally provides a framework to structure the flow of data releases in real time to trace out which variables drive the evolution of the nowcast. Drawing on Bańbura and Modugno (2014), the decomposition of the source of nowcast revision can be written formally. Let Ω_{t-1} and Ω_t denote two consecutive data releases whose difference is I_{new}, that is, the observations for the j^{th} variable in Ω_t, missing in Ω_{t-1}, that the new release made available[14]. The interest is in inspecting how the release of I_{new} affected the updated nowcast of the ECB monetary policy stance. Given the orthogonality between I_{new} and Ω_t, it is possible to write the expression for the nowcast revision as the difference between the new and the old now-cast:

$$\underbrace{\mathbb{E}(Y_t^{k_z,n}|I_{new})}_{\text{news}} = \underbrace{\mathbb{E}(Y_t^{k_z,n}|\Omega_t)}_{\text{new forecast}} - \underbrace{\mathbb{E}(Y_t^{k_z,n}|\Omega_{t-1})}_{\text{old forecast}} \,. \tag{6}$$

[12] Appendix C provides details on how I turned a continuous variable into a categorical one.

[13] The set of regressors are: core inflation, industrial production, M3, EONIA, 10Y BUND-BTP spread, trade-weighted exchange rate and forecasts of inflation and GDP.

[14] To be precise, Ω_{t-1} and Ω_t can differ for two reasons: the first reason is that new observations for the j^{th} variable become available and the second one is that some of the past data might have been revised. I abstract from the latter case.

Note that in equation (6), $\mathbb{E}(Y_t^{k_z,n}|I_{new}) \neq 0$ only if $I_{new} = y^j - y^j_{t|\Omega_{t-1}} \neq 0$, that is, the nowcast is updated only if the new observations in Ω_t are different from the forecasts made at time $t-1$. Therefore, it is not the new release that leads to the revision of the forecast but the unexpected component, the *news*, from the latest release. Abstracting from the problem of parameter uncertainty, the expression for the revision can be developed further as follows:

$$\mathbb{E}(Y_t^{k_z,n}|I_{new}) = \mathbb{E}(Y_t^{k_z,n} I'_{new})\, \mathbb{E}(I_{new} I'_{new})^{-1} I_{new} \;, \tag{7}$$

where, for the state space model written as in equations (3) and (4):

$$\begin{aligned} \mathbb{E}(Y_t^{k_z,n} I'_{new}) &= \Lambda^{z,\cdot}\, \mathbb{E}[(F_t^{k_z,n} - F^{k_z,n}_{t|\Omega_{t-1}})(F_t^{k_z,n} - F^{k_z,n}_{t|\Omega_{t-1}})']\Lambda_j^{z,\cdot'} \\ \mathbb{E}(I_{new} I'_{new}) &= \Lambda_j^{z,\cdot}\, \mathbb{E}[(F_{tj}^{k_z,n} - F^{k_z,n}_{tj|\Omega_{t-1}})(F_{tj}^{k_z,n} - F^{k_z,n}_{tj|\Omega_{t-1}})']\Lambda_j^{z,\cdot'} + \sum j \;, \end{aligned} \tag{8}$$

where $\Lambda^{z,\cdot}$ and $\Lambda_j^{z,\cdot}$ are the factor loadings of the observation equation with the rows corresponding, respectively, to the ECB monetary policy stance variable and the j^{th} variable and $\sum_j$ is a diagonal matrix filled by j^{th}'s data. The expectations are estimated using the Kalman filter and smoother. As a result, it is possible to find a vector of weights $A = (a_1, \ldots, a_j)$ such that:

$$\underbrace{\mathbb{E}(Y_t^{k_z,n}|\Omega_t)}_{\text{new forecast}} - \underbrace{\mathbb{E}(Y_t^{k_z,n}|\Omega_{t-1})}_{\text{old forecast}} = A(\underbrace{y^j - y^j_{t|\Omega_{t-1}}}_{\text{news}}) \;. \tag{9}$$

Equation (9) states that the nowcast revision can be decomposed as a weighted average of the news in the most recent release. This relationship allows me to compute the weighted marginal contribution of each variable to the updated nowcast. More precisely, since in the dataset used for the empirical application there is simultaneous release of several variables, the results of the decomposition have been aggregated into groups of series that share approximately the same release day and economic characteristics as outlined in Table 1.

3.2. Benchmark Model

In the empirical section, I will compare the performance of the DFM with that of a forecast-based policy rule with contemporaneous and forward-looking measures of inflation and output gap (Jansen and De Haan, 2009) where the dependent variable, rather than being the nominal interest rate, is the monetary policy stance of the ECB (ECB_t) derived in Section 2.1. The choice of a forecast-based policy rule compared to a simple outcome-based one is led by the fact that the ECB is found to set interest rates in a forward-looking manner (see Gerlach, 2007; Gorter et al., 2008). I then estimate the model in the following specification:

$$ECB_t = \alpha + \phi_\pi(\pi_t - \pi^*) + \phi_\gamma(\gamma_t - \gamma_t^*) + \phi_{\pi,h}\mathbb{E}_t\pi_{t+h} + \phi_{\gamma,h}\mathbb{E}_t\gamma_{t+h} + \epsilon_t \ , \quad (10)$$

where α is just a constant, $\pi_t - \pi^*$ indicates the difference between the growth rate of core inflation and the ECB inflation target, $\gamma_t - \gamma_t^*$ denotes the difference between the growth rate of output and the output gap, π_{t+h} and γ_{t+h} are respectively the inflation and output growth ECB staff forecasts four quarters ahead ($h = 4$)[15]; ϕ_π, ϕ_γ, $\phi_{\pi,h}$, $\phi_{\gamma,h}$ indicate the ECB's response parameters to be estimated and ϵ_t is a monetary policy shock capturing deviations from the systematic policy response to output and inflation.

4. Empirical Results

The DFM is parametrized with one factor per frequency and variable type and one lag[16]. It is then estimated recursively on the time frame ranging from 1 January 2005 to 31 December 2020. The bridge equation is also estimated as part of the entire econometric framework. The performance of this model is thus assessed from different perspectives[17]. First, Figure 4 reports the in-sample estimates of the actual vs fitted values of ECB monetary policy stance (first column) as well as the in-sample fit of the bridge equation (second column).

[15] Constant four-quarter ahead forecasts are derived as follows: December, $y_{t+4} = y_{t+4}$; March, $y_{t+4} = 0.75y_{t+3} + 0.25y_{t+7}$; June, $y_{t+4} = 0.5y_{t+2} + 0.5y_{t+6}$; September, $y_{t+4} = 0.25y_{t+1} + 0.75y_{t+5}$. The series have then been interpolated to fit the regression model. I used industrial production as a measure of output and derived the output gap using the Hodrick-Prescrott (HP) filter.

[16] This choice is motivated by simplicity and by the fact that results based on two factors and/or more lags are qualitatively similar to those based on one factor.

[17] Since the DFM is built to nowcast a variable that I developed, in Appendix D I test the "external" validity of the model nowcasting variables related to the ECB's monetary policy stance but not developed by myself. The last exercise carried out in Section 4 provides further evidence on the robustness of the DFM.

Figure 4: In-Sample Results

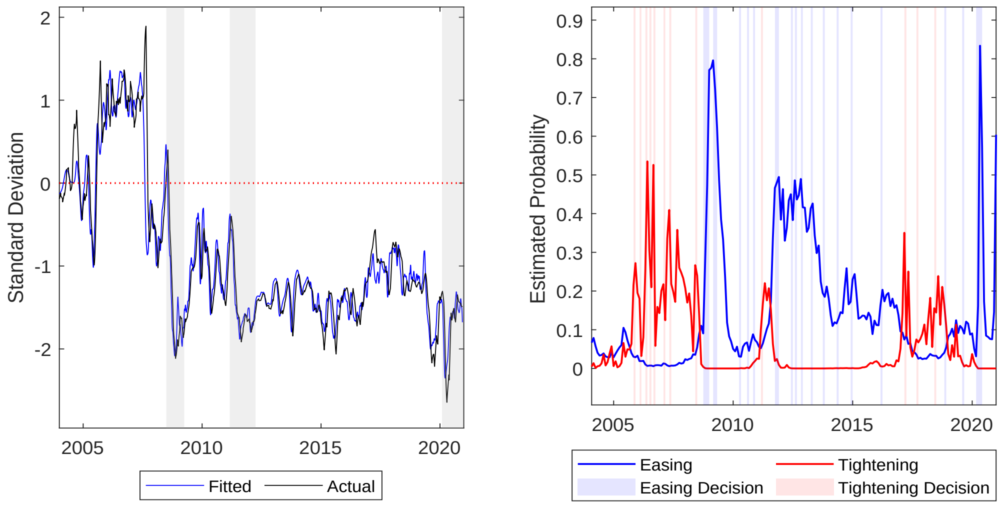

Note: The left figure shows the fitted values (blues) against the actual values (black) of the ECB monetary policy stance, while the figure on the right displays the estimated probability of an easing (blue) and a tightening (red) against actual monetary policy decisions to tighten (shaded red) and ease (shaded blue).

As the figure shows, the fitted values track closely the actual index of monetary policy stance. Besides, the estimated probability of tightening and easing significantly increase in correspondence of introductory statements where such decisions were taken. More precisely, the bridge equation correctly predicts 86% of tightening decisions and 90% of easing decisions[18].

Second, I evaluate the average precision of the nowcasts by running a recursive out-of-sample exercise that produces nowcasts for the nearest press conference. The model is updated every week over the evaluation sample 2005–2020 from the day after the press conference to the day before the press conference and replicates at each point of the forecast evaluation sample the "pseudo" real-time data availability to that point in time. The model is re-estimated each time in order to take into account parameter uncertainty. In particular, I focus on point forecast evaluation using Root Mean Squared Forecast Error (RMSFE) statistics. The accuracy of the DFM-based predictions are then compared to the benchmark model described in Section 3.2. Since the frequency of the ECB press conferences is approximately monthly until 2014 before moving to a six-week cycle, I calculate the RMSFE for each subperiod. Figure 5 reports the result.

[18] To understand how these percentages were calculated see Appendix C.

Figure 5: Root Mean Squared Forecast Error (RMSFE)

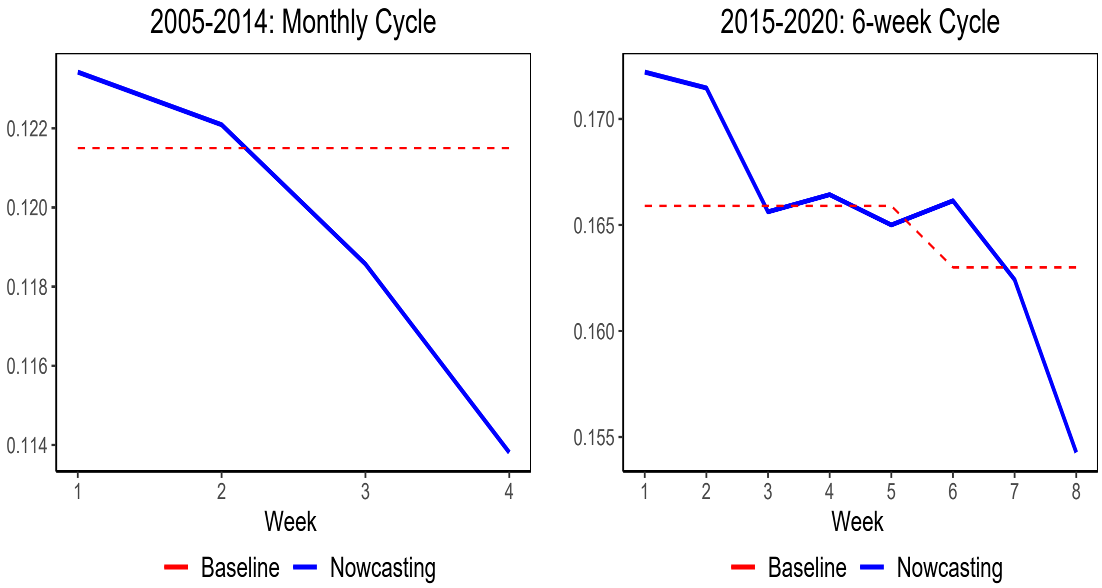

Note: The figure shows the RMSFE of the nowcasting model compared to the baseline model outline in Section 3.2. The first column displays the RMSFE for the period 2005-2014, while the second one exhibits the RMSFE from 2015 to 2020.

In either periods, the RMSFE of the DFM displays a downward-sloping behaviour, that is, the more the data become available, the more the accuracy of the prediction improves. In contrast, the RMSFE of the baseline model outlined in Section 3.2 remains constant in the monthly cycle and reduces only in the fifth week of the six-week cycle in correspondence of the release of real activity data. Although the baseline model outperforms the DFM at the beginning of the forecasting period, it tends to lag behind as the weeks progress. The DFM, therefore, maintains an informational advantage compared to the benchmark model due to the ablity to process mixed-frequency data in a unified framework.

To complement this analysis, I show the most important variables driving the nowcast of the ECB monetary policy stance. Figure 6 documents the average absolute weekly impact of each block on the updated nowcast computed in real time. In other words, Figure 6 sheds light on the reaction function of the ECB.

Figure 6: Average Absolute Contribution

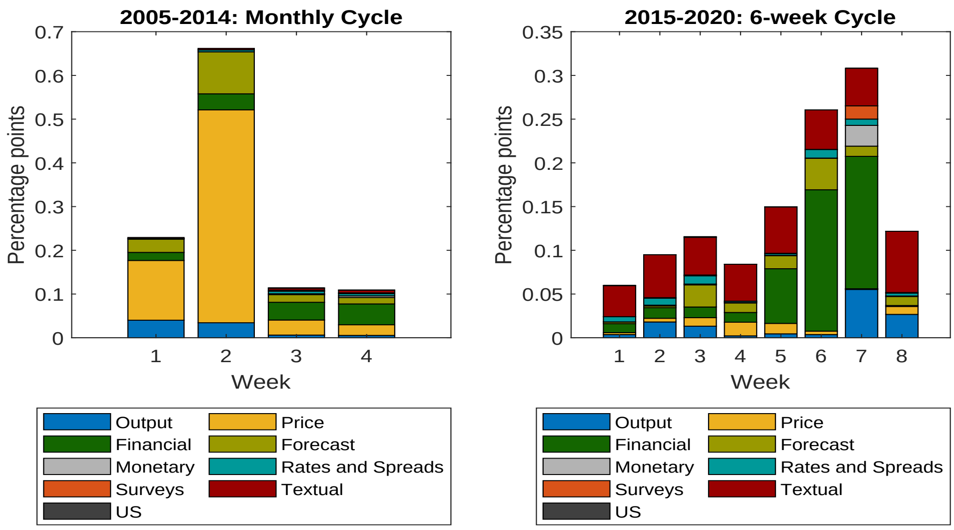

Note: The figure reports the average (absolute) weekly impact of each series, grouped by category, computed in real time over the evaluation sample 2005–2020. The first column shows the monthly-cycle (2005-2014), while the second one the six-week cycle (2015-2020).

From this figure, three features can be highlighted. First, while in the period 2005-2014 the main contributors to changes in nowcasts are, in order of importance, price, forecast, financial, output and us data, in the period 2015-2020, the most prominent blocks are financial, textual, forecast, output, prices and rates and spread data. Second, the monthly-cycle period displays a bell shape indicating that the most useful information for the nowcast arrives in the second week of the nowcasting period, when most of the hard data first become available. Conversely, the six-week cycle approximates a bimodal distribution with the third and seventh week being respectively the minor and major mode. In absolute terms, the most informational weeks are the sixth and seventh week where the second release of hard data become available. Instead, news move the nowcast relatively less at the beginning, where signals are likely to be weak, and at the end of the period, where the room for improvement has significantly reduced. Last, the contribution of textual variables have been large in the period 2015-2020, bringing new evidence to the usefulness of text-data in forecasting purposes.

Third, to further test the model performance I replicate a real-time nowcast for two historical "easing" announcements - the Outright Monetary Transactions (OMT) on August 2, 2012 and the Asset Purchase Programme (APP) on January 22, 2015[19] - and two historical "tightening"

[19] The announcement of the OMTs was preceded by then-President Mario Draghi's "whatever it takes" at the

ones - the 25 basis points (bps) interest rise on July 3, 2008 and April 7, 2011. Figure 7 and Figure 8 show the evolution of the nowcast for the events of interest: the red line indicates the point forecasts of the monetary policy stance, while the black line represents the nowcast of the expected probability of an easing decision.

Figure 7: Nowcasting Historical Episodes - Easing

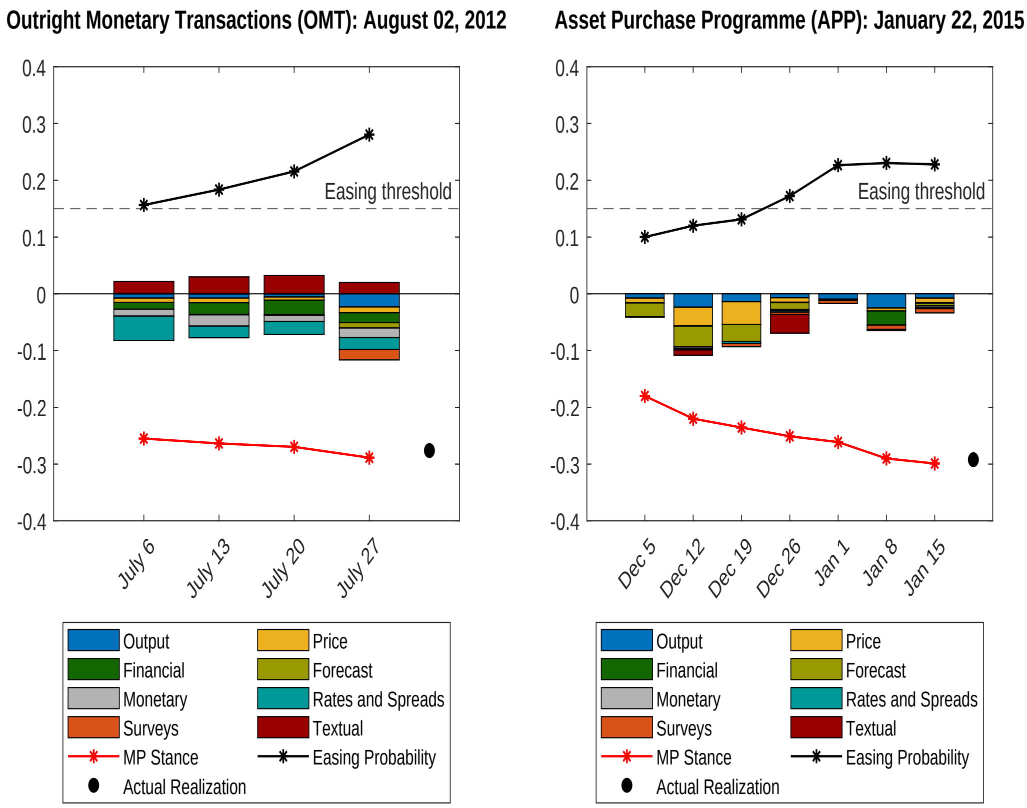

Note: The figure shows the evolution of the nowcast for the announcements of the OMTs and the APP. While the black line represents the nowcast of the expected probability of an easing decision, the red line indicates the nowcast of the monetary policy stance. News contribution have been rescaled for better visualization.

Focusing on the announcement of the OMTs, the nowcast of the ECB's monetary policy stance (red line) starts with a significant dovish realization before gradually turning more negative approximating the actual value of the press conference (black circle). Setting 15% as a threshold for a monetary policy easing (see Appendix C), the expected probability of an easing decision (black line) is robustly predicted four weeks ahead of the actual Governing Council meeting. This nowcast is mainly driven by the significant downward surprises from financial,

Global Investment Conference in London on July 26 2012.

monetary and rates and spread data, indicating that the evolution of the nowcast is due to a tightening of financial and monetary conditions as well as a widening of sovereign rate differentials[20]. Turning to the APP, while the monetary policy stance (red line) becomes more accurate as the weeks progress, the model anticipates a monetary easing (black line) four weeks in advance of the press conference. For this event, the nowcasts are mostly driven by lower-than-expected forecast, price and output data[21].

Moving to tightening decisions, at the peak of the financial crisis, on July 3, 2008, the Governing Council of the ECB decided to increase the key interest rates by 25 bps motivating the decision as follows: "This decision was taken [· · ·] to counteract the increasing upside risks to price stability over the medium term. HICP inflation rates have continued to rise significantly since the autumn of last year"[22]. Coherently, the first column of Figure 8 shows that the downward pressure stemming from spread and rates, output and survey is mainly counteracted by higher-than-expected price-related variables and to a lesser extent by the forecast block. Additionally, considering a tightening threshold of 8% (see Appendix C), the model predicts a hike four weeks in advance of the actual press conference. Turning to the 25 bps rate increase on April 7, 2011 "warranted in the light of upside risks to price stability that we have identified in our economic analysis"[23], the second column of Figure 8 displays a mildly hawkish monetary policy stance driven by the price and forecast blocks. Moreover, although the conditional probability of a tightening decision declines after the third week, it remains always above the tightening threshold throughout the nowcasting period.

[20] The ECB Governing Council justified the OMTs stating that "the severe malfunctioning in the price formation process in the bond markets of euro area countries [· · ·] need to be addressed in a fundamental manner". *Introductory statement to the press conference (with Q&A)*, Mario Draghi, President of the ECB, Vítor Constâncio, Vice-President of the ECB, Frankfurt am Main, 2 August 2012.

[21] The ECB decided to launch the APP as a result of "inflation dynamics [· · ·] [being] weaker than expected [· · ·] [and as a result of] a further fall in market-based measures of inflation expectations over all horizons and [of] the fact that most indicators of actual or expected inflation stand at, or close to, their historical lows". *Introductory statement to the press conference (with Q&A)*, Mario Draghi, President of the ECB, Frankfurt am Main, 22 January 2015.

[22] *Introductory statement to the press conference (with Q&A)*, Jean-Claude Trichet, President of the ECB, Vítor Constâncio, Vice-President of the ECB, Frankfurt am Main, 3 July 2008.

[23] *Introductory statement to the press conference (with Q&A)*, Jean-Claude Trichet, President of the ECB, Vítor Constâncio, Vice-President of the ECB, Frankfurt am Main, 7 April 2011.

Figure 8: Nowcasting Historical Episodes - Tightening

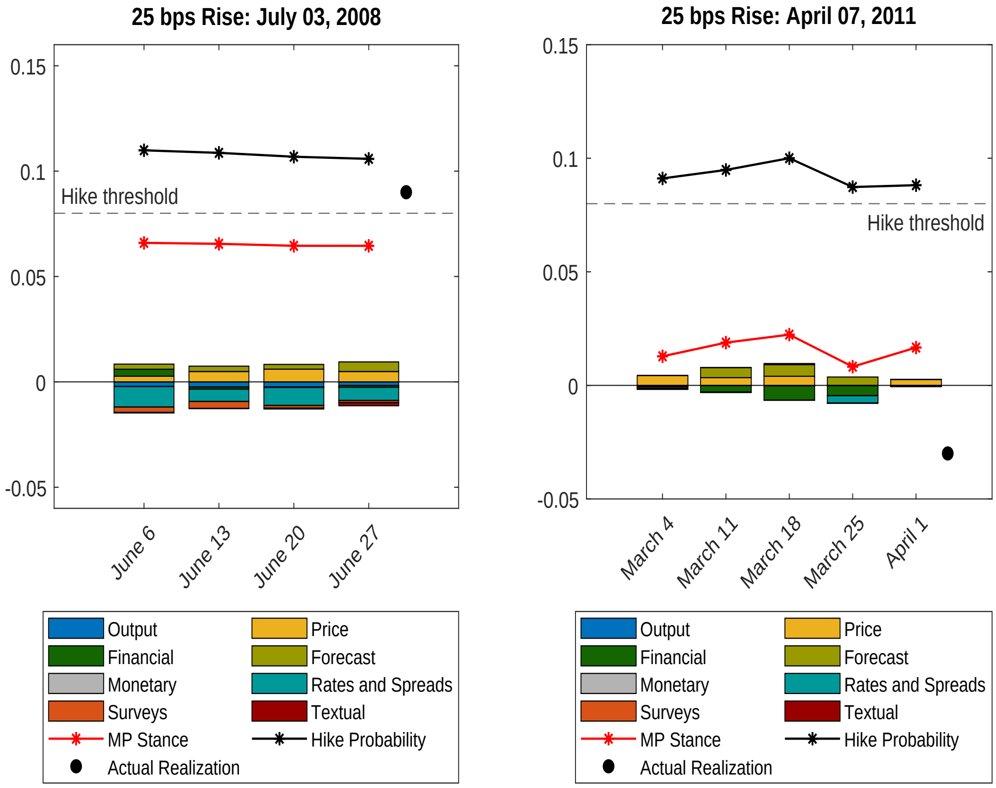

Note: The figure shows the evolution of the nowcast for the announcements of the 25bps rate increase on July 3, 2008 and on April 7, 2011. While the black line represents the nowcast of the expected probability of an easing decision, the red line indicates the nowcast of the monetary policy stance. News contribution have been rescaled for better visualization.

The final exercise I propose is to nowcast, out-of-sample, the evolution of the EONIA rate from January 2008 to December 2009 exploiting the forecasts of the daily factor of ECB monetary policy stance, output and inflation-related variables[24]. This period is particularly important since the European overnight swap rates for the first time approximated the (then-perceived) zero lower bound as a result of the intensification of the 2007-2009 global financial crisis. Figure 9 graphs what would have been the forecast for the EONIA rates conditional on the information available at the time of the analysis.

[24] EONIA is assumed to be unknown during each month even if its data are released daily. Output variables include industrial production, public and private GDP forecast, while inflation variables refer to core inflation, public and private HICP forecast.

Figure 9: Forecasting EONIA

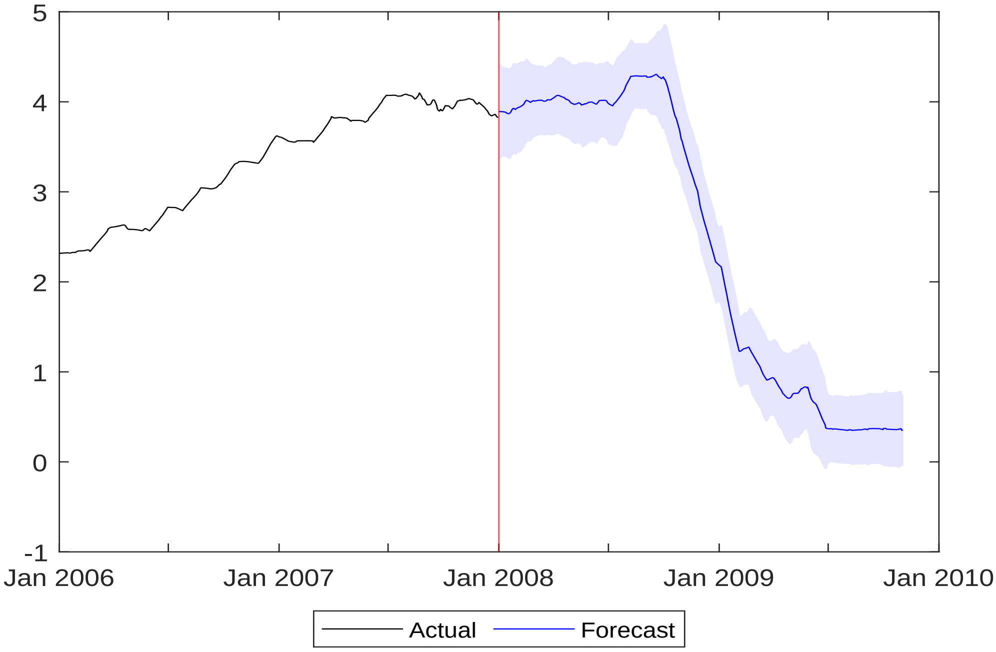

Note: The figure shows the actual values of EONIA rates (black line) and the out-of-sample nowcast with its 90% confidence intervals (blue line).

As the figure shows, the model proves to be valuable also in tracking the evolution of the EONIA rates in a period of high volatility and uncertainty. In fact, since the end of 2008, the forecasts of the EONIA sharply decline reflecting the validity of the forecasting strategy.

5. Conclusion

This paper proposes an econometric framework to track in real time the monetary policy stance and decisions of the ECB exploiting conventional and unconventional data at a sampling frequency higher or equal than monthly that become available between two consecutive press conferences. This mixed-frequency dataset is modelled as a dynamic factor model estimated at daily frequency and augmented with a bridge equation that takes the specification of a multinominal logit. Such a framework outputs three pieces of information: the nowcast of the ECB monetary policy stance, the forecast of the conditional probability that the ECB will actually take a monetary policy decision and the block of variables that drives the revision of each nowcast at every point in time.

The empirical results are the following: first, that the model provides higher forecast accuracy than a benchmark model approximated by a "pseudo" Taylor rule; second, the inclusion of textual variables in the dataset contributes considerably to the improvement of the forecasting performance; third, the model is accurate in tracking the monetary policy stance and decisions of the ECB at key historical announcements; last, the DFM proves to be useful in forecasting the EONIA rates from January 2008 to December 2009.

APPENDIX

A. Text Mining: Data and Techniques

Section A.1 provides details on the estimation of the LDA model; Section A.2 compares the indexes derived from the press conferences with alternative dictionaries; Section A.3 describes the textual dataset and Section A.4 sheds additional light on the quantification of the time series extracted from the textual dataset.

A.1. Estimating the Latent Dirichlet Allocation (LDA) Model

This paper implements the LDA model developed by Blei et al. (2003) and follows the estimation algorithm described by Griffiths and Steyvers (2004). For a full explanation of the mechanics of the model I refer to Heinrich (2009); here I will only deploy the essential features.

Formally, let D denote the entire corpus composed by d documents where $N = \sum_{d=1}^{D} N_d$ is total number of words, K the number of latent topics and V the size of the vocabulary. A corpus D has a distribution of topics given by $\boldsymbol{\theta_d}$ and, in turn, each topic has a distribution of words denoted by $\boldsymbol{\varphi_k}$, with both $\boldsymbol{\theta_d}$ and $\boldsymbol{\varphi_k}$ assumed to have conjugate Dirichlet distributions with hyper parameters α and β. Note that bold-font variables denote vectors. Each document i in the corpus D is an iterated choice of topics $\mathbf{Z_{d,n}}$ and words $\mathbf{W_{d,n}}$ drawn from the multinomial distribution using $\boldsymbol{\theta_d}$ and $\boldsymbol{\varphi_k}$. Moreover, let t be a term in V and denote $P(t|z = k)$, the mixture component, one for each topic, by $\boldsymbol{\Phi} = \{\boldsymbol{\varphi_k}\}_{k=1}^{K}$. Finally, let $P(z|D = d)$ define the topic mixture proportion for document d, with one proportion for each document $\boldsymbol{\Theta} = \{\boldsymbol{\theta_d}\}_{d=1}^{D}$. The aim of the algorithm is then to approximate the distribution:

$$P(\mathbf{Z}, \mathbf{W}; \alpha, \beta) = \frac{P(\mathbf{W}, \mathbf{Z}; \alpha, \beta)}{P(\mathbf{W}; \alpha, \beta)} \tag{A.1}$$

using Gibbs simulations, where α and β are the (hyper) parameters controlling the prior conjugate Dirichlet distributions for $\boldsymbol{\varphi_k}$ and $\boldsymbol{\theta_d}$, respectively.

With these definitions, the probability of the model can be written as follows:

$$P(\mathbf{Z}, \mathbf{W}, \boldsymbol{\Theta}, \boldsymbol{\Phi}; \alpha, \beta) = \prod_{k=1}^{K} P(\boldsymbol{\varphi_k}; \beta) \prod_{i=1}^{I} P(\boldsymbol{\theta_d}; \alpha) \prod_{n=1}^{N} P(\mathbf{Z_{d,n}}|\boldsymbol{\theta_d}) P(\mathbf{W_{d,n}}|\boldsymbol{\varphi_k}) \tag{A.2}$$

Integrating out the parameters $\boldsymbol{\varphi}$ and $\boldsymbol{\theta}$:

$$P(\mathbf{Z},\mathbf{W};\alpha,\beta) = \int_{\Theta}\int_{\Phi} P(\mathbf{Z},\mathbf{W},\boldsymbol{\Theta},\boldsymbol{\Phi};\alpha,\beta)d\boldsymbol{\Phi}d\boldsymbol{\Theta}$$

$$= \int_{\Phi}\prod_{k=1}^{K} P(\boldsymbol{\varphi_k};\beta)\prod_{i=1}^{I}\prod_{n=1}^{N} P(\mathbf{W_{d,t}}|\boldsymbol{\varphi_{z_d,t}})d\boldsymbol{\Phi} \int_{\Theta}\prod_{d=1}^{I} P(\boldsymbol{\Theta_k};\alpha)\prod_{n=1}^{N} P(\mathbf{Z_{d,t}}|\boldsymbol{\Theta_k})d\boldsymbol{\Theta} \quad \text{(A.3)}$$

In equation (A.3), the first integral does not include $\boldsymbol{\theta}$ nor the second integral contains $\boldsymbol{\varphi}$. As a result, $\boldsymbol{\varphi}$ and $\boldsymbol{\theta}$ can be solved separately. Thus, drawing on the properties of the conjugate Dirichlet distribution it can be shown that:

$$\int_{\Phi}\prod_{k=1}^{K} P(\boldsymbol{\varphi_k};\beta)\prod_{i=1}^{I}\prod_{n=1}^{N} P(\mathbf{W_{d,t}}|\boldsymbol{\varphi_{z_d,t}})d\boldsymbol{\Phi} = \frac{\Gamma(\sum_{k=1}^{K}\alpha_k)}{\prod_{k=1}^{K}\Gamma(\alpha_k)}\frac{\prod_{k=1}^{K}\Gamma(n_d^{(k)}+\alpha_k)}{\Gamma(\sum_{k=1}^{K}n_d^{k}+\alpha_k)} \quad \text{(A.4)}$$

and

$$\int_{\Theta}\prod_{d=1}^{I} P(\boldsymbol{\Theta_k};\alpha)\prod_{n=1}^{N} P(\mathbf{Z_{d,t}}|\boldsymbol{\Theta_k})d\boldsymbol{\Theta} = \prod_{k=1}^{K}\frac{\Gamma(\sum_{t=1}^{V}\beta_t)}{\prod_{t=1}^{V}\Gamma(\beta_t)}\frac{\prod_{d=1}^{D}\Gamma(n_k^{(t)}+\beta_t)}{\sum_{t=1}^{V}\Gamma n_k+\beta_t} \quad \text{(A.5)}$$

where here $n_d^{(k)}$ denotes the number of word tokens in the d^{th} document assigned to the k^{th} topic, and $n_k^{(t)}$ is the number of times the t^{th} term in the vocabulary has been assigned to the k^{th}th topic.

Since in equation (A.1) $P(\mathbf{W};\alpha,\beta)$ is invariant to Z, the conditional distribution $P(\mathbf{Z}|\mathbf{W};\alpha,\beta)$ can be derived from $P(\mathbf{W},\mathbf{Z};\alpha,\beta)$ directly using Gibbs simulation and the conditional probability:

$$P(Z_{(d,n)}|\mathbf{Z}_{\text{-}(d,n)},W;\alpha,\beta) = \frac{P(Z_{(d,n)},\mathbf{Z}_{\text{-}(d,n)},\mathbf{W};\alpha,\beta)}{P(\mathbf{Z}_{\text{-}(d,n)},\mathbf{W};\alpha,\beta)} \quad \text{(A.6)}$$

where $Z_{(d,n)}$ denotes the hidden variable of the n^{th} word token in the d^{th} document, and $Z_{\text{-}(d,n)}$ denotes all Zs but $Z_{(m,n)}$. Denoting the index of a word token by $i = (m,n)$, and using the expressions in (A.4) and (A.5), cancellation of terms (and some extra manipulations exploiting the properties of the gamma function) yields:

$$P(Z_i = k|\mathbf{Z}_{(i)}, \mathbf{W}; \alpha, \beta) \propto (n_{d,\_i}^{(k)} + \alpha_k) \frac{n_{k,\_i} + \beta_t}{\sum_{t=1}^{V} n_{k,\_i}^{(t)} + \beta_t} \tag{A.7}$$

where the counts $n_{\cdot,\_i}^{(\cdot)}$ indicate that token i is excluded from the corresponding document or topic. Therefore, sampling topic indexes using equation (A.7) for each word in a document and across documents until convergence allows us to approximate the posterior distribution given by (A.1).

The model can be estimated as described once the researcher sets three parameters: the number of topics K and the two parameter vectors of the Dirichlet priors α and β. While I already discussed in the paper how to optimally estimate K, α and β are defined as a function of the number of topics and unique words (Griffiths and Steyvers, 2004) as $\alpha = \frac{50}{K}$ and $\beta = \frac{200}{N}$.

A.2. Alternative Tone's Measures

Figure A.1 compares the monetary policy stance index derived in Section 2.1 with alternative measures available in the literature.

Figure A.1: Alternative Tone's Measures

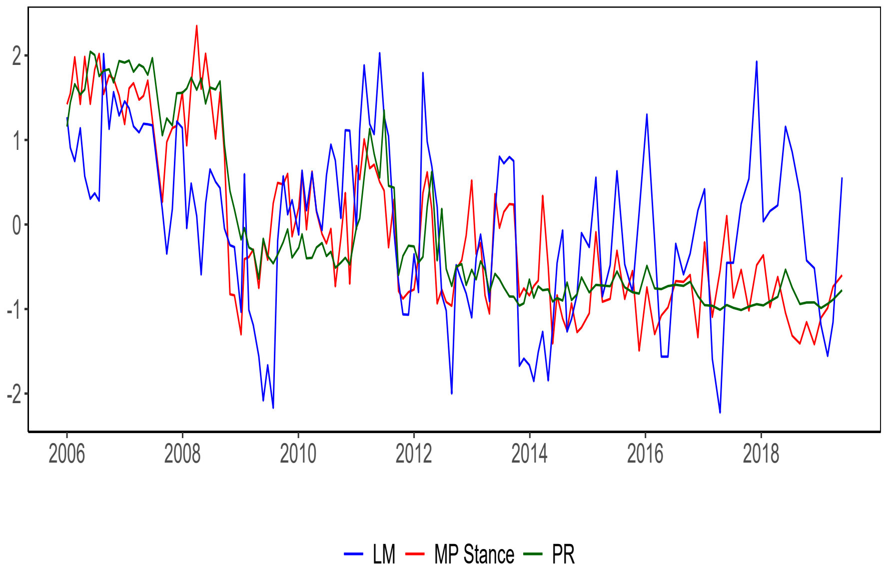

Note: The figure shows the evolution of the monetary policy stance (red) index against alternatives specifications available in the literature.

The figure shows a strong correlation (0.90) with the more sophisticated index developed by Picault and Renault (2017), while it departs from the index based on Loughran and McDonald

(2011)'s dictionary with a correlation of 0.42.

A.3. The Textual Dataset

In Figure A.2, I provide some descriptive statistics of the textual dataset collected in December 2020 from *Precise*.

Figure A.2: Descriptive Statistics of Textual Data

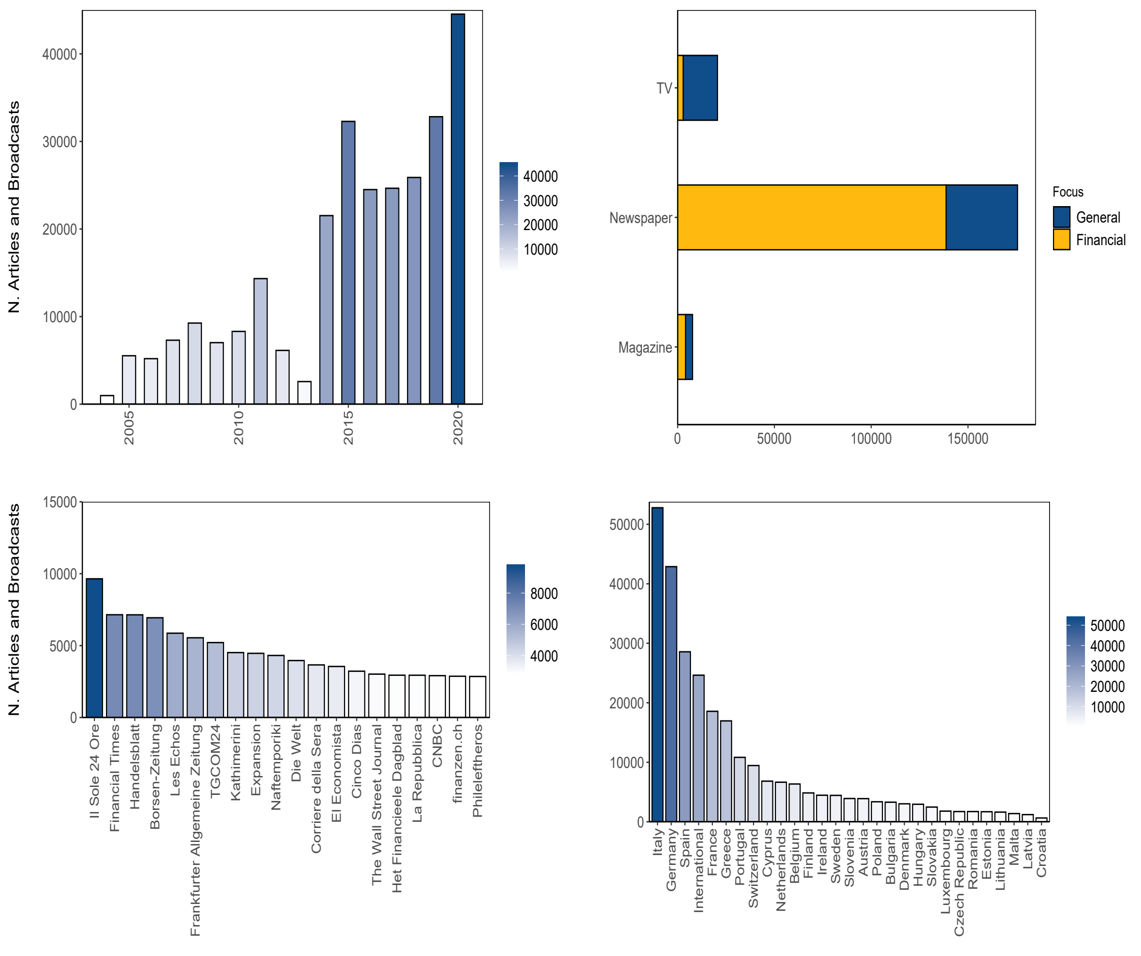

Note: The figure shows descriptive statistics for the textual dataset. Starting with the first row, the first column presents the number of documents per year, while the second column the type of document and the field of specialization. In the second row, the first column details the number of documents by outlet, while the second column the number of documents by country of origin.

It includes around 300,000 documents whose breakdown by year is not constant as shown by the figure in the the first column of the first row. In fact, from 2004 to 2013, the average number of documents per year stands around 5,000, while from 2014 on, the average is tilted significantly upwards to approximately 33,000 documents per year. The figure in the second column of the first row shows the composition of the documents where the large majority of them is composed of specialized newspaper articles. As for the second row, the first column breaks the number of articles into outlets' names and the second one into the country of origin of the outlet. To control for any bias, in the analysis I filtered out the time series regressing them on countries' dummies.

A.4. Quantifying Precise's Time Series

As mentioned in the paper, I applied LDA to the Precise's dataset with $K = 80$. Out of eighty topics, only sixty were clearly identifiable and I list them in the first column of the table below. These identified topics were further scaled down to forty following the criterion of monetary policy relevance (second column). Finally, these topics were manually grouped into seven macro topics (third column) on the basis of the conceptual similarity between them. Therefore, the table below provides a summary of how this exercise of dimensionality reduction was carried out.

Identified Topics	**Relevant Topics**	**Macro Topics**
Equity developments		
French politics		
European Commission President		
Eastern EU politics		
EU crisis	X	Financial crisis
Bank of England policy		
Core countries growth	X	Economic outlook
Banking crisis	X	European banks
Macro uncertainty	X	Economic outlook
Exchange rates	X	Economic outlook
Lending conditions	X	Monetary Policy
Stress tests	X	European banks
Eurexit	X	Eurexit
NPLs	X	European banks
Italian Politics		
Greek bonds	X	Financial crisis
Greek bailout	X	Financial crisis
Banknotes		
Rating Agencies	X	Financial crisis
Mario Draghi		
ECB Board Members		
Italian Banks	X	European banks

Greek Banks	X	European banks
Spanish Banks	X	European banks
German Banks	X	European banks
Digital Currency		
Lagarde		
FGCJ Ruling	X	Monetary Policy
IMF	X	Financial crisis
Banking supervision	X	European banks
Fixed-Income Developments		
Sovereign Bonds	X	Financial crisis
Financial stress	X	Financial crisis
Spreads	X	Financial crisis
Sovereign-debt crisis	X	Financial crisis
Fiscal rules	X	Fiscal policy
Fiscal stance	X	Fiscal policy
German politics		
Inflation dynamics	X	Inflation outlook
Deflation	X	Inflation outlook
TLTROs	X	Monetary policy
Negative rates	X	Monetary policy
Forward guidance	X	Monetary policy
Asset purchases	X	Monetary policy
Fed policies		
US politics		

Trade war	X	Economic outlook
US economy		
NGEU	X	Fiscal policy
Market expectations		
European growth	X	Economic outlook
Core inflation	X	Inflation outlook
Brexit	X	Eurexit
Bank profitability	X	European Banks
Liquidity support	X	Monetary policy
Bank strategies	X	European banks
Monetary statistics	X	Monetary Policy
Climate change		
Cryptocurrency		
Bank of Japan policies		

These seven topics were treated differently. In fact, while "financial crisis", "eurexit" and "European banks" were left as topic probabilities, "monetary policy", "economic outlook", "inflation outlook" and "fiscal policy" needed a further quantification step. In quantifying these topics I followed the same steps as Section 2.1. However, in Section 2.1 the field-specific dictionaries were based on the ECB press conferences that have much narrower number of words compared to the *Precise*'s dataset. To inspect such a large dataset and look for polarized words relevant to the already existing field-specific dictionaries, I relied on word embedding (Stoltz and Taylor, 2019) which is a tool for identifying similarities between words that occur within a "context window". The algorithm is based on the Word Mover's Distance (WMD) that attempts to find the "closest neighbour" for each word so that the "cost" of moving all the words in one collection (sentences, subsection, etc.) to the positions of all the words in another collection is minimized. Thus, collections sharing many semantically similar words should have smaller distances than collections with very different words. Formally, the WMD algorithm finds the values of a matrix $\mathbf{T}$ that minimize "moving" one document, D, to the other, D':

$$WMD_{ij} = \min_{\mathbf{T}\geq 0} \sum_{i,j=1}^{n} T_{ij} c(i,j) \tag{A.8}$$

where $c(i,j)$ indicates the cosine distance between the i^{th} and j^{th} word in the n-dimensional embedding space and under the constraints such that:

$$\sum_{i=1}^{n} T_{ij} = d_i \ , \forall_i \in \{1,\ldots,n\} \tag{A.9}$$

$$\sum_{j=1}^{n} T_{ij} = d'_j \ , \forall_i \in \{1,\ldots,n\} \tag{A.10}$$

where equation (A.9) says that the sum of row word i in $\mathbf{T}$ is equal to the relative frequency of i in document D (after any word removal), d_i, and equation (A.10) that the sum of column word j in $\mathbf{T}$ is equal to the relative frequency of j in D' (again after word removal), d'_j. An example can clarify: if the word "fiscal" had a relative frequency of 0.25 in a document after any word removal, then the sum of the row and column with the word "fiscal" must each sum to 0.25. The intuition is to weight each ij cosine distance by "how much" of the relative frequency of i in D will move to j in D'.

To measure "how close" certain words are to a focal concept, however, an additional step, called Relaxed Word Mover's Distance (RWMD), is required. With RWMD, the flow matrix weighting for each i, j pair is solved twice: once with just the constraint from equation (A.9) removed, and then once with just the constraint from equation (A.10) removed. The RWMD for each i, j pair is then calculated twice following equation (A.8), and the final reported RWMD score for the i, j pair is:

$$RWMD_{ij} = \max(\min_{\mathbf{T} \geq 0} \sum_{i,j=1}^{n} T_{ij} c(i,j), \min_{\mathbf{T} \geq 0} \sum_{i,j=1}^{n} T'_{ij} c(i,j)) \tag{A.11}$$

In the paper, for example, I applied equation (A.11) to find the closest words to the focal concepts of "inflation", "economy", "monetary", and "fiscal", where the focal concepts simply correspond to the topics of the field-specific dictionaries. What this equation provides is a list of words that are closely related to each focal point. Once these lists are derived, the last steps consist in finding which words are not present in the dictionaries derived in Section 2.1. These words are then manually classified as "positive" or "negative" using discretion with regard to the topic of interest.

B. Total Dataset

The table below lists all the variables used in the model with its respective transformation, publication lag, block and frequency. In the transformation column: 1 corresponds to the variable being left in level, 2 to first difference, 3 to log-difference and 4 to percentage change; in the frequency column, D stands for daily, M for monthly and I for irregular; finally, in the column that indicates the publication lag, the numbers denote months.

Series	**Transformation**	**Publication Lag**	**Frequency**
Real Activity			
Industrial Production Total	3	1	M
Industrial Production (construction)	3	1	M
Industrial Production: manufacturing	3	1	M
Unemployment (labour force)	4	1	M
Harmonised Index of Consumer Prices (HICP)	3	1	M
HICP (HICP: services	3	1	M
HICP: excl. energy and food	3	1	M
HPPI: Industry excluding construction	3	1	M
Oil Prices (Brent)	3	1	M
Consumer Price Index (CPI)	3	1	M
CONSUMER PRICES OF PETROLEUM PRODUCTS (PUMP PRICES)	3	1	M
HWWI COMMODITY PRICES FOR THE EURO AREA: Total Excluding Energy	3	1	M
Surveys			
EUROCOIN	1	0	M
Composite Output EA (PMI)	1	0	M
Manufacturing (PMI)	1	0	M
Services (PMI)	1	0	M
Construction (PMI)	1	0	M
Capacity Utilization (PMI)	1	0	M
Euro Area Big 2 (PMI)	1	0	M
Employment expectations	1	0	M
Unemployment expectations	1	0	M
Italy CLIFS	1	0	M
Spain CLIFS	1	0	M
Portugal CLIFS	1	0	M
Business Climate Indicator	1	0	M
Economic sentiment indicator	1	0	M

Series	Transformation	Publication Lag	Frequency
Financial			
DJ U.S. Total Stock Market Index	3	0	D
SP 500 Index	3	0	D
NASDAQ Composite	3	0	D
STOXX Euro	3	0	D
EUSTOXX Banks	3	0	D
VIXEU	3	0	D
CDS Italy	3	0	D
CDS Spain	3	0	D
CDS Portugal	3	0	D
CDS France	3	0	D
CDS Germany	3	0	D
BARCLAYSMBSFIXXRATE	2	0	D
FTSEMIB	2	0	D
IBEX 30	2	0	D
CAC 40	2	0	D
DAX 30	2	0	D
Rates and Spreads			
10Y GDP-Weighted Nominal Yields	2	0	D
5Y5YINFSWAP rates	2	0	D
2Y2YINFSWAP rates	2	0	D
Goldman Sachs Financial Condition Index	2	0	D
EURINFLSWAP0COUP	2	0	D
EONIA	2	0	D
EURIBOR 3M	2	0	D
LIBOR-OIS	2	0	D
OISEONIA	2	0	D
Repo funds DE	2	0	D

Series	**Transformation**	**Publication Lag**	**Frequency**
Repo funds IT	2	0	D
Repo funds FR	2	0	D
EUFCI	2	0	D
EUDESPREAD	2	0	D
EUSWAP10Y	2	0	D
IT10Y	2	0	D
DE10Y	2	0	D
1W SWAP	2	0	D
1Y1Y EONIA Forward	2	0	D
EU10Y	2	0	D
FR10Y	2	0	D
DEIT10Y	2	0	D
1Y1YSPRD	2	0	D
3MEOIS	2	0	D
3MGLBR	2	0	D
3MGLOIS	2	0	D
DE2Y10Y	2	0	D
DEIT2Y	2	0	D
DEIT5Y	2	0	D
DEGR10Y	2	0	D
DESP10Y	2	0	D
DEPT10Y	2	0	D
EUREONIA	2	0	D
EUROIS3M	2	0	D
FR2Y10Y	2	0	D
GBPOIS1M	2	0	D
DE2Y10Y	2	0	D
ES2Y10Y	2	0	D

Series	**Transformation**	**Publication Lag**	**Frequency**
FR2Y10Y	2	0	D
PT2Y10Y	2	0	D
IT2Y10Y	2	0	D
Forecasts			
Bloomberg Weighted Average Private GDP Forecast	4	0	D
Bloomberg Weighted Average Official GDP Forecast	4	0	D
Bloomberg Weighted Average Private IP Forecast	4	0	D
Bloomberg Weighted Average Official IP Forecast	4	0	D
Bloomberg Weighted Average Private CPI Forecast	4	0	D
Bloomberg Weighted Average Private Unemployment Forecast	4	0	D
Textual Newspaper			
ECB Monetary Policy Stance (Newspaper)	1	0	D
Economic Outlook Index (Newspaper)	1	0	D
Inflation Outlook Index (Newspaper)	1	0	D
Eurexit Index	1	0	D
Financial Crisis Index	1	0	D
European Banks Index	1	0	D
Fiscal Policy	1	0	M
Textual PC			
Monetary Policy Stance Index	1	0	I
Economic Outlook Index	1	0	I
Inflation Outlook Index	1	0	I
Monetary			
M3	2	1	M
MFI loans to euro area residents excluding MFIs and general government	2	1	M
Securities other than shares issued in euro by non-MFI corporations	2	1	M
Total assets	2	1	M
Lending to euro area credit institutions in euro	2	1	M

Series	Transformation	Publication Lag	Frequency
Main refinancing operation	2	1	M
Longer-term refinancing operations	2	1	M
Securities of euro area residents in euro	2	1	M
Deposit facility rate	2	1	M
Main refinancing rate	2	1	M
Marginal lending facility rate	2	1	M
Loans to insurance corporations and pension funds	2	1	M
Loans to other financial intermediaries (including investment funds)	2	1	M
Loans to nonfinancial corporations	2	1	M
Households	2	1	M
Lending for house purchase	2	1	M
Interest Rates New Deposits from households (overnight)	2	1	M
New deposits: Deposits from nonfinancial corporations	2	1	M
New deposits: repos	2	1	M
Mixed			
ZEW FINANCIAL MARKET SURVEY: current macro	4	0	M
ZEW FINANCIAL MARKET SURVEY: expected macro	4	0	M
Sentix overall economic index	4	0	M
Sentix break-up economic index	4	0	M
Euro Area Composite Financial Conditions Index	2	0	M
Inflation factor	3	1	M
Uncertainty factor	3	1	M
Sovereign CISS (weighted)	2	1	M
Exchange Rates			
Japanese yen	3	0	M
UK pound	3	0	M
US dollar	3	0	M
Harmonized indicators against 38 trading partners	3	0	M

Series	Transformation	Publication Lag	Frequency
Nominal effective exchange rate of the euro: EER-19	3	0	M
Real (CPI) effective exchange rate of the euro: EER-19	3	0	M
US			
JP Morgan Global PMI	3	0	M
PMI North America	3	0	M
Inflation Expectations (1Y), Michigan University	3	1	M
Inflation Expectations (5Y), Michigan University	3	1	M
Federal funds (effective)	2	0	D
Treasury yields 1Y	2	0	D
Treasury yields 10Y	2	0	D

C. State Space Representation and Estimation

Section C.1 explains the temporal aggregation implemented in the Dynamic Factor Model, Section C.2 provides details of the Expectation Maximization (EM) algorithm and Section C.3 sheds light on the bridge equation.

C.1. Temporal Aggregation

Letting bold-font variables indicate the vector version of the variables, I write the model in equations (3) and (4) compactly as follows:

$$\mathbf{Y_t} = \mathbf{\Lambda}_t \mathbf{F_t} + \mathbf{e_t} \tag{C.12}$$

$$\mathbf{F_t} = \mathbf{A_t} \mathbf{F_{t_1}} + \mathbf{u_t} \tag{C.13}$$

with $\mathbf{F_t} = [\mathbf{F_t}^{k,i} \mathbf{F_t}^{k,m} F_t^d]'$. The last element in the vector $\mathbf{F_t}$ is the scalar F_t^d and can be interpreted as the latent common daily monetary policy stance index. F_t^d is the only scalar in $\mathbf{F_t}$ since the other elements are vectors containing aggregator variables used to handle the mixed-frequency property of the model. In mixed-frequency models, lower frequency variables are usually treated as the highest frequency series, i.e. daily in this case, with missing observations (see Foroni and Marcellino (2013) for an overview). For a generic variable Y_t^k, time aggregation from higher to lower frequency is restricted as follows:

$$\begin{aligned} y_t^k &= \log(v_{1,t}^k) - \log(v_{1,t-k}^k) \\ &\approx \log(\sum_{i=0}^{k-1} v_{1,t_i}^k) - \log(\sum_{i=k}^{2k-1} v_{1,t-i}^k) \\ &\approx \sum_{i=0}^{k-1} \log(v_{1,t_i}^k) - \sum_{i=0}^{2k-1} \log(v_{1,t-i}^k) \\ &= \sum_{i=0}^{2k-2} \omega_i^k y_{1,t-i} \; , t = k, 2k, \dots , \end{aligned} \tag{C.14}$$

where y_t^k is the observed low-frequency growth rate, v_t^k its level, and $\omega_i^k = i+1$ for $i = 1, \dots, k-1$; $\omega_i^k = 2k-i-1$ for $i = k, \dots, 2k-2$; and $\omega_i^k = 0$ otherwise. It follows from equation (C.14) that imposing a common factor structure to y_t^k gives:

$$y_t^k = \sum_{i=0}^{2k-2} \omega_i^k y_{1,t-i} = \sum_{i=0}^{2k-2} \omega_i^k (\Lambda F_{t-i}^d + e_{t-i}) \tag{C.15}$$

However, equation (C.15) makes the inference rather challenging since it significantly increases the number of state variables in the model. To sort this issue and limit the size of the state vector, I employ the double cumulator approach in Bańbura et al. (2013) where the temporal aggregator variables are recursively updated such that at the end of each respective period we have:

$$F_t^k = \sum_{i=0}^{2k-2} \omega_i^k F_{t-i} \quad t = k, 2k, \ldots \tag{C.16}$$

As shown in Bańbura et al. (2013), this expression can be computed recursively with the help of two (additional) state variables. In particular, by introducing the auxiliary variable $\bar{F}_t^k$ and denoting $R(\cdot, k)$ the positive remainder of the division by k, $\tilde{F}_t^k$ can be obtained recursively as follows:

$$\tilde{F}_t^{k,f} = \begin{pmatrix} \bar{F}_t^{k,f} \\ F_t^{k,f} \end{pmatrix} = \begin{cases} \begin{pmatrix} \bar{F}_t^{k,f} + \omega_{k-1}^{k,f} F_t \\ 0 \end{pmatrix}, & t = 1 + k + 1, 2k + 1, \ldots \\ \begin{pmatrix} F_{t-1}^{k,f} + \omega_{R(k-t,k)}^{k,f} F_t \\ 0 \end{pmatrix}, & \text{otherwise.} \end{cases} \tag{C.17}$$

For the stock variables only single aggregator variable is necessary:

$$F_t^{k,s} = \begin{cases} \omega_{k-1}^{k,s} F_t \ , & t = 1 + k + 1, 2k + 1 \\ F_{t-1}^{k,s} + \omega_{k-1}^{k,s} F_t \ , & \text{otherwise.} \end{cases} \tag{C.18}$$

This is implemented via the transition equation (4) with the following weight vector $\mathcal{W}_t^k$ and selection matrix $\mathcal{I}_t^k$:

$$\mathcal{W}_t^{k,f} = \begin{cases} \begin{pmatrix} -\omega_{k-1}^{k,f} \\ 0 \end{pmatrix}, & t = 1+k+1, 2k+1, \ldots \\ \begin{pmatrix} -\omega_{R(k-t,k)}^{k,f} \\ -\omega_{R(k-t,k)+k}^{k,f} \end{pmatrix}, & \text{otherwise.} \end{cases} \tag{C.19}$$

$$\mathcal{I}_t^{k,f} = \begin{cases} \begin{pmatrix} 0 & I_r \\ 0 & 0 \end{pmatrix}, & t = 1+k+1, 2k+1, \ldots \\ I_{2r}, \text{otherwise}, \end{cases} \tag{C.20}$$

and for the stock variables simply: $\mathcal{W}_t^{k,s} = -\omega_{R(k-t,k)}^{k,s}$ and $\mathcal{I}_t^{k,s} = \begin{cases} 0, & t = 1, k+1, \ldots \\ I_r, & \text{otherwise.} \end{cases}$.

In general, equations (C.17)-(C.20) can handle temporal aggregation from higher to lower frequencies for a range of k values. In the model specification outlined in Section 3, both $k = k_i, k_m, k_d$ are considered, where the k refers to the (average) number of days in the irregular, monthly and daily frequencies, respectively.

To handle different number of days per month or irregular, I follow Bańbura et al. (2013) in approximating the flow variables as follows:

$$z_t^k = \frac{k}{k_t} \sum_{i=0}^{k_t-1} z_{t-i} \, , t = k_1, k_1 + k_{k_1+1}, \ldots, \tag{C.21}$$

where k_t is the number of business days in the period (month or irregular) that contains day t and k is the average number of business days per period over the sample. As a result, $\gamma_t^k = z_t^k - z_{t-k}^k$ becomes:

$$\gamma_t^k = k \left(\sum_{i=0}^{k_t-1} \frac{i+1}{k_t} \gamma_{t-i} + \sum_{i=k_t}^{k_t+k_{t-k_t}-2} \frac{k_t+k_{t-k}-i-1}{k_{t-k}} \gamma_{t-i} \right) , t = k_1, k_1 + k_{k_1+1}, \ldots, \tag{C.22}$$

Hence, this results in time-varying weights and the formulas above should be updated with: $\omega_{t,i}^{k,f} = k\frac{i+1}{k_t}$ for $i = 0, 1, \ldots, k_{t-1}$, $\omega_{t,i}^{k,f} = k\frac{k_t+k_{t-k_t}-i-1}{k_{t-k_t}}$ for $i = k_t, k_{t+1}, \ldots, k_t + k_{t-k_t} - 2$ and

$\omega_{t,i}^{k,f} = 0$ otherwise.

C.2. The Expectation Maximization Algorithm

To describe the Expectation Maximization (EM) algorithm for the daily model in equations (3) and (4), I follow once again Bańbura and Modugno (2014). Assuming one lag in the factor VAR, I can group the parameters as follows: $\theta = (\mu, \Lambda, \Phi, \sum_E, \sum_U)$, where the only restriction is that $\sum_E$ is diagonal. Let $T_v = \max_n T_v(n)$ also denote the time index of the most recent observation in Ω_v. It is then possible to write the log-likelihood in terms of $Y_t^{k_z,n} = (Y_1^{k_z,n}, Y_2^{k_z,n}, \ldots, Y_{Tv}^{k_z,n})$ and $F_t^{k_z,n} = (F_1^{k_z,n}, F_2^{k_z,n}, \ldots, F_{Tv}^{k_z,n})$ as $l(Y_t^{k_z,n}, F_t^{k_z,n}; \theta)$. After that an initial estimate of the parameters $\theta(0)$ is computed on a sample of data where missing observations are handled using splines, the EM algorithm proceeds in two steps as follows:

$$\begin{aligned} E - step &: L(\theta, \theta(j)) = E_\theta[l(Y_t^{k_z,n}, F_t^{k_z,n}; \theta)|\Omega_v], \\ M - step &: \theta_{j+1} = \underset{\theta}{\operatorname{argmax}}\, L(\theta, \theta(j)). \end{aligned} \tag{C.23}$$

The new parameter estimates in the M-step can be obtained in two steps, first $\Lambda(j+1)$ and $\Phi(j+1)$ are given by:

$$\begin{aligned} vec(\Lambda^{z,\cdot}(j+1)) &= \left(\sum_{t=1}^{Tv} E_{\theta(j)}[Y_t^{k_z,n} F_t^{k_z,n'}|\Omega_v] \otimes \mathcal{S}_t\right) \left(\mathcal{S}_t \sum_{t=1}^{Tv} E_{\theta(j)}[Y_t^{k_z,n} F_t^{k_z,n'}|\Omega_v]\right)^{-1}, \\ \Phi(j+1) &= \left(\sum_{t=1}^{Tv} E_{\theta(j)}[F_t^{k_z,n} F_{t-1}^{k_z,n'}|\Omega_v]\right) \left(\sum_{t=1}^{Tv} E_{\theta(j)}[F_{t-1}^{k_z,n} F_{t-1}^{k_z,n'}|\Omega_v]\right)^{-1}. \end{aligned} \tag{C.24}$$

where $\mathcal{S}_t$ is a selection matrix, that is, a diagonal matrix that takes the value of one for non-missing observations in $Y_t^{k_z,n}$ and zero otherwise. $\mathcal{S}_t$ allows the estimation to work with arbitrary pattern of missing data. $\Lambda^{z,\cdot}$ are estimated blockwise, by frequency and by stock or flow type, using the corresponding block of $Y_t^{k_z,n}$ and aggregator variable.

Second, given the new estimates of $\Lambda^{z,\cdot}$ and Φ, the covariance matrices can be obtained as follows:

$$\sum_E(j+1) = diag(\sum_{t=1}^{Tv}(\mathcal{S}_t Y_t^{k_z,n} Y_t^{k_z,n'} \mathcal{S}_t' - \mathcal{S}_t Y_t^{k_z,n} E_{\theta(j)}[F_t^{k_z,n'}|\Omega_v]\Lambda^{z,\cdot}(j+1)'\mathcal{S}_t$$
$$- \mathcal{S}_t \Lambda^{z,\cdot}(j+1)E_{\theta(j)}[F_t^{k_z,n}|\Omega_v]Y_t^{k_z,n'}\mathcal{S}_t + \mathcal{S}_t\Lambda^{z,\cdot}(j+1)E_{\theta(j)}[F_t^{k_z,n}F_t^{k_z,n'}|\Omega_v]\Lambda^{z,\cdot}(j+1)'\mathcal{S}_t$$
$$+ (I_N - \mathcal{S}_t)\sum_{E(j)}(I_N - \mathcal{S}_t))), \tag{C.25}$$

and

$$\sum\nolimits_U(j+1) = \frac{1}{T}\left(\sum_{t=1}^{Tv} E_{\theta(j)}[F_t^{k_z,n}F_t^{k_z,n'}|\Omega_v] - \Phi(j+1)\sum_{t=1}^{Tv} E_{\theta(j)}[F_t^{k_z,n}F_t^{k_z,n'}|\Omega_v]\right) \tag{C.26}$$

where the expectations $E_{\theta(j)}(\cdot)$ are obtained via the Kalman filter and smoother, and the estimates for Φ and $\sum_U$ follow from taking the elements from the conditional covariances of the state vector corresponding to $F_t^{k_z,n}$.

C.3. The Bridge Equation

The bridge equation is introduced to equip the model with the possibility to forecast the probability that the ECB will take a monetary policy decision given some incoming information. It owes its definition to its function, that is, it "bridges" the forecast variables from the DFM with the dependent variable. In this case, the specification of the bridge equation follows a multinomial logit. To create the categorical variable of monetary policy decisions, I then followed some basic principles. First, a tightening decision to be classified as such must contain one (or a combination) of the following announcements: an increase in interest rates, a reduction in asset purchases and a hawkish revision of the forward guidance; second, an easing decision to be classified as such must instead include one (or a combination) of the following announcements: an interest rate cut, the launch of an asset purchase programme or a type of long-term refinancing operations (LTROs), an increase in asset purchases, a dovish revision of the forward guidance and an extension of the collaterals eligible for repos; last, monetary policy remains constant if no decision in the first or second group is announced. As a result, from January 2002 to December 2020 tightening decisions total 13, while easing one amount to 33. Figure A.2 displays the tightening monetary policy decisions (shaded red) and the easing ones (shaded blue).

Figure C.1: Out-of-Sample Application

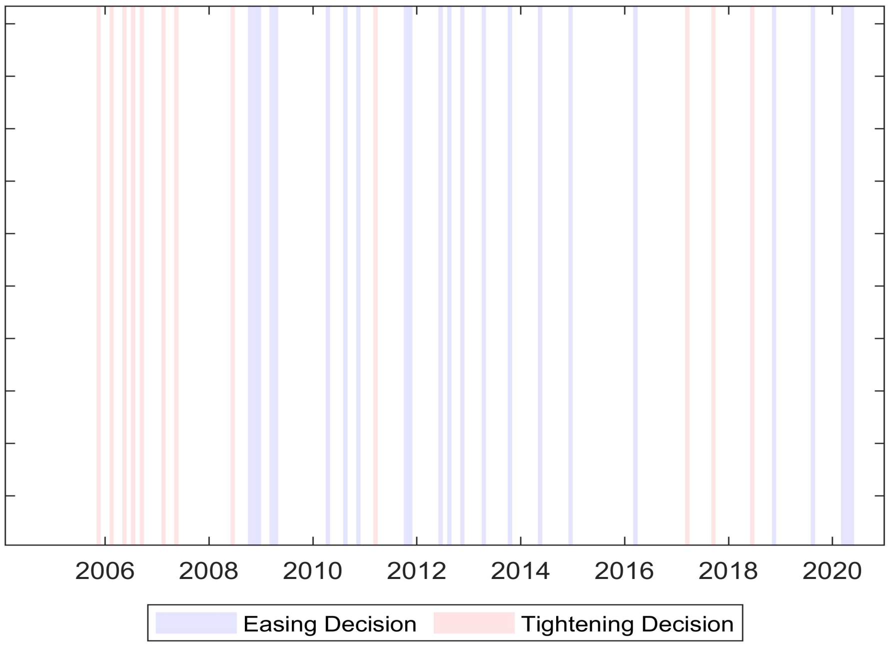

Note: The figure shows actual tightening (shaded red) versus easing decisions (shaded blue).

The performance of the bridge equation is assessed by two statistics: the mean absolute forecast squared error (MAFE) and the pseudo R^2 statistics. The MAFE is simply given by the mean absolute difference between the model-implied probabilities and the actual outcome of the categorical monetary policy variable. Therefore, once the probabilities have been estimated is straightforward to compute:

$$MAE = \frac{1}{T} \sum_{t=1}^{T} (|y_t - P(y_t = j)|) \tag{C.27}$$

On the other hand, pseudo R^2 is a measure of fit defined as the weighted sum of fractions of correctly identified easing, tightening and constant decisions. Formally, it can be written:

$$R^2 = \frac{n_{tt}}{n_t} \cdot \frac{n_t}{n} + \frac{n_{ee}}{n_e} \cdot \frac{n_e}{n} + \frac{n_{cc}}{n_c} \cdot \frac{n_c}{n} = \frac{n_{tt}+n_{ee}+n_{cc}}{n} \tag{C.28}$$

where n is the number of press conferences for which forecasts have been computed, n_t, n_e and n_c are, respectively, the number of tightening, easing and constant decisions and n_{tt}, n_{ee} and n_{cc} are the number of correctly identified tightening, easing and constant decisions. It is important to note that the pseudo R^2 relies on a conversion of the model-implied recession probabilities into a categorical variable. That is, given an estimated probability of recession

$\hat{P}(y_t = j)$ for some time t, a mapping $D : [0,1] \rightarrow \{0,1\}$, $\hat{y}_t = D[P(y_t = j)]$, needs to be implemented to convert probabilities into alleged monetary policy decisions. Usually, such a decision rule depends on a threshold c^* such that $\hat{y}_t = j$, if $\hat{P}(y_t = j) > c^*$. Hence, the choice of the threshold c^* is key. Setting $c^* = 0.5$ is often too conservative a criterion, especially when the overall fraction of tightening and easing decisions in the sample is relatively small (see Green, 2012 on discrete choice models). A more natural choice is to set c^* equal to the frequency of easing and tightening decisions measured over a long period. For this reason, I set $c^* = 0.15$ for easing decisions and $c^* = 0.08$ for tightening ones. While in the out-of-sample exercises, this threshold could be adjusted period by period, I treat it as constant so that results are independent of such time variation.

D. Validating the DFM

This section attempts to validate the DFM by nowcasting the EONIA, the EU shadow rate (Wu and Xia, 2016) and a textual measure of monetary policy stance derived using Loughran and McDonald (2011)'s dictionary. Such a validation exercise aims to test whether the DFM performs well on time series related to the ECB's monetary policy stance that are different from the one I developed in the paper.

For every series, I display the out-of-sample nowcasting of January, April and August 2005. Figure D.1 shows the results for, respectively, EONIA, EU shadow rates and the alternative textual measure of ECB monetary policy stance[25]. Figure 9 should also be considered part of this exercise.

[25] The textual measure should not be taken as representative of the ECB monetary policy stance since it suffers from the flows outlined in Section 2.2

Figure D.1: Validating the DFM with Out-of-sample Nowcasts

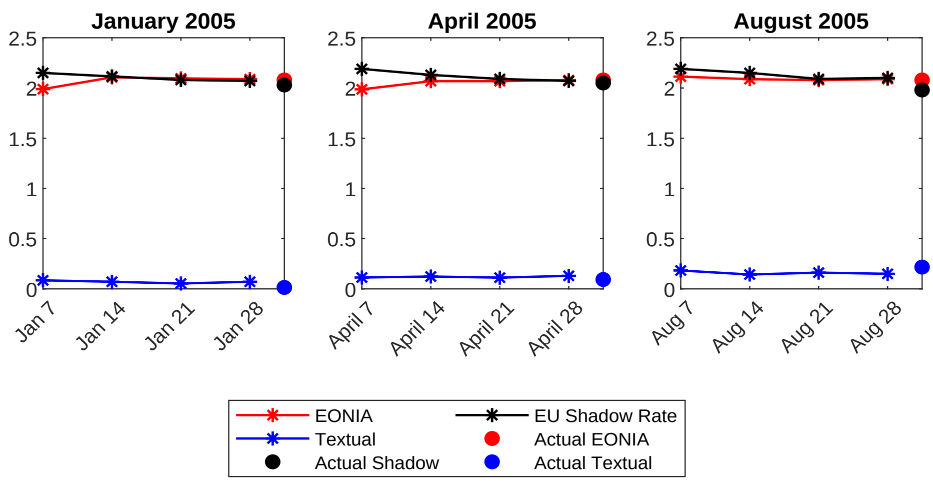

Note: The figure shows the out-of-sample nowcasts of EONIA (red line), EU Shadow rate (black line) and an alternative textual measure of ECB monetary policy stance (blue line) for January, April and August 2005.

The validation exercise, along with the results shown in Figure 9, sheds light on the validity of the DFM beyond the application to the textual measure of ECB monetary policy stance derived in this paper. The DFM, in fact, tracks accurately the actual realizations for the variables of interest.

References

Apel, M. and M. Blix Grimaldi (2012, April). The Information Content of Central Bank Minutes. Working Paper Series 261, Sveriges Riksbank (Central Bank of Sweden).

Bańbura, M., D. Giannone, M. Modugno, and L. Reichlin (2013). Now-casting and the real-time data flow. In C. G. G. Elliott and A. Timmermann (Eds.), *Handbook of Economic Forecasting*, Volume 2, Chapter 0, pp. 195–237. Elsevier.

Bańbura, M. and M. Modugno (2014). Maximum Likelihood Estimation of Factor Models on Datasets with Arbitrary Pattern of Missing Data. *Journal of Applied Econometrics 21*(1), 1–21.

Baranowski, P., H. Bennani, and W. Doryń (2021). Do the ECB's introductory statements help predict monetary policy? Evidence from a tone analysis. *European Journal of Political Economy 66*, 101964.

Bennani, H., N. Fanta, P. Gertler, and R. Horvath (2020). Does central bank communication signal future monetary policy in a (post)-crisis era? The case of the ECB. *Journal of International Money and Finance 104*(19), 102167.

Bennani, H. and M. Neuenkirch (2017). The (home) bias of European central bankers: new evidence based on speeches. *Applied Economics 49*(11), 1114–1131.

Blei, D. M., A. Y. Ng, and M. I. Jordan (2003). Latent Dirichlet Allocation. *Journal of Machine Learning Research 3*, 993–1022.

Bok, B., D. Caratelli, D. Giannone, A. M. Sbordone, and A. Tambalotti (2018). Macroeconomic Nowcasting and Forecasting with Big Data. *Annual Review of Economics 10*(1), 615–643.

Cao, J., T. Xia, J. Li, Y. Zhang, and S. Tang (2009). A density-based method for adaptive LDA model selection. *Neurocomputing 72*(7-9), 1775–1781.

Cimadomo, J., D. Giannone, M. Lenza, A. Sokol, and F. Monti (2020, August). Nowcasting with large Bayesian vector autoregressions. Working Paper Series 2453, European Central Bank.

Dempster, A. P., N. M. Laird, and D. B. Rubin (1977). Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society. Series B (Methodological) 39*(1), 1–38.

Foroni, C. and M. Marcellino (2013, February). A survey of econometric methods for mixed-frequency data. Working Paper 2013/06, Norges Bank.

Gerlach, S. (2007). Interest Rate Setting by the ECB, 1999-2006: Words and Deeds. *International Journal of Central Banking 3*(3), 1–46.

Giannone, D., L. Reichlin, and D. Small (2008). Nowcasting: The real-time informational content of macroeconomic data. *Journal of Monetary Economics 55*(4), 665–676.

Gorter, J., J. Jacobs, and J. De Haan (2008). Taylor rules for the ECB using expectations data. *Scandinavian Journal of Economics 110*(3), 473–488.

Green, W. H. (2012). *Econometric Analysis* (7th Edition ed.). Upper Saddle River.

Griffiths, T. L. and M. Steyvers (2004). Finding Scientific Topics. *Pnas 101*(SUPPL.1), 5228–5235.

Heinrich, G. (2009). Parameter estimation for text analysis (based on Gibbs sampling). *Technical Report, Fraunhofer IGD*.

Jansen, D.-J. and J. De Haan (2009). Has ECB communication been helpful in predicting interest rate decisions? An evaluation of the early years of the Economic and Monetary Union. *Applied Economics 41*(16), 1995–2003.

Loughran, T. and B. McDonald (2011). When is a liability not a liability? Textual Analysis, Dictionaries, and 10-Ks. *Journal of Finance 66*(1), 35–65.

Mariano, R. S. and Y. Murasawa (2003). A new coincident index of business cycles based on monthly and quarterly series. *Journal of Applied Econometrics 18*(4), 427–443.

Marozzi, A. (2021). The ECB and the Cost of Independence. Unearthing a New Doom-Loop in the European Monetary Union. BAFFI CAREFIN Working Papers 21152, BAFFI CAREFIN, Centre for Applied Research on International Markets Banking Finance and Regulation, Universita' Bocconi, Milano, Italy.

Modugno, M. (2013). Now-casting inflation using high frequency data. *International Journal of Forecasting 29*(4), 664–675.

Picault, M. and T. Renault (2017). Words are not all created equal: A new measure of ECB communication. *Journal of International Money and Finance 79*, 136–156.

Rinker, T. W. (2019). *sentimentr: Calculate Text Polarity Sentiment.* Buffalo, New York. version 2.8.0.

Stoltz, D. S. and M. A. Taylor (2019). Concept Mover's Distance: measuring concept engagement via word embeddings in texts. *Journal of Computational Social Science 2*(2), 293–313.

Sturm, J. E. and J. de Haan (2011). Does central bank communication really lead to better forecasts of policy decisions? New evidence based on a Taylor rule model for the ECB. *Review of World Economics 147*(1), 41–58.

Taylor, J. B. (1993). Discretion versus policy rules in practice: two critical points. A comment. *Carnegie-Rochester Confer. Series on Public Policy 39*(C), 215–220.

Thorsrud, L. A. (2020). Words are the New Numbers: A Newsy Coincident Index of the Business Cycle. *Journal of Business and Economic Statistics 38*(2), 393–409.

Wieland, V. and M. Wolters (2013). Forecasting and policy making. In C. G. G. Elliott and A. Timmermann (Eds.), *Handbook of Economic Forecasting*, Volume 2, Chapter 0, pp. 239–325. Elsevier.

Wu, J. C. and F. D. Xia (2016). Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound. *Journal of Money, Credit and Banking 48*(2-3), 253–291.