

# The Central Bank Crystal Ball: Temporal information in monetary policy communication\*

David Byrne<sup>†</sup> Robert Goodhead<sup>‡</sup>  
Michael McMahon<sup>§</sup> Conor Parle<sup>¶</sup>

First Draft: May 6, 2021

This Draft: July 11, 2021

## Abstract

Effective central bank communication provides information that the public does not have but wants to learn. Rather than arising from private information per se, we argue that the public's information deficit reflects three main channels: (i) an updated *evaluation* of the state of the economy based on past data, (ii) a different *projection* from the state to the likely future evolution of the economy, and (iii) differences in the reaction to the economic outlook. This paper provides a methodology to quantify the temporal information in the text of central bank communication. Both the backward-looking *evaluation* and forward-looking *projection* channels significantly improve the ability to explain the news in yields around ECB announcements compared with topic-based text analysis. This paper also shows that speeches generate greater news in yields the larger the information deficit remaining after the previous ECB Governing Council press conference.

**Keywords:** Monetary Policy, Communication, Machine Learning

**JEL Codes:** E52, E58, C55

---

\*McMahon gratefully acknowledges financial support from the European Research Council (Consolidator Grant Agreement 819131). We thank seminar participants at the University of Oxford for helpful comments. The views presented in the paper are those of the authors alone and do not represent the official views of the Central Bank of Ireland or the European System of Central Banks. Any errors remain ours alone.

<sup>†</sup>Central Bank of Ireland. Email: david.byrne@centralbank.ie

<sup>‡</sup>Central Bank of Ireland. Email: robert.goodhead@centralbank.ie

<sup>§</sup>University of Oxford, CEPR, CfM (LSE). Email: michael.mcmahon@economics.ox.ac.uk

<sup>¶</sup>Central Bank of Ireland and Trinity College Dublin. Email: Conor.Parle@centralbank.ie

# 1 Introduction

A burgeoning literature in empirical macroeconomics focuses on the impact monetary policymakers have on asset prices when they release statements to the press, or enact other forms of communication. By studying the responses of asset prices in the window around the communication event, studies such as Kuttner (2001), Gürkaynak et al. (2005) and Gertler and Karadi (2015) isolate the surprise component of central bank communication. From these studies, we learn that market participants update their beliefs in response to central bank communication. Such movements provide strong evidence of market participants having an “information deficit” relative to policymakers, in the sense that market participants are unable to fully anticipate what a central banker will say, prior to their statements. The asset price responses provide evidence that this information deficit is, at least partly, filled by central bank statements from which useful information for the formation of expectations is garnered.

While such empirical approaches have provided a wealth of information regarding the consequences of central bank communication, they have not isolated the specific communication transmission mechanism. They do not model or assess the exact verbal and numerical information contained within central bank statements and map this to the asset price responses.

This study aims to address this gap by studying new measures of both the communication transmission mechanism, and the information deficit. It is part of a complementary empirical approach that begins with the content of statements and tries to relate these to the asset price movements. By directly quantifying the extent of any information deficit between market participants and policymakers, this approach can foster a greater understanding of the communication transmission mechanism. This is important to policymakers who may be trying to design communication to have complement their policy actions, and certainly want to avoid adverse effects on investor beliefs.

Specifically, in this paper we contribute to the literature in four main ways. First, we develop an algorithm to quantify the temporal dimensions of the central bank communication. This allows us to focus on the role and importance of temporal information in the the communication transmission mechanism. While there has been a growing emphasis on the importance of communicating the future intentions of monetary policy makers as part of the management of expectations, studies into the explicit importance of temporal information in textual data are, surprisingly, rare to date. Our methodology adapts new tools from the Natural Language Processing (NLP) literature in order to best suit our central bank communication context.

*A priori*, one might anticipate that information regarding the future might be espe-

cially important for financial markets, relative to information about the past. However, we present a generic description of the process for monetary policy decision making, and the markets anticipation of decisions, that shows that two important sources of the information deficit are likely to be the central banks assessment of economic conditions and their likely evolution. This assessment relies on both *evaluation* of the state of the economy based on past data, and *projection* from the state to the future.

Using our newly-constructed temporal information, our second contribution is to show that the importance of temporal information. We find that information regarding the future has an important role at explaining asset price movements, and we interpret this as evidence in favour of the idea that markets react to information about central bank forecasting procedures. Moreover, we show that even backward-looking data is extremely informative. This, we argue, captures important contextualisation of the data. We conclude that markets react to information regarding the processes by which a central bank interprets data.

Notwithstanding the progress we have made in measuring the temporal dimensions, our third contribution is to provide evidence that there are limits to the extent that the NLP approach can fully uncover the effect of communication on belief updating. There is an identification problem which is difficult to overcome and therefore results from such studies should be carefully interpreted.

Finally, our fourth contribution is to develop a novel measure of the information deficit using the questions raised by journalists after the opening statement of the press conference. These are, to the best of our knowledge, the first direct measures of the information deficit. Since speeches are frequently used by policymakers to address and clarify issues that are not covered in sufficiently during press conferences, that policymakers believe markets wish to hear about. We show that where a speech is closely related to the questions asked after the most-recent statement, the speech is well placed to address the information deficit between the ECB and financial markets, and therefore leads to a stronger market reaction (greater news and belief updating).

## 2 Monetary Policy Surprises and Information Deficits

In this section we show that the complexity of the monetary policy process can be grouped into three broad aspects. Each is a potential source of news for market participants. Although monetary policy is inherently forward-looking, we show that the source of news that generates market surprises could as reasonably come from backward-looking assessments of the current data context as from the forward-looking forecast.

## 2.1 Monetary Policymaking Process

Consider a monetary policymaker in month  $m$  who has to decide the interest rate,  $i_m$ . They have access to a large amount of data capturing macroeconomic trends, surveys, market prices, and other relevant information. We represent this as a high dimensional vector  $X_m^{CB}$ . The decision-making process can be summarised in two broad steps:

1. The policymaker maps the data into a vector of beliefs about the current state of the economy:  $\Omega_m^{CB} = g_m(X_m^{CB})$  where  $g_m(\cdot)$  captures the analysis of contemporaneous data, as well as the forecasting exercise including the judgement applied.
2. The policymaker selects the appropriate interest rate as a function of this state:  
$$i_m = f_m(g_m(X_m^{CB}))$$

We have allowed both the analysis/forecasting function, and the reaction function to vary each meeting. This description is generic. In a typical forward-looking reaction function, the state consists of  $h$ -period ahead forecasts of inflation deviations from target and the output gap ( $\Omega_m^{CB} = [\hat{\pi}_{m,m+h}^e, \hat{y}_{m,m+h}^e]$ ), and  $f_m(\cdot)$  is linear and typical does not vary with time. In a standard DSGE model, there is also a pure monetary policy shock added linearly to the endogenous reaction terms; in our framework,  $f_m(g_m(X_m^{CB})) - f_{m-}(g_m(X_m^{CB}))$  would capture the same kind of shift in policy for a given state of the economy.

The area that typically has received less attention in macro models is  $g_m(\cdot)$ . Those who have worked in a central bank will be aware that this function captures the very heart of the interest rate process. This function, broadly, captures two important analytical steps:

**Evaluation:** This involves following the developments across a range of domestic and international markets using macroeconomic, financial and other data. Recent data is analysed as well as put in context in order to interpret its movements (e.g. is a recent increase in inflation likely to be transitory?). This analysis is largely backward-looking in the sense that it uses the latest available data pertaining mostly to recent developments. However, it may also draw on more historical data for context and within statistical analyses of the data.

**Projection:** Monetary policy is inherently forward-looking – it is made on the basis of forecasts of key variables, and the risks around them. Forecasts typically use a suite of models as well as judgement informed by the economic analysis.

## 2.2 Information Deficits and Market Surprises

Market participants form expectations of the central bank policy choice and these are reflected in market prices of various securities. We denote the average expectation just before the announcement of the monetary policy decision at time  $m$  as  $\mathbb{E}[i_t | \mathcal{I}_{m-}^{mkt}] = \tilde{f}_{m-}(\tilde{g}_{m-}(X_{m-}^{mkt}))$  where  $m-$  indicates the moment right before the month  $m$  decision,  $\mathcal{I}_{m-}^{mkt}$  is the information set at that moment,  $X_{m-}^{mkt}$  is the data that the market is looking at, and  $\tilde{f}_{m-}(\cdot)$  and  $\tilde{g}_{m-}(\cdot)$  are the market beliefs about the central bank equivalent functions described earlier.

The monetary event, an announcement of policy and the associated communications, provides information from the central bank. The new, updated information is  $\mathcal{I}_m^{mkt}$  and the expectation is  $\tilde{f}_m(\tilde{g}_m(X_m^{mkt}))$ . The market surprise, as measured in the empirical literature is, therefore,  $\varepsilon_m^i = \mathbb{E}[i_m | \mathcal{I}_m^{mkt}] - \mathbb{E}[i_{m-} | \mathcal{I}_{m-}^{mkt}]$ . This surprise could be generated by changes in any of the three broad aspects of the central bank decision:

1.  $X_m^{CB}$  could be revealed to include new information not in  $X_m^{mkt}$ .
2. The central bank could simply provide more details about how it assesses/forecasts the economy, or reveal that it has re-evaluated how it is assessing the state of economy and forecasting its evolution as captured by  $g_m(\cdot) \neq \tilde{g}_{m-}(\cdot)$
3. The central bank may choose to react to a given state of economy more or less aggressively than previously;  $f_m(\cdot) \neq \tilde{f}_{m-}(\cdot)$ .

To understand the drivers of each of these channels further, we can approximately decompose the market surprise  $\varepsilon_m \equiv \tilde{f}_m(\Omega_m^{mkt}) - \tilde{f}_{m-}(\Omega_{m-}^{mkt})$  using a first-order Taylor series expansion of  $\tilde{f}_{m-}(\Omega_{m-}^{mkt})$  around the pre-announcement view of the economic state  $\Omega_{m-}^{mkt}$ .<sup>1</sup>

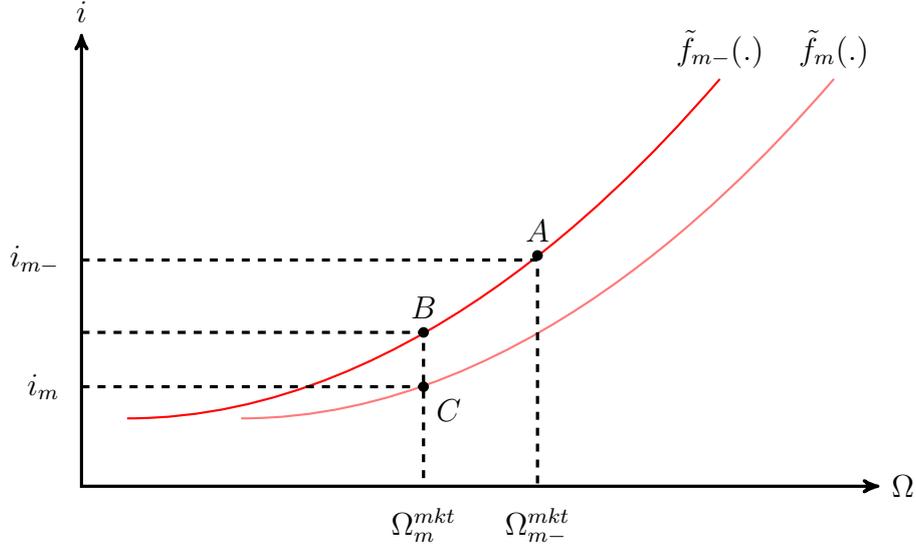
$$\varepsilon_m \approx \underbrace{\tilde{f}_m(\Omega_m^{mkt}) - \tilde{f}_{m-}(\Omega_m^{mkt})}_{\text{Updated Reaction function}} - \underbrace{(\Omega_m^{mkt} - \Omega_{m-}^{mkt})\tilde{f}'_{m-}(\Omega_m^{mkt})}_{\text{Updated State}}$$

To see this decomposition visually, consider Figure 1. Point A is the pre-announcement expectation for the interest rate ( $i_{m-} = \tilde{f}_{m-}(\Omega_{m-}^{mkt})$ ); Point C captures the post-announcement expectation ( $i_m = \tilde{f}_m(\Omega_m^{mkt})$ ). The effect of the updated state of the economy is captured by the move from A to B while B to C reflects the impact of the update reaction function at the new assessment of the economy.<sup>2</sup>

<sup>1</sup> $\tilde{f}_{m-}(\Omega_m^{mkt}) \approx \tilde{f}_{m-}(\Omega_{m-}^{mkt}) + (\Omega_m^{mkt} - \Omega_{m-}^{mkt})\tilde{f}'_{m-}(\Omega_m^{mkt})$ . Subtracting this from  $\tilde{f}_m(\Omega_m^{mkt})$  yields the expression in the text.

<sup>2</sup>Of course, other decompositions will give slightly different weight to each component, especially where the reaction function is non-linear.

**Figure 1: Market Surprise Decomposition**



We can further decompose the effect of the update of the economic situation using a first-order Taylor series expansion of  $g_m(X_m)$  around existing data  $X_{m-}$ :

$$\begin{aligned}
 \Omega_m^{mkt} - \Omega_{m-}^{mkt} &= \tilde{g}_m(X_m^{mkt}) - \tilde{g}_{m-}(X_{m-}^{mkt}) \\
 &= [\tilde{g}_m(X_{m-}^{mkt}) - \tilde{g}_{m-}(X_{m-}^{mkt})] + [\tilde{g}_m(X_m^{mkt}) - \tilde{g}_m(X_{m-}^{mkt})] \\
 &\approx \underbrace{[\tilde{g}_m(X_{m-}^{mkt}) - \tilde{g}_{m-}(X_{m-}^{mkt})]}_{\text{Reassessment}} + \underbrace{(X_m^{mkt} - X_{m-}^{mkt})\tilde{g}'_m(X_{m-}^{mkt})}_{\text{Effect of New Info}}
 \end{aligned}$$

Putting this all together, we have:

$$\begin{aligned}
 \varepsilon_m &\approx \underbrace{\tilde{f}_m(\Omega_m^{mkt}) - \tilde{f}_{m-}(\Omega_m^{mkt})}_{\text{Updated Reaction function}} \\
 &\quad - \underbrace{[\tilde{g}_m(X_{m-}^{mkt}) - \tilde{g}_{m-}(X_{m-}^{mkt})]}_{\text{Reassessment}} \tilde{f}'_{m-}(\Omega_m^{mkt}) \\
 &\quad - \underbrace{(X_m^{mkt} - X_{m-}^{mkt})\tilde{g}'_m(X_{m-}^{mkt})\tilde{f}'_{m-}(\Omega_m^{mkt})}_{\text{Effect of New Info}} \quad (1)
 \end{aligned}$$

An earlier literature has emphasised an important role for private information in monetary policy. Central banks have *some* information that is not publically available, such as regulatory data from financial firms, or sometimes have early access to data. Nonetheless, markets have access to almost all of the underlying data, and it would be rare for the central bank announcement to formally reveal whatever private information there is. If  $X_m^{CB} \simeq X_{m-}^{mkt}$ , genuinely new information (the 3rd element) would not be an important driver of the surprises.

What is revealed, and what could be the source of perceived informational advantage, is a differential assessment of economic conditions. The central bank often has greater analytical resources in the form of large teams of economists and multiple models. This is captured by  $g_m(\cdot)$  and reassessments of the economic state are often communicated. This could be being putting more or less weight on different sources of data to because it better explains the data we observe, or an updated forecast having reassessed the nature and pattern of recent forecast errors. This dimension could entail both forward-looking *projection* as well as backward-looking *evaluation* of the data.

This reassessment information makes up the majority of communicated announcements. For example, consider the Opening Statement following the September 12th 2019 Governing Council Meeting delivered by then-ECB-President Mario Draghi (Draghi 2019). He states that “Today’s decisions were taken in response to the continued shortfall of inflation with respect to our aim. In fact, incoming information since the last Governing Council meeting indicates a more protracted weakness of the euro area economy, the persistence of prominent downside risks and muted inflationary pressures. This is reflected in the new staff projections, which show a further downgrade of the inflation outlook.” This the first of 15 paragraphs (out of a total of 23 substantive paragraphs) providing details of the ECB’s assessment of the economic state.

The first 8 paragraphs describe all the dimensions of the policy decision (interest rates, asset purchases, etc..) which capture the role of the  $f(\cdot)$  function. Of course, much of the policy decision is not news once we condition on the updated assessment of the economy. (This can be seen in the  $f'_{m-}(\Omega_m)$  term in (1)). But sometimes the central bank communicates changes in stance. Draghi (2019) indicates some update saying: “We now expect the key ECB interest rates to remain at their present or lower levels until we have seen the inflation outlook robustly converge to a level sufficiently close to, but below, 2% within our projection horizon, and such convergence has been consistently reflected in underlying inflation dynamics.”

### 3 Measuring the 3rd T: Temporal Communication

The previous section has, we hope, established that monetary policy news may come from both the projection and the evaluation aspects of the central bank decision. While a full estimation of the general model proposed above is beyond the scope of this paper, we wish to examine the extent to which both future and past temporal information is important empirically.

In order to do this, it is necessary to measure the temporal dimension of the information. While existing studies of central bank communication have emphasised two T’s,

Topic and Tone, attention on the 3rd T is limited. One reason for this is that it is not as trivial as you might imagine. In this section, we present the algorithm that we have developed specifically for measuring the temporal dimension in central bank texts. As we will apply this algorithm to ECB monetary policy communications, we first describe the textual data we use.

### 3.1 Textual Data

Our textual data come from two main sources - the ECB press conferences (including the introductory statement, and the associated Q&A session) and the text of speeches delivered by members of the ECB executive board.

The data from the ECB press conferences were manually taken from the website of the ECB.<sup>3</sup> The textual data from the Q&A session contains the questions from journalists – for baseline estimations we cut out the questions, leaving only the responses (which are typically from the ECB President, though occasionally the Vice-President interjects).

Overall, we collected 240 transcripts from the website of the ECB, starting with the first conference of 9th June 1998, and ending with the conference of 10th September 2020.<sup>4</sup>

Policymaker speeches come from the ECB Speeches Dataset, which was created by ECB staff and made available on its website.<sup>5</sup> This is an archive of all speeches by ECB Executive Board members, dating back to February 1997.<sup>6</sup> These data are continually updated by the ECB, we use the version of the dataset that ends with a speech on the 15th September 2020.

Our sample includes 2,203 English speeches.<sup>7</sup> Note that it is often the case that there is more than one speech on a given day. This means the total number of days in which one or more speech occurs is 1,713, i.e. it is lower than the number of speeches. When constructing our measures of temporal orientation for given speeches, we treat speeches given on the same day essentially as one document.

Though the sources of data present the text in a relatively regular manner, in order to use these data we apply some standard cleaning procedures. One important difference is that we need to preserve numerical information associated with dates, whereas numerical

---

<sup>3</sup>See: <https://www.ecb.europa.eu/press/pressconf/html/index.en.html>.

<sup>4</sup>Note that the number of Q&A sessions in our sample is 237, since the questions were not taken during the first three press conferences.

<sup>5</sup>The dataset is here: <https://www.ecb.europa.eu/press/key/html/downloads.en.html>.

<sup>6</sup>The data include speeches delivered by senior officials prior to the formal creation of the ECB in June 1998.

<sup>7</sup>We exclude 159 non-English speeches. There are 16 speeches for which no text is available, which are discarded. There are 34 speeches that merely summarise the title and topic of the speech, and provide a hyperlink to lecture slides – these are also discarded.

information is often jettisoned in other applied computational linguistics studies. We process the statements and the speeches in the same way, though some more cleaning was required in the case of the speeches as discussed in Byrne et al. (2021).

## 3.2 Temporal Tagging

The approach to temporal tagging used in this study categorizes textual data according to three dimensions:

1. categorical references to time - refer to time only in a general sense, and includes references such as “the future”, “in the long-run”, “currently”.
2. numerical references to time - references that can be placed on a calendar such as “next year”, “in the last few months”, as well as more direct numerical references like “1st January 2020” and “2020”.
3. grammatical tense - our tagging algorithm also recognises whether sentences include the use of the present, past, and future tenses.

We shall now explain each in more detail. There is an accompanying technical guide to the algorithms which will provide the interested reader with even more information (Byrne et al. 2021).

### 3.2.1 Tagging Numerical and Categorical Time-References with SUTime

To accurately isolate such references within a large corpus of documents, we employ the SUTime temporal tagger developed in Chang and Manning (2012). SUTime is a rules-based temporal tagger built on regular expression patterns rather than on statistical relationships. Chang and Manning (2012) show that SUTime performs well in comparisons with other temporal taggers in the natural language processing literature across a number of criteria.

A key benefit of SUTime is that it can tag a wide range of representations of time, allowing greater accuracy and flexibility to capture the typical ways in which people speak. SUTime is not limited to absolute date formats such as YYYY/MM/DD, “June 2020” or similar simple references. SUTime is also able to resolve relative date formats, such as “last Friday” or “two months from now”, since the processor takes a reference date as an input.

The outputs of SUTime are temporal tags in the TIMEX3 format (Pustejovsky et al. 2003). Text that can be resolved to a specific date will result in a numerical time tag. For example, “June 2020” or “next June” are numeric dates. SUTime also produces

categorical tags, covering past, present, or future, for more abstract date formats that do not resolve to a specific date.<sup>8</sup>

Although the library of rules in SUTime is large, it is not tailored to the language of monetary policy or central banking. There may be phrases that central bankers use to talk about time that are easily understood by their audiences, but which are not covered by the SUTime rules. For instance, central bankers frequently refer to the “short-term”, “long-run”, and similar constructions. These expressions are, however, not included in the SUTime library. We thus expand SUTime’s library of rules to capture better the ways in which central bankers speak about monetary policy or economics-related topics.

In a similar fashion, central bank communication often refers to dates, times and eras by commonly understood shorthand names. The audience hearing “Great Depression”, “Bretton Woods era” or “Global Financial Crisis”, for instance, is likely to know well to which point in time the speaker is referring. A temporal tagger, on the other hand, would not. We developed a list of relevant economics text-based date expressions and map them to numerical dates, allowing SUTime to process them (discussed fully in Byrne et al. (2021)).

### 3.2.2 Tagging Tense with TMV

Applying computational linguistics to assess whether phrases within a given sentence are in the past, present, or future tenses is a non-trivial task.

Standard computational tools allow one to assign “part of speech” (POS) tags to tokens (words) from corpora of textual data. The tags themselves come from a list of potential word classes for the English language (nouns, verbs, determiners, etc.) However, because POS taggers are applied to tokens, and not verb phrases, their ability to detect tense is limited. To identify the future tense, it is necessary to augment POS taggers with additional tools from the computational linguistics literature.

To identify tense, this study applies the Tense-Mood-Voice tool, introduced by Ramm et al. (2017). This tool is designed to automatically classify verbal complexes (sequences of verbal tokens within a verbal phrase) into their tense, according to a rules-based method. Ramm et al. (2017) distinguish semantic tense from morphosyntactic tense. They give the example of the English sentence “He is leaving at noon”, which is semantically in the future tense, but has the morphosyntactic tense of present progressive. The TMV tool can only provide information about the morphosyntactic tense. The system takes as its

---

<sup>8</sup>SUTime will also identify if the text is referring to a range of dates from one point to another, or a duration such as “for three months”. We do not incorporate information from ranges in our approach, since it is unclear how to resolve such expressions into a single value (one could use the middle of the range, though it is not clear whether such expressions should be treated in the same manner as dates, so we prefer to omit these cases).

argument individual sentences. It then identifies verbal complexes from these sentences, before applying a sequence of around 32 rules to these verbal complexes. For example, the system understands that, for the simple future tense, the modal auxiliary “will” (or “shall”) precedes the infinitive form of the verb, so “I will go” is correctly identified as the future tense.

The system assigns verbal complex to four forms of the present tense (present, present progressive, present perfect, present perfect progressive), four form of the past tense (past, past progressive, past perfect, past perfect progressive), and four forms of future tense (two respective forms of the future and future progressive tenses are identified). We do not distinguish between different forms of present, future or past tenses. We assign the four future tenses to a general future tense category, and likewise for the present and past tenses. We assign the two conditional tenses that are about the past to the past tense category, and we assign the two conditional future tenses to the future category. We do not consider non-finite verbal complexes.

TMV classifies sentences according to their tense but we additionally, ex-post, assign certain expressions in the present tense to the future. Some of these are particularly prevalent in central bank speak. For example, central bankers frequently make statements using expressions such as “we expect”. “we forecast”, or “we predict”. The full list of verb forms we additionally assign to the future are reported in Byrne et al. (2021).

### 3.2.3 Temporal Tags: Example and Some Issues

To recap, by applying the SUTime and TMV approaches, we identify numerical and categorical time-references (SUTime), as well as past, present, and future verbal complexes (TMV). Table 1 shows two example sentences from our text corpus.

The blue highlighted text captures future tagged content while the orange highlights capture the past references. Phrases marked **Numerical** or **Categorical** are tagged using the SUTime tool (or our additional central bank time references), while phrases marked with **Tense** are tagged using the TMV tool. Phrases marked with **Tense\*** are tagged as present tense using the TMV tool, but coded as future using a bespoke dictionary of present tense phrases that evoke future considerations.

The first sentence from the speech at Sintra is clearly about the future, and the second is clearly referring to the past but providing context to their recent decision making. Our argument is that such context, especially the symmetric nature of their objective function, is useful and important information for markets trying to predict future interest rates.

The sentences from the Press Conference highlight another important aspect of the measures. The first of these sentences shows that we capture important forward-looking information with the adjusted tense reference as well as the categorical SUTime measure;

**Table 1:** Example of Our Temporal Parsing Approach

Past Tag	Textual Data	Future Tag
	<p>“In the absence of improvement, such that the sustained return of inflation to our aim is threatened, additional stimulus <b>will be</b> required. In our recent deliberations, the members of the Governing Council <b>expressed</b> their conviction in pursuing our aim of inflation close to 2% in a symmetric fashion.</p> <p>ECB President Draghi, Speech at Sintra, 18th June 2019.</p>	Tense
	<p>“Over the <b>medium term</b> underlying inflation <b>is expected</b> to increase, supported by our monetary policy measures, the ongoing economic expansion and robust wage growth. . . . This assessment is also broadly reflected in the <b>September 2019</b> ECB staff macroeconomic projections for the euro area, which <b>foresee</b> annual HICP inflation at 1.2% in <b>2019</b>, 1.0% in <b>2020</b> and 1.5% in <b>2021</b>.”</p> <p>ECB President Draghi, Press Conference, 12th September 2019.</p>	Categorical Tense* Tense* Numerical Numerical

Notes: Phrases marked Numerical and are tagged as future/past using the SUTime tool. Phrases marked Categorical are tagged using the SUTime tool, with an additional bespoke dictionary of future words specific to the language of central banking (for example, “medium-run”). Phrases marked with Tense are tagged as future/past tense using the TMV tool. Phrases marked with Tense\* are tagged as present tense using the TMV tool, but coded as future using a bespoke dictionary of present tense phrases that evoke future considerations, designed for use with central bank communications (for example, “expect”, “foresee”).

standard tense analysis would either miss this reference, or even classify it as past tense!

The second of these sentences highlights that the numerical references may sometimes blur the measure of temporal orientation coming from the other measures. Our baseline measures described below classify future and past according to sentences tagged by *any* of the measures. We also conduct our analysis using topics constructed according to each temporal tagger separately; this more disaggregated approach will generally generate even stronger results though the qualitative nature of the baseline results is the same.

### 3.3 Measures of Document Temporal Orientation

Once sentences are tagged appropriately, we are in a position to create measures of temporal orientation from our textual data. To do this we first identify all sentences in our corpora that contain *at least one* reference to the future, be it a numerical future reference, a categorical future reference, or the use of the future tense. We then identify

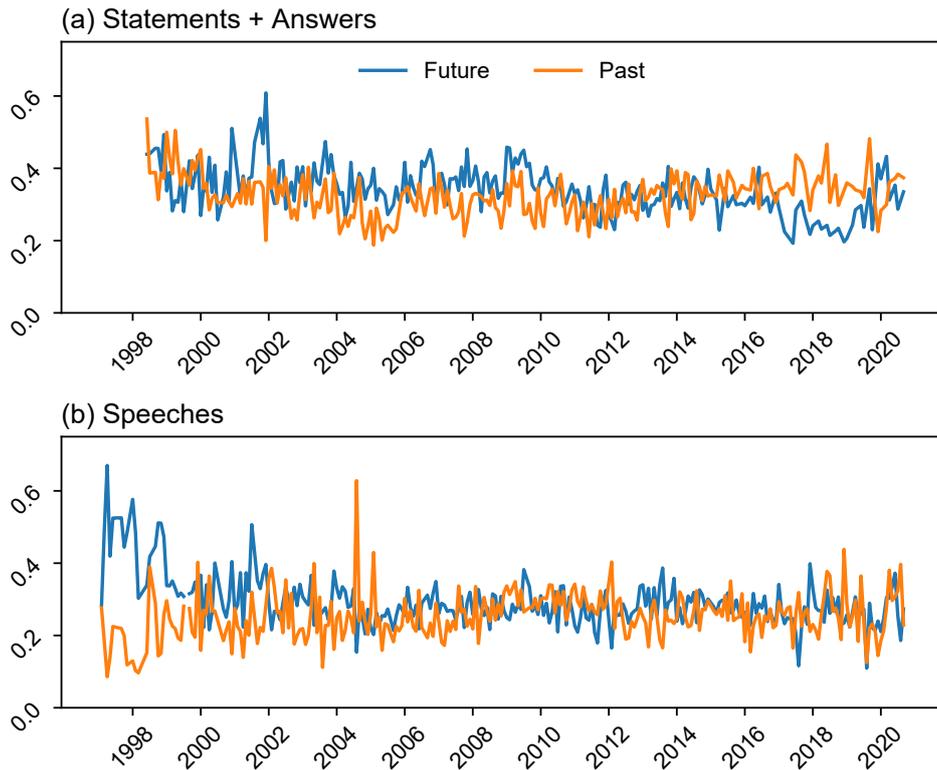
all sentences in our corpora that contain at least one reference to the past (according to any form of tag).

Formally, consider a sentence  $s \in (1, N^d)$  found in document  $d$ . Let  $T_s^{fut} = 1$  if and only if sentence  $s$  contains at least one temporal expression relating to the future (be it numerical, categorical, or the use of the future tense). Define  $T_s^{pst}$  analogously as an indicator variable that equals 1, if and only if sentence  $s$  contains at least one reference to the past. Note that a sentence can be tagged as *both* about the future, and about the past, according to this scheme.

The document level temporal orientation measures,  $p_d^{fut}$  for future and  $p_d^{pst}$  for past, are defined as:

$$p_d^{fut} = \frac{\sum_{s=1}^{s=N^d} T_s^{fut}}{N^d}, \quad (2)$$

$$p_d^{pst} = \frac{\sum_{s=1}^{s=N^d} T_s^{pst}}{N^d}. \quad (3)$$



**Figure 2:** The Time-Series of Future Orientation Across Corpora

Notes: The figure displays  $p_d^{fut}$  across time for the speeches, as well as the statements and answers. The time-series are averaged at monthly frequency.

Figure 2 shows the time-series of the future and past measures, for both the statements

and answers corpus, and the speeches corpus. Note that the share of sentences marked as future and past fluctuates around 40% in the statements and answers corpus, and 20% in the speeches corpus. We conclude that the statements and answers corpus contain a greater proportion of temporal sentences, on average. The future and past shares do appear to be inversely correlated, which makes sense, given their construction. We have constructed our temporal taggers to allocate phrases or words to the mutually exclusive categories of past, present, and future. *Ceteris paribus*, an increased number of sentences tagged only as future will lead to a reduced past share. However, it is still possible that the future and past shares can co-move positively, in the sense that such movements are possible in the case that the share of sentences marked only as present falls. Another important point about these figures is that both series are monthly averages. In the case of the speeches, the figures therefore understate the level of variation from speech to speech.

From Figure 7, several stylised facts emerge about temporal information in ECB communication. Panel (a) displays information from the statements and answers corpus. In the earliest years of the sample we observe a slow decline in the share of references to the past, from the period 1999 to 2004. We do not observe a clear rise in the measure of futureness, meaning that share of future references was kept at a generally similar level during this period (with some higher-frequency fluctuations). We observe a slow decline in the level of futureness from the 2009 period, with a particularly sharp fall in the 2017 to 2019 period. We have found that these dynamics are driven by the categorical and tense measures of futureness.

We observe a clear increase in the measure of future orientation in the pre-1999 period, before a slow decline in this measure until around 2005. This is driven largely by movement in the tense and numerical indicators of futureness, and may reflect a large number of references to the 1st of January 1999 in the earliest speeches, as well as a general discussion of a the future prospects of the monetary union. These trends are reflected (with opposite sign) in the measure of past orientation. In the post-2005 period, the measures derived from the speeches do not vary much at business cycle frequency, though there us a degree of higher-frequency variation. Note that the use of aggregate measures obscures variation that may occur at the level of individual topics, as we will examine later in this study.

### **3.3.1 Construction of Temporal Topics**

While these document-level future and past orientation measures may provide reasonable summary measures of the overall temporal orientation of a speech, or press conference, one concern is that such an approach would amalgamate information from a fairly diverse

range of subjects. It is also of interest to study whether the temporal orientation of given topics within a speech changes the way market participants respond to given speeches. For example, a discussion of the future path of interest rates in a given speech may be more relevant for market participants than a discussion of past interest rate choices (even though the context of how the past is shaping current decisions may also be informative).

To capture the temporal orientation of given topics, we need a way to combine our temporal information with measures of topic from textual data. To measure topics, we use Latent Dirichlet Allocation (LDA) following Blei et al. (2003) and applied to central bank communication in Hansen et al. (2017). LDA is applied at the sentence level of our corpus.

One complication is that our overall corpus includes information from several sources, in the sense that we wish to extract topics from both press conferences (which are largely focussed on monetary policy) and from speeches (which may include references to a more diverse range of subjects, including law, history, for example). Given that the number of speeches (2,203) greatly exceeds the number of press conference (240), treating all of our textual data as one corpus would tend to favour the extraction of topics that explain the speeches, potentially to the detriment of model fit for the press conferences.

We choose to estimate our topic model on a corpus comprised of the introductory statement during press conferences, and the responses to questions during the Q&A session (but not the questions themselves). The speeches are therefore *not* included in the estimation of the topic model. Given a topic model estimated on the press conference textual data, we extrapolate our topic model to the textual data from the speeches (*querying* in the language of information retrieval).

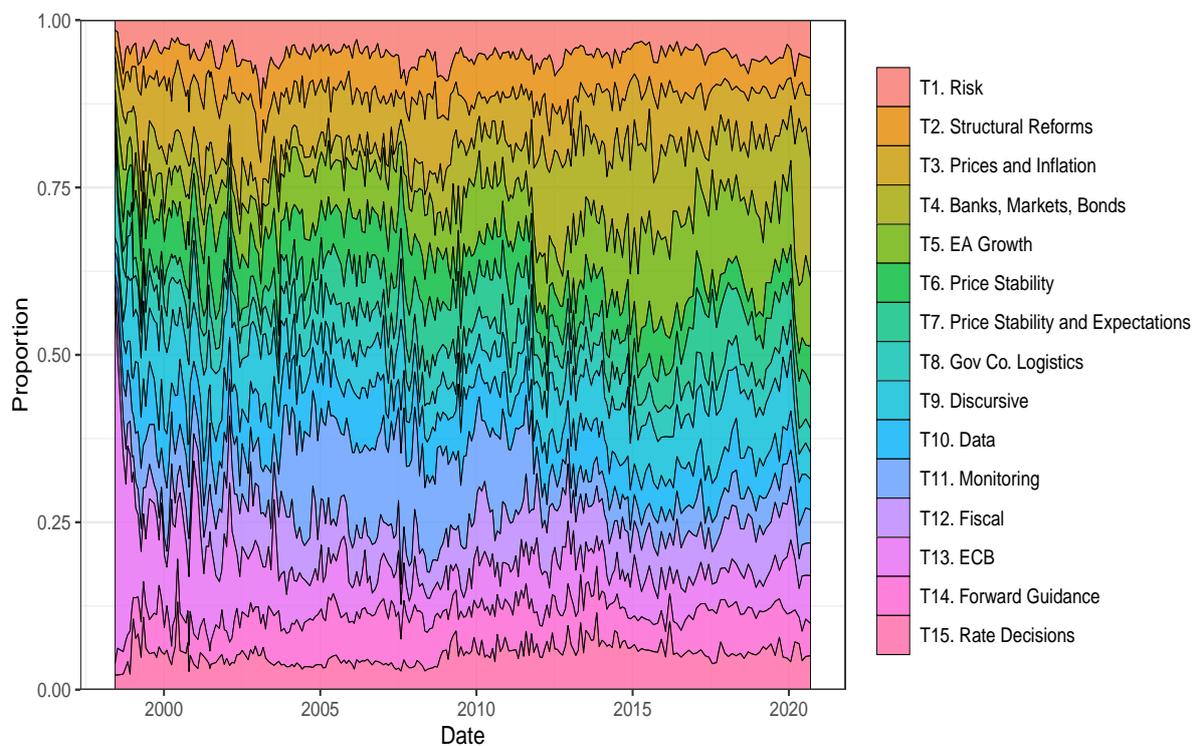
A number of pre-processing steps are taken prior to the fitting of the topic model. Note that these pre-processing steps are applied to the corpus prior to the application of the LDA model, but not prior to the application of the SUTime and TMV tools. These steps are largely standard, and include the removal of numbers and punctuation, the removal of standard English stopwords (e.g. “the”), the conversion of words to lower case and the use of a stemmer to reduce words to their root form (so “developing” and “developed” are both stemmed to the same root, “develop”).

The model was estimated using a Gibbs sampling approach, with a burn in of 1000 iterations and a total of 2000 iterations for the fitting process. We set the number of topics to 15, which led to a distinct and interpretable series of topics.

In turn, the fitted LDA model of 15 topics was extrapolated out of sample to the speeches dataset using a second Gibbs sampling approach with identical numbers of burn in and iterations in total. A minor concern is the existence of words that may be in the speeches data but not the press conferences. Reassuringly, in total these account for

only 4.3% of the total count of words that occur in the speeches dataset, and as such their exclusion from the training data is unlikely to have a fundamental impact on our goodness of fit.

The fifteen generated topics are shown in Table **tab:topicfull**. The fifteen generated topics each have meanings that are broadly interpretable in line with certain aspects of ECB communication. A number of topics are directly related to the primary mandate of price stability (Topics 3, 6, and 7), while other groups are more directly related to monetary policy actions (Topics 14 and 15). In addition, two topics are clearly related to structural parts of the statements, such as Topics 2 and 12, which are directly related to structural reforms, fiscal policy and the stability and growth pact, while Topic 9 appears to be related to broad analysis of risks that can be seen in the balance of risks segment of the press conference. A time series of the relevant topics is shown in Figure ??



**Figure 3:** Time series of the LDA topic proportions in the statements and questions dataset

Notes: This figure shows the distribution of the probabilities from a LDA topic model fit on the ECB statements and answers over time. Each section of the area chart reflects the implied proportion of the overall statement and answers dedicated to that given topic at the relevant point in time.

The distribution of the topics over time suggests that while on the whole the structure of the document has been relatively homogeneous, there are changes over time. First, Topic 4, which relates to Banks, Markets and Bonds increases in importance in the

2010s, and increases further in 2020. This likely reflects the increased importance of asset purchases during this time point. Similarly, Topic 13, which deals with a number of ECB specific administrative concepts is at its peak towards the start of the sample, reflecting the discussion of the institutional formation of the ECB at that time point. Other items of note are the fact that Topics 1,2 and 15 are relatively static in the sample over time. This reflects the fact that they represent some structural components of the statements that are fixed over time (the balance of risks, discussions of structural reforms and rate decisions respectively).

Given this estimation of our topic model, the next stage is to create our measures of temporal topic. Expressed in words, the future (past) temporal topic for a given document represents the average topic share for that topic, when we restrict the corpus to only those sentences that contain a future (past) reference. Formally, for a given document  $d$  with  $N_d$  sentences, and a given topic  $k$ ,  $k \in (1, \dots, K)$  with  $K = 15$ , we denote the statement (or speech) future oriented topic measure ( $\theta_d^{fut}(k)$ ) and the past orientation measure ( $\theta_d^{pst}(k)$ ) respectively according to:

$$\theta_d^{fut}(k) = \frac{1}{N_d^{fut}} \sum_{s=1}^{s=N_d} T_s^{fut} \theta_{k,s}, \text{ where } N_i^{fut} \equiv \sum_{i=1}^{i=N_d} T_i^{fut} \quad (4)$$

$$\theta_d^{pst}(k) = \frac{1}{N_d^{pst}} \sum_{s=1}^{s=N_d} T_s^{pst} \theta_{k,s}, \text{ where } N_i^{pst} \equiv \sum_{i=1}^{i=N_d} T_i^{pst}. \quad (5)$$

where  $\theta_{k,s}$  is the topic share for topic  $k$  associated with the sentence  $s$ .

To demonstrate some of the properties of our new measures of temporal orientation, Panel (a) of Figure 4 displays the evolution of measures relating to topic 7 and topic 12 over time. These measures are constructed for the statements corpus, and therefore vary at the frequency of the Governing Council meetings. Topic 7 is broadly related to price stability, and inflation expectations, this can be seen from Panel (b). Topic 12 relates to fiscal policy and government, as can be seen from Panel (c).

Figure 4 displays several important properties of our measures graphically. The first point to make is that, when studying the overall topic shares, for both topics, we see evidence of both higher-frequency changes in topic shares between statements, as well as more persistent movements. For example, the overall measure of the price stability and expectations topic increases during the period of the financial crisis period, decreases during the resolution of the European sovereign debt-crisis, before rising again. The fiscal policy topic is elevated in the post-2008 crisis period. Of course, time-varying topic shares have been estimated for central bank statements in a number of significant studies to date, and the observation that these shares evolve over time is not novel to this paper.

Part of the contribution of this study is to separately estimate temporal topic shares, and their evolution is also plotted in Figure 4. We can observe several features of the shares immediately. The first is that they are correlated with the overall topic share. This is unsurprising, given that they are calculated essentially as “sub-corpora” of the overall corpus. The measures also appear more volatile than the overall measure. Again, this was perhaps to be anticipated, given that they are averages applied to a smaller number of sentences (i.e. those only those sentences that are tagged as future or past, respectively). This makes it likely that the new measures exhibit a somewhat greater degree of noise. However, what is also clear from the Figure is that the new temporal topics are not reducible to the overall topic measure, and that they can display interesting dynamics that are less pronounced in the overall measure. For example, in the case of Topic 7, we note that there is a clear increase in the future topic share during the 2008 crisis period. There is also a larger increase during the period of 2015 to 2019, during which the ECB was conducting purchases as part of its PSPP programme. This was also the period during which the ECB President linked net asset purchases to a “sustained adjustment in the path of inflation”, meaning the future evolution of inflation received a great deal of emphasis in discussion. For the case of topic 12, there is also evidence of a decline in the past temporal topic in the post-2014 period, perhaps indicating the diminishing quantity of analysis and reflection devoted to fiscal matters, after the resolution of the European sovereign debt crisis.

## 4 Analysis of the Yield Curve News

The framework presented in section 2 suggests that central bank communication will cause market news when market participants, in aggregate, update their beliefs because the communication fills an information deficit that they had. We argued that, while some of the information deficit would relate to the central bank’s outlook for the economy (the projection dimension), reassessment and contextualisation of current data should be important too.

Armed with our measures of the temporal dimension of central bank communication, we can provide an empirical assessment of the drivers of market surprises. If the information deficit comes only from the projection dimension of policy making, the inclusion or exclusion of past information should have no bearing on the extent to which we can explain the asset price news. On the other hand, if interpretation matters, it should, at least on average, result in past temporal topics being able to explain some of the asset price news.



sponding observations of a vector of  $p$  potential predictor variables  $\{\mathbf{x}_i\}_{i=1}^N$ , where  $\mathbf{x}_i$  is of dimension  $(p \times 1)$ . The following minimisation problem is:

$$\min_{\beta \in \mathbb{R}^p} \left\{ \frac{1}{2} \sum_{i=1}^N (y_i - \mathbf{X}_i \beta)^2 + \lambda \left[ \frac{1}{2} (1 - \alpha) \|\beta\|_2^2 + \alpha \|\beta\|_1 \right] \right\},$$

for some  $\lambda \geq 0$  and  $\alpha \in [0, 1]$ , where  $\|u\|_p \equiv \sum_{j=1}^N (|u_j|^p)^{1/p}$  is the  $l_1$ -norm. Here  $y = (y_1, \dots, y_N)$  denotes an  $N$ -vector of responses of interest, and  $\mathbf{X}$  is a  $N \times p$  matrix of independent variables. Note that when  $\alpha = 1$  this specification.

We define the parameter  $\alpha$  to be equal to 0.99. We estimate the parameter  $\lambda$  by 10-fold cross-validation.

One complicating factor when using LASSO estimation is that, in the case of highly correlated independent variables, parameters can be selected by the routine in a somewhat arbitrary manner.<sup>9</sup> To account for this, we tend to summarise the results from our LASSO estimation routine using a non-parametric bootstrap. We draw with replacement from our dataset  $M$  observations, and do this 5000 times, and store the distribution of estimated parameters. Note that for each bootstrapped dataset, we estimate a different value of  $\lambda$  via 10-fold CV. To compute the distribution of adjusted  $R^2$  we adopt the following procedure. For a given bootstrapped dataset, having estimated a LASSO specification, we then re-estimate our prediction equation via OLS, conditional on the subset of variables that were assigned non-zero coefficients by the LASSO algorithm. Our measure of adjusted  $R^2$  for this draw is taken from this OLS regression.

## 4.2 Effects of Temporal Communication on Market Yield News

A growing literature in the empirical macro-economics literature focusses on the extraction of monetary policy surprises from policy announcements, by isolating changes in asset prices in narrow windows around the announcement. Early papers studied uni-dimensional monetary policy surprises (Kuttner (2001)). To study the impact of our textual measures of ECB Press Conference communication on market yields, we use intra-daily data from the Euro Area Monetary Policy Event Database (EA-MPD) of Altavilla et al. 2020. These data are the change in asset prices in response to the ECB statements, recorded as the difference in the price of assets before and after a narrow window (around 45 minutes) of the ECB conference. The assumption is that monetary policy news should be the most important source of variation for these asset prices, given

---

<sup>9</sup>This issue is widely known in the literature surrounding LASSO specifications. Taddy (2017) cautions that (for the case of LASSO) cross-validation “can lead to over-fit for unstable algorithms whose results change dramatically in response to data jitter”. See also Gentzkow et al. (2019) for discussion of this issue.

the tightness of the time-span of the window. In this case our measures of the signal financial markets receive (in response to information from the press conference) should not be systematically related to other signals (for example signals about aggregate demand that do not come from discussions within the press conference).

We measure news as the absolute value of the change in yields; this gives a measure of the update in beliefs without direction mattering. Table 2 presents the bootstrapped mean Adjusted  $R^2$  values from post-LASSO OLS regressions of the absolute value of changes in yields on sets of variables selected by the LASSO estimation. The main message is that the temporal dimension is very relevant for to explain the news and both future and past temporal topics matter. This suggests that communication of both interpretation and projection seems important to address the markets' information deficit.

More specifically, regardless of the communication measure used, communication has its strongest explanatory power for news in yields of 1- to 3-year maturity. Explanatory power is lowest at the shorter and longer ends of the yield curve. This hump shape is present across all combinations of topic regressors (rows of the table).

The most basic measures of communication content, the topics  $\theta_k$ , capture some of the news systematically (Adjusted  $R^2$  between 0.15 and 0.3). This suggests that simply knowing what the press conference discussion is about is useful to pick up systematic variation in the market reaction.

Separately adding future and past topics increases the explanatory power (rows 2 and 3); the variation that the topics capture rises by about a quarter so for the 2-year OIS the explanatory power goes from 0.28 on adjusted  $R^2$  to 0.37 (future added) or 0.36 (past added). Nearer the short end, the past topics actually capture more of the variation whereas further out the yield curve, the future topics typically explain more.

The strongest effect comes from allowing the LASSO to select over topics from both past and future dimensions. Adding both the disaggregated temporal topics explains even more again suggesting that these temporal topics are not simply picking up the same thing. Moreover, perhaps because of the differences in references that each tagger is picking up, using separate temporal topics for each tagger gives the LASSO an even greater disaggregated set of topics to select from and allows us to capture even more of the variation (measured with adjusted  $R^2$  which penalises simply selecting more topics).

### 4.3 Effects of Temporal Communication on Risk Premiums

Several recent papers have drawn attention to the ways in which central bank statements appear to affect the risk-premium (e.g. Hansen et al. (2019), Leombroni et al. (2021)). In order to analyse the temporal effect on risk premiums, we follow an approach to their calculation that is broadly similar to Leombroni et al. (2021).

**Table 2:** Adjusted R-Squared of yield curve by specification (union)

Specification	OIS 1M	OIS 1Y	OIS 2Y	OIS 3Y	DE 5Y	DE 10Y
Topics Only	0.25	0.29	0.28	0.23	0.19	0.15
Topics and Future	0.32	0.36	0.37	0.31	0.25	0.19
Topics and Past	0.35	0.37	0.36	0.29	0.22	0.20
Topics, Future and Past	0.41	0.43	0.43	0.36	0.27	0.24
Topics, Future and Past*	0.59	0.62	0.63	0.59	0.56	0.47
Forecasts & Revisions	Yes	Yes	Yes	Yes	Yes	Yes

**Notes:** This table shows the bootstrapped mean Adjusted R-Squared values from post-LASSO Ordinary Least Squares regressions of the absolute value of changes in yields on sets of variables selected by the LASSO estimation. The bootstrap procedure is non-parametric with 500 draws and estimates an Adjusted R-Squared for each draw. Changes in yields are calculated in an intra-daily window around the ECB’s press conference. “Specification” refers to the set of variables available to be selected in the LASSO estimation. Each of Topics ( $\theta_k$ ), Future ( $\theta_k^{FUT}$ ) and Past ( $\theta_k^{PAST}$ ) contains 15 variables. Each specification includes forecasts of current year and one year ahead annual GDP growth and inflation, and revisions to these forecasts where applicable. The forecast variables are constrained to be always included in the set of selected variables.

These authors take the first principal component of a collection of movements in interest rate futures around press conferences of the ECB. They then regress the movements in equities provoked by press conferences onto their principal component, and take the residual as a measure of a “risk-premium shock”. The idea is that such a movement represents a component of the effect of central bank of communication on market appraisal of risk, since it is orthogonal (by construction) to the effects of central bank communication on the risk-free part of the yield curve (since it is orthogonal to OIS contract movements).

While the shock we create in this study is highly related to that of Leombroni et al. (2021), we differ slightly in our implementation of their procedure. In our data and sample we found that merely regressing equities on the first principal component of OIS contracts could lead to very low  $R^2$ , with even negative adjusted  $R^2$  in evidence. In the case of such low  $R^2$ , there would be little reason to use an orthogonalised risk-premium shock, as opposed to the direct movements in equities themselves. We achieved slightly higher  $R^2$  when we regressed movement equities on a group of five OIS contracts, and so adopt this procedure to achieve a risk-premium surprise.<sup>10</sup>

Table 3 shows the analysis of the risk premium. As before, the regression results suggest that (i) central bank messaging matters for the market news, (ii) the temporal dimension is important, and (iii) both the forward and the backward-looking information plays an important role. This suggests that the interpretation and context given

<sup>10</sup>The 5 contracts we use are one-month, three-month, 6-month, 1-year, 2-year.

to decisions, even if it covers information that is widely known and understood, is an important dimension of the central bank narrative.

**Table 3:** Adjusted R-Squared of Leombroni et al. (2021) Risk Premium surprise by specification

Specification	Risk Premium
Topics Only	0.17
Topics and Future	0.25
Topics and Past	0.29
Topics, Future and Past	0.36
Topics, Future and Past*	0.56
Forecasts & Revisions	Yes

**Notes:** This table shows the bootstrapped mean Adjusted R-Squared values from post-LASSO Ordinary Least Squares regressions of the absolute value of a risk premium surprise on sets of variables selected by the LASSO estimation. The bootstrap procedure is non-parametric with 500 draws and estimates an Adjusted R-Squared for each draw. The surprise is constructed from changes in equities and yields in an intra-daily window around the ECB’s press conference, following the identification scheme of Leombroni et al. (2021). The “Specification” refers to the set of variables available to be selected in the LASSO estimation. Each of Topics ( $\theta_k$ ), Future ( $\theta_k^{FUT}$ ) and Past ( $\theta_k^{PAST}$ ) contains 15 variables. The specification marked by \* indicates the use of disaggregated Future and Past measures where  $\theta_k^{FUT}$  and  $\theta_k^{PAST}$  contain 45 variables each. Each specification includes forecasts of current year and one year ahead annual GDP growth and inflation, and revisions to these forecasts where applicable. The forecast variables are constrained to be always included in the set of selected variables.

## 5 Identification Problem

### 5.1 The Identification Problem Facing Markets

The preceding analysis shows that both backward- and forward-looking information seems to convey important information. However, an important point to make is that though we can conclude that the information is important, we cannot conclude exactly what beliefs are updated as a result. As discussed in Hansen and McMahon (2018), there is a clear identification problem deriving from the high-dimensionality of the problem. This is why we do not focus particularly on whether the future or past topics are monetary policy or inflation or something else.

To see this, consider a market participant who knows the most-recently selected  $i_t$  but is unsure of the specific central bank assessment of the economy  $\Omega_t$  nor their specific reaction function  $f_t(\cdot)$ . The identification problem is the fact that it is not possible to pin down the elements of the assessment and reaction function. Based on a particular market belief of the central bank assessment,  $\tilde{\Omega}_t$ , and given the selected interest rate, there is an

implied reaction function,  $\tilde{f}_t(\cdot)$ , such that  $i_t = \tilde{f}_t(\tilde{\Omega}_t)$ . Alternatively, if they have a strong view about  $\tilde{f}_t(\cdot)$ , then it will imply a belief about the state of the economy  $\tilde{\Omega}_t$ .

If the central bank were to subsequently provide more clarity on their assessment of the state of the economy, perhaps through a backward-looking description of their conjunctural analysis, this backward-looking information would provide information on  $\tilde{\Omega}_t$  which also requires an update of  $\tilde{f}_t(\cdot)$ . For example, if markets learn that policymakers perceive the state of the economy is better than markets had expected, then the only justification for the current interest rate is that policymakers are more dovish than previously assumed.

In fact, while the news on the state of the economy is what was communicated, the impact from the updated view of the reaction function might be even more important. This would follow if we perceived reaction function preferences as more persistent than the state of the economy. If we assumed, as an extreme example, that the reaction function shifted according to a unit root, then this update would have important implications for the entire yield curve whereas the economic news may only have relevance for the state of the economy over the next few quarters.

This means that while we can explore the effect of communicated information on market beliefs, we cannot assume that backward looking information only impacts beliefs about the state of the economy. In other words, our approach allows us to say something about the nature of information, but it does not allow us to conclusively understand what beliefs that information caused to update. To see this, and to see that backward-looking does, in fact, seem to lead beliefs on monetary reaction functions to update, we shall now extend our earlier analysis to market news that has been decomposed into different constituent parts.

## 5.2 Monetary Policy Surprise Analysis

### 5.2.1 Data

Recent contributions in the monetary event study literature have followed Gürkaynak et al. (2005) and sought to decompose asset price movements into multiple forms of surprise, according to structural criteria. This study employs a number of ECB monetary policy surprise series from leading recent papers in the literature.

The first set of surprises used in this study are the four surprise series of Altavilla et al. (2019): the “target” surprise, the “timing” surprise, the “forward guidance” surprise and the “QE surprise”. The first surprise is extracted only from asset price movements around the press release of the ECB, which is posted on the ECB website around 30 minutes before the press conference itself. Altavilla et al. (2019) argue that this surprise reflects a

conventional surprise to interest rates. The latter three surprises (timing/forward guidance/QE) are extracted from the press conference window. Of these three surprises, the forward guidance surprise and the QE surprise are assumed not to load on short-term interest rate futures contracts. A further assumption is made that the QE surprise explains minimal variation in the post-crisis period. The timing surprise is allowed to load on any asset price, however it has only a weak relation to the shorter and longest-term rates. Altavilla et al. (2019) argue that this surprise represents a form of update about the timing of ECB decisions in the near future, as opposed to what we might understand by intentional rate guidance about the medium run.

Because our data sample for our empirical estimation runs until December 2019, we update the shock series of Altavilla et al. (2019). We do so using the codes provided on their website, in conjunction with data from the EA-MPD.

The second approach to identifying monetary policy surprises used in this paper follows the approach Jarociński and Karadi (2020) use to separate the monetary policy surprise and the “information surprise”. These authors also use information from high-frequency movements in asset prices around policy announcement days (in their case, they study both the US and euro area cases). These authors are interested in controlling for a component of the announcement that may relate to the communication of information about macro-economic variables and their expected progression. To do this the authors use a sign-restrictions approach, restricting true monetary policy surprises to raise interest rates, but reduce equities. An information surprise, on the other hand, is constrained to raise interest rates, but also to raise equities (in that sense an information surprise is “good news” when it raises rates, and “bad news” when it lowers rates).

It is important to emphasise that in this study we do not strictly use the exact surprise series of Jarociński and Karadi (2020). The reason is that Jarociński and Karadi (2020) identify their policy surprises in a one-step estimation procedure, within a monthly Bayesian VAR model, which incorporates zero-restrictions on the relation between high-frequency and low-frequency data. Because there were occasions where there was more than one policy statement within a month, and because Jarociński and Karadi (2020) cumulate their surprises prior to applying their identification procedure, an exact mapping of the monthly surprise series provided in the replication materials of Jarociński and Karadi (2020) and the meetings of the ECB does not exist for all months in the sample. We therefore prefer to apply the sign-restrictions of Jarociński and Karadi (2020) to information that is derived solely from the high-frequency data, without applying zero-restrictions to macro-economic series. We also use information from the EA-MPD, whereas Jarociński and Karadi (2020) created a novel dataset of their own. We do use the same two contracts as Jarociński and Karadi (2020), namely the two-month OIS rate and

the Eurostoxx50 index. However, the sign-restrictions we employ are essentially the same as Jarociński and Karadi (2020), and therefore we refer to these surprises as “Jarociński and Karadi (2020) style” shocks throughout the paper.

### 5.2.2 Analysis of Decomposed Surprises

Ignoring the identification problem, you might expect that the backward-looking information, when applied to the Altavilla et al. (2019) decomposition, should have relatively more importance than forward-looking information for the Target factor, and likely less importance for Forward Guidance factor. Timing and Quantitative Easing (QE) factors would be somewhere in between. Table 4 shows the results from extending our Bayesian bootstrap analysis to the Altavilla et al. (2019) decomposition.

The basic result of the earlier analysis remains - both forward and backward looking temporal topics play an important role in explaining news variation. However, consistent with the difficulty of mapping information to specific beliefs being updated, we find that past topics are relatively *least* important for the target factor. This suggests, though does not necessarily prove, that the context may indeed be providing information on the reaction function of the central bank to market participants.

**Table 4:** Adjusted R-Squared of Altavilla et al. (2019) surprises by specification

Specification	Target	Timing	Forward Guidance	Quantitative Easing
Topics Only	0.23	0.28	0.29	0.21
Topics and Past	0.29	0.36	0.33	0.26
Topics and Future	0.37	0.36	0.34	0.27
Topics, Future and Past	0.42	0.42	0.37	0.32
Topics, Future and Past*	0.59	0.62	0.63	0.59
Forecasts & Revisions	Yes	Yes	Yes	Yes

**Notes:** This table shows the bootstrapped mean Adjusted R-Squared values from post-LASSO Ordinary Least Squares regressions of the absolute value of the surprises constructed in Altavilla et al. (2019) on sets of variables selected by the LASSO estimation. The bootstrap procedure is non-parametric with 500 draws and estimates an Adjusted R-Squared for each draw. The surprises are constructed from changes in yields in an intra-daily window around the ECB’s press conference. For further discussion, see Altavilla et al. (2019). The “Specification” refers to the set of variables available to be selected in the LASSO estimation. Each of Topics ( $\theta_k$ ), Future ( $\theta_k^{FUT}$ ) and Past ( $\theta_k^{PAST}$ ) contains 15 variables. The specification marked by \* indicates the use of disaggregated Future and Past measures where  $\theta_k^{FUT}$  and  $\theta_k^{PAST}$  contain 45 variables each. Each specification includes forecasts of current year and one year ahead annual GDP growth and inflation, and revisions to these forecasts where applicable. The forecast variables are constrained to be always included in the set of selected variables.

We can provide further evidence of this by exploring the effect on Jarociński and Karadi (2020) surprises by specification. Table 5 presents the results. As before, both

temporal topics systematically contain information to explain each shock, and the joint inclusion captures the most variation. In line with a reasonable prior belief that context is mostly about the economic state, backward-looking information is most informative about identified information shocks. Nonetheless, they also play an important role in explaining the news in identified policy shocks. Forward looking information is very important for both of the decomposed shocks too.

**Table 5:** Adjusted R-Squared of Jarociński and Karadi (2020) surprises by specification

Specification	Information	Monetary
Topics Only	0.21	0.23
Topics and Past	0.32	0.28
Topics and Future	0.36	0.32
Topics, Future and Past	0.45	0.37
Topics, Future and Past*	0.56	0.65
Forecasts & Revisions	Yes	Yes

**Notes:** This table shows the bootstrapped mean Adjusted R-Squared values from post-LASSO Ordinary Least Squares regressions of the absolute value of an information surprise and a monetary policy surprise on sets of variables selected by the LASSO estimation. The bootstrap procedure is non-parametric with 500 draws and estimates an Adjusted R-Squared for each draw. The surprises are constructed from changes in equities and yields in an intra-daily window around the ECB’s press conference, following the identification scheme of Jarociński and Karadi (2020). The “Specification” refers to the set of variables available to be selected in the LASSO estimation. Each of Topics ( $\theta_k$ ), Future ( $\theta_k^{FUT}$ ) and Past ( $\theta_k^{PAST}$ ) contains 15 variables. The specification marked by \* indicates the use of disaggregated Future and Past measures where  $\theta_k^{FUT}$  and  $\theta_k^{PAST}$  contain 45 variables each. Each specification includes forecasts of current year and one year ahead annual GDP growth and inflation, and revisions to these forecasts where applicable. The forecast variables are constrained to be always included in the set of selected variables.

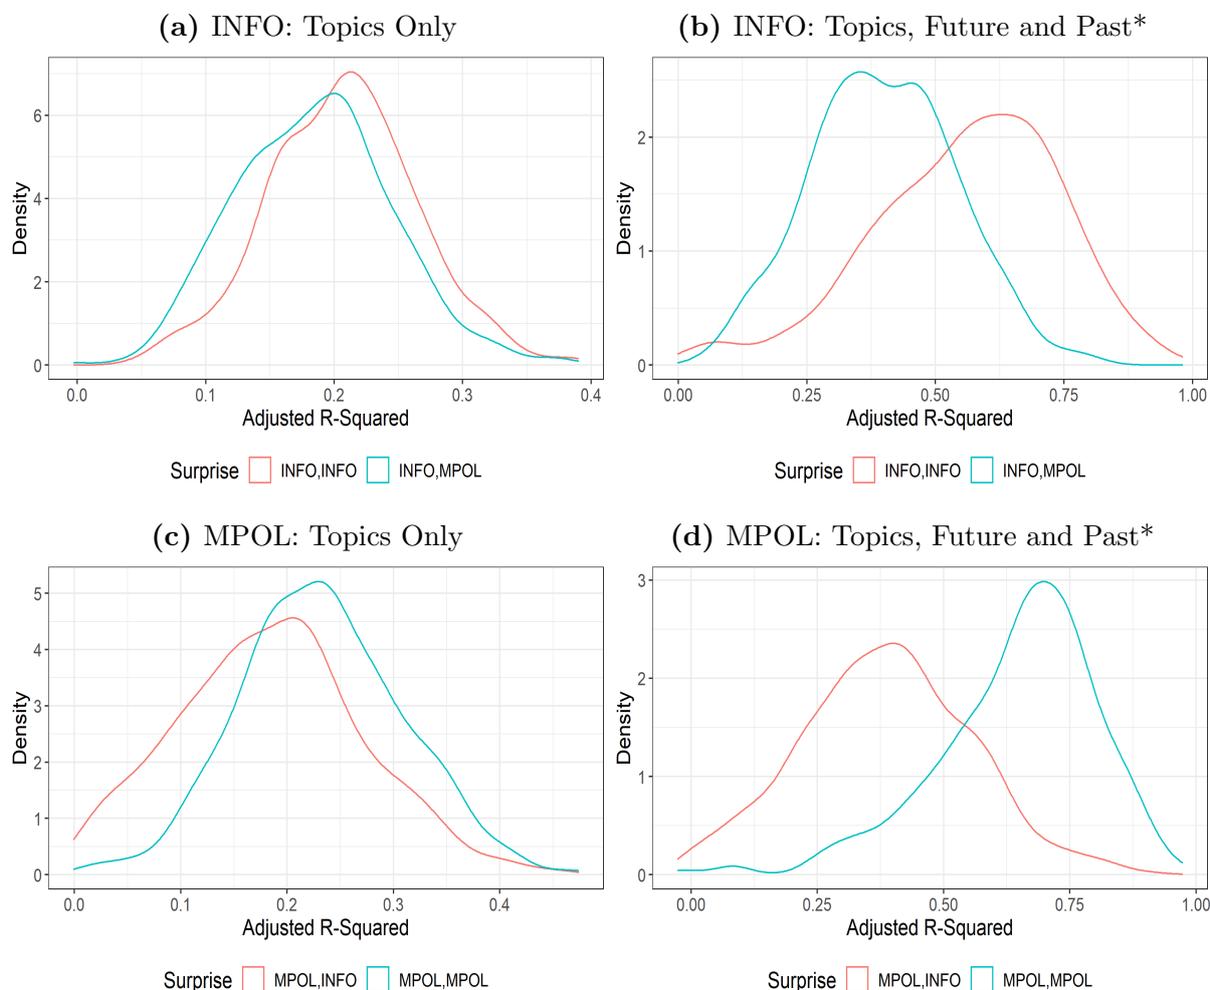
Figure 6 plots the full distributions of Adjusted R-Squared values from Post-LASSO Ordinary Least Squares regressions that underpin Table 5. The left column shows the regressions for the exercise in which only topics without any temporal dimension are included in the topic set (“Topics Only” row of the table), while the right column shows the most disaggregate results with both future and past temporal topics (“Topics, Future and Past\*” row).<sup>11</sup> The top row, Figures 7(a) and 7(b), show the results for the information surprise, “INFO”, and each figure includes the findings from regressing on the set of variables selected by LASSO for the information surprise itself and the set selected for the monetary policy surprise (“MPOL”) for the same bootstrap draw. Figures 7(c) and 7(d) repeat these results but where the monetary policy surprise (“MPOL”) is the dependent variable.

Two main results stand out. The first is that the topics, while capturing broad themes,

<sup>11</sup>Forecast control variables are constrained to be always included in the set of selected variables.

are too broad to capture a lot of the market news. This is reflected in both the fact the Adjusted  $R^2$  is relatively low at around 0.2 (as in the Table). The second is that there is little extra explanatory power from using the topics selected by LASSO for one surprise in explaining the other surprise – the distributions are very similar. When we instead use the disaggregated temporal shocks, the Adjusted  $R^2$  jumps up to around 0.6-0.7 for the own shocks, and the specificity of the selected is much greater – using the topics selected for the other shock, almost halves the explanatory power. This shows that the communication that moves markets is best captured in high-dimensional measures of the messaging rather than simply broad themes.

**Figure 6:** Contribution of temporal variables to identifying distinct information sets



**Notes:** Figure 7(a) shows the distributions of Adjusted R-Squared values from Post-LASSO Ordinary Least Squares regressions of the information surprise (“INFO”) on the set of variables selected by LASSO for the information surprise itself and the set selected for the monetary policy surprise (“MPOL”) for the same bootstrap draw. The set of variables available to be selected includes the Topics only ( $\theta_k$ ). Figure ?? shows the distributions of Adjusted R-Squared across the bootstrap draws when the set of variables available to be selected has Topics ( $\theta_k$ ), Future ( $\theta_k^{FUT}$ ) and Past ( $\theta_k^{PAST}$ ). Each specification includes forecasts of current year and one year ahead annual GDP growth and inflation, and revisions to these forecasts where applicable. The forecast variables are constrained to be always included in the set of selected variables.

## 6 Speeches Analysis

As a final step, we extend our analysis to ECB executive board member speeches described in Section 3.1. The speeches themselves are not as functionally important as the Governing Council press conference, and can cover a wide array of non-monetary policy topics. Nonetheless, some speeches, such as Mario Draghi’s “whatever it takes” speech have had large scale market impact. Further to this, speeches can play a role in clarifying comments made in the previous press conference, or indeed can be used to steer the conversation ahead of upcoming monetary policy decisions. The temporal aspect of the speeches is thus of particular importance, owing to both the clarification and guiding role that these can play.

### 6.1 Speech Yields Data

When studying the effect of the speeches on financial market variables, we use daily (end-of-day) series, downloaded from Bloomberg. There are several reasons for this. The first is that many of the speeches are given outside market trading hours, meaning the construction of an intra-daily movement is impossible. The second is that we do not know exactly when the information contained within the speeches became generally available to markets, since this information is not recorded in the dataset. For these reasons we use a two-day window around the speeches. This means that if a speech occurs after trading hours on a given day, the signal from this speech is still recorded (a one-day window would preclude this). However, for the speeches that occur during trading hours we also use a two-day window, to ensure the window is consistent for different speeches.<sup>12</sup> For our empirical specifications, we drop 90 speech-day observations that fall on Saturdays or Sundays, and we also drop 150 speech-day observations that fall on Governing Council meeting days, the day before, or the day after. This means the total number of speech-day observations available for analysis is 1,473.

### 6.2 Basic speech regressions that back up our previous findings

Table 2 repeats the analysis of Table 2 and 3 but applied to the speeches. The basic findings hold though the explanatory power is generally much weaker. This likely reflects two drivers. First, using the 2-day window introduces more non-speech news in the surprise measures and, therefore, the speech content should be unrelated to these movements limiting the amount of variation the information can explain. The second is

---

<sup>12</sup>Note that for a speech delivered on a Friday, we employ a window from market-close on Thursday to market-close on Monday. We account for potential heterogeneity in the treatment of speeches across days using day-of-the-week fixed effects in our empirical specifications.

that if speeches are, on average, less central to market monetary-policy belief formation. This may be because they are not always about monetary policy. Or, even when they are about monetary policy, they are less likely to introduce new information and may often just restate the opening statement which has already been communicated (and hence may not get the same amount of attention as a result).

**Table 6:** Adjusted R-Squared of Yield Curve by Specification – ECB Speeches

	OIS 1M	OIS 1Y	OIS 3Y	DE 5Y	DE 10Y	RP
Topics	0.177	0.260	0.263	0.221	0.115	0.062
Topics and Future	0.190	0.267	0.269	0.227	0.126	0.076
Topics and Past	0.190	0.268	0.272	0.232	0.127	0.078
Topics, Future and Past	0.203	0.276	0.278	0.237	0.138	0.091
Topics, Future and Past*	0.245	0.323	0.326	0.287	0.207	0.158

Notes: This table shows the bootstrapped mean Adjusted R-Squared values from post-LASSO Ordinary Least Squares regressions of the absolute value of changes in yields on sets of variables selected by the LASSO estimation. The bootstrap procedure is non-parametric with 500 draws and estimates an Adjusted R-Squared for each draw. Changes in yields are calculated in a two-day window around ECB speeches. Specification refers to the set of variables available to be selected in the LASSO estimation. Each of Topics ( $\theta_k$ ), Future ( $\theta_k^{fut}$ ) and Past ( $\theta_k^{pst}$ ) contains 15 variables. The specification marked “Topics, Future and Past\*” does not aggregate information from numerical, categorical, and tense temporal references. It creates individual temporal topics for each of these three forms of temporal reference. Therefore this specification allows for 90 ( $3 * (2 * 15)$ ) potential temporal topics, as well as 15 overall topic shares. Each specification includes a series of control variables, detailed in the main text. Control variables are constrained to always be included in the set of selected variables.

### 6.3 Speeches Addressing the Information Deficit

As already discussed, the impact of any communication is going to be related to its newsworthiness. A speech which addresses market participants’ *information deficit* is more likely to generate news and lead beliefs to update. This may be clarification of things which the market is unclear about, or updates to the latest thinking of the central bank on the state of the economy or the reaction of monetary policy. However, it is very hard to specifically identify what the market doesn’t know in order to see whether this is the case.

Using the Q&A part of the ECB press conferences, we propose a novel measure of this information deficit of speeches based on the similarity of the information contained within speeches and the questions from the press conference held immediately previous to the speech. Our measure is based on the assumption that the questions, albeit asked

by financial journalists, are addressed to the information deficit. That is, the questions are assumed to highlight the issues that journalists wish for clarification on. While there are always things that could be made clearer, we are also implicitly assuming that the most pressing gaps are addressed first.

For our speech-specific measure, we measure the issues that make up the information deficit following a given Governing Council meeting using the content of the questions in the associated press conference. We then will measure whether Governing Council members' speeches address the same issues. Of course, ECB officials can address these questions immediately through direct answers in the press conference so we will also assess the extent to which they do this. (Though, of course, answering a question on a prevalent topic does not necessarily exhaust all interest in the issue which be multi-faceted or very complicated.)

Specifically, we compute our similarity measure using the document term matrix, where a document is a speech or the questions.<sup>13</sup> Terms are weighted by the term frequency-inverse document frequency (TF-IDF) score. Using this weighted document-term matrix, and letting  $dtm_i$  be a vector corresponding to the row of the weighted document term matrix associated with document  $i$ , we calculate the similarity between two documents  $i$  and  $j$  as:

$$Sim_{i,j} \equiv \frac{\sum_{k=1}^n (dtm_{ik} \times dtm_{jk})}{\sqrt{\sum_{k=1}^n dtm_{ik}^2} \times \sqrt{\sum_{k=1}^n dtm_{jk}^2}} \quad (6)$$

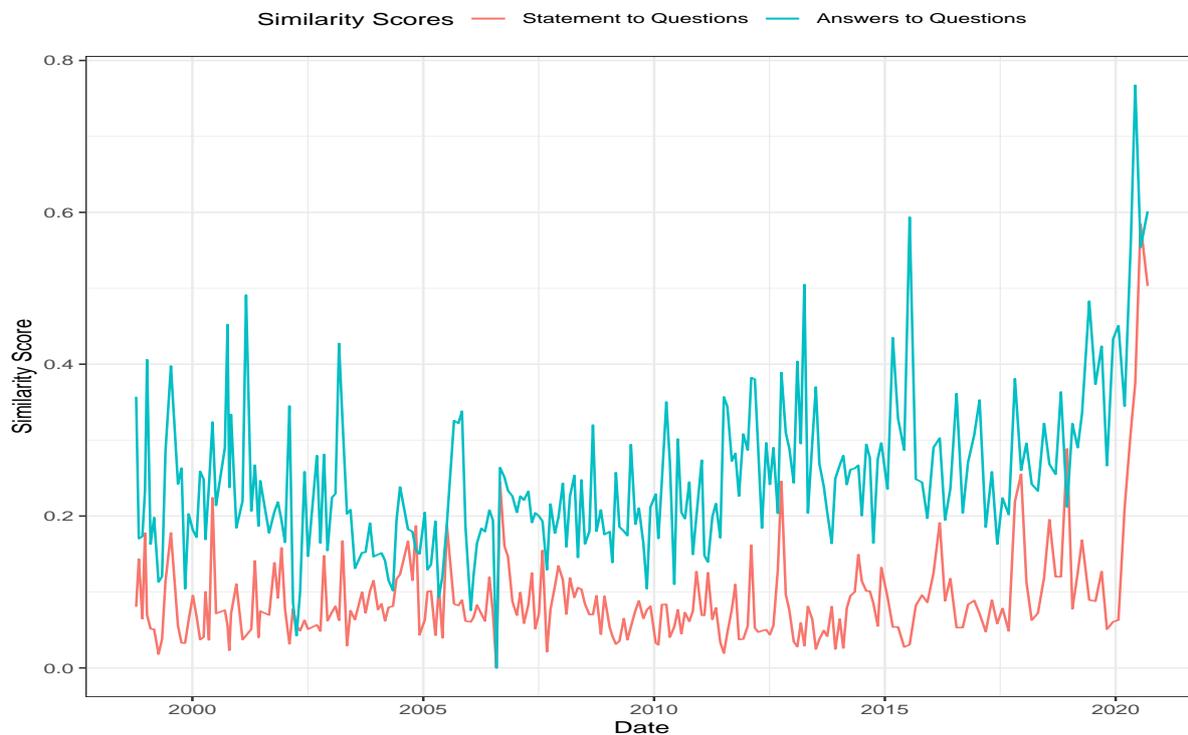
where  $n$  is the total number of terms in the document term matrix, and  $dtm_{ik}$  refers to the  $k$ th term in the vector corresponding to document  $i$ .

Figure 8 show the similarity scores for statements ( $Sim_{S,Q}$ ) and answers ( $Sim_{A,Q}$ ) when both are compared to the questions in associated press conference. A number of trends emerge. First, the similarity of the questions to the answers is consistently higher than all other measures. This provides an internal validity check for the measure as the answers should, in theory, directly address the topics in the questions. Second, we note a sharp increase towards the end of the sample. This is primarily driven by the onset of

---

<sup>13</sup>Prior to calculation, we clean steps are taken in line with the steps for the topic model, excluding the creation of the set of n-grams. In addition, due to the large sparse nature of the weighted document term matrix, an additional step is taken to remove the most rare and most frequent terms (those that occur in less than 1% of the documents and greater than 50% of the documents respectively). Speeches that occur on the same day are treated as one long speech. This is due to the fact that the windows of observation in our event study are two-day windows, and as such any movements can be attributed to any speech on that given day. All speeches that occur on the same day as a Governing Council press conference are dropped from the sample for the calculation of similarity indices.

the Covid-19 crisis in 2020, where the topics of all discussions became polarised towards one specific topic.

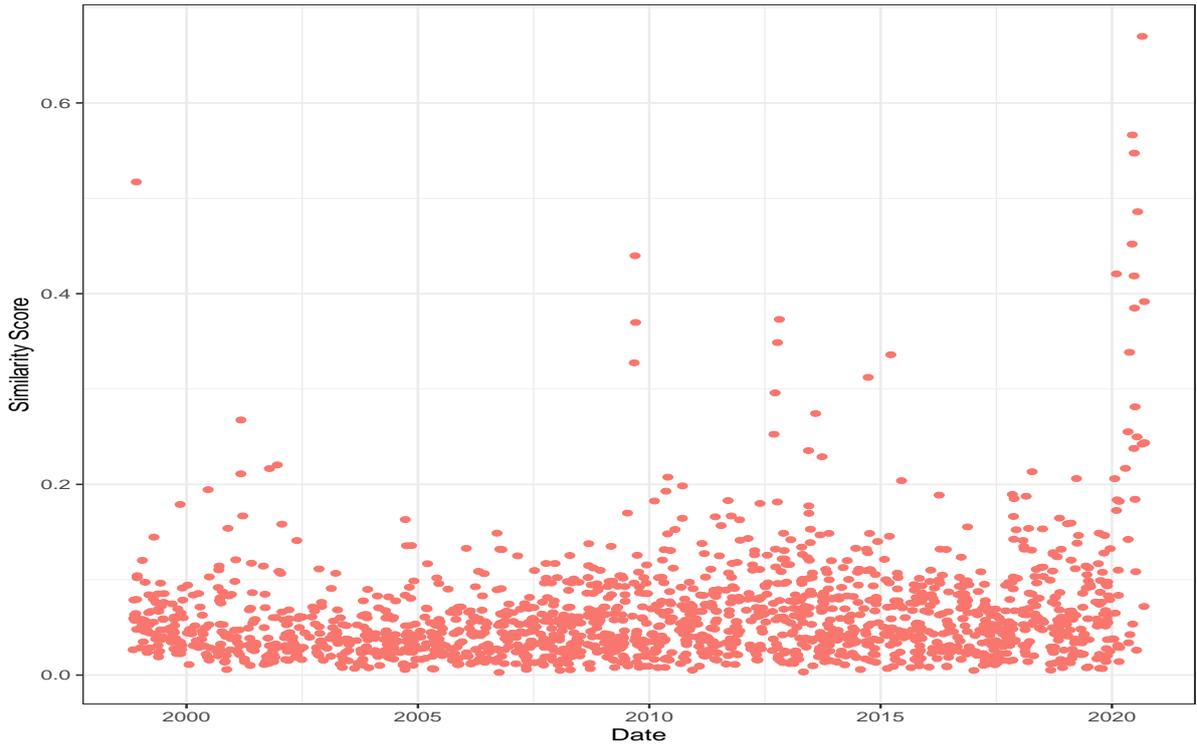


**Figure 8:** Similarity Scores for Press Conference Questions compared to Statements and Answers

Notes: This figure shows a plot of the similarity measures as shown in Equation 6 comparing the content of the questions in a given ECB Governing Council press conference with the text of the accompanying answers and statements

Figure 9 plots the similarity for speeches compared to the questions from the previous press conference ( $Sim_{sp,Q}$ ). There is a measure for each individual speech day but they are stacked on the date of the Governing Council meeting from which the questions are derived. The Covid effect is also present in speeches in 2020. Moreover, other clusters of increased similarity are seen across the sample. A notable example is in late 2009, when a large number of speeches were made discussing the path forward of ECB monetary policy in the context of the crisis measures implemented during the financial crisis. In addition, it is noteworthy that speeches that contain references to “hearing” and “European Parliament” in their titles have, on average, a much larger recorded similarity score (0.102 compared to 0.062). This highlights the extent to which such meetings can be used to bridge any information deficit.

We now use our measures of similarity in a regression analysis. The idea is to regress asset price news associated with a speech, the same dependent variable in Table 6, on



**Figure 9:** Similarity Scores for Speeches compared to Press Conference Questions

Notes: This figure shows a plot of the similarity measures as shown in Equation 6 comparing the content of each speech by an ECB board member with the questions from the previous Governing Council press conference

whether the speech addresses the information deficit measured using our similarity measures. The hypothesis is that speeches addressing the topics of the questions should be more likely to give rise to news. We also include the similarity of the press conference answers to the press conference questions to capture the already provided information, as well as its interaction with  $Sim_{Sp,Q}$ .

Table 7 presents the regression results and supports the idea that speeches addressing the information deficit give rise to more market news. However, the extent of this news may be reduced when speeches are similar to the answers given in the press conference. These effects are strongest at the 1- to 5-year maturity range.

Table 8 shows the main results are robust to including similarity of statements to questions ( $Sim_{S,Q}$ ), as well as the interaction of this variable with  $Sim_{Sp,Q}$ .

## 7 Conclusion

Over recent decades, central banks have increased the depth, range and frequency of their communication with the public. Central bankers give speeches to the public, respond to

**Table 7:** Speech-Question similarity and Information Deficit

	OIS 1M	OIS 1Y	OIS 3Y	DE 5Y	DE 10Y
$Sim_{Sp,Q}$	0.17 (5.18)	12.47** (5.32)	18.90** (7.51)	18.47** (8.77)	14.70* (8.57)
$Sim_{A,Q}$	0.93 (1.47)	1.70 (1.51)	2.78 (2.21)	5.11** (2.49)	4.92** (2.44)
$Sim_{Sp,Q} \times Sim_{A,Q}$	-10.5 (18.65)	-42.7** (19.16)	-63.5** (26.81)	-71.1** (31.55)	-56.7* (30.84)
Constant	-9.79 (7.27)	4.84 (7.47)	1.28 (10.44)	4.92 (12.29)	9.40 (12.02)
Speaker FE	Yes	Yes	Yes	Yes	Yes
Macro controls	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Day of Week FE	Yes	Yes	Yes	Yes	Yes
Topics, Future and Past	Yes	Yes	Yes	Yes	Yes
N	1246	1246	1113	1249	1249
Adj. R-Squared	0.22	0.34	0.30	0.15	0.07

**Notes:** This table shows OLS estimates from regressions of the absolute value of the two day change in yields around a speech on measures of the speech similarity with the press conference questions and the Information Deficit. “ $Sim_{Sp,Q}$ ” and “ $Sim_{A,Q}$ ” denote the similarity of a speech and of the press conference answers to the questions, respectively. Fixed effects are included for the Executive Board member, for the year and for the day of the week in which the speech was given. Six macroeconomic surprises are included, five for the euro area and one for the U.S. The statistical significance level is displayed as \* $p < 0.1$ ; \*\* $p < 0.05$ ; \*\*\* $p < 0.01$ .

queries from the media, interpret recent economic developments, project their intended policy path, and provide forecasts of key economic variables. A number of studies have shown that central bank communication can measurably surprise, or generate news. For communication to be effective in this way, it must, at least in part, fill some information deficit on the part of the public. The exact nature of the information deficit has been less studied in the literature.

Earlier studies supposed that the information deficit was related to private information on the part of the central bank. However, markets have broadly the same access to data that the central bank has, reducing the scope for substantively new information of this type to be released by the central bank. In this paper, we argue instead that the source of the information advantage that generates the news in central bank communication comes from the central bank updating either its assessment of the current state of the economy, or its mapping of the state into the appropriate monetary policy stance.

**Table 8:** Speech-Question similarity and Information Deficit - Robustness to Statements

	OIS 1M	OIS 1Y	OIS 3Y	DE 5Y	DE 10Y
$Sim_{Sp,Q}$	0.44 (5.05)	9.06* (5.17)	17.71** (7.23)	15.17* (8.47)	14.42* (8.29)
$Sim_{A,Q}$	1.65 (1.53)	2.33 (1.56)	4.18* (2.26)	6.51** (2.56)	5.60** (2.50)
$Sim_{S,Q}$	-1.0 (2.41)	-5.1** (2.46)	-5.1 (3.39)	-10.8*** (4.03)	-7.8** (3.95)
$Sim_{Sp,Q} \times Sim_{A,Q}$	-21.5 (19.99)	-44.0** (20.47)	-75.2*** (28.43)	-79.4** (33.53)	-57.8* (32.81)
$Sim_{Sp,Q} \times Sim_{S,Q}$	25.54 (26.02)	31.73 (26.64)	39.13 (35.68)	53.70 (43.64)	17.71 (42.70)
Constant	1.47** (0.69)	5.33*** (0.70)	3.41*** (0.94)	5.73*** (1.15)	4.59*** (1.13)
Speaker FE	Yes	Yes	Yes	Yes	Yes
Macro controls	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Day of Week FE	Yes	Yes	Yes	Yes	Yes
N	1247	1247	1114	1250	1250
Adj. R-Squared	0.21	0.34	0.31	0.16	0.08

**Notes:** This table shows OLS estimates from regressions of the absolute value of the two day change in yields around a speech on measures of the speech similarity with the press conference questions and the Information Deficit. “ $Sim_{Sp,Q}$ ”, “ $Sim_{A,Q}$ ” and “ $Sim_{S,Q}$ ” denote the similarity of a speech, of the press conference answers and of the Introductory Statement to the questions, respectively. Fixed effects are included for the Executive Board member, for the year and for the day of the week in which the speech was given. Six macroeconomic surprises are included, five for the euro area and one for the U.S. The statistical significance level is displayed as \* $p < 0.1$ ; \*\* $p < 0.05$ ; \*\*\* $p < 0.01$ .

While central banks may have no advantage purely in terms of methodology or training than markets, they often have an advantage in terms of the resources: the number of staff or models available to generate the assessment of the state. Hence the central bank can devote resources to *Interpretation* of economic data - choosing the weights to put on different data sources at each point in time and how to place developments in the data into context. This process is inherently backward-looking.

Central banks also communicate forward-looking information, given the inherent nature of monetary policy. They engage in *Projection* exercises - providing some outlook for the expected evolution of a variable or for their policy reaction to the state. This is another potential source of news - the central bank could surprise markets by indicating

a more or less aggressive policy response in the future. We argue that forward-looking communication by central banks is multi-faceted and is not reducible to the numerical forecast data that they also publish.

To capture these Interpretation and Projection features of central bank communication requires measuring the temporal dimension of central bank communication. In the previous literature, attention has chiefly been paid to the Topic and/or Tone of the communication. This paper provides a methodology to capture the Third “T” - Time. We include an algorithm to measure the time orientation of speech using Natural Language Processing methods. Using an event study methodology, we show that our temporal measures significantly increase the explicability of measures of news around the press conference following ECB Governing Council meetings. Measures of past and future content contribute in roughly similar amounts and are not substitutes. This supports the view that both the Interpretation and Projection dimensions matter - central bank communication is not just informative through forward-looking information. We also show that these text-based measures contribute beyond numerical-only forecast data, indicating the importance of a multi-faceted communication approach rather than solely publishing forecasts without accompanying context.

To support the view that central bank communication is informing by filling an information deficit on the part of the public, we measure the similarity of the text of ECB speeches, and of the answers provided by the President, to questions from the media during the press conference. We measure the post-press conference information deficit through the distance between the questions asked and the answers received. The ECB would have an opportunity to fill this deficit through speeches in the weeks following the Governing. We find that the more closely a speech then matches the media’s questions, the greater the news embedded in yield changes in a window around that speech. However, the contribution of a speech is diminished the smaller the extant information deficit after the press conference, supporting the view that communication matters when it is informative on topics about which the public would like to learn.

## References

- Carlo Altavilla, Luca Brugnolini, Refet S. Gürkaynak, Roberto Motto, and Giuseppe Ragusa. Measuring euro area monetary policy. *Journal of Monetary Economics*, 108: 162–179, 2019.
- David M. Blei, Andrew Y. Ng, and Michael I. Jordan. Latent dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022, March 2003. ISSN 1532-4435. URL <http://dl.acm.org/citation.cfm?id=944919.944937>.
- David Byrne, Robert Goodhead, Michael McMahon, and Conor Parle. The 3t’s of central bank communication: A technical guide. Working paper, Mimeograph, May 2021.
- Angel X. Chang and Christopher D. Manning. SUTIME: A Library for Recognizing and Normalizing Time Expressions. Technical report, 8th International Conference on Language Resources and Evaluation, 2012.
- Mario Draghi. Introductory statement, press conference, 12 september 2019. <https://www.ecb.europa.eu/press/pressconf/2019/html/ecb.is190912~658eb51d68.en.html>, 2019.
- Matthew Gentzkow, Bryan Kelly, and Matt Taddy. Text as Data. *Journal of Economic Literature*, 57(3):535–574, September 2019.
- Mark Gertler and Peter Karadi. Monetary policy surprises, credit costs, and economic activity. *American Economic Journal: Macroeconomics*, 7(1):44–76, 2015.
- Refet S Gürkaynak, Brian Sack, and Eric Swanson. Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements. *International Journal of Central Banking*, 1(1), May 2005.
- Stephen Hansen and Michael McMahon. *How central bank communication generates market news*, chapter 15, pages 71–162. Handbook of Macroeconomics. VoxEU, February 2018.
- Stephen Hansen, Michael McMahon, and Andrea Prat. Transparency and Deliberation within the FOMC: A Computational Linguistics Approach. *The Quarterly Journal of Economics*, 2017. Forthcoming.
- Stephen Hansen, Michael McMahon, and Matthew Tong. The long-run information effect of central bank communication. *Journal of Monetary Economics*, 108:185–202, 2019.

- Marek Jarociński and Peter Karadi. Deconstructing monetary policy surprises?the role of information shocks. *American Economic Journal: Macroeconomics*, 12(2):1–43, April 2020. doi: 10.1257/mac.20180090. URL <https://www.aeaweb.org/articles?id=10.1257/mac.20180090>.
- Kenneth N. Kuttner. Monetary policy surprises and interest rates: Evidence from the fed funds futures market. *Journal of Monetary Economics*, 47(3):523–544, June 2001.
- Matteo Leombroni, Andrea Vedolin, Gyuri Venter, and Paul Whelan. Central Bank Communication and the Yield Curve. *Journal of Financial Economics*, forthcoming, 2021.
- James Pustejovsky, José Castaño, Robert Ingria, Roser Saurí, Robert Gaizauskas, Andrea Setzer, Graham Katz, and Dragomir Radev. TimeML: Robust Specification of Event and Temporal Expressions in Text. Technical report, AAAI, 2003.
- Anita Ramm, Sharid Loaíciga, Annemarie Friedrich, and Alexander Fraser. Annotating Tense, Mood and Voice for English, French and German. Technical report, Proceedings of ACL 2017, System Demonstrations, 2017.
- Matt Taddy. One-Step Estimator Paths for Concave Regularization. *Journal of Computational and Graphical Statistics*, 26(3):525–536, 2017.
- Hui Zou and Trevor Hastie. Regularization and Variable Selection via the Elastic Net. *Journal of the Royal Statistical Society*, 67:301–320, 2005. Part 2.

**Table 9:** Intra-Daily Monetary Policy Surprises and their Relation to the Yield Curve

	OIS 1M	OIS 3M	OIS 6M	OIS 1Y	OIS 2Y	OIS 3Y	DE 5Y	DE 10Y	Equities
Target	1.00	0.75	0.63	0.50	0.37	0.34	0.17	0.02	-0.04
Timing	0.33	0.83	1.00	1.17	1.02	0.75	0.67	0.39	-0.00
FG	0.00	0.17	0.41	0.75	1.00	0.91	0.97	0.48	-0.02
QE	0.00	0.02	-0.02	0.02	0.22	0.48	0.89	1.11	-0.01
MPOL	0.30	1.10	1.49	1.95	2.06	1.80	1.73	1.00	-0.29
INFO	0.33	1.04	1.28	1.66	1.73	1.41	1.35	0.71	0.36
RP	0.00	-0.00	-0.00	-0.00	-0.00	-0.32	-0.20	-0.01	1.00

Notes: The timing, forward guidance (FG) and quantitative easing (QE) surprises are updated versions of those used in Altavilla et al. (2019). The monetary policy (MPOL) and information (INFO) are derived from sign-restrictions equivalent to those applied in Jarociński and Karadi (2020). The risk-premium surprise is created as the component of the response of the movement in Eurostoxx50 that is orthogonal to 6 OIS movements (between one month and two years). The dependent variables are taken the EA-MPD dataset of Altavilla et al. (2019). For the case of the target shock, the dependent variables are from the press-release window of the EA-MPD. For all other shocks, the dependent variables are from the press-conference window.

**Table 10:** Topics from LDA topic model fit on the statements and answers corpus, terms ordered by weight within topic

1. Risk	2. Structural Reforms	3. Prices and Inflation	4. Banks, Markets, Bonds	5. EA Growth	6. Price Stability	7. Price Stability and Expectations	8. Gov Co-Logistics
risk	euro_area	price	bank	euro_area	price_stabil	price_stabil	meet
uncertainiti	structur_reform	increas	measur	contin	monetari_polic	inflat_expect	govern_council
econom	market	inflat	market	support	govern_council	medium_term	ecb
euro_area	competit	effect	bond	growth	risk	close	confer
outlook	economi	expect	monetari_polic	expect	monitor	line	presid
relat	increas	year	programm	demand	develop	inflat	outcom
financi_market	countri	month	certain	loan	decis	deliv	introductori
develop	need	oil_price	differ	remain	close	mandat	decis
high	growth	inflat_rate	liquid	recoveri	stanc	maintain	staff
growth	product	energi	effect	credit	assess	continu	project
downsid_risk	improv	reflect	actual	ongo	continu	necessari	today
global	employ	project	use	improv	inform	level	next
remain	strengthen	higher	credit	domest	medium_term	anchor	pleas
impact	order	come	countri	sector	appropri	medium	attend
downsid	reform	current	risk	household	account	aim	last

9. Discursive	10. Data	11. Monitoring	12. Fiscal	13. ECB	14. Forward Guidance	15. Rate Decisions
well	year	import	countri	ecb	rate	rate
go	quarter	situat	govern	central_bank	continu	interest_rate
chang	growth	observ	fiscal	govern_council	remain	decid
comment	euro_area	market	growth	european	monetari_analysi	key
thing	data	regard	stabil	euro	growth	basi
last	indic	cours	euro_area	member	under	ecb_interest_rate
certain	last	economi	fiscal_polic	institut	low	oper
ask	confirm	europ	implement	presid	monetari	point
actual	assess	level	pact	decis	money	loan
discuss	inform	present	need	govern	confirm	accord
year	recent	respons	progress	bank	interest_rate	period
reason	survey	decis	econom	nation	liquid	estim
never	expect	differ	import	work	expans	annual_growth
word	start	certain	polic	respons	euro_area	increas
happen	forecast	world	commit	treati	broad	eurostat