

## Should Central Banks Care About Text Mining?

### A Literature Review

Jean-Charles Bricongne, Raquel Caldeira & Baptiste Meunier<sup>1</sup>

June 2024, WP #950

#### ABSTRACT

As text mining has expanded in economics, central banks appear to also have ridden this wave, as we review use cases of text mining across central banks and supervisory institutions. Text mining is a polyvalent tool to gauge the economic outlook in which central banks operate, notably as an innovative way to measure inflation expectations. This is also a pivotal tool to assess risks to financial stability. Beyond financial markets, text mining can also help supervising individual financial institutions. As central banks increasingly consider issues such as the climate challenge, text mining also allows to assess the perception of climate-related risks and banks' preparedness. Besides, the analysis of central banks' communication provides a feedback tool on how to best convey decisions. Albeit powerful, text mining complements – rather than replaces – the usual indicators and procedures at central banks. Going forward, generative AI opens new frontiers for the use of textual data.

**Keywords:** Text Mining, Sentiment Analysis, Central Banking, Generative AI, Language Models

**JEL classification:** C38, C55, C82, E58, L82

---

<sup>1</sup> [jean-charles.bricongne@banque-france.fr](mailto:jean-charles.bricongne@banque-france.fr): Banque de France, Paris 1 University and LEO (Laboratory of Economics of Orleans); [raquel.mail@gmail.com](mailto:raquel.mail@gmail.com): previously Paris 1 University; [Baptiste.Meunier@ecb.europa.eu](mailto:Baptiste.Meunier@ecb.europa.eu): ECB (seconded from Banque de France) and AMSE. We are very grateful to M. McMahon, O. de Bandt, A. de Gaye, M. Menaa, and participants to internal seminars for useful comments. The views expressed in this paper are those of the authors, and do not necessarily represent those of the Banque de France, the LEO, Paris I University, the European Central Bank, or the AMSE

## NON-TECHNICAL SUMMARY

Over the last years, text mining has expanded in economics, in particular in relation with central banks as shown by figure 1 that indicates the number of occurrences in articles registered in Google Scholar of “text analysis” combined with “central bank” keywords. Among other players, central banks and supervisory institutions have mobilized it for multiple uses, in relation with their missions.


Text mining is a polyvalent tool to gauge the economic outlook in which central banks operate, notably as an innovative way to measure inflation expectations. Beyond inflation, text mining also contains significant information for forecasting key variables such as GDP, private consumption, or employment. Research has shown that text-based indicators retain significance for forecasting even after controlling for traditional indicators, suggesting that text captures information that would not otherwise be reflected by the usual indicators.

This is also a pivotal tool to assess risks to financial stability, both at the macro (economy-wide) and micro (for individual financial institution) levels. The economic literature shows a clear link between text-based sentiment analysis and financial markets, demonstrating how capturing investors’ sentiment via text can help predict market returns.

Beyond financial markets, text mining can also help supervising individual financial institutions. As central banks increasingly consider issues such as the climate challenge, text mining also allows to assess the perception of climate-related risks and banks’ preparedness.

The analysis of central banks’ communication provides a feedback tool on how to best convey decisions. Central banks’ communication delivers information outside central banks and may be studied, using text analysis, to assess its impact on the financial markets. Communication can indeed be an effective tool for central banks as it can impact financial markets, increase the predictability of monetary policy actions, and help achieving macroeconomic goals. In addition to financial markets, central banks increasingly target non-experts. For example, tweets from central banks announcing the launch of new coins and banknotes, and those related to monetary policy decisions, seem to be associated with a higher public engagement.

**Figure 1. Number of occurrences of “text analysis” AND “central bank” keywords in Google Scholar**



Sources: Google Scholar, authors’ calculations

It is important to notice that, albeit powerful, text mining complements – rather than replaces – the usual indicators and procedures at central banks. The literature generally concludes that a combination of text-based and traditional indicators is more adequate. Beyond technical problems at the pre-processing stage (e.g., multilingual text), text-mining techniques can exhibit interpretative issues. Indeed, one word may have several meanings and can be interpreted in multiple ways: such an ambiguity leads to noise in the data. Moreover, sentiment analysis often relies on a binary classification of positive and negative words but several typical nuances such as negation, irony, ambiguity, idioms, and neologisms, make such a binary classification likely too simplistic.

Going forward, generative AI also opens new frontiers for the use of textual data, for example for ideation and feedback, writing, background research, data analysis, coding, or mathematical derivations. However, the literature also acknowledges challenges associated with bias, interpretability, and reproducibility which seem important when considering LLMs into monetary policy.

---

## Les banques centrales doivent-elles s'intéresser à l'analyse textuelle ? Une revue de la littérature

### RÉSUMÉ

Alors que l'analyse textuelle s'est développée en économie, il semble que les banques centrales aient également investi dans ce domaine, comme nous le constatons dans les cas d'utilisation de l'analyse textuelle parmi les banques centrales et les institutions de supervision. L'analyse textuelle est un outil polyvalent pour évaluer la conjoncture économique dans laquelle opèrent les banques centrales, notamment comme moyen innovant de mesurer les anticipations inflationnistes. C'est également un outil essentiel pour évaluer les risques pour la stabilité financière. Au-delà des marchés financiers, l'analyse textuelle peut également aider à superviser les institutions financières individuelles. À mesure que les banques centrales prennent de plus en plus en compte des problèmes tels que le défi climatique, l'analyse textuelle permet également d'évaluer la perception des risques liés au climat et d'évaluer la préparation des banques. En outre, l'analyse de la communication des banques centrales fournit un outil de rétroaction sur la meilleure manière de transmettre les décisions. Bien que puissante, l'analyse textuelle complète – plutôt que remplace – les indicateurs et procédures habituels des banques centrales. À l'avenir, l'IA (intelligence artificielle) générative ouvre de nouvelles frontières pour l'utilisation des données textuelles.

Mots-clés : analyse textuelle, analyse de sentiment, banques centrales, IA générative, modèles de langage

Les Documents de travail reflètent les idées personnelles de leurs auteurs et n'expriment pas nécessairement la position de la Banque de France. Ils sont disponibles sur [publications.banque-france.fr](https://publications.banque-france.fr)

## Introduction

Recent years have seen an exponential usage of text mining usages in economics. This has been made possible by the increasing quantity of text available to the public. One popular example of text used as data had been the volume of Google searches related to some keywords (Choi and Varian, 2012; Varian, 2014; Scott and Varian, 2014, 2015). Beyond the quantity, the nature of available text has changed: the rise of social media has made it easier for people to express publicly their opinions, sentiment, ideas, and beliefs (Sudarsa et al., 2018). Twitter, where conciseness is required, has proved a highly tractable source of interpretable text (Buono et al., 2017). The expansion of data sources has come together with new tools to exploit it. Text data, high dimensional by nature (Gentzkow et al., 2019), has benefitted from the development of machine learning techniques geared towards such high-dimensional datasets.

Text mining is an umbrella term for different methods aimed at extracting meaning from strings of letters.<sup>1</sup> In that sense, “text mining” is to machines what “reading” is to humans. And while machines can process far greater quantities of text, the main challenge is to make machines capable of capturing the meaning of words as well as a human could do. In that vein, a central topic in text mining is sentiment analysis, which consists in extracting the sentiment (generally positive *vs.* negative) from text.<sup>2</sup> The importance of such a sentiment analysis in

---

<sup>1</sup> The term first appeared in Nasukawa and Yi (2003), but research started earlier (Tong, 2001; Morinaga et al., 2002; Pang et al., 2002). The Oxford English Dictionary defines text mining as the process or practice of examining large collections of written resources to generate new information, typically using specialized computer software. It is a subset of the larger field of data mining. The goal of text mining is to discover relevant information in text by transforming the text into data that can be used for further analysis through the use of a variety of analysis methodologies; natural language processing (NLP) is one of them. There is another nuance between NLP and text analysis in the sense that the former uses text but also other forms of communication (speeches, images, signs... for example) whereas text analysis logically focuses on text only. When messages are conveyed orally, there is room to also detect emotions, beyond the sentiment in policy texts, as shown by Gorodnichenko et al. (2023).

<sup>2</sup> Such a measurement of sentiment is connected to “Natural Language Processing” (NLP). More broadly, NLP research focuses on showing how, in a downstream task, an algorithm can improve our ability to model and understand language. By contrast, many use cases in central banks (and more generally in economics) are rather doing economic research with text – but not NLP research. In addition, the use of text in economics often relies on simple methods like counting the number of occurrences of some keywords in a corpus of text – as those methods are easier to implement and their results easier to validate. Nevertheless, when comparing accuracy across different techniques, simple word counts are found to perform very well (Ahrens et al., 2024). Finally, simple word counts are also often used in central banks to provide evidence on the latest trends in the economy – see, for example, the

economics has a theoretical background as multiple studies have shown the pivotal role of “sentiment” in shaping anticipations, financial markets, housing prices, and inflation expectations (Case and Shiller, 2003; Fischer and Stamos, 2013; Bayer et al., 2021).

The literature on text mining has skyrocketed in recent years (**Figures 1a and 1b**). While Bholat et al. (2015) wrote in 2015 that “*text mining has been historically less used as a technique in economics*”, the state of the literature is largely different by now – and central banks have followed the trend. As noted in Gentzkow et al. (2019), text as data is applied on various occasions, from measuring economic uncertainty to predicting financial market returns. Why? A first reason is that *qualitative* text contains a similar – if not higher – informative content than *quantitative* figures (Diaz-Sobrinó et al., 2020; Consoli et al., 2021) since text can convey a deep range of nuanced information.<sup>3</sup> A second reason is the rapid expansion of text available for analysis (newspaper, social media, Google trends, etc.). Most notably, social media have provided economists with a vast amount of unfiltered information about households’ sentiment.<sup>4</sup> Such insights on common people’s views were before largely inaccessible to economists, notably central bankers, who rather relied on surveys of dedicated specialists or firms (e.g., PMI surveys based on purchasing managers, SPF based on professional forecasters).

Against this background, this paper reviews the use cases of text mining across central banks, as well as supervisory institutions. One strand of the literature that drew a large attention is the analysis of central banks’ communication. Beyond this already well documented literature, this paper aims at taking a holistic view covering all aspects of central banks’ mandates as well

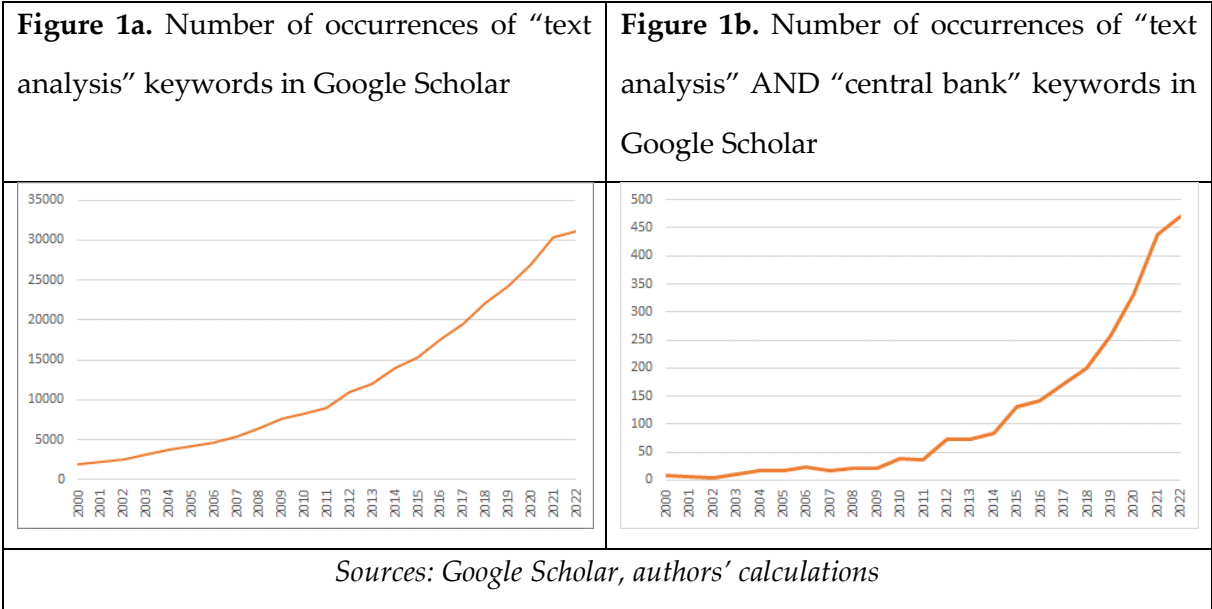
---

appetite of firms for “re-shoring” in Attinasi et al. (2024) in the broader context of trade fragmentation (Attinasi et al., 2023a; 2023b; 2023c).

<sup>3</sup> Diaz-Sobrinó et al. (2020) and Consoli et al. (2021) show for example that a word count of some specific terms in the *Quarterly Economic Bulletin* of Banco de España enables the forecaster to mimic very closely the GDP forecasts done by Banco de España. This means that the *qualitative* narrative embedded in the text contains similar information to that conveyed by the *quantitative* forecasts.

<sup>4</sup> An interest of news-based data compared to social media is however that the former can have a longer timespan: for example, the newspaper-based Financial Stress Indicator (FSI) of Püttmann (2018) spans from 1889 to 2016.

as some of their emerging missions.<sup>5</sup> In that respect, the two key pillars of central bank mandates are the control of inflation and financial stability. The latter should be taken in a broad sense, not only covering financial assets, but also housing markets and international capital flows. Central banks activities are understood in this paper in a broad sense, also including banking supervisory activities – that are generally undertaken by specific public institutions that are very close to central banks. Prudential policies (micro and macro) are also covered in this study. Finally, many central banks have recently started to address climate change risks (Dikau and Volz, 2021), which this paper also covers.



Text mining is a polyvalent tool for gauging the macroeconomic outlook in which central banks operate. It can help to assess the inflation outlook and monitor the formation of inflation expectations on the households’ side. Beyond inflation, text mining also contains significant information for forecasting key variables such as GDP, private consumption, or employment (Levenberg et al. 2014; Ellingsen et al., 2021; Barbaglia et al., 2022). Research has shown that text-based indicators retain significance for forecasting even after controlling for traditional

<sup>5</sup> Even though some use cases do not apply to some central banks depending on their mandates, this paper aims at covering all potential applications – reflecting the diversity of the missions covered by central bankers across the world (Carstens, 2019).

indicators, suggesting that text captures information that would not otherwise be reflected by the usual indicators. While the timeliness of text-based indicators is particularly relevant for near-term forecasting (*nowcasting*), research has shown that it retains some predictive power for the short-term (see de Bandt et al. (2023) on inflation) and medium-term outlook – the horizon which central banks typically target. An important body of the literature has also measured economic uncertainty based on text (Baker et al., 2016) with these indicators now an integral part of central banks' dashboard.

Besides controlling inflation, text mining can also help central banks in assessing the risks to financial stability, both at the macro (economy-wide) and micro (for individual financial institution) levels. The economic literature shows a clear link between text-based sentiment analysis and financial markets, demonstrating how capturing investors' sentiment *via* text can help predict market returns. More relevant for central banks in relation with their financial stability mandate, text mining has also been shown to be a potent predictor for volatility and stress in financial markets. Text also captures risks in the housing market – where sentiment among homebuyers plays a key role in shaping the real-estate cycle. As a key advantage, text mining can provide very granular indicators to track the highly fragmented housing market. Finally, text mining can be applied to individual financial institutions to assess idiosyncratic risks, for example by analysing CEO letters, annual reports, or earning calls. This practice can help for the supervision of financial institutions. If applied across all banks, this can also assess systemic risks in the financial system. In that sense, text mining is a helpful tool for *micro-* and *macro-*prudential policies.

Going forward, textual analysis allows central banks to approach emerging missions and measure the impact of their own communication. Text mining is a tool to assess the preparedness of banks to climate-related risks and, more broadly, can also measure the prevalence of climate uncertainty. Researchers highlighted the role of the latter on financial markets – demonstrating the importance of the topic for central banks. In addition, a large strand of the literature has been dedicated to the analysis of central banks' communication. These papers can be used as a feedback mechanism for central banks to understand how their

communication (and even the tone of their voice as shown in Gorodnichenko et al., 2023) affect the economy. Finally generative AI and notably Large Language Models (LLMs) entail a large potential for economics in general – and central banks in particular – across different areas such as writing, coding, analysis, and forecasts.

While numerous papers have addressed how text can be used in economics in general (e.g., Gentzkow et al., 2019), this paper differs by focusing explicitly on central banks. In that sense, it updates and extends Bholat et al. (2015). It differs however from Bholat et al. (2015) whose goal was to give step-by-step explanations of text mining techniques; by contrast, our paper joins the literature at a more mature stage and rather focuses on how central banks can use text mining across various applications, from forecasting to risk monitoring.

The rest of the paper is organised as follows: **section 1** focuses on how text mining can help monitoring the macroeconomic outlook, **section 2** shows examples of text mining usage for ensuring financial stability. **Section 3** discusses the emerging missions of the central banks and **section 4** addresses the complementarity of text data with usual indicators.

## **Section 1. Text mining is a polyvalent tool for gauging the macroeconomic outlook in which central banks operate**

### *1. Measuring uncertainty*

Text-based indicators have been widely used to measure uncertainty and have become an integral part of the indicators considered by central banks. In a seminal paper, Baker et al. (2016) developed an index of economic policy uncertainty (EPU) for the US and major economies. Their index combines the frequency of the references to “uncertainty” in newspapers with the variance of anticipated inflation rates – for instance, the EPU rises around events such as Gulf Wars, 9/11 attacks, and big fiscal policy disputes. Following the same idea, others have constructed complementary risk indicators such as Caldara and Iacoviello (2022) with the geopolitical risk (GPR) index – by taking the share of articles mentioning adverse



geopolitical events in global newspapers.<sup>6</sup> They find a role of the GPR in business fluctuations and show that it captures complementary information not reflected in the EPU. The index has also been extended to specific areas such as trade policy uncertainty (Caldara et al., 2020), financial stress (Püttmann, 2018), and climate policy uncertainty (Gavriilidis, 2021). As these text-based uncertainty indices became increasingly popular in the literature, the EPU has joined the list of indicators monitored by central banks (Broto et al., 2021).

The importance of text-based indicators for uncertainty in central banks relates to the negative impact of uncertainty shocks on the macroeconomic environment. EPU indices have also been shown useful to anticipate business cycle fluctuations by Rogers and Xu (2019). Several studies found empirically a detrimental effect of the EPU on activity – e.g., Meinen and Roehle (2017) for large European countries, Fontaine et al. (2017) for China, and Colombo (2013) for the euro area. In the case of Spain, Ghirelli et al. (2019) construct an EPU index and find a significant dynamic relationship between their index and the macroeconomic outlook: specifically, an unexpected rise in uncertainty leads to a significant reduction of GDP, consumption, and investment. In the same vein, central banks have used such uncertainty indicators to assess the impact on economic conditions around the Covid-19 pandemic (Gieseck and Rujin, 2020) or the Russian aggression of Ukraine (Bobasu and de Santis, 2022).<sup>7</sup> Particularly relevant for central banks' mandates, de Santis and van der Veken (2022) documented the inflationary effects of uncertainty shocks.

Central banks may also apply text mining technics to some of their own productions to produce “in-house” uncertainty indicators. This may be the case for example when business surveys include some comments on which text mining technics may be implemented. This is the case for example for the Banque de France business survey (**Figure 2**). It enables to capture

---

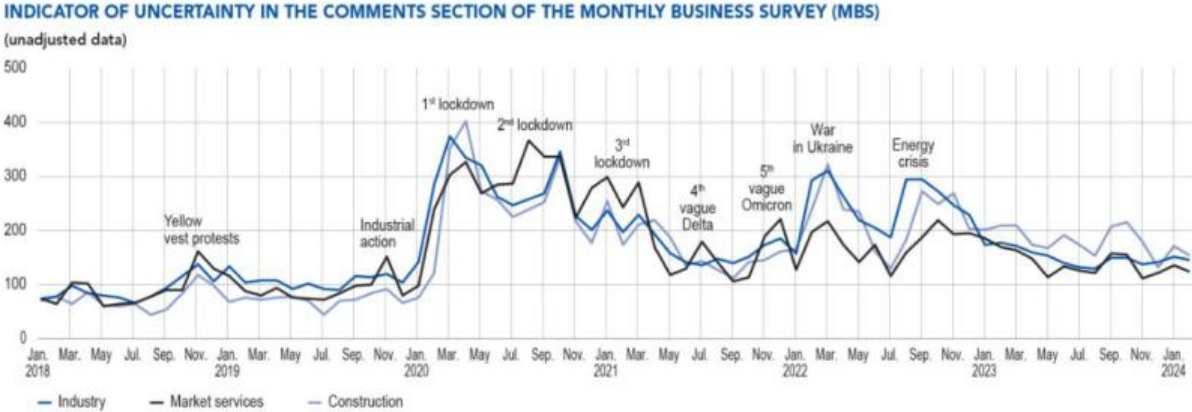
<sup>6</sup> They define geopolitical risk as “*the threat, realization, and escalation of adverse events associated with wars, terrorism, and any tensions among states and political actors, that affect the peaceful course of international relations*”. On top of Baker et al. (2016), their approach also draws on Saiz and Simonsohn (2013) who show that document frequencies in large decentralized textual databases can capture the cross-sectional variation in the occurrence frequencies of social phenomena. The GPR shares some patterns with the military spending news variable of Ramey (2011).

<sup>7</sup> In a similar vein, Baker et al. (2020b) have assessed the impact of Covid-related uncertainty on the contraction of output (In the case of Covid crisis).

on a monthly basis the impact of major events as mentioned in the firms' comments (social movements, Covid-19, war in Ukraine...), differentiated by sector (industry, services or construction). This can thus be more "tailor-made" than standard uncertainty indicators derived from Baker et al. (2016) for example.

**Figure 2.** Indicator of uncertainty in comments of monthly business survey

Source: Banque de France business survey



Note: The reference value is set at 100 and corresponds to the value around which the indicator fluctuates in normal periods.

In a more direct way, uncertainty is found to directly affect policy making at central banks. Studying the US Fed’s monetary policy, Cieslak et al. (2023) distinguish between Fed-managed uncertainty *vis-à-vis* uncertainty emanating from the economy. Quantifying how they affect the policy stance, they find that higher inflation uncertainty strongly predicts a more hawkish policy stance – even after controlling for expectations and other sources of uncertainty. They argue the effect of uncertainty on Fed’s decisions reflects concerns with maintaining credibility on the inflation anchor.

**2. Nowcasting<sup>8</sup> and forecasting economic conditions**

The first interest of text-based data for nowcasting relates to timeliness as such indicators are generally available in almost real-time. Multiple papers exploiting this timeliness have

<sup>8</sup> Nowcasting relates to the prediction of the present or very near future. While a real-time assessment of the business cycle is key for policymakers (Consoli et al., 2021), national statistics are generally released with a large

reported improvements in nowcasting accuracy from text-based indices. Text-based indicators are available with little delay, sometimes in near real-time. This makes them particularly relevant for nowcasting, as they contain early information about the current state of the economy – and therefore appear potent to predict macroeconomic variables that are otherwise released with considerable delay like GDP which is published around two months after the end of the quarter. In that vein, Aprigliano et al. (2021) have proposed high-frequency trackers of the Italian economy based on textual data; they show that they track adequately short-term fluctuations in activity. More broadly, Woloszko (2020) combines Google searches and neural networks to nowcast GDP in real-time growth across many economies and reports a high accuracy of these trackers.

The second interest of text-based data for forecasting relates to their capacity to capture some information that is otherwise not measured by conventional statistics. The literature has first started showing that text contains the predictive information contained in usual indicators: a proof of concept was documented by Levenberg et al. (2014) who show that it is possible to predict the direction of US employment using only web-based text data. This suggested that there is a large amount of predictive information inherent to text. Beyond this, several studies have shown that text-based data retain a significant predictive power, even after controlling for conventional economic variables. This strongly suggests that text contains information not captured by these usual economic variables. This is the case of Ferrara and Simoni (2022) when nowcasting GDP with Google searches: they empirically find that text-based indices increase the accuracy, even after controlling for official variables. Barbaglia et al. (2022) have evaluated the informational content of news-based sentiment indices for forecasting GDP and similarly find that the predictive power of text-based data is robust to the inclusion of conventional

---

delay. In this context, nowcasting has become increasingly popular. Giannone et al. (2008) have popularized the use of dynamic factor models for that purpose, with this tool currently being used at central banks such as the Federal Reserve Bank of Atlanta (GDPNow), the Federal Reserve Bank of New York (Bok et al., 2017), and the ECB (Angelini et al., 2011). On top of nowcasting their own economies, central banks have also been using factor models to assess economic conditions at the global level, notably at Banque de France (Ferrara and Marsilli, 2019; Jarret and Meunier, 2022). The nowcasting literature recently extended to machine learning methods and non-traditional data sources (for example, Bricongne et al., 2021; d'Aspremont et al., 2023; Bricongne et al., 2023; and Chinn et al., 2023).

macroeconomic indices. Ellingsen et al. (2021) have conducted a more extensive real-time out-of-sample forecasting horserace across several target variables (GDP, output, consumption, investment) and a large scope of explanatory variables (using the FRED-MD of Mc Cracken and Ng, 2016). They report a significant predictive power of text-based indicators, and their results suggest that news-based data are especially useful for anticipating consumption – as news-based data seems to be very adequate on capturing households' sentiment. On a different perspective, analysis of central banks communication can prove useful when macroeconomic forecasts are not released to the public. This is the case for the People's Bank of China: Lin et al. (2023) embed text-based indices into forecast models and find that the predictive information from communication texts improves the real-time out-of-sample prediction performance.

Several papers have exploited the qualities of text-based data for nowcasting and report significant improvements in accuracy. The literature has shown accuracy gains from nowcasting GDP with news-based sentiment indices in the UK (Kalamara et al., 2022), Norway (Thorsrud, 2020), or a panel of EU economies (Ashwin et al., 2021). Beyond GDP, Caldara and Iacovello (2022) showed that the GPR (a news-based index, see **section 1.1**) forecasts business investment, while Rambacussing and Kwiatkowski (2020) find that news data perform well at predicting output, inflation, and unemployment. As regards the latter, Levenberg et al. (2014) have demonstrated the ability of text-based indices to predict US non-farm payrolls, a key indicator of the health of the American economy. While these studies use sentiment indicators or word counts, topics are also used to forecast macroeconomic variables: Larsen and Thorsrud (2019) have shown that using newspaper topics can help forecasting economic changes in US, Japan, and EU.<sup>9</sup> Beyond the macroeconomic outlook, text-based indices have also been used for financial indices: Tobbäck et al. (2018) extend Baker et al. (2016)'s method with machine learning and show that their revised EPU index has a high predictive power to forecast changes in yields, spreads of sovereign bonds, and spreads on credit default swaps.<sup>10</sup> On a

---

<sup>9</sup> Ardia et al. (2019) have applied a similar framework for predicting US economic development.

<sup>10</sup> More specifically, they show that Baker et al. (2016) method is likely to lead to false positive / negative and improve their methodology by using a support vector machines (SVM) classification.

more original note, Lehrer et al. (2021) capture the emotional content of emojis in text to predict the evolution of cryptocurrencies with a horizon of a few days.

The interest of text-based data for nowcasting could however differ depending on the state of the economy. Ferrara and Simoni (2022) show greater accuracy gains from text-based indicators during periods of downturn. This is in line with Thorsrud (2020) reporting a similar asymmetry, with greater accuracy gains at turning points of the business cycle. Such an asymmetry could however pave the way for the use of text-based data in risk monitoring. Barbaglia et al. (2022) exploit this idea and show that text-based sentiment can explain the tail risks for several variables.

The interest of text-based data for nowcasting appears to also depend on the time of the forecast – with larger value added at the beginning of the quarter when few other indicators are available. Ashwin et al. (2021) show that their daily text data (based on SMS) significantly enhance nowcasts of quarterly GDP for EU economies in the first half of the quarter, but much less afterwards. This is in line with the value added of timely data being higher when few other indicators are available – as is generally the case for high-frequency data (Bricongne et al., 2020).

It should finally be noted that the interest of text-based data is not confined to nowcasting, but also applies to forecasting several months ahead. Consoli et al. (2021) find, in the case of Spain, that the narrative text-based indicator from Banco de España is informative not only at short-term horizons, but also at one- to two-year ahead. This is the case for predicting housing prices: Çepni and Khorunzhina (2023) show that text-based sentiment indices about the housing market is a strong predictor of future house prices up to three quarters ahead. Beyond forecasting macroeconomic variables, Ferrari Minesso et al. (2023) propose a methodology based on text analysis to forecast US recessions which outperforms standard yield curve-based forecast at medium horizon (up to 8 months ahead).

### *3. Monitoring the inflation outlook and the formation of inflation expectations*

Multiple studies have shown that text-mining based on press articles or Twitter provide an adequate real-time indicator for inflation and inflation expectations. Larsen et al. (2021) construct inflation indicators from press articles in the US and find high predictive power of such indicators for both inflation and inflation expectations. Gabrielyan et al. (2019) and Kalamara et al. (2022) construct similar inflation indicators from press data in the UK and echo the potential of text mining to predict inflation expectations.<sup>11</sup> Angelico et al. (2022) use Italian tweets and machine learning to construct real-time indicators of consumer inflation expectations in Italy: they find a strong correlation between their Twitter-based indicators and the usual metrics for inflation expectations, both monthly surveys and daily market-based expectations. On top of being a good real-time proxy, they also show that their Twitter-based index provides complementary information with respect to usual metrics of inflation expectations (e.g., the ECB Survey of Professional Forecasters). In the same spirit, de Bandt et al. (2023) show that both newspapers and Twitter data are well correlated with official inflation figures for France. Their results confirm the explanatory power of text-based data for nowcasting and forecasting of inflations expectations, even after controlling for oil prices and alternative inflation expectations metrics based on survey or markets (e.g., consensus forecast, inflation linked swaps), but using a semi-supervised approach.<sup>12</sup>

### *4. Anticipating international capital flows*

When taking a global perspective, text-mining can help predict the behaviour of foreign investors and therefore help anticipating international capital flow movements. Fraiberger et al. (2021) construct a media sentiment index for a range of countries. They find a sharp contrast

---

<sup>11</sup> By contrast, Rambacussing and Kwiatkowski (2020) do not find that the British press has a significant predictive power for UK inflation. Kalamara et al. (2022) however show this depends on the method: combining word counts with supervised machine learning delivers the highest forecasting improvements. Other text-based methods – such as the ones employed in Rambacussing and Kwiatkowski (2020) – perform less well.

<sup>12</sup> A theoretical underpinning of why news and other text can propagate into inflation expectations can be found in the epidemiological model of Carroll (2001), and further developed in other papers mentioned in Carroll and Wang (2022). In this approach, consumers' expectations of inflation stem from exposure to (common) news media sources and can be "infected" by what they report.

between effects of local *vs.* global news: while optimism (pessimism) in local news predicts a small and transitory increase (decrease) in local stock returns, global news has a larger and persistent impact. They also find that news affects local stocks mainly through the investment decisions of foreign – rather than local – investors. Sentiment from reports of international organizations such as the International Monetary Fund (IMF) can influence expectations of market participants and triggers movements in international capital markets. Couharde et al. (2021) find that IMF’s Regional Economic Outlook have significant repercussions on bond yields. This sheds light on the importance of IMF communication for guiding and managing markets’ expectations, including downwards in risky periods.

## **Section 2. Text mining can also help central banks ensuring financial stability at both micro- and macro-prudential levels**

### ***1. Monitoring macro risks to financial stability ...***

There is an extensive literature on how sentiment drives financial markets (Kindleberger and Aliber, 2005; Akerlof and Schiller, 2009); it should therefore be no surprise that sentiment analysis has gained a lot of interest for monitoring them. This has been done with different text sources, from press articles to social media. Multiple studies have used sentiment analysis on Twitter to forecast trends in the Dow Jones Industrial (Ranco et al., 2015; Bollen et al., 2011; Mittal and Goel, 2012). Facebook has also been used by Siikanen et al. (2018) to anticipate stock returns. Before Twitter, Antweiler et al. (2004) used posts on Raging Bull and Yahoo! Finance and measured the bullishness of the messages.<sup>13</sup> They found that an increase in the number of messages is a predictor of an increase in trading volume, and that trades rise when messages are more bullish. Traditional press has also proved useful, with for instance Schumaker and Chen (2006) using financial news articles and stock quotes to assess the evolution of stock prices twenty minutes after the news article is released.

---

<sup>13</sup> Launched in 1997, *RagingBull.com* is a website focused on financial literacy and day trading.

A key interest for central banks is the capacity of text-based indicators to measure financial volatility, and more generally stress in the financial markets. Sprenger et al. (2014) make a broad usage of Twitter data, and find relations between volatility and tweet disagreement, between trading volume and volume of tweets, as well as between stock returns and tweet sentiment. Similarly, Fernandez et al. (2021) show a correlation between market volatility, sovereign risk, foreign exchange rate volatility, and their Twitter-based sentiment index for the Mexican stock market. Using articles from the Wall Street Journal, Tetlock (2007) had already assessed interactions between media and stock market. He has found that abnormally high or low media pessimism predicts high market trading activity, and high media pessimism forecasts downward pressure on market prices. Nyman et al. (2021) analyse how narratives in financial news drive markets developments: they find that changes in the emotional content of news are highly correlated across different data sources and show the formation (and subsequent collapse) of exuberance in news prior to the global financial crisis. Such metrics can help central banks to be warned about impending financial system stress.<sup>14</sup> More broadly, textual data can be used to derive metrics of overall financial risk as in Al-Haschimi et al. (2023) where the authors rely on topical modelling to decompose financial risk in China into its thematic drivers.<sup>15</sup>

## **2. ... *Notably housing risks***

Housing risks, at the crossroads of the financial and real sectors, are key for central banks. A sizeable fraction of financial crises has originated in the housing sector (Reinhart and Kenneth, 2009; Gerlach, 2012) and recessions coinciding with housing busts are, on average, longer and more pronounced (Claessens et al., 2008). Fluctuation in housing prices and rents is also a source of economic and social concern for policymakers (Dietzel, 2015).

---

<sup>14</sup> For instance, during the pandemic, Baker et al. (2020a) and Croce et al. (2020) have used news or Twitter to quantify the propagation of information about epidemic conditions towards the international financial markets. They find a sizeable impact of such news on the stock market.

<sup>15</sup> The authors argue in addition that their index can identify major episodes of overall heightened financial risks in China, which cannot be consistently captured using other financial data. This suggests the importance of text data to capture risks in a context where official statistics would be otherwise scarce.



Sentiment is a particularly potent factor in the housing market: not only this market has been shown to be prone to the influence of irrational sentiment, but also social media can lead to over-reaction from market participants. A large body of the literature has explained volatility of housing prices by consumer psychology and irrational behaviour. Case and Shiller (2003) argued that the “animal spirits” of investors play an important role in the volatility in housing prices. Similarly, the “noise trader” approach (De Long et al., 1990) provides a theoretical background for sentiment analysis in the housing market. This approach divides investors between fully rational and imperfectly rational groups. These unexplained “irrational” factors, driven by “noise traders” or “irrational investors” are market sentiment. More generally, the housing market is characterized by high market segmentation, information asymmetry, and lack of short-selling mechanisms, all of which make participants vulnerable to sentiment-induced pricing errors (Case and Shiller, 2003; Fischer and Stamos, 2013). In those markets, social media are rapidly replacing newspapers, increasing interactions among market participants (Da et al., 2015) and being an increasingly important source of information for individual home buyers (Bailey et al., 2018; Cerchiello and Nicola, 2018; Bayer et al., 2021). Social media spread market sentiment faster and can amplify the reactions (De Fusco et al., 2018). All of these suggest that social media can be considered a valuable and innovative source of real-estate market sentiment and provide reliable indicators (Sinyak et al., 2021).

While the role of sentiment in housing market is well documented, its measurement is not straightforward – but can be enhanced by text mining. Common measures of market sentiment have been surveys (Case et al., 2012) or other macroeconomic indicators (Baker and Wurgler, 2007). However, there are important caveats in using such metrics for the housing market, known to be highly fragmented into a myriad of local markets, with researchers often facing issues in implementation of surveys or in the availability of data (Lambertini et al., 2013). By contrast, text mining can help to construct such local indicators at a minimal cost. Soo (2018) develops metrics of housing sentiment for 34 US cities by quantifying the qualitative tone of local news. He finds a significant predictive power of this index for future house prices up to two years ahead. He also finds that media sentiment has a greater effect in markets in which

speculative investors are more prevalent and where home buyers appear to be less informed. Such findings are confirmed by Hausler et al. (2018), Ruscheinsky et al. (2018), or Beracha et al. (2019). Even in countries with more constraints as regards freedom of speech, indicators based on social media can provide valuable information. For instance, Li et al. (2022) show that a market sentiment indicator based on social media brings additional explanatory power to predict house prices in 30 large and medium-sized cities in China.

Central banks communication can also influence households' expectations in the field of housing, not only due to the speech wording, but also through the speaker's voice tone and body language, as shown by Binder et al. (2023). Text analysis enables thus to control for the first factor. In the case of the US, consumers interpret heterogeneously Chair Powell's voice tone and body language at the press conference, which significantly influence their house price expectations

### ***3. A micro-prudential tool: supervising individual financial institutions***

Text-mining can also be used to assess risks to individual banks. Sentiment analysis of social media focused on individual firms has long been used to predict company's stock returns and volatility (Yu and Yuan, 2011). Nopp et al. (2015) look at how such sentiment analysis can be used by banking supervisors to evaluate risk. They find that CEO letters and banks' annual reports contain valuable forward-looking information for the near-term. They also find a strong correlation between the Tier 1 capital ratio evolution over time and the uncertainty measured by text-based indicators. Such analysis, when performed on all systemic banks, can inform on systemic risks. For instance, Cerchiello et al. (2017) use tweets to assess sentiment about large Italian banks, assessing *in fine* the systemic risk in the Italian banking sector.

Besides risks, text analysis can also be used by regulators to assess the compliance of financial institutions with regulations. Moreno-Bernal and Caminero-García (2022) use text mining to track the compliance with the environmental, social and governance (ESG) regulations. Their textual analysis is based on Pillar 3 reports that significant banks under European supervision are now required to release. Their results indicate that although there is a higher awareness of

ESG risks among banks, the preparedness is generally low. More broadly, the use of text data enables a more systematic scrutiny of banks' reports by supervisors. For instance, the ECB performed in 2020 an analysis of the comprehensiveness of climate-related and environmental risk disclosures for around 120 financial institutions in a "non-automatic" way. It ended with a similar conclusion as Moreno-Bernal and Caminero-García (2022) that the preparedness of financial institutions to such risks was generally low and that significant efforts were needed to promote transparency on the climate-related and environmental risks (ECB, 2020). Another exercise has been performed by Comunale and Manera (2024), who review the literature on the effects of Artificial Intelligence (AI) adoption and the ongoing regulatory efforts concerning this technology in different fields, including financial stability. They find that, across countries, regulations differ widely in scope and approaches and face difficult trade-offs.

### **Section 3. Textual analysis allows central banks to approach their emerging missions and better understand the impact of their own communication**

#### ***1. Climate change and natural disasters***

Over the last few years, central banks have started to consider the climate issue that can affect their capacity to meet their monetary and financial stability objectives (Batten et al., 2016). For instance, climate change was part of ECB's strategy review (ECB, 2021). The underlying rationale is that climate change can affect financial stability through physical and transition risks,<sup>16</sup> by triggering the materialization of existing risks pertaining to credit, operational, market, migration, credit spread, real estate, or strategic dimensions. Climate change has also a wide-ranging impact in terms of business activities notably for agriculture, forestry, fishery,

---

<sup>16</sup> *Physical* risks refer to the occurrence of a climate-related natural disaster (floods, droughts, etc.) while *transition* risks relate to the impact of climate-related transition policies (quota of carbon emissions, ban of polluting products, etc.).

health, energy, mining, transportation, and tourism (ECB, 2020). The impact can be a loss of profitability of firms or a devaluation of their assets.

Economic research has used textual analysis to demonstrate and quantify the effect of climate change on asset prices and market returns. Using text-based indicators for physical and transition risks, Bua et al. (2022) show the existence of climate-related risk premiums in equity markets in the euro area. Meinarding et al. (2020) detect transition risk shocks by combining abnormal financial returns with textual analysis on newspapers. They find that these shocks worsen credit conditions, leading *in fine* to financial instability. Climate-sensitive industries, like oil, are particularly affected.<sup>17</sup> In banking supervision, assessing the impact of climate change on the solvability of financial institutions however requires granular information on banks' geo-spatial exposures, to be combined with estimates of physical risks (ESRB, 2021).

Beyond the materialization of climate risks, climate uncertainty can itself affect financial stability and growth. Textual analysis has been used to gauge climate uncertainty in the spirit of Baker et al. (2016). Choi et al. (2020) identify how the public realizes and responds to global warming. They construct first a monthly Global Warming Search Volume Index measuring public interest and find that when the weather is unusually warm, stocks of carbon-intensive firms tend to underperform. They find that the cause is that retail investors – but not institutional investors – become sellers of their stocks of carbon-intensive firms. Gavriilidis (2021) constructs a Climate Policy Uncertainty index, based on major US newspapers. This index spikes near major climate-related events, such as new emissions legislation and global strikes about climate among other. Finally, Baker and Bloom (2013) use textual analysis to show how uncertainty on natural disasters – like the 2011 earthquake in Japan – affects GDP growth.

---

<sup>17</sup> Based on the idea that climate change affects financial assets, Engle et al. (2020) have proposed an approach to construct portfolios hedged against climate risk: their strategy is based on two indices measuring how climate change is discussed in the news. The first index is based on the correlation between the text content of the Wall Street Journal each month and a fixed climate change vocabulary. The second index measures the volume of negative climate news by applying sentiment analysis to articles about climate change.

## 2. *Feedback tool to analyse the impact of central banks' decisions and communication*

Central banks' communication delivers information outside central banks and may be studied, using text analysis, to assess its impact on the financial markets. Blinder et al. (2008) showed that communication can be an effective tool for central banks as it can impact financial markets, increase the predictability of monetary policy actions, and help achieving macroeconomic goals. Using a dictionary-based approach, Petropoulos and Siakoulis (2021) get to the same conclusion that the sentiment in central banks' speeches can forecast the future evolution of financial markets.<sup>18</sup> Several studies have highlighted that the communication of central banks can predict the future path of interest rates (Chague et al., 2013) and the markets' expectations about future monetary policy (Handlan, 2024). This includes long-term effects: as shown in Hansen et al. (2019), long-run interest rates respond to central bank communication.<sup>19</sup> Topics do matter in communication: Istrefi et al. (2023) show that a higher topic intensity or negative tone is associated with more monetary policy accommodation than implied by the state of the economy, when their index for speech tone is added to a standard Taylor rule. Yet, communication by monetary authorities relating to financial stability can unsettle markets (Born et al., 2010). For instance, Bulir et al. (2014) found evidence that clearer central bank communication can, at times, help to reduce noise in financial markets. Finally, while most of the literature finds that central bank communication acts primarily *via* reduced uncertainty, Ahrens et al. (2023) show how monetary policy news affect financial volatility and tail risk *via* implied changes in the forecasts of GDP, inflation, and unemployment. Central banks' communication should be understood in a broader sense than just speeches. Istrefi et al. (2022) point that, for the ECB, communication outside of regular monetary policy meeting days has a significant effect on daily movements on financial markets (Eonia rates, market-based inflation expectations and sovereign bond rates). In addition, Bennani et al. (2020) looked at *ad*

---

<sup>18</sup> In a similar vein, Baumgärtner and Zahner (2021) use a dictionary approach on over 23,000 documents from 130 central banks. They use a language model trained using machine learning techniques to locate words and documents in a multidimensional vector space. They apply these techniques with several examples in the fields of monetary policy surprises, or financial uncertainty, confirming the predictive power in the embeddings.

<sup>19</sup> The authors show that the signals from Bank of England's Inflation Reports drive long-run interest rates through the term premium, but do not affect short-run rates.

*hoc* communication from the ECB and find that communication about conventional measures particularly helps understanding how the ECB will modify interest rates in the future. Their findings demonstrate the importance of transparent communication in understanding the future direction of monetary policy. Finally, besides only words, the tone of the voice of central bankers during their speeches might matter in shaping financial markets (Gorodnichenko et al., 2023). Overall, this suggests that central banks' communication is a broad topic, beyond well-oiled speeches in monetary meetings. Going further, text analysis of central banks releases can also assess whether the impact of monetary policy could be improved with further information disclosure: Fischer et al. (2023) show that if FOMC meeting transcripts—were accessible to the public in real time – instead of with a five-year lag, market policy expectations could substantially improve forecasting accuracy.

In addition to financial markets, central banks increasingly target non-experts. This is the approach used by Ehrmann and Wabitsch (2021), who analyse English and German Twitter traffic about the ECB to understand whether its communication is received by non-experts and how it affects their views. It shows that Twitter traffic is responsive to ECB communication, also for non-experts. For several ECB communication events, Twitter constitutes primarily a channel to relay information: tweets become more factual, and the views expressed more moderate and homogeneous. Other communication events, such as former President Draghi's "*whatever it takes*" statement, trigger persistent traffic and a divergence in views. Thus, Twitter also serves as a platform for controversial discussions. The findings suggest that central banks manage to reach non-experts. Masciandaro et al. (2024) confirm that social media can be a useful tool to target a larger audience. Focusing on the new tools that central banks employ in their communication, they show that tweets from central banks announcing the launch of new coins and banknotes, and those related to monetary policy decisions, are associated with a higher public engagement.

Textual methods can also assess central banks' efficiency in more conceptual frameworks. For instance, Handlan and Gáti (2024) have conceptualized a framework to estimate central banks' communication rules – as a mapping between central banks' expectations on the economy and

the words used in speeches. They find evidence for systematic communication rules – though they can vary over time. In the same vein, Carretta et al. (2015), focusing on the national culture framework provided by Hofstede et al. (2010), tried to define which supervisory culture is most effective in preserving the stability of European banks. To do that, they observe to what extent public speeches by heads of national supervision authorities reflect the national cultural values of the Hofstede framework. They find that a supervisory culture oriented to “collectivism” and “uncertainty avoidance” helps safeguarding financial stability.<sup>20</sup> Another possible use is to apply text analysis to legal texts related to central banks to qualify or compare them (e.g., in terms of independence) such as in AlAjmi et al. (2023).

#### **Section 4. Text analysis to enhance traditional indicators and methods, notably with the advent of generative AI**

##### ***1. Text as a complement to traditional indicators***

Some papers highlight the comparative advantages of text mining over traditional indicators. Bholat et al. (2015) show four strengths of text-based uncertainty metrics over usual measures: *(i)* while most uncertainty measures are based on the volatility of financial assets, textual analysis can capture uncertainty beyond financial markets; *(ii)* the cause of uncertainty is not always evident in traditional indicators whereas the source of uncertainty is frequently indicated directly in the text; *(iii)* far longer time series can be acquired based on newspapers than those based on financial option;<sup>21</sup> and *(iv)* cross-country measures of uncertainty may be easily derived from national media. Beber et al. (2015) extract from news real-time measures of inflation, output, employment, and sentiment at various intervals and frequencies. They show their method provides real-time, daily, unbiased, and objective reading of the state of the macroeconomy. They find that their method delivers more accurate and timelier

---

<sup>20</sup> In a similar effort, Goldsmith-Pinkham et al. (2016) extract textual information from concerns raised by Federal Reserve supervisors to measure supervision and its interaction with the other Basel pillars.

<sup>21</sup> Püttmann (2018) for example builds an indicator of financial stress from 1889 for the US using newspapers.

predictions of changes in economic conditions compared to forecasting approaches based on usual indicators. In the same vein, Thorsrud (2016) finds better relative performance of news-based indicators compared with conventional methods or professional forecasters.

The literature generally concludes that a combination of text-based and traditional indicators is more adequate. Beyond technical problems at the pre-processing stage (e.g., multilingual text), text-mining techniques can exhibit interpretative issues. Gaikwad et al. (2014) highlight the ambiguity issue: although research has been conducted on the sense of words (whether it is negative or not, e.g., in Loughran and McDonald, 2011), the problem of ambiguity cannot be eliminated. One word may have several meanings and can be interpreted in multiple ways: such an ambiguity leads to noise in the data. Talib et al. (2016) highlight the importance of domain knowledge integration for specific fields to improve interpretations. Moreover, sentiment analysis often relies on a binary classification of positive and negative words but as underlined by Jurafsky and Manning (2012), several typical nuances such as negation, irony, ambiguity, idioms, and neologisms, make such a binary classification likely too simplistic. Therefore, a combination of both text-mining and traditional measures is often appropriate as shown by Schumaker and Chen (2006) whose model containing both article terms and stock price at the time of article release had the best performance in closeness to the future stock price. This is confirmed by Li et al. (2022) for housing, who show that market sentiment complements standard variables to explain housing prices.<sup>22</sup>

## 2. *Large Language Models*

Generative AI, notably Large Language Models (LLMs) like *ChatGPT*, are widely considered as a deeply transformative technology with the potential to affect wide areas of the society. For example, Noy and Zhang (2023) show that ChatGPT can make writers 40% faster while improving output quality, while Peng et al. (2023) show that LLM-based coding assistant allows programmers to code 50% faster. In economics, Korinek (2024) insists on the potential

---

<sup>22</sup> On a more transversal topic, the combination of text with contextual data (e.g., situational characteristics) can enhance the user experience and effectiveness of scholarly recommender system interfaces – for example *rScholar* (Champiri et al., 2023) for Google Scholar.



of generative AI for research along six areas: ideation and feedback, writing, background research, data analysis, coding, and mathematical derivations. However, the literature also acknowledges challenges associated with bias, interpretability, and reproducibility (Korinek, 2023) – which seem important when considering LLMs into monetary policy. In this spirit, the European Central Bank (Moufakkir, 2023) and the US Federal Reserve (Horwich, 2023) have both acknowledged using AI – and planning to open new uses cases, as a promising area for the future.

Specifically for text analysis, LLMs provide a state-of-the-art and efficient tool to perform broad tasks in NLP – since the primary goal of LLMs is to process and generate human-like text. For example, ChatGPT has been shown to outperform smaller-scale deep learning models like FinBERT by 35% in sentiment classification (Fatouros et al., 2023), highlighting the strong capacity of LLMs to perform various tasks of sentiment analysis. More specifically on monetary policy, LLMs can be used to classify central banks' announcements (hawkish / dovish) in a way that aligns with target and path surprise factors (Smales, 2023). The same applies for FedSpeak in Hansen and Kazinnik (2023) who show that LLMs significantly improve classification performance over other methods. The authors also demonstrate that ChatGPT can identify macroeconomic shocks using a narrative approach. LLMs are also used to predict macroeconomic variables and anticipate future monetary policy decisions. For instance, Woodhouse and Charlesworth (2023) test whether LLMs can offer predictive qualities of future interest rate decisions. Using speeches from the Bank of England Monetary Policy Committee (MPC) members, they find that ChatGPT can predict future interest rate decisions. In the same vein, Faria e Castro and Leibovici (2023) use LLMs to produce conditional inflation forecasts and show that LLM-based forecasts perform better than the Survey of Professional Forecasters (SPF). In the same vein, Bybee (2023) uses a LLM to generate a survey of economic expectations based on news articles; he finds that the results are in line with standard surveys of economic analysts, suggesting that generative AI can simulate the human cognitive process. LLMs may also be used to compare central bank communication, as in Evdokimova et al. (2023) who find emerging market central banks' communication to

outperform that of the Federal Reserve and the ECB. Finally, LLMs can also be used to revisit the classification of central banks' communication, as in Pfeifer and Marohl (2023) who construct *CentralBankRoBERTa*, an economic agent classifier that identifies the emotional content of sentences in central bank communications.

## **Conclusion**

This paper shows the role that text mining can take in central banking, from monitoring the economic outlook to supervising risks in the financial system or in individual financial institutions. Text mining has proven useful to gauge inflation expectations, at the heart of central banks' mandates. The polyvalence of text mining also helps tackling a wide array of risks across financial markets, housing markets, or related to individual financial institutions. Besides these missions, text mining is also pivotal to approach climate-related risks, not only to measure it but also to assess the preparedness of the financial system. Finally, text mining is also key for a central bank to understand the impact of its own communication – a crucial role given that “*monetary policy is 98 percent talk and 2 percent action*” (Bernanke, 2022).

Both the availability of text and the methods to exploit it are expanding rapidly, so the importance of text in central banking applications is expected to continue to grow. The review of use cases in this paper suggests several areas where innovation could proceed rapidly. For instance, text analysis might move from *ad hoc* dictionary methods (which impose beliefs of the forecasters) to more sophisticated and data-driven methods. Richer representations, such as word embeddings and linguistic models that draw on natural language processing tools have seen a large success in the general literature, and one can see great potential for their application in central banking. Large Language Models, notably, appear as a potent tool to expand the use of text-based models in central banks.

## References

- Acosta, M. (2023). "A New Measure of Central Bank Transparency and Implications for the Effectiveness of Monetary Policy", *International Journal of Central Banking*, 19(3), 49-97.
- Ahrens, M., Erdemlioglu, D., McMahon, M., Neely, C. J., & Yang, X. (2023). "Mind Your Language: Market Responses to Central Bank Speeches", *unpublished manuscript*.
- Ahrens, M., Gorduza, D., & McMahon, M. (2024). "Benchmarking NLP for Economics and Finance", *unpublished manuscript*.
- Akerlof, G. A., & Shiller, R. J. (2010). "Animal spirits: How human psychology drives the economy, and why it matters for global capitalism", Princeton University Press.
- AlAjmi, K., Deodoro, J., Khan, A. & Moriya, K. (2023). "Predicting the Law: Artificial Intelligence Findings from the IMF's Central Bank Legislation Database", *Working Paper*, No. 2023/241, International Monetary Fund.
- Al-Haschimi, A., Apostolou, A., Azqueta-Gavaldon, A., & Ricci, M. (2023). "Using machine learning to measure financial risk in China", *Working Paper Series*, No 2767, European Central Bank.
- Andersen, T. G., Bollerslev, T., Diebold, F. X., & Vega, C. (2003). "Micro effects of macro announcements: Real-time price discovery in foreign exchange", *American Economic Review*, 93(1), 38-62.
- Angelico, C., Marcucci, J., Miccoli, M., & Quarta, F. (2022). "Can we measure inflation expectations using Twitter?", *Journal of Econometrics*, 228(2), 259-277.
- Angelini, E., Camba-Mendez, G., Giannone, D., Reichlin, L., & Ruenstler, G. (2011). "Short term forecasts of euro area GDP growth", *Economic Journal*, 14, C25–C44.
- Antweiler, W., & Frank, M. Z. (2004). "Is all that talk just noise? The information content of Internet stock message boards", *The Journal of Finance*, 59(3), 1259-1294.

Aprigliano, V., Emiliozzi, S., Guaitoli, G., Luciani, A., Marcucci, J., & Monteforte, L. (2021). "The power of text-based indicators in forecasting the Italian economic activity", *Temi di discussione*, No 1321, Banca d'Italia.

Ardia, D., Bluteau, K., & Boudt, K. (2019). "Questioning the news about economic growth: Sparse forecasting using thousands of news-based sentiment values", *International Journal of Forecasting*, 35(4), 1370-1386.

Arias, M., Arratia, A., & Xuriguera, R. (2014). "Forecasting with Twitter data", *ACM Transactions on Intelligent Systems and Technology*, 5(1), 1-24.

Ashwin, J., Kalamara, E., & Saiz, L. (2021). "Nowcasting euro area GDP with news sentiment: a tale of two crises", *Working Paper Series*, No. 2616, European Central Bank.

d'Aspremont, A., Ben Arous, S., Bricongne, J-C., Lietti, B., & Meunier, B. (2023). "Satellites Turn "Concrete": Tracking Cement with Satellite Data and Neural Networks", *Working papers*, No 916, Banque de France.

Attinasi, M. G., Boeckelmann, L., Hespert, L., Linzenich, J., & Meunier, B. (2024). "Global trade in the post-pandemic environment", *Economic Bulletin Boxes*, 1, European Central Bank.

Attinasi, M-G., Boeckelmann L., and Meunier, B. (2023a). "Friend-shoring global value chains: a model-based analysis", *Economic Bulletin Box*, 2, European Central Bank.

Attinasi M-G., Boeckelmann, L., and Meunier, B. (2023b). "The economic costs of supply chain decoupling", *Working Paper Series*, No 2839, European Central Bank.

Attinasi, M-G., Boeckelmann L., and Meunier, B. (2023c). "Unfriendly friends: Trade and relocation effects of the US Inflation Reduction Act", *VoxEU*, 7 July 2023.

Bailey, M., Cao, R., Kuchler, T., & Stroebel, J. (2018). "The economic effects of social networks: evidence from the housing market", *Journal of Political Economy*, 126(6), 2224-2276.

Baker, S. R., & Bloom, N. (2013). "Does Uncertainty Reduce Growth? Using Disasters as Natural Experiments", *NBER Working Paper*, No. 19475, National Bureau of Economic Research.

Baker, S. R., Bloom, N., & Davis, S. J. (2016). "Measuring Economic Policy Uncertainty", *The Quarterly Journal of Economics*, 131(4), 1593-1636.

Baker, S. R., Bloom, N., Davis, S. J., Kost, K. J., Sammon, M. C., & Viratyosin, T. (2020a). "The unprecedented stock market impact of COVID-19", *NBER Working Paper*, No. 26945, National Bureau of Economic Research.

Baker, S. R., Bloom, N., Davis, S. J., & Terry, S. J. (2020b). "Covid-Induced Economic Uncertainty", *NBER Working Paper*, No. 26983, National Bureau of Economic Research.

Baker, M., & Wurgler, J. (2007). "Investor Sentiment in the Stock Market", *Journal of Economic Perspectives*, 21(2), 129-152.

de Bandt, O., Bricongne, J.-C., Denes, J., Dhenin, A., de Gaye, A. & Robert, P.-A. (2023). "Using the Press to construct a new indicator of inflation perceptions in France", *Working paper*, No 921, Banque de France.

Barbaglia, L., Consoli, S., & Manzan, S. (2022). "Forecasting with economic news", *Journal of Business & Economic Statistics*, 1-12.

Baumgärtner, M., & Zahner, J. (2021). "Whatever it takes to understand a central banker – Embedding their words using neural networks", *Marburg Centre for Institutional Economics Paper Series*, No. 2021/01.

Batten, S., Sowerbutts, R., & Tanaka, M. (2016). "Let's talk about the weather: the impact of climate change on central banks", *Working paper*, No 603, Bank of England.

Bayer, P., Mangum, K., & Roberts, J. W. (2021). "Speculative fever: investor contagion in the housing bubble", *American Economic Review*, 111(2), 609-651.

Beber, A., Brandt, M. W., & Luisi, M. (2015). "Distilling the macroeconomic news flow", *Journal of Financial Economics*, 117(3), 489-507.

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

Beracha, E., Lang, M., & Hausler, J. (2019). "On the relationship between market sentiment and commercial real estate performance—a textual analysis examination", *Journal of Real Estate Research*, 41(4), 605-638.

Bernanke, B. (2022). *The Fed from the Great Inflation to COVID-19* [speech audio recording], <https://www.brookings.edu/events/ben-bernanke-the-fed-from-the-great-inflation-to-covid-19/>

Bholat, D., Hansen, S., Santos, P., & Schonhardt-Bailey, C. (2015). "Text Mining for Central Banks", *European Economics: Macroeconomics & Monetary Economics eJournal*.

Binder, C., Kuang, P., & Tang, L. (2023). "Central Bank Communication and House Price Expectations", *NBER Working Paper*, No. 31232, National Bureau of Economic Research.

Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). "Latent Dirichlet Allocation", *Journal of machine Learning research*, 3, 993-1022.

Blinder, A. S., Ehrmann, M., Fratzscher, M., De Haan, J., & Jansen, D. J. (2008). "Central Bank Communication and Monetary Policy: A Survey of Theory and Evidence", *Working Paper Series*, No 898, European Central Bank.

Bobasu, A., & de Santis, R. (2022). "The impact of the Russian invasion of Ukraine on euro area activity via the uncertainty channel", *Economic Bulletin Box*, 4, European Central Bank.

Bollen, J., Mao, H., & Zeng, X. (2011). "Twitter mood predicts the stock market", *Journal of computational science*, 2(1), 1-8.

Bok, B., Caratelli, D., Giannone, D., Sbordone, A. M., & Tambalotti, A. (2017). "Macroeconomic Nowcasting and Forecasting with Big Data", *Staff Reports*, No 830, Federal Reserve Bank of New York.

Born, B., Ehrmann, M., & Fratzscher, M. (2010). "Macroprudential Policy and Central Bank Communication" in: Bank for International Settlements (ed.), *Macroprudential regulation and policy*, 60, 107-110.

Bricongne J.-C., Coffinet J., Delbos J.-B., Kaiser V., Kien J.-N., Kintzler E., Lestrade A., Meunier B., Mouliom M., & Nicolas T. (2020). "Tracking the economy during the Covid-19 pandemic: The contribution of high-frequency indicators", *Bulletin de la Banque de France*, No 231.

Bricongne J.-C., Meunier B., & Pical, T. (2021). "Can satellite data on air pollution predict industrial production?", *Working Papers*, No. 847, Banque de France.

Bricongne, J.-C., de Bandt, O., de Gaye, A., Denès, J., Hubert, P., & Robert, P.-A. (2022). "New indicators of perceived inflation in France based on media data", *Banque de France Notepad*, No. 299.

Bricongne, J.-C., Meunier, B., & Pouget, S. (2023). "Web-scraping housing prices in real-time: The Covid-19 crisis in the UK", *Journal of Housing Economics*, 59(PB).

Broto, C., Alonso, J., Cáceres García, E., Fernandez Lafuerza, L. G., Galán, J., Lamas, M., Martínez Resano, J. R., Melnychuk, M., Merino, S., Stupariu, P., Tejedor, M. L., Alonso Alvarez, I., Alloza, M., Burriel, P., Garcia, S., Gil, M., Molina, L., Rodriguez-Moreno, M., Villanueva, E., Alonso, A., Caminero, T., Gonzalez, C. I., Gonzalez-Perez, M. T., Nalda Tomas, M., & Martínez-Valero, P. J. (2021). "How Do Central Banks Identify Risks? A Survey of Indicators", *Occasional Paper*, No. 2125, Banco de España.

Bua, G., Kapp, D., Ramella, F., & Rognone, L. (2022). "Transition Versus Physical Climate Risk Pricing in European Financial Markets: A Text-Based Approach", *Working Paper Series*, No. 2677, European Central Bank.

Bulir, A., Cihak, M., & Jansen, D. J. (2014). "Does the Clarity of Inflation Reports Affect Volatility in Financial Markets", *Working Paper*, No. 14/175, International Monetary Fund.

Buono, D., Mazzi, G. L., Kapetanios, G., Marcellino, M., & Papailias, F. (2017). "Big data types for macroeconomic nowcasting", *Eurostat Review on National Accounts and Macroeconomic Indicators*, 1, 93-145.

Bybee, L. (2023). "Surveying Generative AI's Economic Expectations", *arXiv pre-print*.

Caldara, D., & Iacoviello, M. (2022). "Measuring Geopolitical Risk", *American Economic Review*, 112(4), 1194-1225.

Caldara, D., Iacoviello, M., Molligo, P., Prestipino, A., & Raffo, A. (2020). "The economic effects of trade policy uncertainty", *Journal of Monetary Economics*, 109, 38-59.

Carretta, A., Farina, V., Fiordelisi, F., Schwizer, P., Stentella, L., & Francesco, S. (2015). "Don't Stand So Close to Me: The role of supervisory style in banking stability", *Journal of Banking and Finance*, 52(C).

Carroll, C. (2001). "The Epidemiology of Macroeconomic Expectations," *NBER Working Paper Series*, No. 8695, National Bureau of Economic Research.

Carroll, C., & Wang, T. (2022). "Epidemiological expectations", *NBER Working Paper Series*, No. 30605, National Bureau of Economic Research, October.

Carstens, A. (2019). "The new role of central banks", speech to the Financial Stability Institute's 20<sup>th</sup> anniversary conference, *A cross-sectoral reflection on the past, and looking ahead to the future*.

Case, K. E., Shiller, R. J., & Thompson, A. K. (2012). "What have they been thinking? Homebuyer behavior in hot and cold markets", *Brookings Papers on Economic Activity*, 1, 265-315.

Case, K. E., & Shiller, R. J. (2003). "Is there a bubble in the housing market?", *Brookings Papers on Economic Activity*, 2, 299-362.



- Çepni, O., & Khorunzhina, N. (2023). "Geography of Housing Sentiment over Business Cycles", *mimeo*, <http://dx.doi.org/10.2139/ssrn.4350783>
- Cerchiello, P., Giudici, P., & Nicola, G. (2017). "Twitter data models for bank risk contagion", *Neurocomputing*, 264, 50-56.
- Chague, F. D., de Losso, R., Giovannetti, B. C., & Manoel, P. (2013). "Central Bank Communication Affects Long-Term Interest Rates", *Working Paper*, No. 2013\_07, University of São Paulo.
- Champiri, Z. D., Fisher, B. D., Sinnamon, L. S. (2020). "rScholar: An Interactive Contextual User Interface to Enhance UX of Scholarly Recommender Systems", *Proceedings of the 22<sup>nd</sup> HCI International Conference*.
- Chen, M., Deininger, M., Lee, S. J., & Sicilian, M. J. (2020). "Identifying Financial Crises Using Machine Learning on Textual Data", *31<sup>st</sup> Annual congress of the European Economic Association Congress*.
- Chinn, M. D., Meunier, B., & Stumpner, S. (2023). "Nowcasting world trade with machine learning: a three-step approach", *Working Paper Series*, No 2836, European Central Bank.
- Choi, H., & Varian, H. (2012). "Predicting the present with Google Trends", *Economic Record*, 88(s1), 2-9.
- Cieslak, A., Hansen, S., McMahon, M., & Xiao, S. (2023). "Policymakers' Uncertainty", *Working Paper*, No. 31849, National Bureau of Economic Research.
- Claessens, S., Kose, A., & Terrones, M. E. (2008). "Financial Cycles: What? How? When?", *Working Paper*, No. 11/76, *International Monetary Fund*.
- Colombo, V. (2013). "Economic policy uncertainty in the US: Does it matter for the euro area?", *Economics Letters*, 121(1), 39-42.

Comunale, M., & Manera, A. (2024). "The Economic Impacts and the Regulation of AI: A Review of the Academic Literature and Policy Actions", *Working Paper*, No. 2024/065, International Monetary Fund.

Consoli, S., Reforgiato Recupero, D., & Saisana, M. (2021). *"Data Science for Economics and Finance. Methodologies and Applications"*, Springer.

Couharde, C., Bennani, H., & Wallois, Y. (2021). "Do IMF Reports Affect Market Expectations? A Sentiment Analysis Approach", *EconomiX Working Papers*, No. 2021-6, University of Paris Nanterre.

Croce, M. M., Farroni, P., & Wolfskeil, I. (2020). "When the Markets Get COVID: COntagion, Viruses, and Information Diffusion", *CEPR Discussion Paper*, No. 14674.

Da, Z., Engelberg, J., & Gao, P. (2015). "The sum of all FEARS investor sentiment and asset prices", *The Review of Financial Studies*, 28(1), 1-32.

Deerwester, S., Dumais, S. T., Furnas, G. W., Landauer, T. K., Harshman, R. (1990). "Indexing by latent semantic analysis", *Journal of the American society for information science*, 41(6), 391-407.

De Fusco, A., Ding, W., Ferreira, F., & Gyourko, J. (2018). "The role of price spillovers in the American housing boom", *Journal of Urban Economics*, 108, 72-84.

De Long, J. B., Shleifer, A., Summers, L. H., & Waldmann, R. J. (1990). "Noise trader risk in financial markets", *Journal of Political Economy*, 98(4), 703-738.

De Santis, R., & van der Veken, W. (2022). "Deflationary financial shocks and inflationary uncertainty shocks: an SVAR Investigation", *Working Paper Series*, No 2727, European Central Bank.

Diaz-Sobrino, N., Ghirelli, C., Hurtado, S., Perez, J. J., & Urtasun, A. (2020). "The narrative about the economy as a shadow forecast: an analysis using bank of Spain quarterly reports", *Working Paper*, No. 2042, Banco de España.

Dietzel, M. (2015). "Sentiment-based predictions of housing market turning points with Google trends", *International Journal of Housing Markets and Analysis*, 9, 108-136.

Dikau, S., & Volz, U. (2021). "Central bank mandates, sustainability objectives and the promotion of green finance", *Ecological Economics*, 184, 107022.

ECB (2020). "Guide on climate-related and environmental risks / Supervisory expectations relating to risk management and disclosure", *Technical report*.

ECB (2021). "Climate change and monetary policy in the euro area", *Occasional Paper Series*, No. 271, European Central Bank.

Eckley, P. (2015). "Measuring economic uncertainty using news-media textual data", *MPRA Paper*, No. 64874.

Ehrmann, M., & Wabitsch, A. (2021). "Central bank communication with non-experts: a road to nowhere?", *European Central Bank Working Paper*, No. 2594, October.

Ellingsen, J., Larsen, V. H., & Thorsrud, L. (2022). "News media versus FRED-MD for macroeconomic forecasting", *Journal of Applied Econometrics*, 37(1), 63-81.

Engle, R. F., Giglio, S., Kelly, B., Lee, H., & Stroebel, J. (2020). "Hedging Climate Change News", *The Review of Financial Studies*, 33, 1184-1216.

Evdokimova, T., Nagy-Mohacsi, P., Ponomarenko, O. & Ribakova, E. (2023). "Central Banks and Policy Communication: How Emerging Markets have Outperformed the Fed and ECB", *Peterson Institute for International Economics Working Paper*, No. 23-10, November.

Faria e Castro, M., & Leibovici, F. (2023). "Artificial Intelligence and Inflation Forecasts", *Working Paper Series*, No 2023-015C, Federal Reserve Bank of Saint Louis.

Fatouros, G., Soldatos, J., Kouroumali, K., Makridis, G., & Kyriazis, D. (2023). "Transforming sentiment analysis in the financial domain with ChatGPT", *Machine Learning with Applications*, 14, 100508.

- Fernandez, R., Palma Guizar, B., & Rho, C. (2021). "A sentiment-based risk indicator for the Mexican financial sector", *Latin American Journal of Central Banking*, 2(3).
- Ferrara, L., & Marsilli, C. (2019). "Nowcasting global economic growth: A factor-augmented mixed-frequency approach", *The World Economy*, 42(3), 846-875.
- Ferrara, L., & Simoni, A. (2022). "When are Google Data Useful to Nowcast GDP? An Approach via Preselection and Shrinkage". *Journal of Business & Economic Statistics*, 1-15.
- Ferrari Minesso, M., Lebastard, L., & Le Mezo, H. (2023). "Text-Based Recession Probabilities", *IMF Economic Review*, 71, 415-438.
- Fischer, E., McCaughrin, R., Prazad, S., & Vandergon, M. (2023). "Fed Transparency and Policy Expectation Errors: A Text Analysis Approach", *FRB of New York Staff Report*, No. 1081.
- Fischer, M., & Stamos, M. Z. (2013). "Optimal life cycle portfolio choice with housing market cycles", *The Review of Financial Studies*, 26(9), 2311-2352.
- Fontaine, I., Didier, L., & Razafindravaosolonirina, R. J. (2017). "Foreign policy uncertainty shocks and US macroeconomic activity: Evidence from China", *Economics Letters*, 155, 121-125.
- Fraiberger, S. P., Lee, D., Puy, D., & Ranci re, R. (2021). "Media sentiment and international asset prices", *Journal of International Economics*, 133.
- Gabrielyan, D., Masso, J., & Uusk la, L. (2019). "Mining news data for the measurement and prediction of inflation expectations", *International Conference on Time Series and Forecasting*, 253-271.
- Gaikwad, S. V., Chaugule, A., & Patil, P. (2014). "Text mining methods and techniques", *International Journal of Computer Applications*, 85(17).
- Gavriilidis, K. (2021). "Measuring Climate Policy Uncertainty", *SSRN Electronic Journal*.
- Gentzkow, M., Kelly, B., & Taddy, M. (2019). "Text as data", *Journal of Economic Literature*, 57(3), 535-574.

Gerlach, S. (2012). *“Housing markets and financial stability”*, speech by Mr. Stefan Gerlach, Deputy Governor of the Central Bank of Ireland at the National University of Ireland.

Ghirelli, C., Hurtado, S., Pérez, J. J., & Urtasun, A. (2021). “New Data Sources for Central Banks” in Ghirelli (ed.), *Data Science for Economics and Finance*.

Ghirelli, C., Pérez, J. J., & Urtasun, A. (2019). “A new economic policy uncertainty index for Spain”, *Economics Letters*, 182, 64-67.

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

Gieseck, A., & Rujin, S. (2020). “The impact of the recent spike in uncertainty on economic activity in the euro area”, *Economic Bulletin Box*, 6, European Central Bank

Goldsmith-Pinkham, P., Hirtle, B., & Lucca, D. (2016). “Parsing the content of bank supervision”, *Staff Report*, No. 770, Federal Reserve Bank of New York.

Gorodnichenko, Y., Pham, T., & Talavera, O. (2023). “The Voice of Monetary Policy”, *American Economic Review*, 113(2), 548-584.

Handlan, A. (2024). “FedSpeak Matters: Statement Similarity and Monetary Policy Expectations”, *unpublished manuscript*.

Handlan, A., & Gáti, L. (2024). “Monetary Communication Rules”, *unpublished manuscript*.

Hansen, S., McMahon, M., & Tong, M. (2019). “The long-run information effect of central bank communication”, *Journal of Monetary Economics*, 108, Elsevier, 185-202.

Hansen, A. L., & Kazinnik, S. (2023). “Can ChatGPT Decipher FedSpeak?”, *unpublished manuscript*, available at SSRN: <https://ssrn.com/abstract=4399406>

Hausler, J., Ruschinsky, J., & Lang, M. (2018). “News-based sentiment analysis in real estate: a machine learning approach”, *Journal of Property Research*, 35, 344-371.

Hofstede, G. J., & Minkov, M. (1991). *Cultures and Organizations: Software of the Mind*, McGraw Hill.

Horwich, J. (2023). "Across the Fed, a mindful exploration of AI is underway", Federal Reserve Bank of Minneapolis, 22 September.

Istrefi, K., Odendahl, F., & Sestieri, G. (2022). "ECB Communication and its Impact on Financial Markets", Working paper, No 859, Banque de France.

Istrefi, K., Odendahl, F., & Sestieri, G. (2023). "Fed communication on financial stability concerns and monetary policy decisions: Revelations from speeches", *Journal of Banking and Finance*, 151(C).

Jardet, C., & Meunier, B. (2022). "Nowcasting world GDP growth with high-frequency data", *Journal of Forecasting*, 41(6), 1181-1200.

Jurafsky, D., & Manning, C. (2012). "Natural language processing", *Instructor*, 212(998), 3482.

Kalamara, E., Turrell, A., Redl, C., Kapetanios, G., & Kapadia, S. (2020). "Making text count: Economic forecasting using newspaper text", *Journal of Applied Econometrics*, 37(5), 896-919.

Kalampokis, E., Tambouris, E., & Tarabanis, K.A. (2013). "Understanding the predictive power of social media", *Internet Research*, 23, 544-559.

Kindleberger, C. P., & Aliber, R. Z. (2005). "Anatomy of a typical crisis" in *Manias, Panics and Crashes*, 21-32, Palgrave MacMillan, London.

Korinek, A. (2024). "Generative AI for Economic Research: Use Cases and Implications for Economists", *Journal of Economic Literature*, forthcoming.

Lambertini, L., Mendicino, C., & Punzi, M. T. (2013). "Expectation-driven cycles in the housing market: evidence from survey data", *Journal of Financial Stability*, 9(4), 518-529.

Larsen, V. H., Thorsrud, L. A., & Zhulanova, J. (2021). "News-driven inflation expectations and information rigidities", *Journal of Monetary Economics*, 117(C), 507-520.

Larsen, V. H. & Thorsrud, L. A. (2019). "Business Cycle Narratives", *Working Paper*, No. 7468, CESifo.

Lehrer, F. S., Xie, T., & Yi, G. (2021). "Do the Hype of the Benefits from Using New Data Science Tools Extend to Forecasting Extremely Volatile Assets?" in Ghirelli (ed.), *Data Science for Economics and Finance*.

Levenberg, A., Pulman, S., Moilanen, K., Simpson, E., & Roberts, S. (2014). "Predicting economic indicators from web text using sentiment composition", *International Journal of Computer and Communication Engineering*, 3(2), 109-115.

Lewis, D. J., Mertens, K., & Stock, J. H. (2020). "Monitoring Real Activity in Real Time: The Weekly Economic Index", *Liberty Street Economics*, Federal Reserve Bank of New York.

Li, J., Wang, Y., & Liu, C. (2022). "Spatial effect of market sentiment on housing price: evidence from social media data in China", *International Journal of Strategic Property Management*, 26(1), 72-85.

Lin, J., Fan, J., Zhang, Y., & Chen, L. (2023). "Real-time macroeconomic projection using narrative central bank communication", *Journal of Applied Econometrics*, 38(2), 202-221.

Loughran, T., & McDonald, B. (2011). "When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks", *The Journal of Finance*, 66(1), 35-65.

Masciandaro, D., Peia, O., & Romelli, D. (2024). "Central bank communication and social media: From silence to Twitter", *Journal of Economic Surveys*, 38, 365-388.

McCracken, M., & Ng, S. (2016). "FRED-MD: A Monthly Database for Macroeconomic Research", *Journal of Business & Economic Statistics*, Vol. 34, No. 4, Special Issue on Big data, 574-589, October.

McLaren, N., & Shanbhogue, R. (2011). "Using Internet search data as economic indicators", *Quarterly Bulletin*, 51(2), Bank of England, 134-140.

Meinen, P., & Roehe, O. (2017). "On measuring uncertainty and its impact on investment: Cross country evidence from the euro area", *European Economic Review*, 92, 161-179.

Meinerding, C., Schüler, Y. S., & Zhang, P. (2020). "Shocks to Transition Risk", *SSRN Electronic Journal*.

Menz, J. O., & Poppitz, P. (2013). "Households' disagreement on inflation expectations and socioeconomic media exposure in Germany", *Discussion Paper*, No. 27, Deutsche Bundesbank.

Mittal, A., & Goel, A. (2012). "Stock prediction using Twitter sentiment analysis", *Stanford University, CS229*, 15, 2352.

Moreno Bernal, Á. I., & Caminero García, T. (2022). "Analysis of ESG disclosures in Pillar 3 reports. A text mining approach", *Occasional Paper*, No. 2204, Banco de España.

Morinaga, S., Yamanishi, K., Tateishi, K., & Fukushima, T. (2002). "Mining product reputations on the web", *Proceedings of the 8<sup>th</sup> ACM SIGKDD international conference on Knowledge discovery and data mining*, 341-349.

Moufakkir, M. (2023). "Careful embrace: AI and the ECB", *The ECB blog*, 28 September.

Nasukawa, T., & Yi, J. (2003). "Sentiment analysis: capturing favorability using natural language processing", *Proceedings of the 2<sup>nd</sup> international conference on Knowledge Capture*, 70-77.

Nopp, C., & Hanbury, A. (2015). "Detecting Risks in the Banking System by Sentiment Analysis", *Conference on Empirical Methods in Natural Language Processing*.

Noy, S., & Zhang, W. (2023). "Experimental evidence on the productivity effects of generative artificial intelligence", *Science*, 381(6654).

Nyman, R., Gregory, D., Kapadia, S, Ormerod, P, Tuckett, D., & Smith, R. (2021). "News and narratives in financial systems: exploiting big data for systemic risk assessment", *Journal of Economic Dynamics and Control*, 127, 104119.



- Pang, B., Lee, L., & Vaithyanathan S. (2002). "Thumbs up? Sentiment Classification using Machine Learning Techniques", *Proceedings of the 2002 Conference on Empirical Methods in Natural Language Processing* 79–86.
- Peng, S., Kalliamvakou, E., Cihon, P., & Demirer, M. (2023). "The impact of AI on developer productivity: Evidence from GitHub Copilot", *unpublished manuscript*.
- Petropoulos, A. & Siakoulis, V. (2021). "Can central bank speeches predict financial market turbulence? Evidence from an adaptive NLP sentiment index analysis using XGBoost machine learning technique", *Central Bank Review*, Vol. 21, Issue 4, 141-153, December.
- Pfeifer, M., & Marohl, V. (2023). "CentralBankRoBERTa: a fine-tuned large language model for central banks communications", *The Journal of Finance and Data Science*, Vol. 9, November.
- Püttmann, L. (2018). "Patterns of Panic: Financial Crisis Language in Historical Newspapers", *SSRN Electronic Journal*.
- Rambaccussing, D., & Kwiatkowski, A. (2020). "Forecasting with news sentiment: Evidence with UK newspapers", *International Journal of Forecasting*, 36(4), 1501-1516.
- Ramey, V. A. (2011). "Identifying Government Spending Shocks: It's All in the Timing", *Quarterly Journal of Economics*, 126(1), 1–50.
- Ranco, G., Aleksovski, D., Caldarelli, G., Grčar, M., & Mozetič, I. (2015). "The effects of Twitter sentiment on stock price returns", *PloS One*, 10(9), 1-21.
- Reinhart, C. M., & Rogoff, K. S. (2009). "The Aftermath of Financial Crises", *American Economic Review*, 99, 466-472.
- Rogers, J. H., & Xu, J. (2019). "How well does economic uncertainty forecast economic activity?", *Finance and Economics Discussion Series*, No. 2019–2085, Board of Governors of the Federal Reserve System.

- Rosa, C., & Verga, G. (2007). "On the consistency and effectiveness of central bank communication: Evidence from the ECB", *European Journal of Political Economy*, 23(1), 146-175.
- Ruscheinsky, J. R., Lang, M., & Schäfers, W. (2018). "Real estate media sentiment through textual analysis", *Journal of Property Investment & Finance*, 36(5), 410-428.
- Saiz, A., & Simonsohn, U. (2013). "Proxying for Unobservable Variables with Internet Document-Frequency", *Journal of the European Economic Association*, 11(1), 137-165.
- Schumaker, R. P., & Chen, H. (2009). "Textual analysis of stock market prediction using breaking financial news: The AZFin text system", *ACM Transactions on Information Systems*, 27(2), 1-19.
- Scott, S. L., & Varian, H. (2014). "Predicting the Present with Bayesian Structural Time Series", *International Journal of Mathematical Modeling and Numerical Optimisation*, 5 (1-2): 4-23.
- Scott, S. L., & Varian, H. (2015). "Bayesian Variable Selection for Nowcasting Economic Time Series", in *Economic Analysis of the Digital Economy*, Goldfarb, A., Greenstein, S. M., & Tucker, C. E. (eds), University of Chicago Press.
- Siikanen, M., Baltakys, K., Kannianen, J., Vatrappu, R., Mukkamala, R., & Hussain, A. (2018). "Facebook drives behavior of passive households in stock markets", *Finance Research Letters*, 27, 208-213.
- Sinyak, N., Tajinder, S., Madhu Kumari, J., & Kozlovskiy, V. (2021). "Predicting real estate market trends and value using pre-processing and sentiment text mining analysis", *Real Estate: Economics, Management*, (1), 35-43.
- Smales, L. A. (2023). "Classification of RBA monetary policy announcements using ChatGPT", *Finance Research Letters*, 58.
- Smith, N. A. (2010). "Text-Driven Forecasting", *Carnegie Mellon University Working Paper*.

- Soo, C. (2018). "Quantifying Sentiment with News Media across Local Housing Markets", *The Review of Financial Studies*, 31(10), 3689-3719.
- Sprenger, T. O., Tumasjan, A., Sandner, P. G., & Welpe, I. M. (2014). "Tweets and trades: The information content of stock microblogs", *European Financial Management*, 20(5), 926-957.
- Sudarsa, D., Pathuri, Siva, K., & Rao, L. (2018). "Sentiment Analysis for Social Networks Using Machine Learning Techniques", *International Journal of Engineering and Technology*, 7, 473-476.
- Sun, A.K., Lachanski, M., & Fabozzi, F.J. (2016). "Trade the tweet: Social media text mining and sparse matrix factorization for stock market prediction", *International Review of Financial Analysis*, 48, 272-281.
- Talib, R., Hanif, M. K., Ayesha, S., & Fatima, F. (2016). "Text mining: techniques, applications and issues", *International Journal of Advanced Computer Science and Applications*, 7(11).
- Tetlock, P.C. (2005). "Giving Content to Investor Sentiment: The Role of Media in the Stock Market", *The Journal of Finance*, 62(3), 1139-1168.
- Thorsrud, L.A. (2020). "Words are the new numbers: A newsy coincident index of business cycles", *Journal of Business & Economic Statistics*, 38, 393-409.
- Thorsrud, L. A. (2016). "Nowcasting using news topics. Big data versus big bank", *Working Paper*, No. 20/2016, Norges Bank.
- Tobback, E., Naudts, H., Daelemans, W., Junqué de Fortuny, E., & Martens, D. (2018). "Belgian Economic Policy Uncertainty Index: Improvement through text mining", *International Journal of Forecasting*, 34(2), 355-365.
- Tong, R.M. (2001). "An Operational System for Detecting and Tracking Opinions in On-Line Discussion", *Proceedings of SIGIR Workshop on Operational Text Classification*.
- Varian, H. (2014). "Big data: New tricks for econometrics", *Journal of Economic Perspectives*, 28(2), 3-28.

Woodhouse, D., & Charlesworth, A. (2023). "Can ChatGPT Predict Future Interest Rate Decisions?", *unpublished manuscript*, available at SSRN: <https://ssrn.com/abstract=4572831>

Wolozko, N. (2020). "Tracking activity in real time with Google Trends", *OECD Economics Department Working Papers*, No 1634.

Yu, J., & Yuan, Y. (2011). "Investor sentiment and the mean–variance relation", *Journal of Financial Economics*, 100(2), 367-381.

## Annex 1: Text-mining techniques<sup>23</sup>

All text mining techniques require the pre-processing of the text. Text as data is naturally high-dimensional, due to the variety of words. While Gentzkow et al. (2019) underlined this to be a strength of text data compared to data usually employed in economics, extracting a low-dimensional information is generally required to make any statistical analysis tractable. Bholat et al. (2015) summarize how to pre-process: define the corpus in scope, transform it into a format amenable to analysis, and break the document into tokens (sequence of characters). Another important step is to remove *stop words* and to make tokens more comparable with each other by lemmatizing, stemming, and/or case folding.<sup>24</sup>

Once data are pre-processed, the several text-mining techniques have been classified by Bholat et al. (2015) as either deductive or abductive.<sup>25</sup> The deductive approaches include Boolean search and dictionary text mining while abductive approaches include Latent Semantic Analysis, Latent Dirichlet Allocation (LDA), and Descending Hierarchical Classification. While deductive approaches are simple and scalable, they ignore words that have not been assessed *a priori* to be informative by the researcher. Conversely, abductive approaches analyse all words with the algorithm taking care of assessing their informative power. Yet, programming is more complex and resource consuming. Nevertheless, all these techniques are based on counting the number of occurrences of words in the text, with the underlying assumption that co-occurrence is a reliable indicator of the topic and sentiment conveyed in the text. Word count can be however inappropriate due to the prominence of a small number of very frequent words. Weighting methods answer this issue. One of the most used methods

---

<sup>23</sup> For a more in-depth study on text mining techniques, the authors point to the excellent handbook of Bholat et al. (2015) who provides detailed and intuitive explanation for the different techniques mentioned in this Annex.

<sup>24</sup> “Stop words” are function (filler) words, with little or no specific meaning in themselves (e.g. “and”, “the”, “a”). “Lemmatization” refers to recast words into their linguistic roots. “Stemming” refers to cutting off affixes and counting just stems. “Case folding” refers to converting all alphabetic tokens to lower case. These three last operations allow similar tokens to be made comparable – for example recasting both “airliner”, “airline” and “airlines” into the same token “airlin”.

<sup>25</sup> “Deduction” begins with a general theory and uses specific datasets to evaluate the theory’s applicability. “Abduction” seeks to derive the most plausible explanation for a specific event from a set of evidence.

is the Term Frequency–Inverse Document Frequency (TF.IDF) which multiplies how many times a word appears in a document by the inverse document frequency of the word across a corpus of documents. In more details on each text-mining technique:<sup>26</sup>

- Boolean search works by looking for the presence or absence of specific terms or expressions defined *a priori* in a document. The algorithm allocates a score of 1 (presence) or 0 (absence) depending on whether the terms / expressions are found in the corpus of text or not.
- Dictionary-based techniques also use a list of keywords (i.e. a dictionary) defined *a priori*. Each document is then represented in terms of the (normalized) frequency of the words in the dictionary. The difference with Boolean search is that it measures the *intensity* rather than the mere *presence* of the keywords in the dictionary. Such dictionaries can be used to assess the sentiment of a text, by assigning words to different sentiments. For example, Tetlock (2007) assesses the tone of the “*Abreast of the Market*” column in the Wall Street Journal by using word lists that reflect positive and negative feelings.<sup>27</sup> The development of acute dictionaries has been a long-standing effort in the literature: criticizing the use of a conventional dictionary in Tetlock (2007), Loughran and McDonald (2011) have proposed a vocabulary tailored for financial terms and shown that it was more accurate than generic dictionaries for predicting asset returns. The rationale is that words can have a pejorative connotation in a broader sense but be neutral in a financial context (for example “tax”, “cost” or “liability”).
- Latent Semantic Analysis (LSA) is an abductive approach, whose key idea is to assume that words are not independent but linked together by underlying yet not observed

---

<sup>26</sup> Comparing across different techniques, Ahrens et al. (2024) find that financial transformer models perform best on text-only classification datasets, but struggle on multimodal economics and finance datasets, highlighting more need for concerted research efforts on fusing multimodal data and pre-training or fine-tuning models for tasks in these domains. They also find that the Loughran and McDonald (2011) dictionary underperforms substantially across all datasets. Finally, they show that simple word count models do very well across all datasets.

<sup>27</sup> Specifically, his analysis uses the Harvard IV-4 dictionaries which contain a wide range of word classifications, for example for pain and pleasure, rituals, and natural processes.

topics – and can be traced back to Deerwester et al. (1990). The main advantage of such a latent variable approach is to connect words to topics.<sup>28</sup> One use case is Acosta (2023) who studies the effect of greater transparency of the U.S. Federal Reserve Open Market Committee (FOMC) meetings.

- Yet, LSA has a weakness: the topics produced are not probabilistic. Latent Dirichlet Allocation (LDA; Blei et al., 2003) rectifies this by assigning probabilities for words and documents to be related to different topics. Larsen and Thorsrud (2019) use LDA to decompose major business newspapers according to the topics they cover. Angelico et al. (2022) build a set of daily measures of inflation expectations on the selected tweets combining LDA with a dictionary-based approach with manually labelled bi-grams and tri-grams.<sup>29</sup>
- Descending Hierarchical Classification, as explained in Bholat et al. (2015), uses a partitioning algorithm to identify statistically significant relationships between words and elementary context units. The algorithm seeks to maximise the Chi-2 values of a contingency table, and tests if further splits improve the Chi-2 values. The iterative process comes to conclusion when a predetermined number of iterations no longer results in statistically significant divisions. Like LSA, Descending Hierarchical Classification is another text mining technique that proceeds on the basis that words are not independent of each other but reflect underlying topics/themes.
- Deep Learning-Based Sentiment Analysis leverages deep learning architectures, such as neural networks, to capture complex patterns and relationships within textual data for sentiment classification. One notable example is Bidirectional Encoder Representations from Transformers (BERT). Transformer-based models aim at capturing contextual information and semantic relationships within language. Unlike traditional models that process text in a sequential manner, BERT and its peers employ

---

<sup>28</sup> A latent variable is a variable that is inferred using models from observed data.

<sup>29</sup> An *n-gram* is a contiguous sequence of *n* words. It is used to predict the occurrence of a word based on the *n* – 1 previous word. When *n* = 2 it is called a bi-gram and when *n* = 3, a tri-gram.

bidirectional attention mechanisms, enabling them to consider both preceding and succeeding words when analysing a given word's context in a sentence. This bidirectionality allows for a more nuanced understanding of language semantics, making these models highly effective in various NLP tasks such as text classification, question answering, and language translation.

- Large Language Models (LLMs) encompass a broader category of models beyond specific architectures like BERT, encompassing expansive neural networks designed to handle vast amounts of textual data. These models, such as GPT-3 (Generative Pre-trained Transformer 3), are characterized by their immense scale, often comprising billions or even trillions of parameters. LLMs can perform diverse natural language tasks without task-specific training, showcasing a degree of versatility. LLMs leverage pre-training on massive text corpus, enabling them to generate coherent and contextually relevant text.