Irving Fisher Committee on
Central Bank Statistics

IFC-Bank of Italy Workshop on "Data science in central banking"

18-20 February 2025

A natural language processing toolbox for the National Bank of Romania[1]

Claudia Voicilă and Daniel Șerbu,
National Bank of Romania

[1] This contribution was prepared for the workshop. The views expressed in this publication are those of the authors and do not necessarily represent the official views of the Committee, its members, or the BIS.

A Natural Language Processing toolbox for the National Bank of Romania

Claudia Voicilă, Daniel Șerbu[1]

Abstract

This paper advocates extending beyond traditional numerical data by enhancing central bank analyses with textual information. We propose integrating Natural Language Processing (NLP) tools into the National Bank of Romania's (NBR) analytical toolkit to align with the EU AI Strategy, national and international digitalization priorities, and Romania's objectives to strengthen financial system supervision in line with OECD standards. To this end, we introduce an NLP toolbox specifically tailored for the NBR, aimed at increasing the efficiency of economists in processing textual sources such as financial news and press releases. The toolbox's primary objective is to align the NBR with the best practices in central banking innovation.

The first part of the analysis examines the use of NLP in processing monetary policy statements, starting with the NBR's adoption of inflation targeting. The dataset was expanded to include statements from the European Central Bank, the U.S. Federal Reserve, and 25 other central banks from emerging economies. A supervised sentiment analysis using built labels was conducted by calibrating FinBERT, a Transformer-based model designed to capture contextual relationships in financial text. This model enables more accurate sentiment classification by identifying domain-specific terms that frequently convey either negative or positive sentiment.

The second part of the paper explores the NBR's daily Romanian financial news archive. Here, a specialized Romanian financial dictionary was developed to build a sentiment index. Sentiment intensity scores were assigned using a Generative AI agent. Additionally, topic modelling via Latent Dirichlet Allocation (LDA) was employed to identify and track the evolution of major news topics, highlighting key themes in the financial news over time.

Keywords Natural Language Processing, Financial Stability, Central Bank Communication

JEL classification: E58, C55

1. Introduction

The domain of Natural Language Processing in Central Banking is still a novelty, different applications emerging constantly, as the desire to keep up with the innovation in commercial banks is high. Despite the level of complexity being lower

[1] The opinions are those of the authors and do not necessarily reflect the views of the National Bank of Romania

than research in computer science areas, economists from central banks around the world gather often to discuss ideas about digital transformation in banking. This paper aims to provide an overview of common NLP analyses on text data from the National Bank of Romania, offering a fresh perspective on a task that has not been explored in depth within the institution. Therefore, the main goal of the present work is to align NBR to the good practice of central banks and get inspired by the rapid technological changes of the last years in terms of text analysis. Building a toolbox for Romanian language is not an easy task as it is not a widely spoken language and since Romanian has many different particularities than English, the language of most NLP tools.

Communication of central banks has an important role in ensuring the effectiveness of monetary policy and supporting the credibility of the institution (on a more general note). In a way, communication can be considered a monetary policy tool (due to the impact it has on inflation expectations and economic behaviour) but can also reveal financial stability concerns. A central bank can communicate in many ways: monetary policy decisions, governor and other board members' speeches, reports (Inflation Report, Financial Stability Report) and presentations of spokespersons etc.

In monetary policy decision statements words are chosen with extreme care, as even the slightest nuance can influence financial markets and shape economic expectations. Usually, a central bank's communication is likely to be efficient only if it is perceived as a credible institution by the public. The relationship between the bank's credibility and effective communication has long been recognized by central bankers, as credibility is an important factor in shaping economic agents' expectations.

This paper proposes a series of custom-made NLP tools for analysing financial news in Romanian and Monetary Policy Statements in both Romanian and English, offering insights into the interaction of central bank communication and financial market perception. Sentiment analysis plays a key role in assessing how the press reacts to monetary policy decisions and that's an essential factor for monetary policy efficiency. At the same time, topic modelling allows us to uncover the most relevant themes in financial news over time, providing a clearer list of subjects of public interest. By leveraging these tools, we aim to enhance the understanding of how monetary policy decisions are communicated, perceived, and ultimately, how they shape financial expectations.

2. English Monetary Policy Statements Database

Monetary policy is the primary responsibility of a central bank, with price stability usually as its foremost objective. Generally, words are chosen with extreme care in monetary policy decision statements, as even the slightest nuance can influence financial markets and shape economic expectations. A central bank's communication is effective only when it is perceived as a credible institution, as credibility is essential for shaping economic agents' expectations and has long been recognized as a key factor by central bankers.

It is possible to get a sense of how efficiently central banks communicate by analysing the clarity, coherence, and influence of their messages on financial markets, economic expectations, and public sentiment. Using NLP techniques, we can also examine how well central banks convey important policy decisions by identifying key events—such as financial crises, policy shifts, and macroeconomic shocks—in textual data. This helps us understand how central banks respond to these events and the broader economic impact of their communication strategies.

The monetary policy statement released by the central bank provides insights into current economic conditions and future outlooks, while also highlighting a range of market risks and vulnerabilities. We employed methods for automated content analysis to extract sentiment from each monetary policy statement since the NBR adopted inflation targeting (2005). By developing a sentiment indicator that measures the frequency of positive or negative words in each board decision published on the website, we constructed credibility indexes for monetary policy based on these statements.

We have web-scraped using a tool called Beautiful soup by Richardson (2007) the monetary policy decisions from the press releases of the Board of the National Bank of Romania that are also published in English (apart from the original Romanian text. The National Bank of Romania released a total of 164 monetary policy statements during 2003 and 2024, out of which 50 were linked to loosening monetary policy (meaning cutting down the reference interest rate), 26 were related to tightening monetary policy (meaning raising the reference interest rate), whereas 88 monetary policy decisions were neutral, meaning they maintained the same level as before for the monetary policy rate. The main factor is the inflation rate, as the purpose of the Romanian central bank is to ensure and maintain price stability, with an inflation target of 2.5% ±1 p.p. since 2013 (Figure 1).

Evolution of the monetary policy rate and the target inflation rate

Figure 1

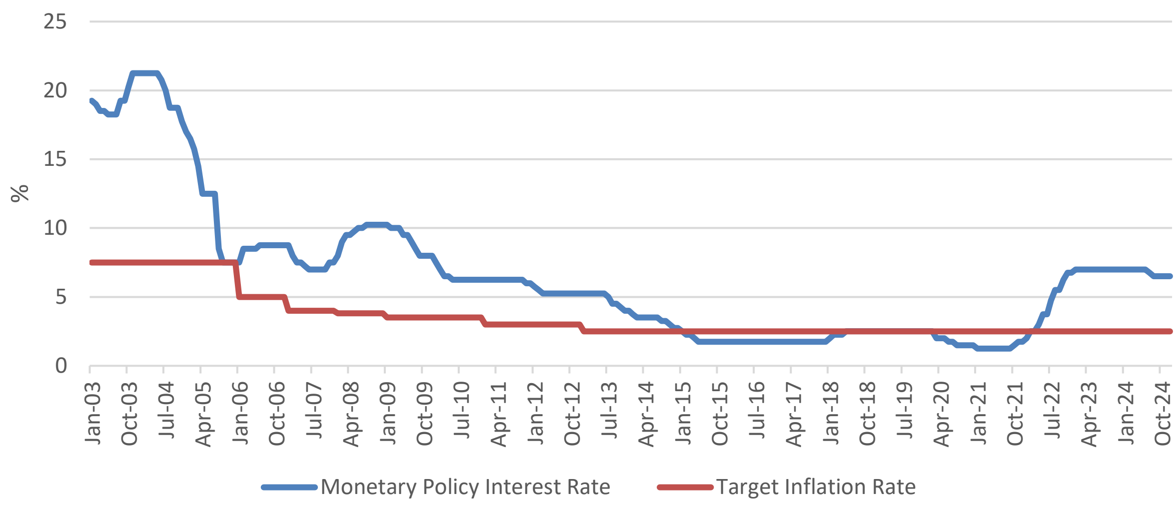

Source: NBR

We enlarged our monetary policy database with monetary policy statements in English from 25 emerging market central banks (Asia: Armenia, Georgia, Israel, India, Kazakhstan, Malaysia, Mongolia, Pakistan, Philippines, South Korea and Thailand; Europe-Asia: Turkey and Russia; Europe: Czech Republic, Hungary, Poland and Ukraine, South America: Brazil, Chile, Colombia and Peru; Africa: Egypt, Nigeria and South Africa; North America: Mexico) along with the Federal Reserves in the United States and the European Central Bank’s press-conference introductory statements that rules over the eurozone countries (Austria, Belgium, Croatia, Cyprus, Estonia, Finland, France, Germany, Greece, Ireland, Italy, Latvia, Lithuania, Luxembourg, Malta, the Netherlands, Portugal, Slovakia, Slovenia and Spain). Using a larger data set we can underline common global trends and better understand the monetary decisions of the National Bank of Romania since they are strongly influenced by the decisions made by European Central Bank (ECB) and the Federal Reserve’s Board (FED) in the USA. The extended database contains 4112 monetary policy statements from 1998 to 2023.

The extended dataset was gathered by Evdokimova et al (2023) and offered by Araujo (2023) that introduced "gingado", an open-source Python library designed to enhance the use of machine learning in the fields of economics and finance. This tool has relevant datasets and integrates seamlessly with major machine learning libraries and offers functionalities that simplify the augmentation of user datasets with authoritative data via the SDMX protocol. Gingado includes a benchmarking feature that optimizes random forest models for robust performance across various scenarios without extensive user input. Furthermore, the library aids in the rigorous documentation of machine learning models, incorporating ethical considerations into the documentation process. The platform is under continuous development, with new features being added to address the specific needs of economic and financial data analysis.

We assigned labels to our monetary policy statements database based on the reference monetary policy interest rate obtained from the Refinitiv database. Based on the monetary policy rate at each decision date, we created three categories: positive sentiment if the interest rate was lowered (loosening monetary policy), negative sentiment if the interest rate was raised (tightening monetary policy), and neutral if the interest rate was unchanged. We then regrouped the neutral category as positive if the previous decision was positive (indicating that the positive economic sentiment continues) and negative if the previous decision was negative (indicating that the negative economic sentiment continues), but only for interest rate levels around 3%, which is considered a low value.

On the extended English Monetary Policy Statements database with labels, we leveraged a pre-trained FinBERT model by ProsusAI, which is specifically designed for financial sentiment analysis. This model, originally trained on financial texts such as SEC filings, earnings calls, and analyst reports, is better suited for monetary policy sentiment classification than general-purpose models like our first option, DistilBERT by Sanh (2019). While DistilBERT is optimized for speed and trained on general sentiment datasets (such as SST-2), it lacks domain-specific knowledge, making it less effective in capturing the nuances of central bank communication.

To enhance sentiment analysis accuracy in monetary policy statements, we fine-tuned FinBERT on a custom dataset of central bank communications, classifying text statements as either positive or negative (removed the neutral sentiment).

Additionally, we extended its application to score individual words based on their sentiment polarity, allowing a deeper exploration of sentiment-laden terms within monetary policy.

The performance of the fine-tuned FinBERT model on the sentiment classification task demonstrates strong accuracy of 90% on the test set (Table 1) Training over 10 epochs resulted in a notable reduction in training loss, particularly in early epochs, signifying effective learning. In terms of classification metrics, the model achieved a precision of 86.5% for negative sentiment and 92.5% for positive sentiment, with recall values that indicate balanced performance across both classes. The F1-scores of 87.8% (negative) and 91.6% (positive) further confirm FinBERT's ability to classify sentiment with high accuracy, ensuring reliable sentiment extraction from monetary policy statements.

Classification report of fine-tuned FinBERT on test sample (10 epochs) Table 1

	Precision	Recall	F1-score	Support
Negative	86.52%	89.19%	87.82%	259
Positive	92.53%	90.60%	91.56%	383
Accuracy			90.03%	642
Macro accuracy	89.53%	89.89%	89.69%	642
Weighted average	90.11%	90.03%	90.05%	642

Source: own computations

After training, we saved the fine-tuned FinBERT model and tokenizer for future use. Given its specialization in financial text analysis, we leveraged this pre-trained model to extract and validate a list of positive and negative words in English, which served as the foundation for our sentiment dictionary. Rather than using the model for document-level sentiment scoring, we focused on identifying sentiment-laden vocabulary relevant to the financial domain. This approach highlights FinBERT's utility not only in contextual sentiment classification but also in supporting the development of fine-grained lexical resources tailored to central bank communication analysis.

Confusion Matrix of FinBERT on test sample from the extended English database Figure 2

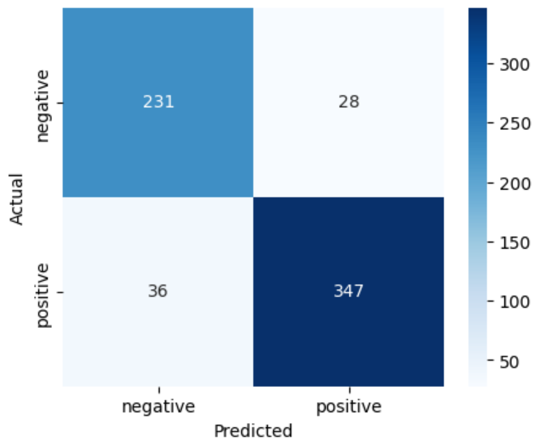

Source: own computations

These results (Figure 2) indicate that FinBERT is probably the most effective choice for financial text analysis, capturing nuanced policy sentiment and providing a robust tool for assessing central bank communication. We used the fine-tuned on monetary policy statements FinBERT model as a starting point for the list of words in the specialized proposed dictionary, however, a tool like this can be very useful in a central bank, as it can analyse sentiment trends over time, assess the impact of policy decisions on financial markets, extract key phrases and topics for explainability, identify comparable past policy statements, classify monetary policies based on sentiment, and even predict market reactions to central bank communications.

3. Romanian Financial News Articles Database

Crafting NLP Tools for Romanian Financial News

The National Bank of Romania has an archive of news starting from 1998, manually collected and selected by employees in the Communication Department, making the text relevant and eliminating the need to filter out unrelated news. The older news, scanned from physical newspapers, requires optical character recognition (OCR). In total, there are 27 years of news, amounting to more than 4,000 PDFs, each containing approximately 30-40 daily news articles. Each day reflects the news from the previous working day. Due to time constraints and limitations in the length of this paper, we will focus on the last 7 years: 2018-2024.

In the preprocessing of Romanian text mainly with NLTK by Loper (2002), several systematic steps were undertaken to ensure the data was appropriately cleaned and prepared for further analysis. The process began with the removal of any watermarks from the text and was followed by the normalization of specific Romanian characters to their correct forms: for example, characters like 'ş' with cedilla were replaced with 'ș' with comma, and 'ţ' with cedilla with 'ț' with comma, ensuring consistency in the representation of Romanian diacritical marks[2]. Next, whitespace normalization was applied to the text, reducing any irregular spacing to single spaces, thereby creating a uniform text structure. Punctuation was then systematically removed using Python's string translation method, which helped in simplifying the text by eliminating non-alphabetic characters that could interfere with tokenization and subsequent analysis.

Following the removal of punctuation, any numerical digits present in the text were filtered out. This step was crucial to focus the analysis on the linguistic content of the text, excluding numeric data which was not relevant for the sentiment and topic modelling tasks at hand. The cleaned text was then processed through Stanza's Romanian language pipeline by Qi et al. (2020), which included tokenization, part-of-speech tagging, and lemmatization. During this stage, the text was converted to lowercase to maintain consistency and reduce redundancy. Lemmatization was employed to reduce words to their base dictionary form, therefore normalizing

[2] There is no uniform approach in the news archive, even though in 2003, a scientific council of the Institute of Linguistics of the Romanian Academy determined that the correct sign is the comma, not the cedilla.

different inflected forms of words to a common base. We used spaCy's NER model to identify named entities labelled as *PERSON* from the Romanian text corpus.

Simultaneously, stop words—common words in Romanian that do not carry significant meaning—were removed from the text. This was done to ensure that the analysis focused on the more meaningful words that contribute to the overall sentiment and thematic content of the text. Additionally, only alphabetic tokens were retained, with non-alphabetic tokens being discarded to further refine the dataset. The resulting tokens, which were now cleaned, lemmatized, and stripped of irrelevant content, were then stored for subsequent analysis. This comprehensive preprocessing ensured that the Romanian text data was in an optimal state for accurate and meaningful linguistic analysis, sentiment scoring, and topic modelling.

One of the most useful NLP tools is topic modelling which can be applied to news such as in our case, but also central bank reports, economic reviews and monetary policy statements to identify key topics and main themes discussed by the policymakers. This tool can help identify emerging trends, priorities and risks in monetary policy. Using this topic modelling analysis, we can track main trends in the news and this type of list can be useful for decision makers at the central bank level.

We constructed a dictionary from the pre-processed Romanian text data using the Gensim library, which systematically mapped each unique term in the corpus to a specific integer ID. By doing so, we established a comprehensive lexicon that enabled us to translate the textual content into a numerical format suitable for computational analysis. To refine this dictionary, we implemented a filtering process that excluded words appearing in fewer than 20 documents and those present in more than 50% of the documents. This step ensured that our analysis focused on terms that were neither too rare nor overly common, thereby raising the model's ability to discover meaningful topics within the text.

Following the dictionary construction and filtering, we generated a Bag-of-Words (BoW) representation for each document. This representation provided a numerical summary of the documents, capturing the frequency of each term in the text as per the dictionary's mappings. Using this BoW format, we applied Latent Dirichlet Allocation (LDA) according to Blei et al.(2003) to the corpus to identify underlying topics. The LDA model was trained with four topics, and we specified parameters to control the training process, including the number of passes over the data and the update frequency. Upon completion of the training, we extracted and examined the topics, each represented by the most significant words associated with it. Additionally, we employed pyLDAvis by Mabey (2018) to visualize these topics, enabling an interactive exploration of the thematic structures present within the documents. This visualization provided a deeper understanding of the distribution and significance of each topic across the corpus.

We identified the five most prominent topics along with their associated keywords, providing a structured overview of the key thematic areas within the corpus for each year (Table 2). In addition to analysing the relative importance of terms within each topic, we also examined the temporal evolution of these themes to assess how their significance has shifted over time (Figure 3). This allowed us to capture emerging trends and changing priorities reflected in the Romanian text data, offering a dynamic perspective on the underlying discourse.

LDA topic analysis – top 10 words per topic for each year

Table 2

	2018	2019	2020	2021	2022	2023	2024
Topic 1- Global Economy, Gold Reserves & Geopolitics	Bancpost, cryptocurrency, UniCredit, Piraeus, Bitcoin, BERD, clause, card, OTP, insolvency	Parliament, Politician, pillar, deputy, greed, IRCC, political-party, treasury, hearing, adviser	Growth, QE, advisor, monetary policy, anticipation, capitalism inflation, monetary	CSALB, ton, insolvency, Treasury, Russian, laundering, grain, governance, corporate, circumstantial	Ruble, grain, ton, oil, Senate, gas, tanker, metal, Gazprom, federation	Minister, political-party, Fidelis, BT, exemption, SME, polish, Israel, OUG	Tezaur, Trump, corporation, ton, Moscow, election, trillion, CEC, Russian, candidate
Topic 2 - Interest Rates, Inflation & Housing Market	Greed, OUG, card, Turkey, quarterly, gov order, revision, normative, pillar, pound	MF, prim-minister, revision, execution, liability, RONEUR, pillar, president, correction	Governor, CSALB, ASF, pillar, file, president, Bitcoin, S&P, polish, reconcile	Dividends, UniCredit, anticipation, adviser, aggregated, absorption, residential, capping, growth	Schengen, admission, Austrian, ASF, pillar, Erste, corporation, SME, FTX, administrator	Suisse, Swiss, UBS, Silicon-Valley, SVB, First Bank, Switzerland, depositor, cryptocurrency	IRCC, governance, capping, instant, ton, poverty, CEC, margin, expensive, mortgage
Topic 3 – Crypto currency, Financial Markets & Digital Assets	Convergence, Italian, poverty, inhabitant, intermediation, circulation, Brussels, insolvency, union, Deutsche	Prim-minister, Chinese, EximBank, RONEUR, ton, polish, Libra, crypto, inhabitant, poverty	UniCredit, user, term, grant, Revolut, FNGCIMM, unemployed, Bitcoin, telephone, authentication	City-Insurance, ASF, court, political-party, RONEUR, blockchain, constitutional, insolvency, deputy	Lombard, HICP, interest rate, Schengen, associate, anticipated, inflation, decrease	ANPC, OTP, Revolut, Tezaur, Intesa, inclusion, ARB, Moscow, SanPaolo, CFA	Trump, Bitcoin, Revolut, flat, cryptocurrency, Fitch, elections, gov decision, amnesty, Alpha
Topic 4 - Social & Economic Trends, Policy & Regulations	Pillar, Parliament, capping, president, political-party, Senate, prim-minister, Bancpost, deputy, saving	Alpha, OTP, adviser, Lagarde, convergence, Christine, deputy-governor, revision, profitable, wealth	Vaccine, pandemic, trillion, S&P, lockdown, associate, provision, rescheduling, Biden, mobility	Ton, OTP, Erdogan, pound, inequality, wheat, wealth, privatization, Turkey, MF	CSALB, revision, postponing, non-reimbursable, execution, MF, court, taxation, duties, suspension	Bitcoin, interbank, lombard, insolvency, crypto currency, decrease, trillion, non-performing, ton, associate	Spontaneous, tradition, correction, abstract, freedom, conduct, decrease, knowledge, western, OTP
Topic 5 - Legal & Financial Disputes, Banking & Consumer Protection	Food, OTP, ASF, repo, vegetables, non-food, expensive, abusive, facility, fuel	CSALB, execution, debt, clause, resolution, unpredictable, abusive, litigation, S&P, forced	Duties, RONEUR, grant, Swiss, fuel, taxpayer, Erste, CFA, EximBank, ton	Banknote, coal, revision, taxpayer, compensation, ANAF, HICP, manufacturing, capping, OPEC	Revolut, user, instant, respondent, crypto, corporation, Swiss, Suisse, Tezaur, pound	CSALB, SIF, Alpha, pillar, privatization, ASF, court, resolution, litigation, solution	CSALB, Moldova, court, resolution, Republican, litigation, Intesa, Alpha, farmer, SanPaolo

[1] Mentions of individuals were automatically identified using Named Entity Recognition (NER) with the spaCy language model. Entities labelled as PERSON were later replaced manually with institutional roles to protect personal data.

[2] The words were translated from Romanian to English for reader's understanding

Source: created by the authors

Heatmap of the Importance of the Main Five Topics Over Timef

Figure 3

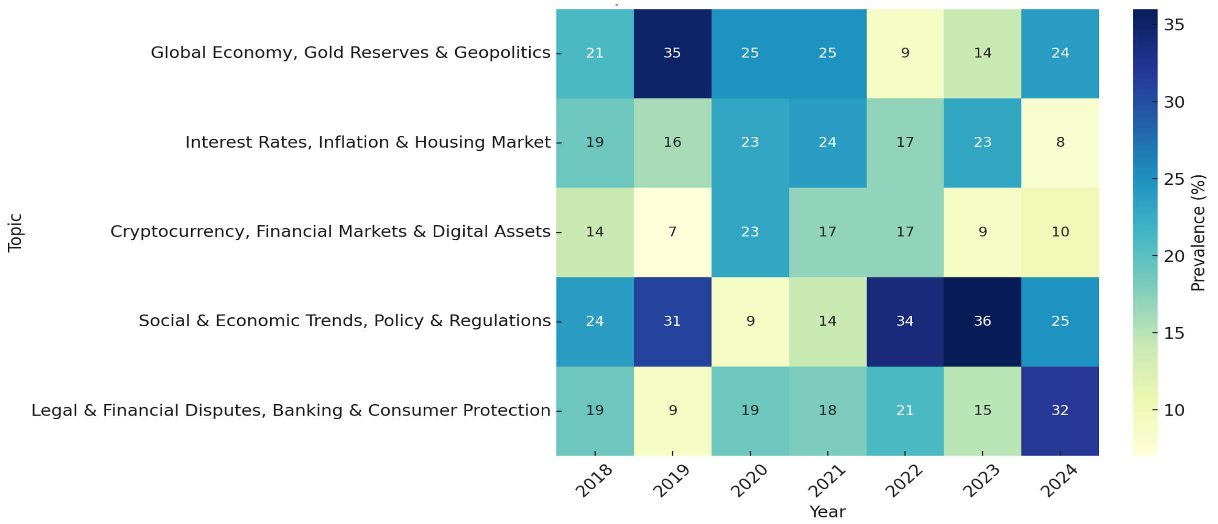

Source: own computations

After constructing the LDA model, we aggregate word frequencies across all topics to create a comprehensive view of the most significant terms within the corpus. In LDA, each word receives a weight within each topic, reflecting its probability of occurrence given that topic. To assess the most thematically significant terms for the corpus, we computed a cumulative score for each word by summing its weights across all topics. Formally, for each word *w*, its overall importance is calculated as:

$$LDA\ Weighted\_Score(w) = \sum_{t=1}^{T} \phi_t(w)$$

Where $\phi_t(w)$ denotes the probability of word w in topic t, and T represents the total number of topics (in our case, T=5). This aggregated score captures the extent to which a word is representative across the entire thematic structure of the corpus, rather than within a single topic alone. We then ranked all terms according to their LDA-weighted scores and extracted the top 10 words per year. These words are visualized and reported as key indicators of the dominant semantic patterns found in the annual financial news review.

An analysis of the top 10 words aggregated across all five LDA topics for the year 2018 points out that the year (Figure 4) was marked by the wide debate on the "tax on greed," which aimed to impose a significant tax on bank profits linked to the level of interest rates, sparking concerns about financial sector stability and economic impact. Additionally, in 2018, Romania experienced some discussions surrounding the country's second-pillar pension system reform, which is a mandatory private pension fund scheme, and Bancpost, one of Romania's major banks, was acquired by Banca Transilvania, marking a significant consolidation in the Romanian banking sector.

Top 10 words based on LDA topic-weighted scores (Year 2018/ 2019) Figure 4/Figure 5

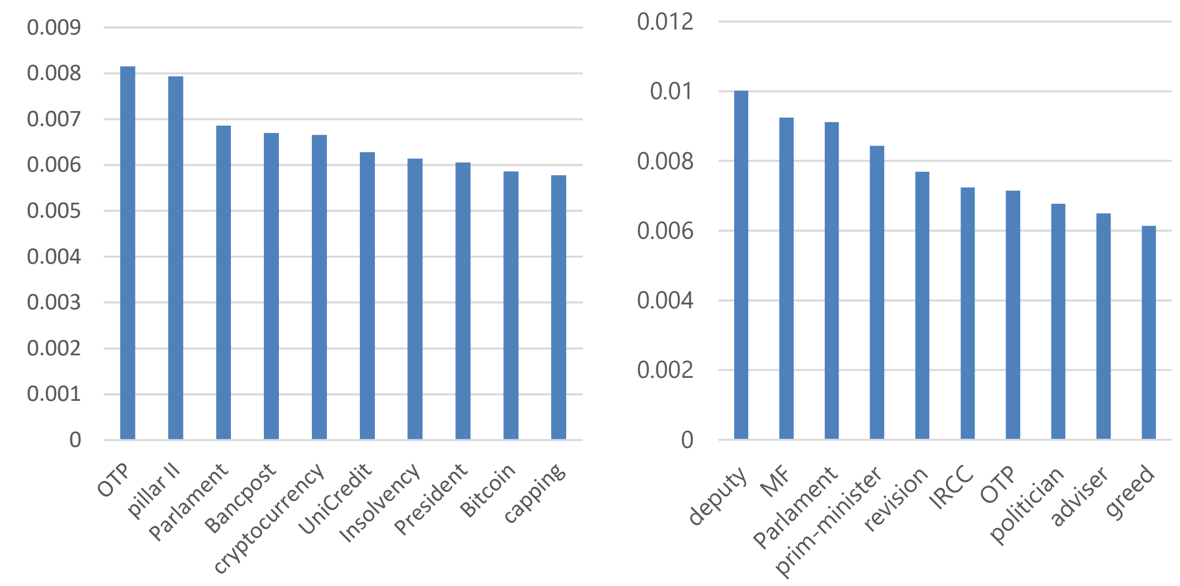

[1] Note: the words were translated from Romanian to English

Source: created by the authors

In 2019 (Figure 5), political discourse was dominated by themes of political instability and leadership challenges, largely shaped by electoral dynamics. Legal controversies involving prominent political figures influenced governance and public trust, while opposition efforts focused on contesting fiscal policies, particularly in the context of budgetary decisions that raised concerns over public finances. At the same time, legislative initiatives aimed at consumer protection in financial services gained traction, reflecting broader discussions on financial stability and regulatory oversight. Overall, the year was marked by intense political debates and policy disputes, highlighting the intersection of governance, economic decision-making and public confidence.

In 2020 (Figure 6), Romania tackled the COVID-19 pandemic with various public health measures and economic interventions, while the National Bank of Romania implemented monetary policies to stabilize the economy amidst the global crisis. Simultaneously, the country navigated the impacts of Brexit, while Revolut expanded its services in Romania, adapting to the changing financial landscape. In 2021 (Figure 7), Romania faced rising inflation, prompting companies to reassess dividend distributions, while the CSALB (Center for Alternative Dispute Resolution in Banking) continued to mediate disputes between consumers and banks. Additionally, there was growing interest in blockchain technology that can be attributed to the significant increase in valuation of the cryptocurrency sector, with increased discussions around its potential impact on the financial sector.

Top 10 words based on LDA topic-weighted scores (Year 2020/ 2021) Figure 6/Figure 7

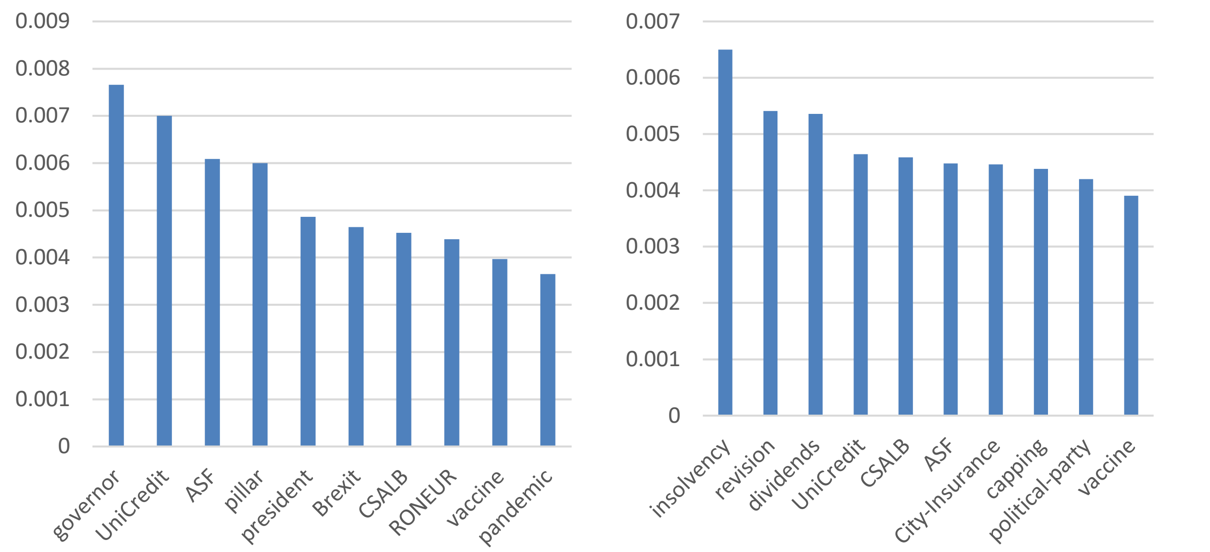

Source: created by the authors

In 2022 (Figure 8), Romania's efforts to join the Schengen Area faced ongoing delays, while the depreciation of the Russian ruble amid global economic instability influenced regional trade dynamics. Additionally, Romania grappled with rising inflation, which impacted purchasing power and economic stability. In 2023 (Figure 9), the turmoil surrounding Credit Suisse's acquisition by UBS and the collapse of Silicon Valley Bank had ripple effects in Romania, contributing to increased market volatility and heightened concerns among investors and financial institutions. The events underscored the interconnectedness of global financial systems, leading to cautious responses from Romanian banks and regulators to ensure stability in the face of external shocks.

In 2024 (Figure 10), the key themes emerging from unsupervised analysis of public discourse in Romania highlight the prominence of electoral events, financial sector developments, and digital asset trends. The U.S. presidential election and Romania's local elections dominated public interest, reflecting the growing impact of political dynamics on economic expectations. At the same time, the rise of cryptocurrencies, particularly Bitcoin, remained a focal point, driven by increased institutional adoption and regulatory discussions. In the financial sector, Revolut's decision to establish a local branch in Romania marked a significant milestone in the country's fintech evolution, enhancing competition in digital banking. Additionally, UniCredit's acquisition of Alpha Bank signalled further consolidation in the banking industry, potentially reshaping market dynamics and credit accessibility. These trends underscore Romania's integration into global financial and technological shifts while navigating the challenges of an election year.

Top 10 words based on LDA topic-weighted scores (Year 2022/2023) Figure 8/Figure 9/ Figure 1010

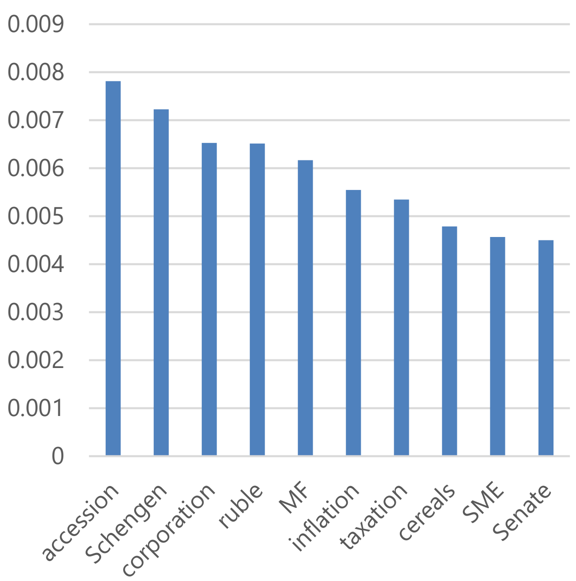

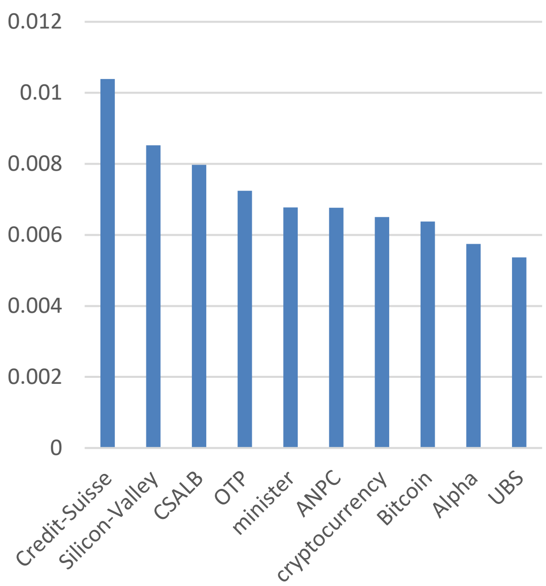

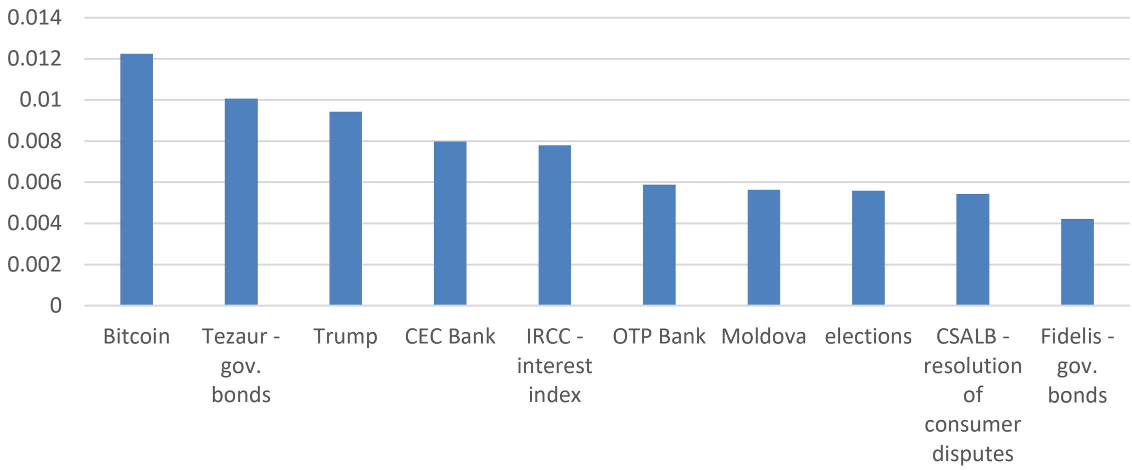

[1] Additional information was added for reader's better understanding

Source: created by the authors

Building the Romanian Financial Stability Dictionary

There are two primary approaches for sentiment analysis that do not use LLMs: supervised learning, which requires many labelled documents, and rule-based, which relies on a previously built dictionary of positive and negative words that can be adapted to the subject of the documents. In our case, we do not have readily available labels and creating them would require an immense effort to go through the daily text due to its sheer size and length, which would also be subjective to the reader to some extent. The best approach, given the limitations, was to build a financial stability dictionary for the Romanian language starting English words identified by FinBERT and fine-tuned with expert-level knowledge in Romanian.

General tools included the lexicons by Hu and Liu (2004), Baccianella et al. (2010) with SentiWordNet 3.0, and VADER by Hutto and Gilbert (2014), which are tailored to consumer reviews and online content but proved inadequate for technical financial texts. Financial-specific dictionaries such as Loughran and McDonald (2011), designed for corporate filings, and the IMF Financial Stability Dictionary by Correa et al. (2017), better captured sentiment but still showed limitations. While the IMF dictionary yielded the most relevant results, its effectiveness was constrained by language, as NBR communications are originally in Romanian. Consequently, the study highlights the need to develop a domain-specific sentiment dictionary in Romanian to improve analysis accuracy in the context of central banking and financial news.

Our positive and negative list of words was draft from FinBERT results in English and was also inspired by the IMF dictionary, the best English option for monetary policy statements. However, the Romanian language has many particularities, so simply machine translating from English would result in poor-quality outcomes. For that reason, we manually translated the terms and ended up with a list of Romanian words that did not perfectly match the English ones therefore we expert added specific Romanian terms. We also investigated RoWordNet by Dumitrescu et al. (2018) to find all possible synonyms and antonyms for our Romanian words on the existing list.

The resulted Romanian Dictionary focuses on Financial Stability issues and can be used to measure the sentiment from the Romanian news database. It contains 150 positive words in lemma form[3] and 182 negative terms in lemma form. We had the list assessed by Gen-AI model (ChatGPT) who was prompt as a Romanian financial stability expert and had to score positive and negative terms on an intensity scale: -0.25 (ambiguous), -0.5 (quite negative), -0.75 (clear negative), -1 (strong negative), and 0.25, 0.5, 0.75, to 1 for positive.

The formula to evaluate the sentiment index is simple and relies on the dictionary with polarities (positive/negative) and their intensity score levels (weights) for each lemma.

For any day and document (t) the simple sentiment index is computed (Figure 11):

$$Sentiment\ Index = \frac{P(t)-N(T)}{P(t)+N(t)} \quad (1)$$

Where:

- P(t) is the list of positive lemmas in the text t that matches the positive words in the dictionary
- N(t) is the list of negative lemmas in the text t that matches the negative words in the dictionary

[3] A lemma is the base or dictionary form of a word, used in linguistic processing to group inflected forms (e.g., "running," "ran," and "runs" all share the lemma "run").

For any day and document (t) the sentiment index using scores (Figure 15):

$$Sentiment\ Score\ Index = \frac{\sum_{i=1}^{P(t)} S_p(i) + \sum_{j=1}^{N(t)} S_n(j)}{\sum_{i=1}^{P(t)} S_p(i) + \sum_{j=1}^{N(t)} |S_n(j)|} \quad (2)$$

Where:

- P(t) is the list of positive lemmas in the text t that matches the positive words in the dictionary
- N(t) is the list of negative lemmas in the text t that matches the negative words in the dictionary
- $S_p(i)$ the score (intensity) of the i-th positive lemma in P(t) according to the positive words dictionary
- $S_n(i)$ the score (intensity) of the j-th positive lemma in N(t) according to the negative words dictionary

Sentiment Index of Monetary Policy using Romanian Financial Stability Dictionary (without and with scores)

Figure 11/ Figure 12

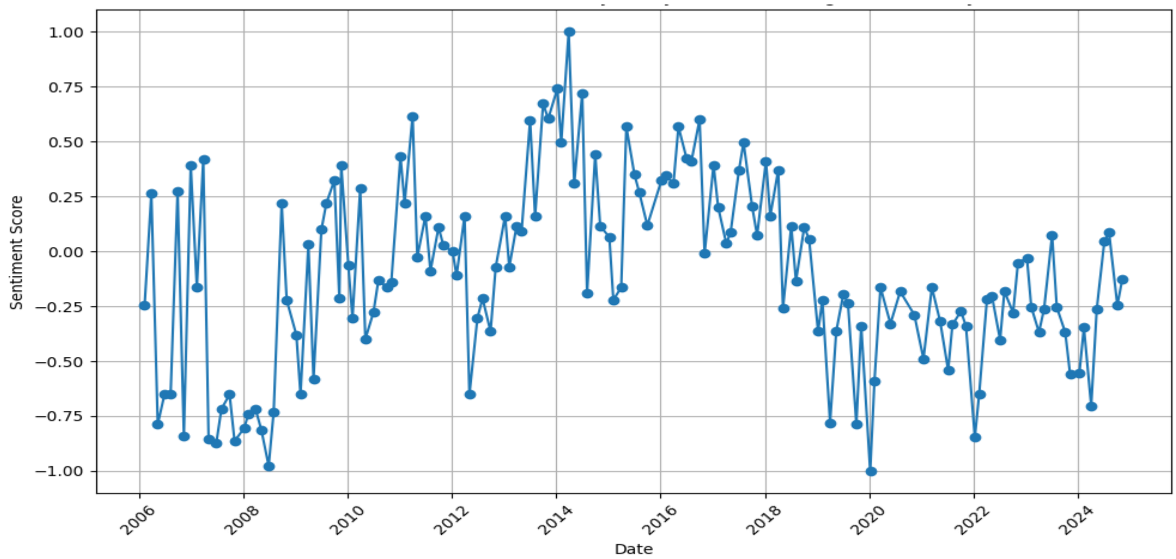

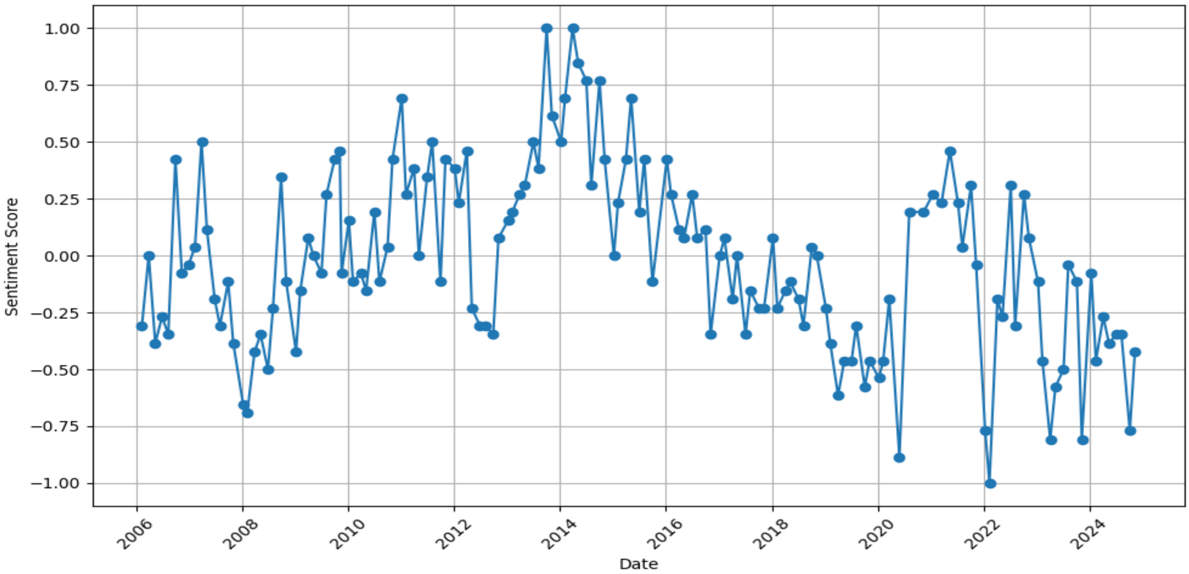

Source: authors' own computations

By leveraging the Romanian Financial Stability Dictionary, we managed to correctly identify trends in monetary policy that the unspecialized method failed to identify, such as the global financial crisis of 2007-2008, the 2012 NBR intervention of raising the monetary policy rate due to rising inflationary pressures related to the Eurozone crisis, the 2014 peak of positive sentiment driven by strong economic growth that boosted consumer confidence, the 2019 drop due to the proposed 'greed tax' for on banks and energy companies, as well as the polarized presidential elections and the 2022 exponential rise in inflation that led to 8 negative interventions by the NBR, etc.

We tested our results against binary labels that we've built starting from the movement of the monetary policy interest rate (series from Refinitiv): positive sentiment if the monetary policy rate was lowered in that meeting, and negative if the monetary policy rate was raised. All neutral decisions, meaning meetings where the monetary policy rate was maintained, were labelled with the same sentiment as the last non-null decision movement; therefore, we considered that the sentiment was maintained until further notice.

The version of the sentiment index with scores outperforms the simple sentiment index by 2% (78% versus 76%), and the final index that we obtained by mixing the two methods matches 92.5% of the labels. This performance is considered excellent, given that neutral decisions were made using the last decision rule and not manually evaluated by a human (Figure 13).

Sentiment Index of Monetary Policy versus Binary NBR Decision Labels Figure 13

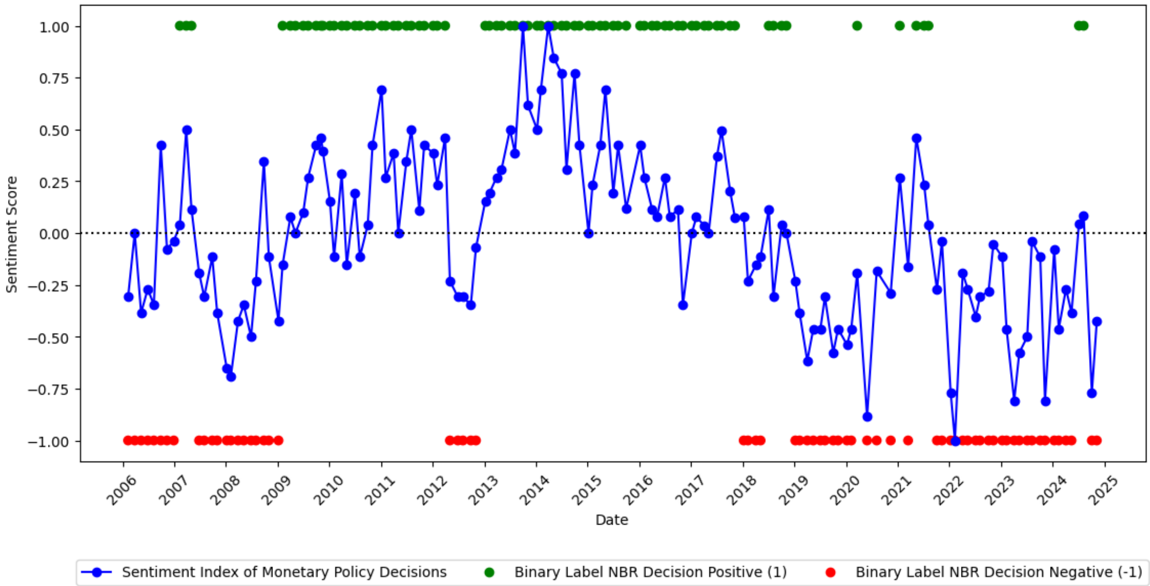

Source: authors' own computations

To assess the performance of our NLP analysis, we examined the correlation between the sentiment index of monetary policy decisions and the most relevant economic variable: the inflation rate, since controlling inflation is the primary goal of monetary policy. We found an expected negative correlation between our computed sentiment and the all-items HICP (Harmonized Index of Consumer Prices) of over 40%, while the benchmark correlation between the monetary policy interest rate and HICP was -60%.

Extracting Sentiment from Financial News Articles

We used the tailor-made dictionary to compute the sentiment in daily financial news using the simple method of computing positive/negative lemmas (Figure 14) and the method that also uses scores (Figure 15).

Daily Sentiment of News using Romanian Financial Stability Dictionary

(Without and with scores) Figure 14/Figure 15

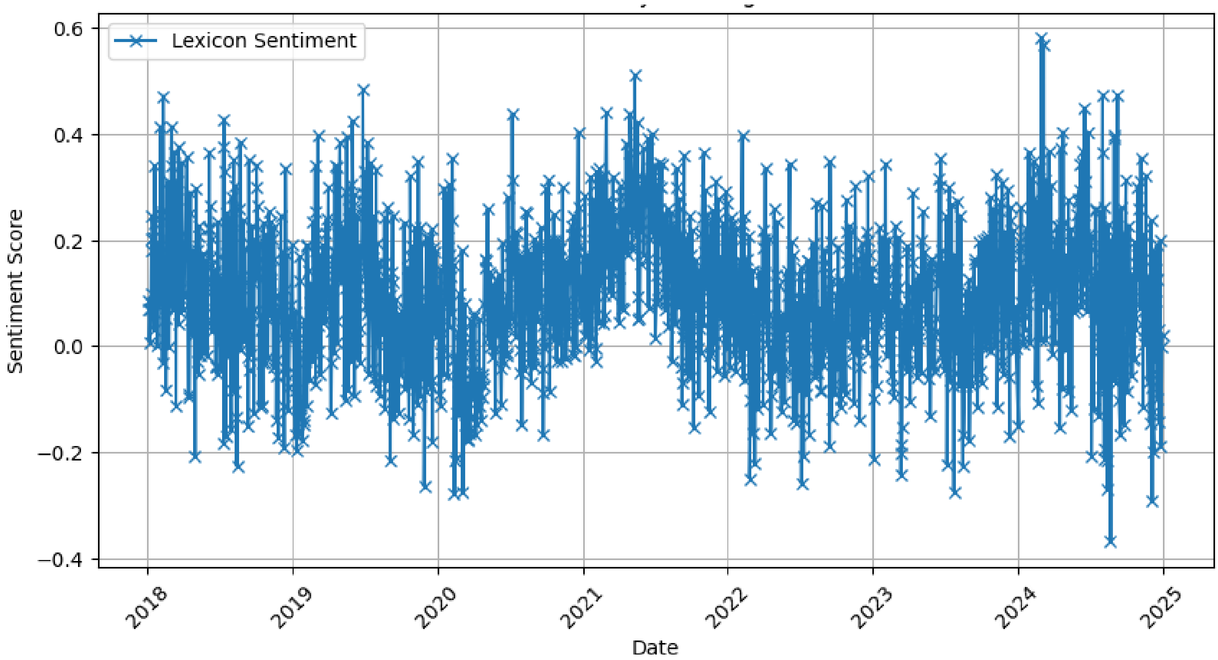

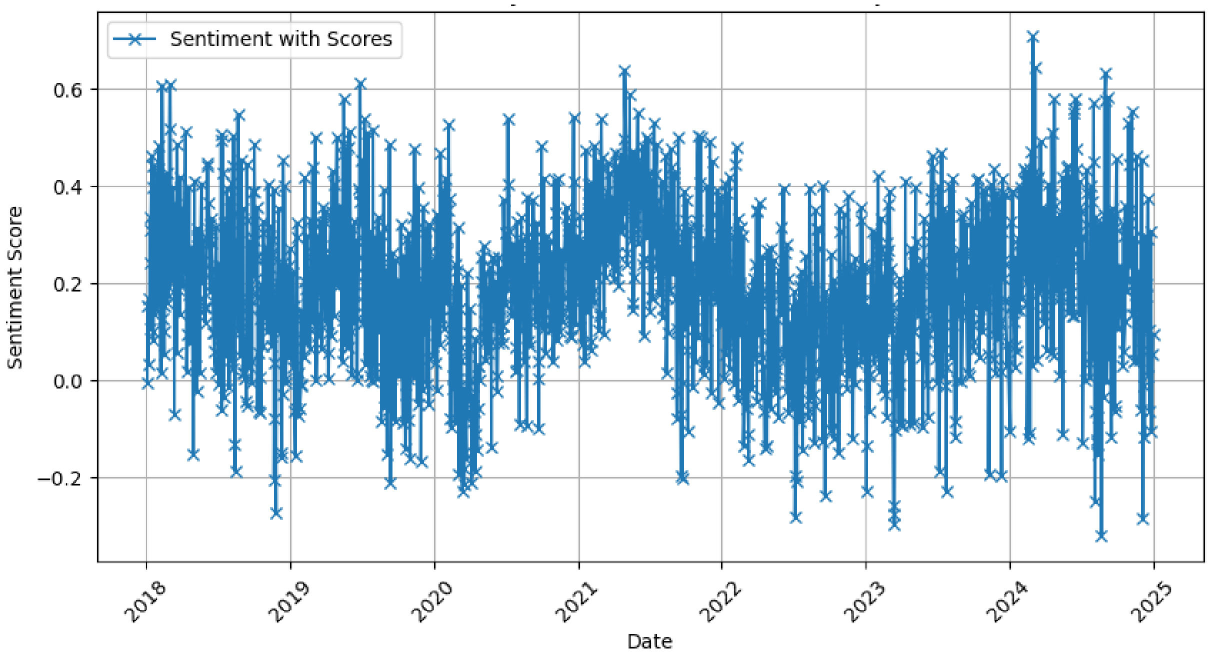

Source: authors' own computations

Sentiment analysis of economic and political discourse in Romania throughout the most recent year, 2024, reveals a predominantly positive outlook in the first half of the year, with sentiment scores remaining above the neutral threshold. However, the second half of the year saw a shift towards negative sentiment, driven by key adverse events. Notably, the annulment of the elections on December 9, 2024, introduced significant political uncertainty, while Fitch's decision to assign a negative outlook to Romania at year-end further weighed on investor confidence. Additionally, the decline in industrial production in August signalled potential economic headwinds, contributing to a deteriorating sentiment. These developments collectively shaped an overall negative sentiment trajectory for the year, highlighting the sensitivity of economic perceptions to both political instability and macroeconomic indicators.

We compared the daily sentiment of financial news with the moments of monetary policy decisions (Figure 16) and showed how the positive sentiment of a monetary policy decision, such as a lowering of the reference interest rate, positively influences the economy. There was a rise in sentiment in mid-2020 despite the COVID-19 pandemic. The opposite pattern can be identified in early 2018 and especially in 2022, where there were 8 negative monetary policy decisions, meaning the NBR raised the interest rate to combat inflation pressures. In just one year, the reference monetary policy rate rose by almost 5 percentage points due to the need to counteract the international inflationary environment. We concluded that there is strong evidence that the press responds to monetary policy decisions and that the communication of the central bank is efficient in influencing the markets.

Daily Sentiment of Financial News and Monetary Policy Decisions — Figure 16

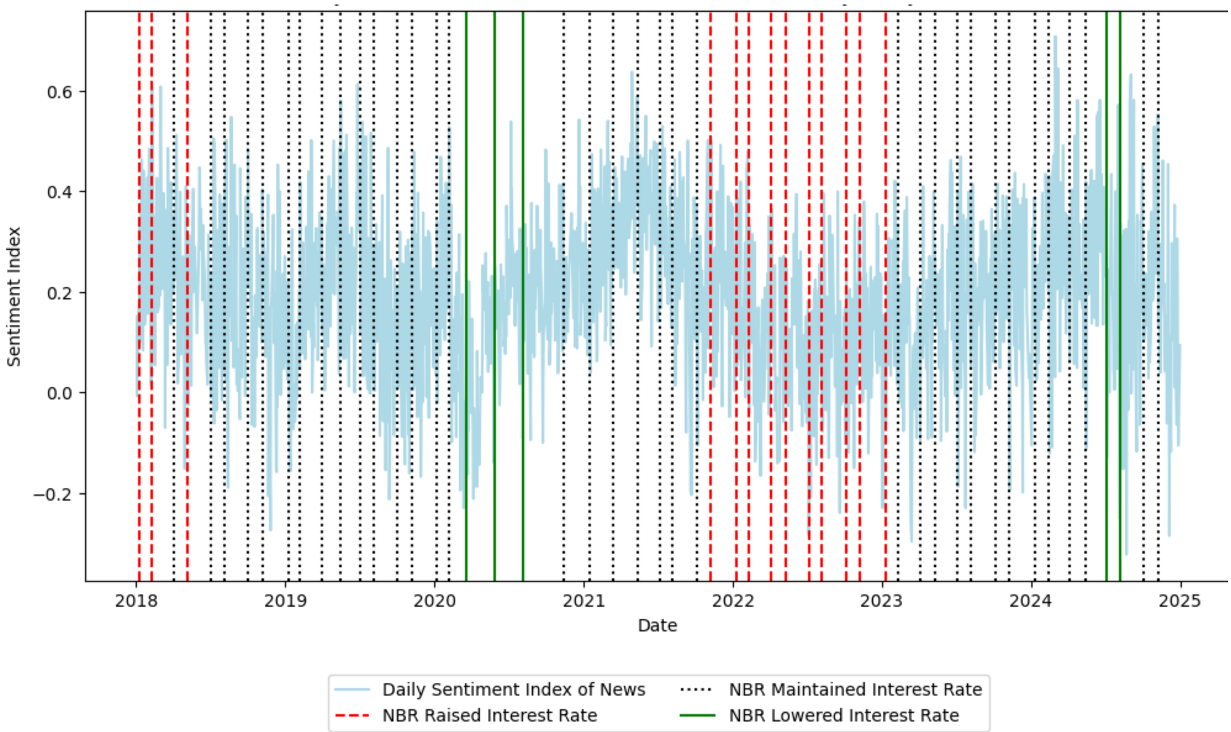

Source: authors' own computations

In conclusion, there is great value in building our own Romanian Dictionary for Financial Stability issues, and we achieved that by combining terms translated from FinBERT in English with manually identified Romanian terms. We computed scores for special terms related to financial stability using Gen-AI. To assess the performance of our daily sentiment indicator, we developed an unsupervised K-Means++ analysis according to Kanungo (2002) on embeddings using the BERT multilingual model by Devlin (2018) after reducing their size using PCA. The results show a 93% overlap between the clusters and the binary sentiments from the simple index computed by counting lemmas, while a slightly lower 96% overlap was observed between the clusters and the binary sentiments from the index computed also using scores.

Daily Sentiment of News and Stock Market Index

Figure 17

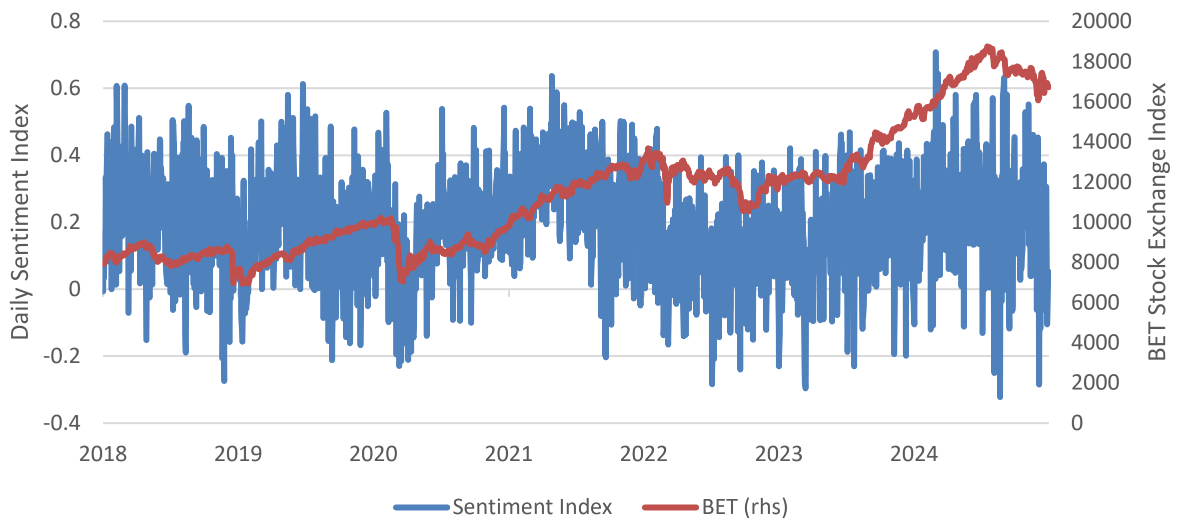

Source: authors' own computations

There is a positive correlation (Figure 17) between our daily sentiment index of financial news and Romanian Stock Exchange Index (BET), but at the quarterly level, the interconnectivity between the two series is most evident, with a correlation degree of 62% (Figure 18).

A negative correlation can be identified (Figure 19) between our sentiment index of financial news and the short-term yields of Romanian government bonds (1 year maturity - indicate the cost of state financing and also reflect inflation pressure), as well as the overnight and 3-month interest rate benchmarks known as the Romanian Interbank Offered Rate (ROBOR3M - part of the variable interest rate households pay for their loans).

Quarterly Sentiment of News and Percentage Change in Stock Market Index Figure 18

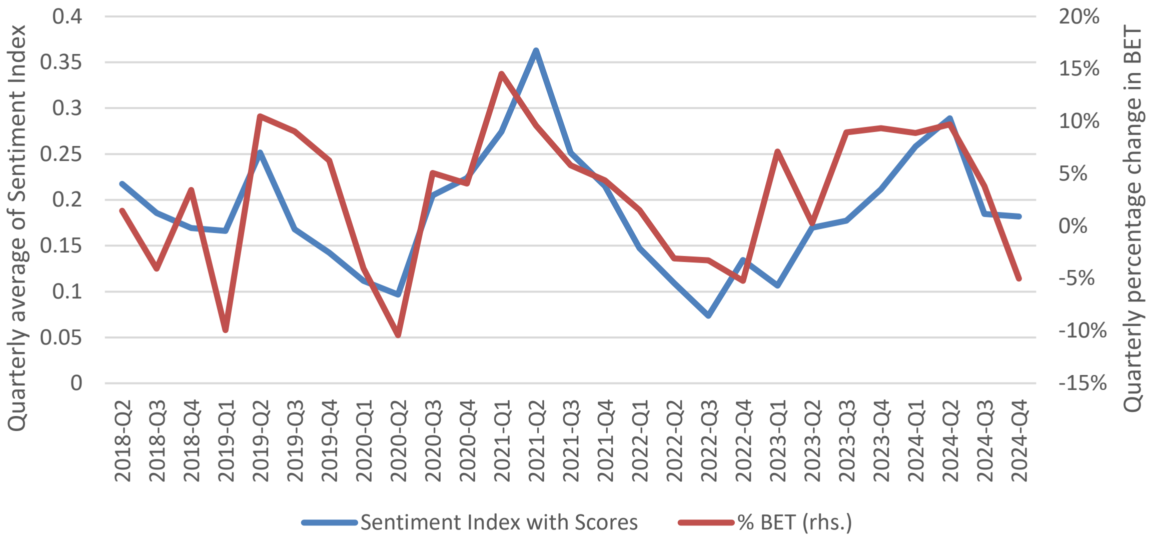

Daily Sentiment of News, Interbank Interest Rates and Yield Rates Figure 19

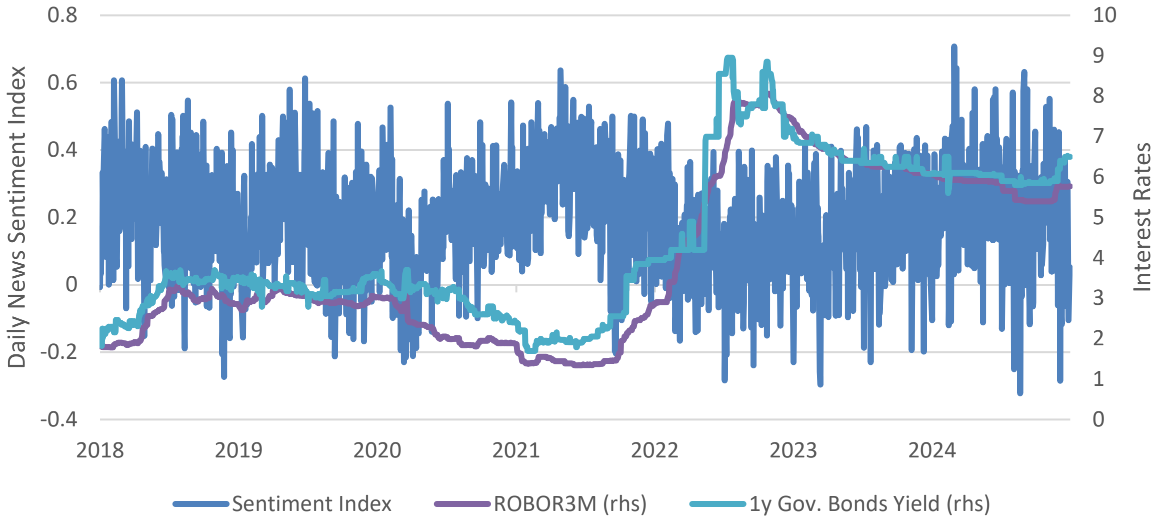

Source: authors' own computations

4. Conclusions

In conclusion, this paper proposes stepping aside from classical numerical, tabular data and investing more time and effort in the use of nontraditional data, not just time series, therefore implementing Natural Language Processing methods at a large scale for some types of text that exist in a central bank. The primary goal is to provide the National Bank of Romania with a guideline for pre-processing Romanian text, building NLP tools, and benchmarking sentiment analysis. Creating a toolbox for

Romanian, when most tools are in English, was challenging due to the need to consider all language particularities.

We believe that text is a rich resource for central banks, and therefore, in this paper, we proposed integrating Natural Language Processing into the National Bank of Romania's set of analyses through a thorough analysis on lexicon-based approaches to sentiment analysis using news and monetary policy decisions. We can only hope that by following these guidelines for text in Romanian, the list of potential use cases for Natural Language Processing will continue to be extended at the National Bank of Romania and that economists will commit to a long-term adoption of text data in their research, as there are many other types of text sources available to them.

The core contribution of this study is the construction of a Romanian Dictionary for Financial Stability, developed by identifying positive and negative terms extracted with FinBERT, fine-tuned on English monetary policy statements. We enriched this dictionary by assigning sentiment intensity scores using a generative AI agent. Our results demonstrate that the tailor-made dictionary outperforms existing solutions in capturing sentiment, with the scoring-enhanced approach yielding superior performance. The monetary policy sentiment index correctly identified sentiment in 92.5% of binary-labelled monetary policy statements. For the unsupervised analysis of daily financial news articles, we applied K-Means++ clustering on multilingual BERT embeddings, achieving overlapping rates of 93% and 96% with our results. Additionally, the daily news sentiment index showed consistent correlations with key macroeconomic variables. Integrating both components of the analysis, we find strong evidence that the press responds to monetary policy decisions and that central bank communication effectively influences financial markets.

5. References

1. Araci, D. (2019). FinBERT: A Pretrained Language Model for Financial Communications. arXiv preprint arXiv:1908.10063.
2. Araujo, Douglas KG. 2023. "Gingado: A Machine Learning Library Focused on Economics and Finance." Working Paper 1122. BIS Working Paper. Bank for International Settlements.
3. Baccianella, S., Esuli, A., & Sebastiani, F. (2010, May). Sentiwordnet 3.0: an enhanced lexical resource for sentiment analysis and opinion mining. In Lrec (Vol. 10, No. 2010, pp. 2200-2204).
4. Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent dirichlet allocation. Journal of machine Learning research, 3(Jan), 993-1022.
5. Correa, R., Garud, K., Londono-Yarce, J. M., & Mislang, N. (2017). Constructing a dictionary for financial stability.
6. Devlin, J. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805.
7. Edison, H., & Carcel, H. (2021). Text data analysis using Latent Dirichlet Allocation: an application to FOMC transcripts. Applied Economics Letters, 28(1), 38-42.

8. 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, (23-10).
9. Hu, M., & Liu, B. (2004, August). Mining and summarizing customer reviews. In Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining (pp. 168-177).
10. Hutto, C., & Gilbert, E. (2014, May). Vader: A parsimonious rule-based model for sentiment analysis of social media text. In Proceedings of the international AAAI conference on web and social media (Vol. 8, No. 1, pp. 216-225).
11. Jolliffe, I. T. (1986). Principal component analysis. Springer Series in Statistics. Springer-Verlag.
12. Kanungo, T., Mount, D. M., Netanyahu, N. S., Piatko, C. D., Silverman, R., & Wu, A. Y. (2002). An efficient k-means clustering algorithm: Analysis and implementation. IEEE transactions on pattern analysis and machine intelligence, 24(7), 881-892.
13. Loper, E., & Bird, S. (2002). Nltk: The natural language toolkit. arXiv preprint cs/0205028.
14. 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.
15. Mabey, B. (2018). pyLDAvis documentation.
16. Qi, P., Zhang, Y., Zhang, Y., Bolton, J., & Manning, C. D. (2020). Stanza: A Python natural language processing toolkit for many human languages. arXiv preprint arXiv:2003.07082.
17. Richardson, L. (2007). Beautiful soup documentation.
18. Sanh, V. (2019). DistilBERT, A Distilled Version of BERT: Smaller, Faster, Cheaper and Lighter. arXiv preprint arXiv:1910.01108.
19. Schmitt, X., Kubler, S., Robert, J., Papadakis, M., & LeTraon, Y. (2019, October). A replicable comparison study of NER software: StanfordNLP, NLTK, OpenNLP, SpaCy, Gate. In 2019 sixth international conference on social networks analysis, management and security (SNAMS) (pp. 338-343). IEEE.
20. Vasiliev, Y. (2020). Natural language processing with Python and spaCy: A practical introduction. No Starch Press.

National Bank of Romania

Rome, February 18th, 2025

A NATURAL LANGUAGE PROCESSING TOOLBOX FOR THE NATIONAL BANK OF ROMANIA

Claudia Voicilă, Senior Economist

Note 1: The opinions are those of the author and do not necessarily reflect the views of the National Bank of Romania

Note 2: Part of my dissertation thesis for the Artificial Intelligence Master's program at the Faculty of Mathematics and Informatics, University of Bucharest.

Agenda

Introduction

Chapter 1. English Monetary Policy Statements Database
Chapter 2. Romanian Financial News Articles Database

Conclusions

Abstract

- This paper introduces a specially designed toolbox aimed at leveraging Natural Language Processing (NLP) to unlock insights for the National Bank of Romania (NBR), enhancing economists' capacity to process text data such as financial news and press releases, exploring areas of Financial Stability and Central Bank Communication.
- We propose implementing scalable Natural Language Processing methods within the National Bank of Romania’s analysis toolkit, focusing on lexicon-based sentiment analysis of daily news and monetary policy decisions. This study seeks to align the institution with the best practices of other central banks and highlights the untapped potential of textual data as a valuable resource in central banking.

Introduction

- The paper is structured in two chapters: the first one addressed monetary policy statements in English from the National Bank of Romania, the European Central Bank, the United States's Federal Reserves and 25 other emerging market countries, with the scope of building a list of positive and negative words specific for central bank text, the starting point of the Romanian Financial Stability Dictionary.
- The second chapter focus on the Romanian news archive database, and we finalize the Financial Stability Dictionary that contains specialized positive and negative words. We augment the method with GenAI scores to measure text-based sentiments in financial texts.

Chapter 1. The English Monetary Policy Statements Database

- The first part focuses on English monetary policy statements starting from the moment NBR adopted inflation targeting. We extended the dataset by including ECB, FED, and 25 other emerging market countries.
- We fine-tuned a Transformer model, FinBERT, on the binary sentiment labels we created for the monetary policy statement based on movements in the monetary policy interest rate of each country (positive sentiment if the interest rate is lowered, negative sentiment if the interest rate is raised).
- After trying unspecialized lexicons, we demonstrated that using specialized Financial Stability Dictionary is crucial for constructing a sentiment index for monetary policy, and we measure the polarity of words using GenAI.

Confusion Matrix of FinBERT on test sample from the extended English database

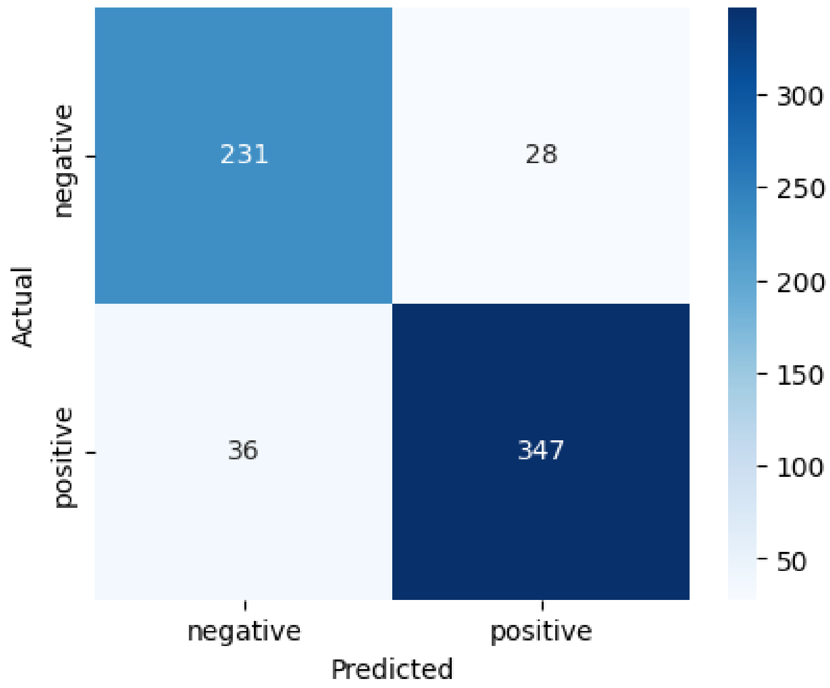

Classification report on test sample (10 epochs)

	precision	recall	F1-score	support
positive	86.52%	89.19%	87.82%	259
negative	92.53%	90.60%	91.56%	383
accuracy			90.03%	642
macro accuracy	89.53%	89.89%	89.69%	642
Weighted average	90.11%	90.03%	90.05%	642

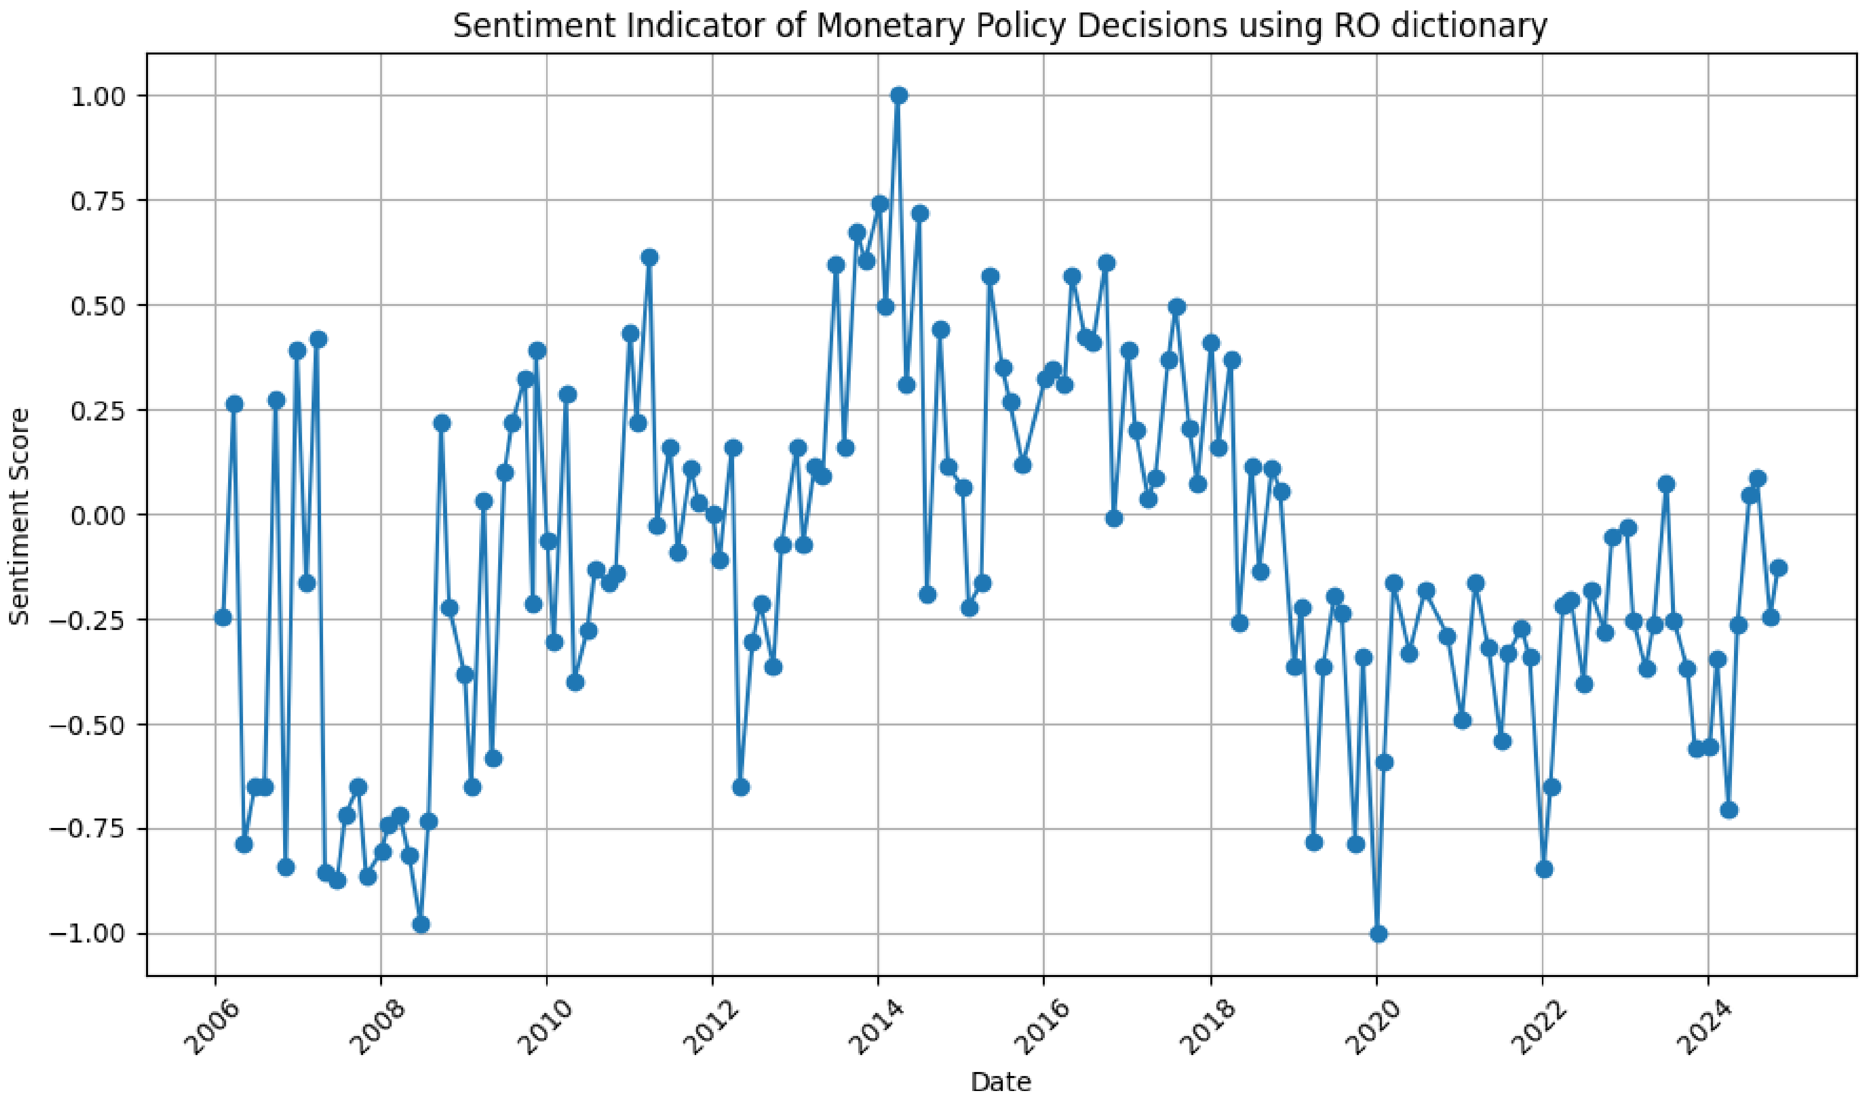
Sentiment Indicator of Monetary Policy Decisions using RO dictionary
1.00
0.75
0.50
0.25
0.00
−0.25
−0.50
−0.75
−1.00
Sentiment Score
2006
2008
2010
2012
2014
2016
2018
2020
2022
2024
Date

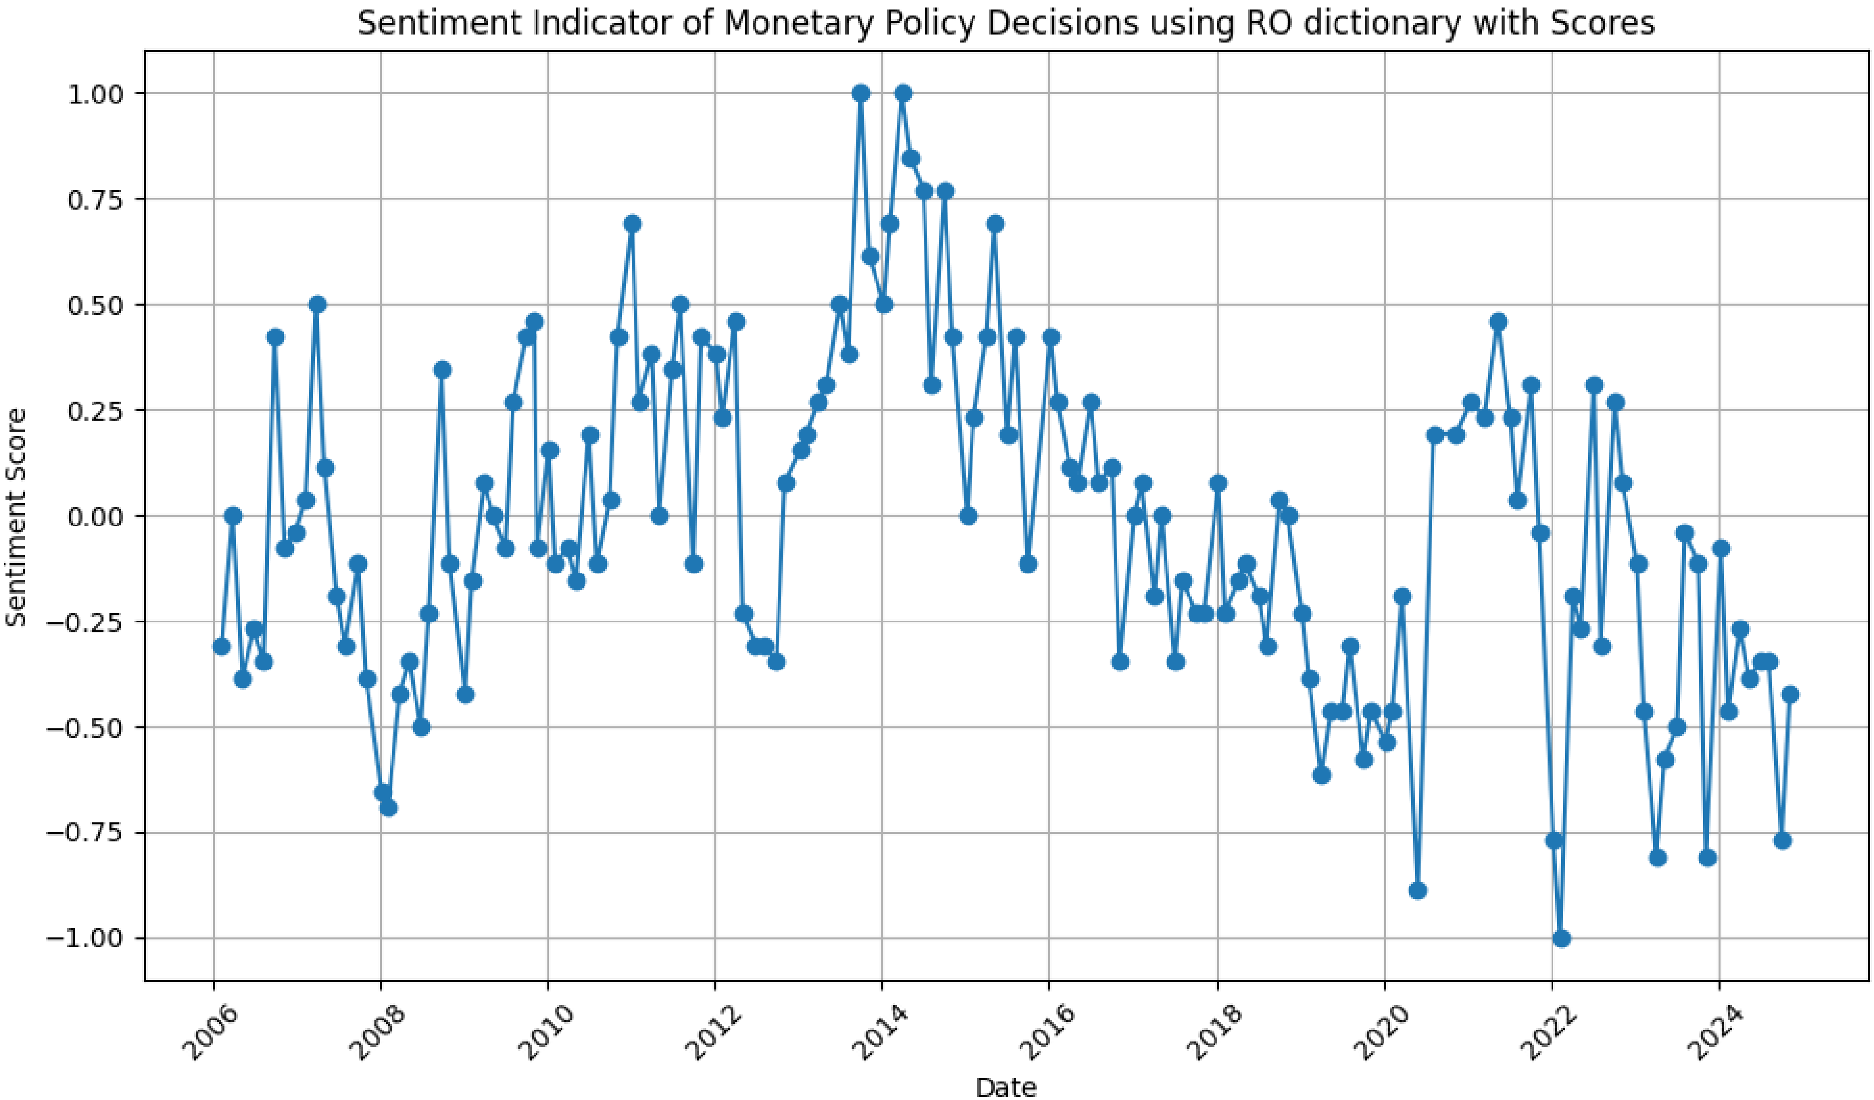
Sentiment Indicator of Monetary Policy Decisions using RO dictionary with Scores
Sentiment Score
Date
1.00
0.75
0.50
0.25
0.00
−0.25
−0.50
−0.75
−1.00
2006
2008
2010
2012
2014
2016
2018
2020
2022
2024

- The version of the sentiment index with scores outperforms the simple sentiment index by 2% (78% versus 76%), and the final index that we obtained by mixing the two methods matches 92.5% of the labels. This performance is considered excellent, given that the neutral decisions were built using the last decision rule (negative sentiment persists if last decision was to raise the interest rate, positive sentiment persists if the last decision was to lower interest rate) and not manually evaluated by a human.

- To assess the performance of our NLP analysis, we examined the correlation between the sentiment index of monetary policy decisions and the most relevant economic variable: the inflation rate, since controlling inflation is the primary goal of monetary policy. We found an expected negative correlation between our computed sentiment and the all-items HICP (Harmonized Index of Consumer Prices) of over -40%, while the correlation between the monetary policy interest rate and HICP was around -60%.

Sentiment Index of Monetary Policy compared to Binary Labels

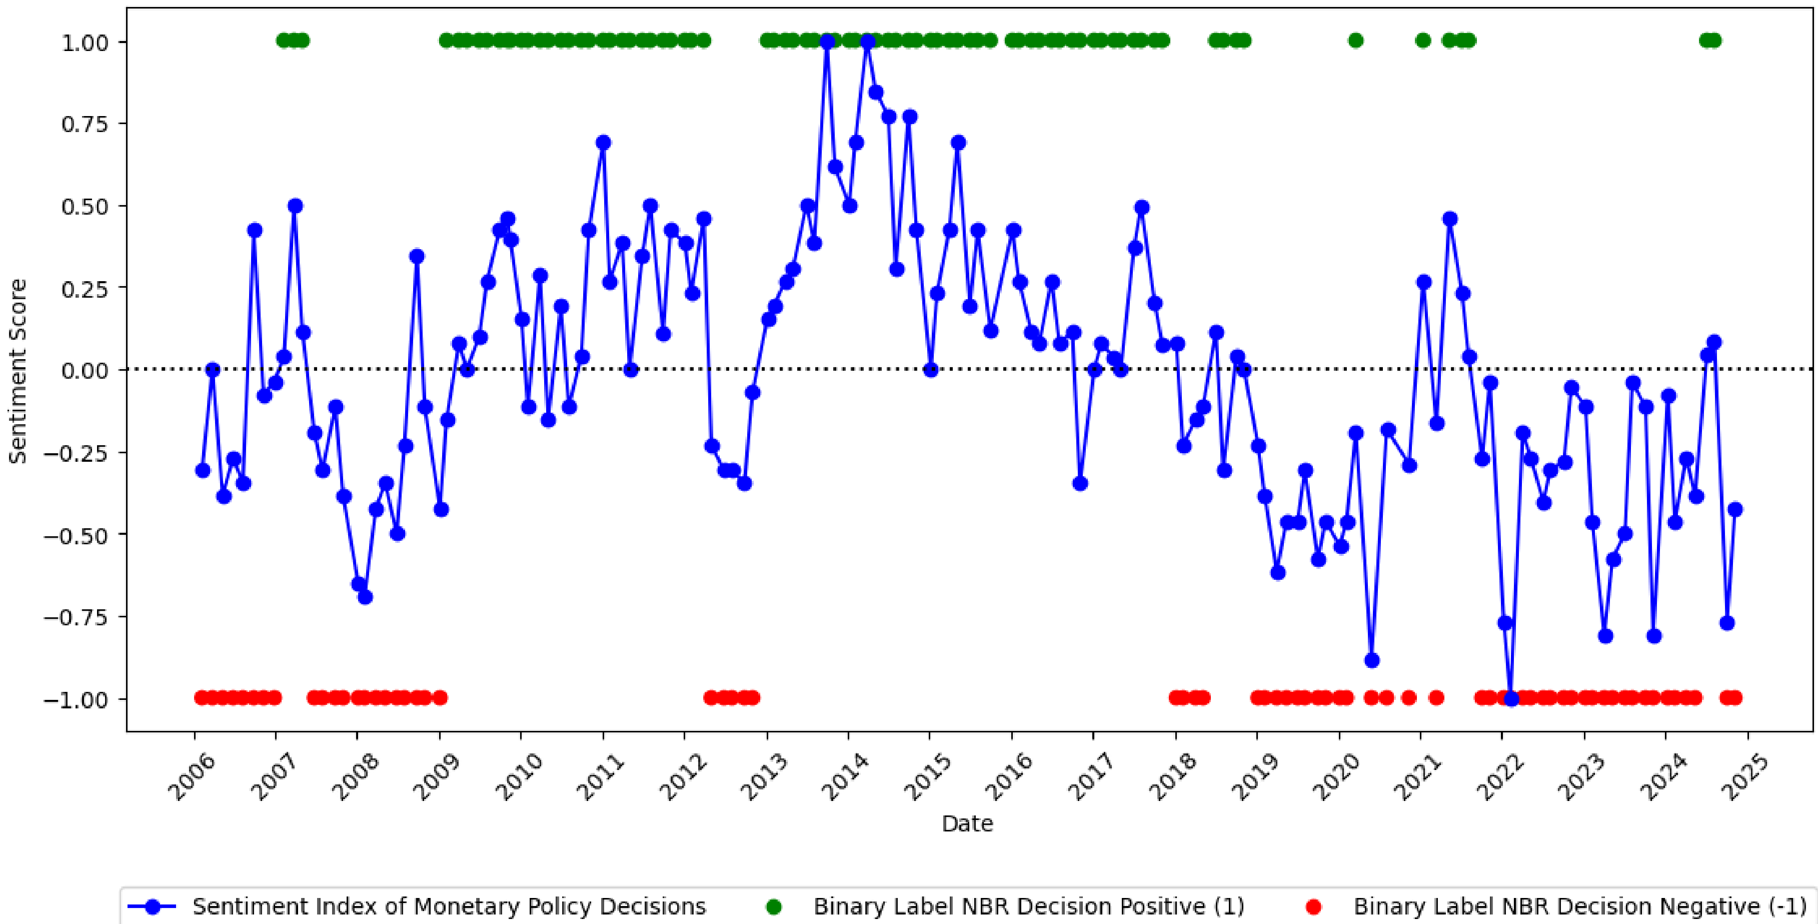

Chapter 2. The Romanian Financial News Articles Database

- The second part addresses the NBR's archive of daily financial news, for which we conduct topic modeling using LDA to track key trends in the news, and propose a tailor-made, Romanian Financial Stability Dictionary to measure sentiments.
- We computed polarity intensity scores by prompting as a financial stability expert a Gen AI agent.
- In absence of labels, we assess the performance by comparing our method with K-Means++ clustering applied to embeddings obtained with multilingual BERT Transformer and reduced with PCA, and by demonstrating correlations with macroeconomic and market variables as well as interactions with monetary policy decisions.

Interactive LDA analysis, Example for year 2023

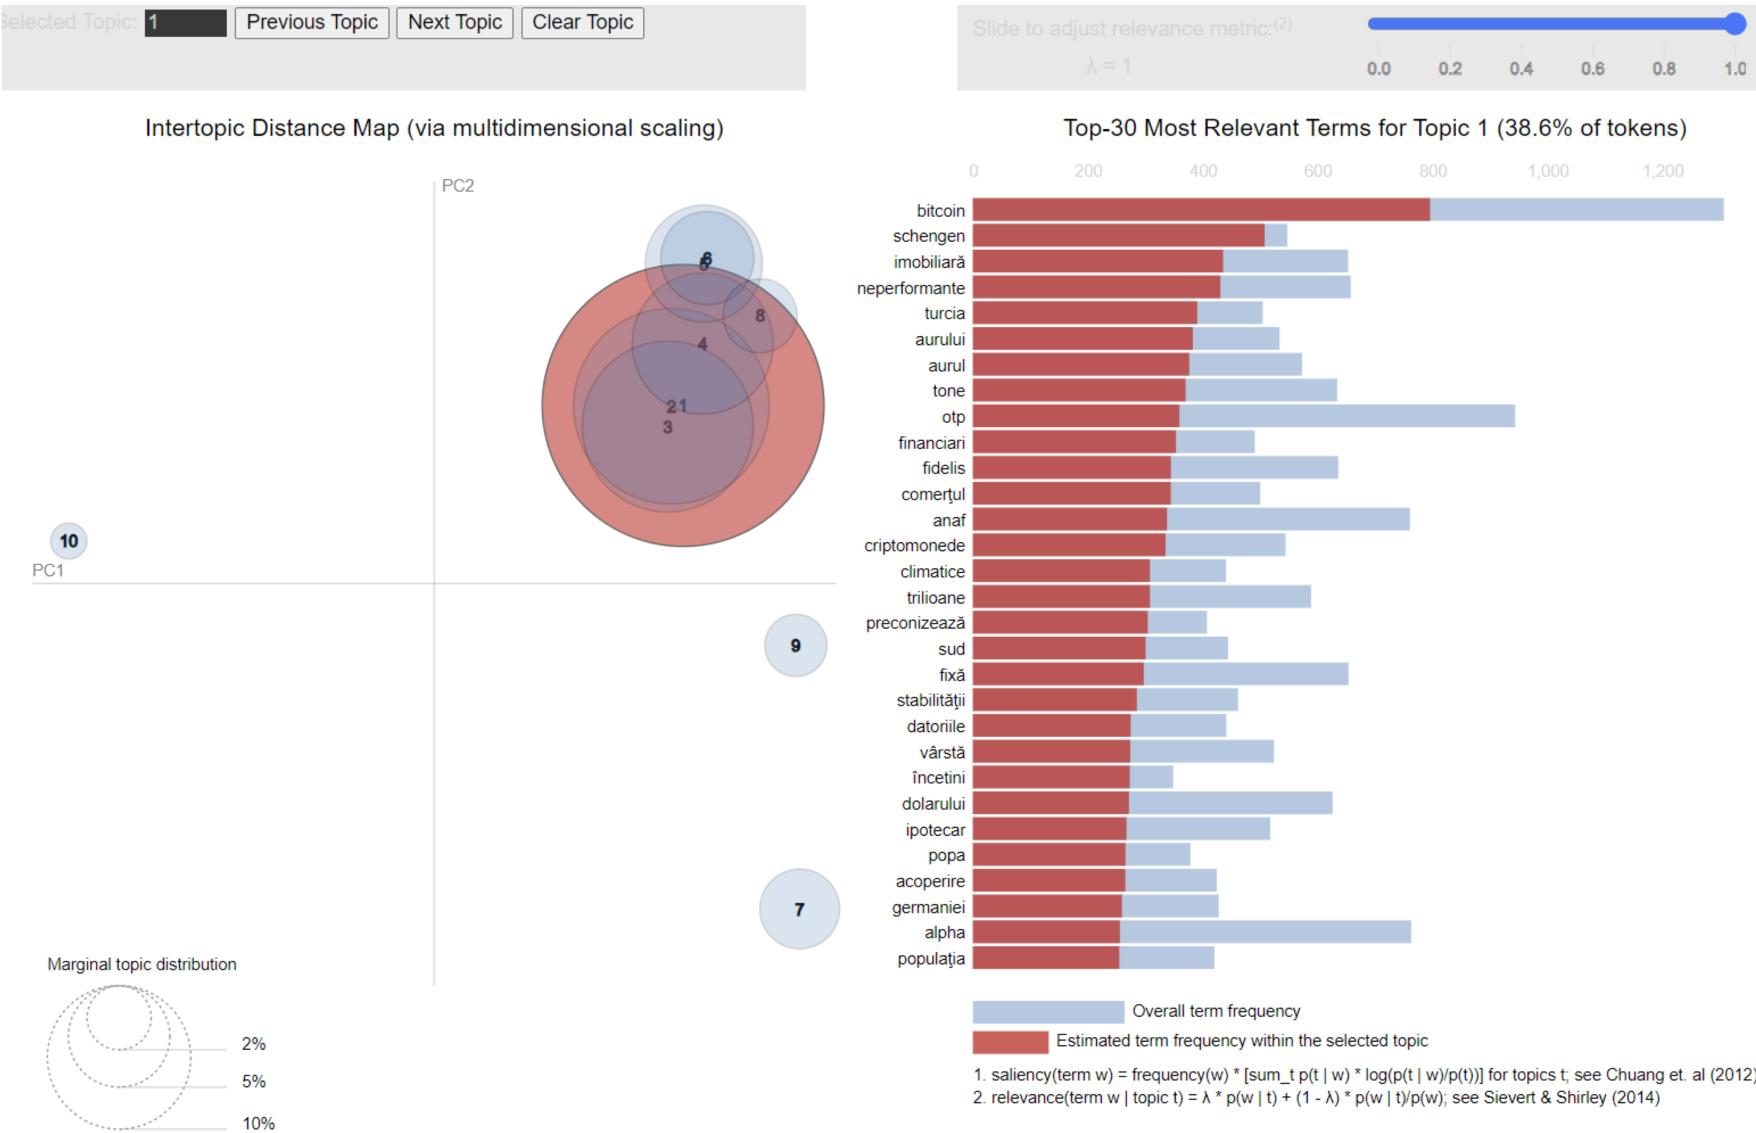

Heatmap of the Importance of the Main Five Topics Over Time

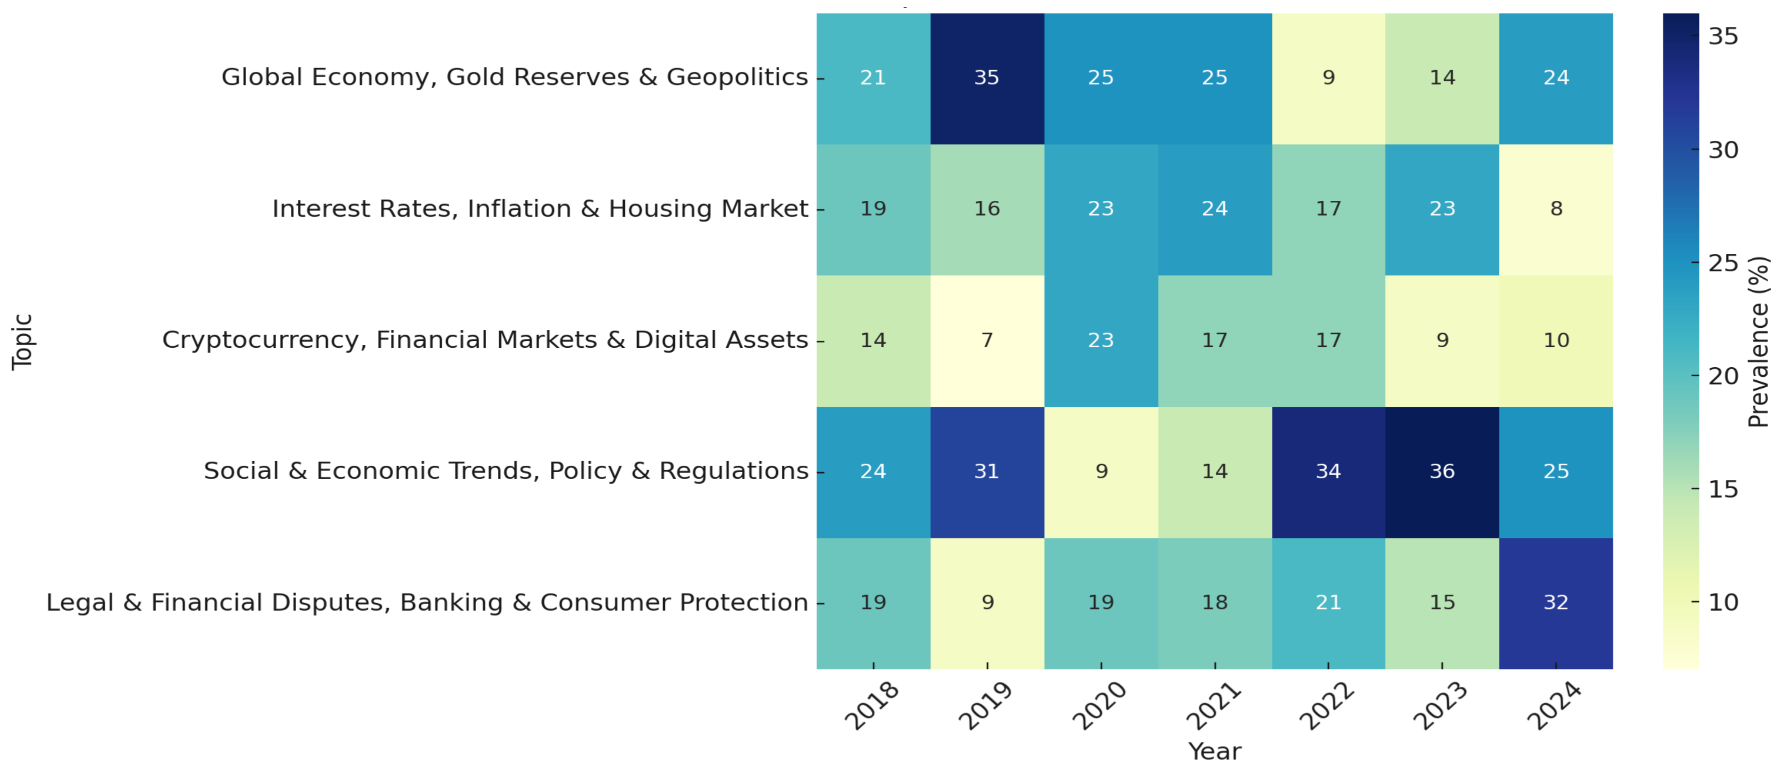

Word Cloud analysis, Example for year 2024

Combined Word Cloud for All Topics

Daily Sentiment of News using Romanian Financial Stability Dictionary

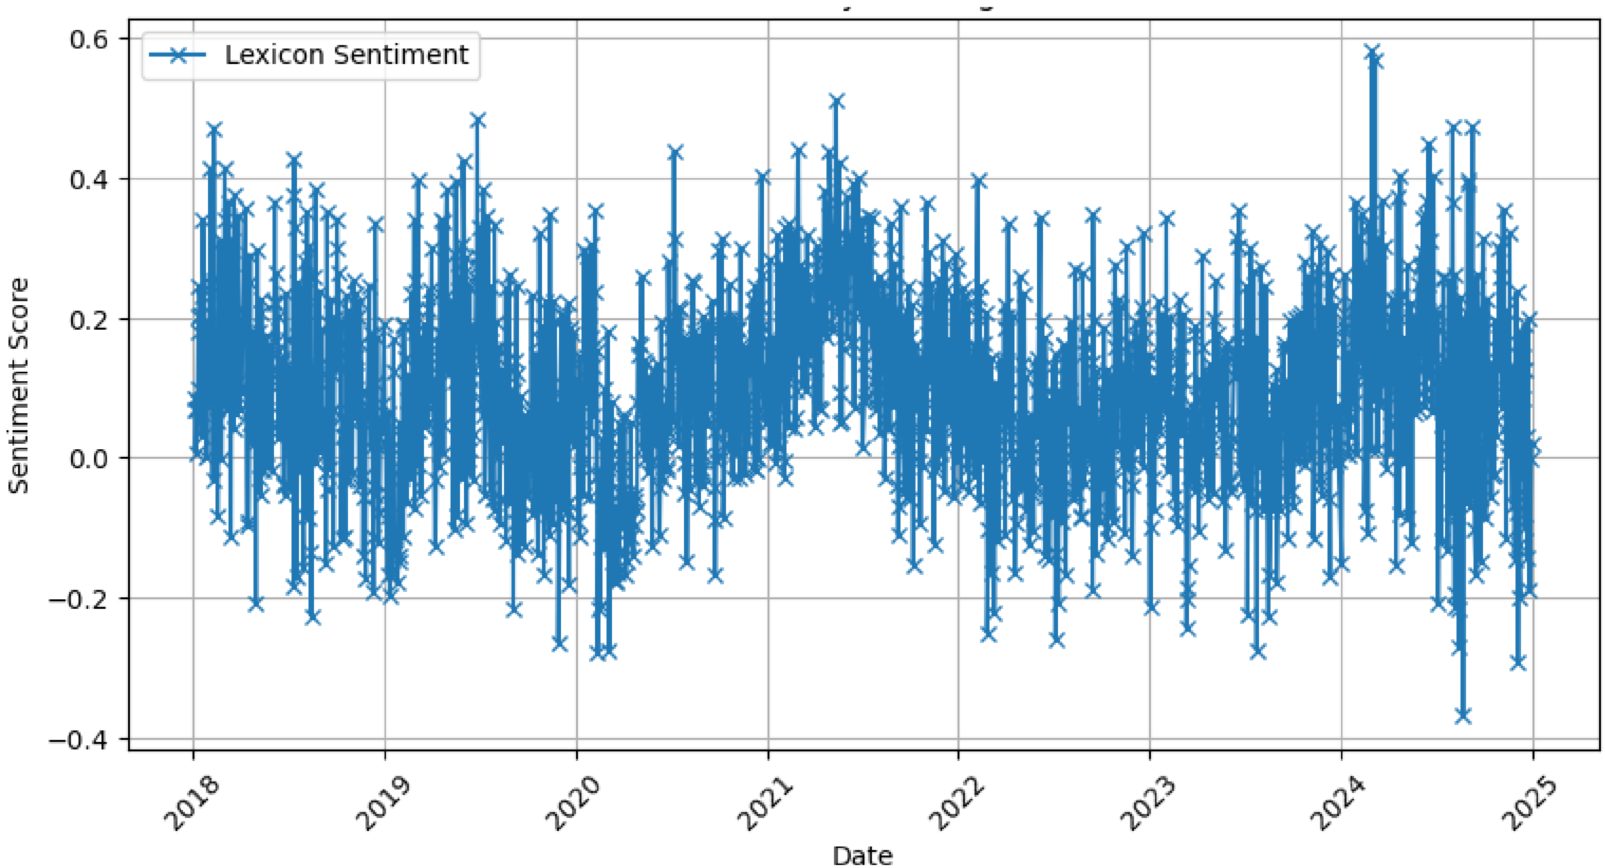

$$Sentiment\ Index = \frac{P(t) - N(T)}{P(t) + N(t)}$$

where P number of positive lemmas and N number of negative lemmas

Daily Sentiment of News using Romanian Financial Stability Dictionary augmented with scores

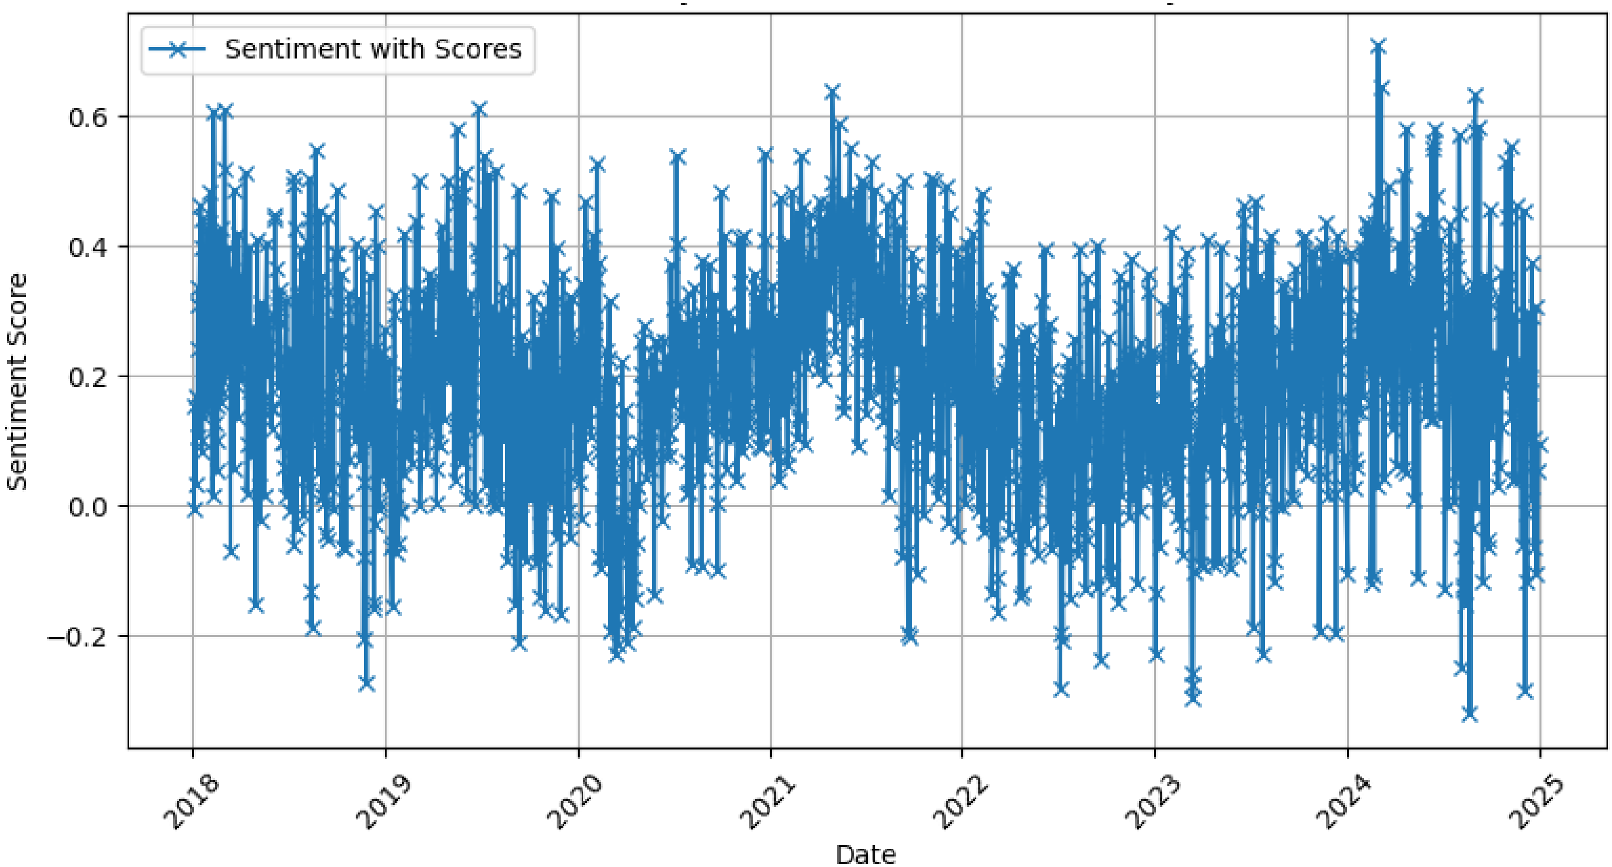

$$Sentiment\ Score\ Index = \frac{\sum_{i=1}^{P(t)} S_p(i) + \sum_{j=1}^{N(t)} S_n(j)}{\sum_{i=1}^{P(t)} S_p(i) + \sum_{j=1}^{N(t)} |S_n(j)|}$$

where Sp scores of positive lemmas and Sn scores of negative lemmas

- To assess the performance of our daily sentiment indicator, we developed an unsupervised K-Means++ analysis according to Kanungo (2002) on embeddings using the BERT multilingual model by Devlin (2018), after reducing their size using PCA. The results show a 93% overlap between the clusters and the binary labels from the simple index, while a slightly higher 96% overlap for the index using scores.

- We compared the daily sentiment of financial news with the moments of monetary policy decisions and showed that there is strong evidence that the press respond to monetary policy decisions and that the communication of the central bank is efficient in influencing the markets.

- We also found a clear correlation with macroeconomic variables: BET, the stock market index, yields of short-term bonds, as well as interbank interest rates ROBOR3M.

Daily Sentiment of Financial News and Monetary Policy Decisions

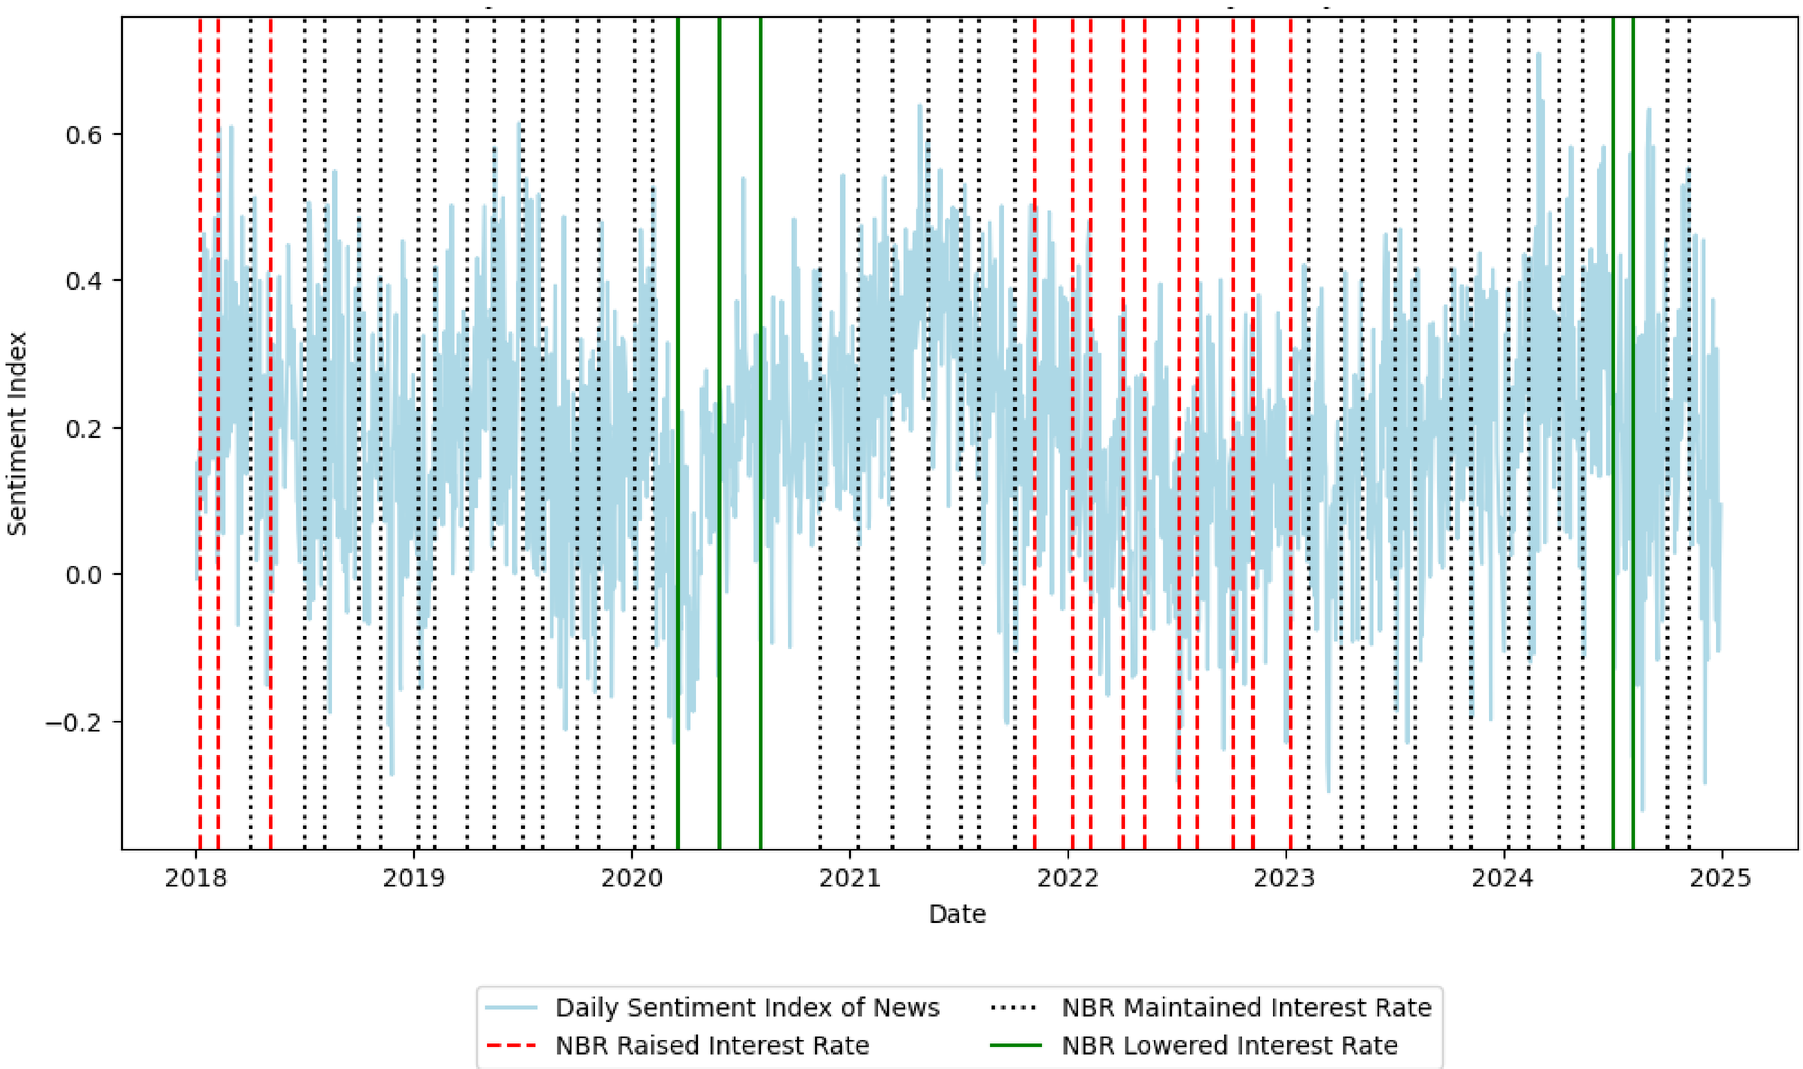

Daily Sentiment of News and Stock Market Index

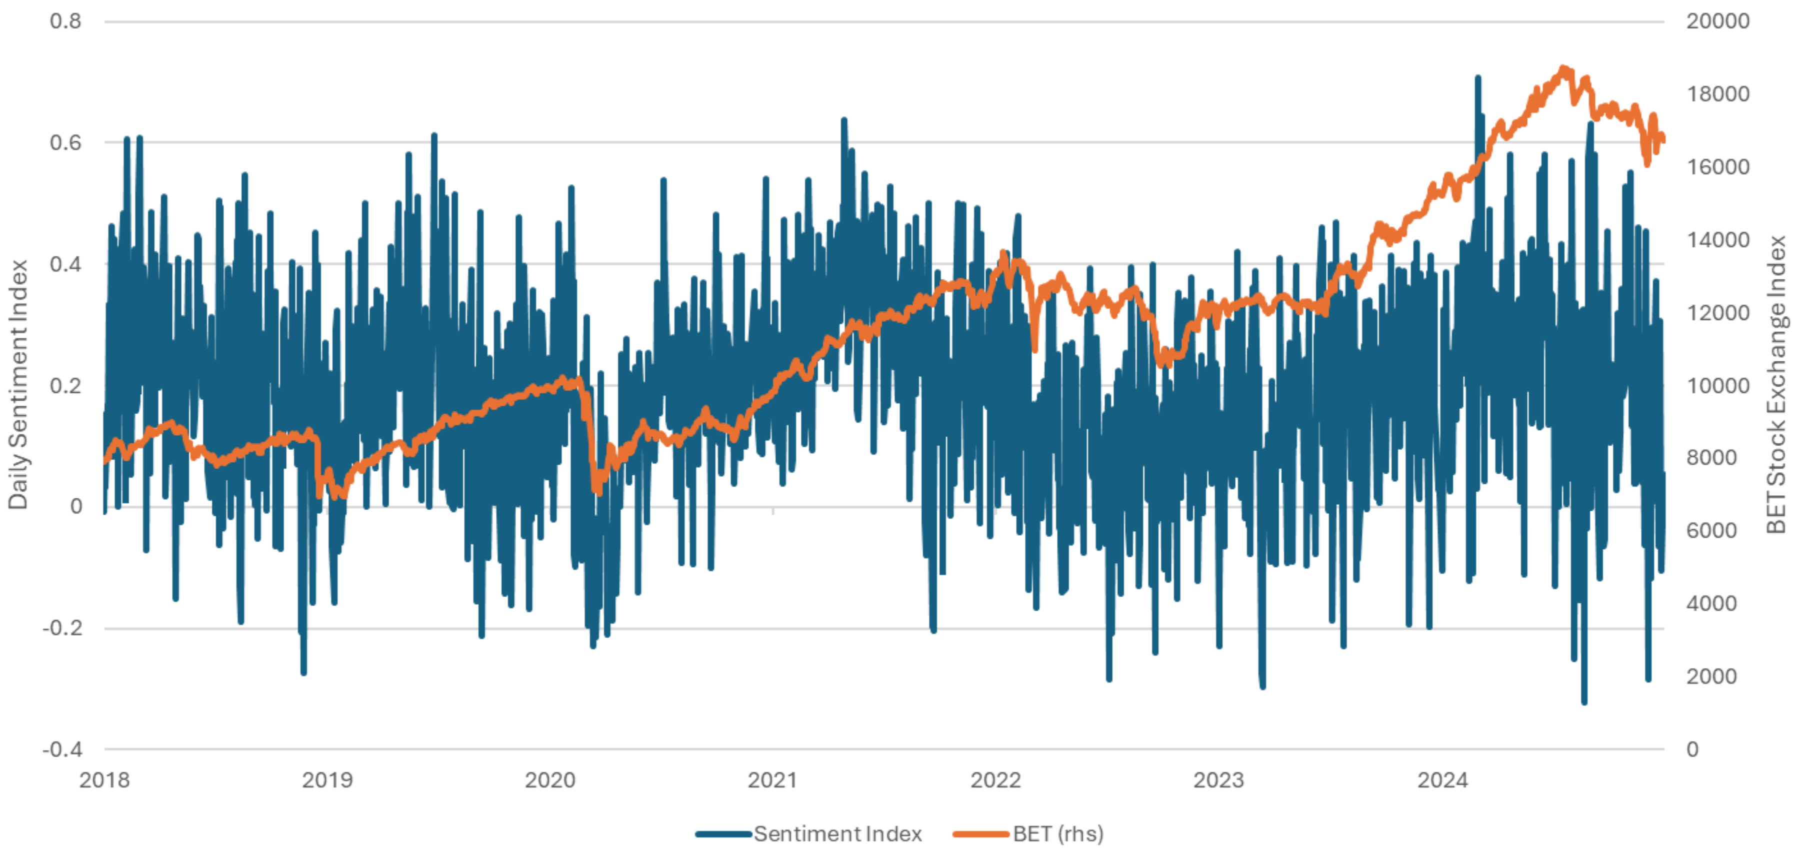

Quarterly Sentiment of News and Percentage Change in Stock Market Index

Daily Sentiment of News, Interbank Interest rates and Sovereign Yield Rates

Conclusions

- This paper proposes stepping aside from classical numerical, tabular data and investing more time and effort in the use of nontraditional data, not just time series, thus implementing natural language processing methods at a large scale for some types of text that exist in a central bank. The primary goal is to provide the NBR with a guideline for pre-processing Romanian text, building NLP tools, and benchmarking sentiment analysis.

- We can only hope that by following this NLP Toolbox for the National Bank of Romania, the list of potential use cases for NLP will continue to be extended at the NBR and that economists will commit to a long-term adoption of text data in their research, as there are many other text sources available.

Thank you!

claudia.voicila@bnro.ro

NATIONAL BANK OF ROMANIA 145

References

Abodayeh, A., Hejazi, R., Najjar, W., Shihadeh, L., & Latif, R. (2023, March). Web Scraping for Data Analytics: A BeautifulSoup Implementation. In *2023 Sixth International Conference of Women in Data Science at Prince Sultan University (WiDS PSU)* (pp. 65-69). IEEE.

Aprigliano, V., Emiliozzi, S., Guaitoli, G., Luciani, A., Marcucci, J., & Monteforte, L. (2023). The power of text-based indicators in forecasting Italian economic activity. *International Journal of Forecasting*, *39*(2), 791-808.

Araci, D. (2019). FinBERT: A Pretrained Language Model for Financial Communications. arXiv preprint arXiv:1908.10063.

Araujo, D., Bruno, G., Marcucci, J., Schmidt, R., & Tissot, B. (2023). Data science in central banking: applications and tools. *IFC Bulletin*, (59).

Araujo, D., Bruno, G., Marcucci, J., Schmidt, R., & Tissot, B. (2023). Machine learning applications in central banking. *Journal of AI, Robotics & Workplace Automation*, *2*(3), 271-293.

Araujo, Douglas KG. 2023. "Gingado: A Machine Learning Library Focused on Economics and Finance." Working Paper 1122. BIS Working Paper. Bank for International Settlements.

Astuti, V., Bruno, G., Marchetti, S., & Marcucci, J. (2023). *News and banks' equities: do words have predictive power?* (Vol. 59). Bank for International Settlements.

Baccianella, S., Esuli, A., & Sebastiani, F. (2010, May). Sentiwordnet 3.0: an enhanced lexical resource for sentiment analysis and opinion mining. In *Lrec* (Vol. 10, No. 2010, pp. 2200-2204).

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

Chong, E., & Ho, S. J. (2023). *Measuring text-based sentiments from monetary policy statements–a Malaysian case study using natural language processing* (Vol. 59). Bank for International Settlements.

.

Chong, E., Ho, C. C., Ong, Z. F., & Ong, H. H. (2022). *Using news sentiment for economic forecasting: a Malaysian case study* (Vol. 57). Bank for International Settlements

Chuang, J., Manning, C. D., & Heer, J. (2012, May). Termite: Visualization techniques for assessing textual topic models. In *Proceedings of the international working conference on advanced visual interfaces* (pp. 74-77).

Correa, R., Garud, K., Londono-Yarce, J. M., & Mislang, N. (2017). Constructing a dictionary for financial stability.

De Bandt, O., Jean-Charles, B., Julien, D., Alexandre, D., De Gaye, A., & Pierre-Antoine, R. (2023). *Utiliser la presse pour construire un nouvel indicateur de perception d'inflation en France*. Banque de France.

Denes, J., Lestrade, A., & Richardet, L. (2022). *Using twitter data to measure inflation perception* (Vol. 57). Bank for International Settlements.

Devlin, J. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

Doerr, S., Gambacorta, L., & Garralda, J. M. S. (2021). Big data and machine learning in central banking. *BIS Working Papers*, (930).

Dumitrescu, S. D., Avram, A. M., Morogan, L., & Toma, S. A. (2018, June). Rowordnet–a python api for the romanian wordnet. In *2018 10th International Conference on Electronics, Computers and Artificial Intelligence (ECAI)* (pp. 1-6). IEEE.

Edison, H., & Carcel, H. (2021). Text data analysis using Latent Dirichlet Allocation: an application to FOMC transcripts. *Applied Economics Letters*, *28*(1), 38-42.

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*, (23-10).

Hu, M., & Liu, B. (2004, August). Mining and summarizing customer reviews. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining* (pp. 168-177).

Hutto, C., & Gilbert, E. (2014, May). Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *Proceedings of the international AAAI conference on web and social media* (Vol. 8, No. 1, pp. 216-225).

Istrefi, K., Odendahl, F., & Sestieri, G. (2022). ECB communication and its impact on financial markets.

Istrefi, K., Odendahl, F., & Sestieri, G. (2023). Fed communication on financial stability concerns and monetary policy decisions: Revelations from speeches. *Journal of Banking & Finance*, *151*, 106820.

Jabbar, M. A., Wibisono, O., Widjanarti, A., & Zulen, A. A. (2023). Machine learning for measuring central bank policy credibility and communication from news. *IFC Bulletins chapters*, *59*.

Jolliffe, I. T. (1986). Principal component analysis. *Springer Series in Statistics*. Springer-Verlag.

Kanungo, T., Mount, D. M., Netanyahu, N. S., Piatko, C. D., Silverman, R., & Wu, A. Y. (2002). An efficient k-means clustering algorithm: Analysis and implementation. *IEEE transactions on pattern analysis and machine intelligence*, *24*(7), 881-892.

Lee, Y., & Seo, B. (2023). Extracting economic sentiment from news articles: the case of Korea. *IFC Bulletins chapters*, *59*.

Loper, E., & Bird, S. (2002). Nltk: The natural language toolkit. *arXiv preprint cs/0205028*.

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.

Mabey, B. (2018). pyLDAvis documentation.

Qi, P., Zhang, Y., Zhang, Y., Bolton, J., & Manning, C. D. (2020). Stanza: A Python natural language processing toolkit for many human languages. arXiv preprint arXiv:2003.07082.

Richardson, L. (2007). Beautiful soup documentation.

Schmitt, X., Kubler, S., Robert, J., Papadakis, M., & LeTraon, Y. (2019, October). A replicable comparison study of NER software: StanfordNLP, NLTK, OpenNLP, SpaCy, Gate. In *2019 sixth international conference on social networks analysis, management and security (SNAMS)* (pp. 338-343). IEEE.

Vasiliev, Y. (2020). *Natural language processing with Python and spaCy: A practical introduction*. No Starch Press.