Banco Central del Ecuador
Revista Cuestiones Económicas
Volumen 35 · Número 2
e–ISSN: 2697–3367

Effectiveness of Sentiment Analysis in Forecasting Ecuador's GDP Annual Growth Rate Using Newspaper Textual Data (2001-2024)

Roberto Páez Plúas

University San Francisco de Quito - School of Economics
Quito, Ecuador

Article Info

Received:
22nd July 2025

Accepted:
28th November 2025

Keywords:
Ecuador
Sentiment analysis
GDP forecasting
LASSO-ARDL
Real-time data

JEL:
C53, E37, C55, E01, O54

DOI:
https://doi.org/10.47550/RCE/35.2.2

Abstract

Ecuador's ongoing crises underscore the need for timely economic forecasting. While business and consumer surveys capture expectations, their monthly frequency limits responsiveness to sudden shocks. This study tests whether sentiment analysis of daily newspaper articles can improve GDP growth forecasts. Through 240.000 articles from 2001-2024, five sentiment indicators were constructed with a customized NLP approach combining TextBlob, VADER, and SpaCy's spanish model. These indicators were incorporated into a LASSO-ARDL model alongside traditional controls and survey-based indicators. Results show that sentiment-enhanced models significantly outperform survey-based benchmark for short-term horizons (1-3 months), reducing RMSE and demonstrating greater stability. At longer horizons, differences become statistically insignificant. The findings highlight the value of integrating real-time, text-derived indicators into forecasting frameworks, offering a scalable complement to traditional methods in emerging and volatile economic contexts.

[1] ORCID: 0009-0002-7710-939X. CRediT: conceptualization, data curation, formal analysis, research, methodology, validation, visualization, writing - original draft, writing - review and editing. Email: rpaezp@alumni.usfq.edu.ec
Copyright © 2025 Paez. Authors retain the copyright of this article. This article is published under the terms of the Creative Commons Attribution Licence 4.0.

Banco Central del Ecuador
Revista Cuestiones Económicas
Volumen 35 · Número 2
e–ISSN: 2697–3367

Eficacia del análisis de sentimiento para pronosticar la tasa de crecimiento anual del PIB de Ecuador mediante datos textuales de periódicos (2001-2024)

Roberto Páez Plúas

Universidad San Francisco de Quito - Escuela de Economía
Quito, Ecuador

Información

Recibido:
22 de julio de 2025

Aceptado:
28 de noviembre de 2025

Palabras clave:
Ecuador
Análisis de sentimiento
Previsión del PIB
LASSO-ARDL
Datos en tiempo real

JEL:
C53, E37, C55, E01, O54

DOI:
https://doi.org/10.47550/RCE/35.2.2

Resumen

Las crisis persistentes en Ecuador ponen de relieve la necesidad de contar con herramientas oportunas de pronóstico económico. Si bien las encuestas empresariales y de consumidores capturan expectativas, su frecuencia mensual limita la capacidad de respuesta ante choques repentinos. Este estudio evalúa si el análisis de sentimiento aplicado a artículos periodísticos diarios puede mejorar los pronósticos del crecimiento del PIB. A partir de 240.000 artículos publicados entre 2001 y 2024, se construyeron cinco indicadores de sentimiento mediante un enfoque de NLP personalizado que combina TextBlob, VADER y el modelo en español de SpaCy. Estos indicadores se incorporaron en un modelo LASSO-ARDL junto con controles tradicionales e indicadores basados en encuestas. Los resultados muestran que los modelos enriquecidos con sentimiento superan significativamente al modelo de referencia basado en encuestas en horizontes de corto plazo (1-3 meses), reduciendo el RMSE y mostrando mayor estabilidad. En horizontes más largos, las diferencias dejan de ser estadísticamente significativas. Los hallazgos resaltan el valor de integrar indicadores en tiempo real derivados de texto en los marcos de pronóstico, ofreciendo un complemento escalable a los métodos tradicionales en contextos económicos emergentes y volátiles.

[1]ORCID: 0009-0002-7710-939X. CRediT: conceptualización, curación de datos, análisis formal, investigación, metodología, validación, visualización, redacción - borrador original, redacción - revisión y edición. Correo electrónico: rpaezp@alumni.usfq.edu.ec.
Copyright © 2025 Paez. Los autores conservan los derechos de autor del artículo. El artículo se distribuye bajo la licencia Creative Commons Attribution 4.0 License.

1. INTRODUCTION

Ecuador encompasses a history of significant challenges, currently facing a rising crime and insecurity crisis, deepening political polarization and economic struggles. Energy shortages and ongoing health system strains have added further complexity to daily life. In this volatile environment, policymakers and private sector leaders need precise insights into public perceptions and trends to make informed decisions and navigate uncertainty effectively.

Business and consumer surveys (BCSs) have been historically essential tools for policymakers and researchers to monitor and forecast the economy. These surveys provide broader insights in people's perception on the current and expected state of economic activity, which is particularly valuable given the often delayed release of macroeconomic indicators. However, their main limitation is inclusiveness: because surveys are generally conducted to mainly representative actors of the economy, they fail to capture the behavior of smaller agents but whom in aggregate still play an in important role in the economic interplay. Thus, restraining efficient answers to rapid, short-term changes in perceptions and shocks to specific sectors of the economy.

An alternative approach is sentiment analysis, which uses natural language processing (NLP) techniques to extract and quantify the tone or emotional context of large volumes of textual data, such as news articles. By converting qualitative expressions into numerical indicators of positive, neutral, or negative sentiment, this method enables researchers to measure public attitudes in real time. Thus, incorporating sentiment analysis offers a valuable approach to addressing Ecuador's complex social and economic challenges by leveraging real-time insights into public perceptions and market conditions. Additionally, the integration of sentiment analysis with econometric models offers a promising solution to the limitations of traditional forecasting. In particular, autoregressive distributed lag (ARDL) models are well suited for combining sentiment indicators with macroeconomic variables, since they can handle variables with different integration orders and capture both short-term dynamics and long-term relationships. This makes them especially relevant in Ecuador's volatile environment, where sudden shocks—such as natural disasters, commodity price fluctuations, or political instability—can quickly reshape public and market perceptions.

The main objective of this study is to evaluate whether sentiment indicators derived from Ecuadorian newspapers between 2001 and 2024 can improve GDP growth forecasts. Specifically, it tests whether incorporating these indicators into an ARDL model with LASSO regularization enhances predictive accuracy compared to a pure survey-based ARDL benchmark. Under this line of research, the paper aims to determine whether high-frequency, text-derived data can complement conventional surveys and provide decision makers with more timely tools for navigating Ecuador's volatile economic environment.

This paper contributes to the literature in three main ways. First, it extends the application of sentiment analysis for macroeconomic forecasting to the case of Ecuador, a small and dollarized economy that remains underexplored in this line of research. Most existing studies have focused on developed countries with deep financial markets and highly institutionalized survey systems, leaving a gap in the understanding of how sentiment analysis can be used in emerging and volatile contexts. Second, the paper develops a set of sector-specific sentiment indicators derived from Ecuadorian newspapers, capturing dimensions such as politics, health, security, and society alongside economic reporting. Third, it evaluates the integration of sentiment into a flexible framework such as LASSO-ARDL, providing evidence on both its advantages and limitations for short versus long-term forecasting horizons. Together, these contributions enrich the growing literature on alternative data sources in economic forecasting and broaden its applicability to the Ecuadorian context.

The paper is organized as follows: section 2 describes the literature review, the data, and key descriptive statistics. Section 3 details the construction of sentiment indicators, the NLP algorithm, and the forecasting methodology and its limitations. Section 4 reports in-sample and out-of-sample forecasting results, incorporating the discussion of findings in light of international studies and policy implications for Ecuador, as well as robustness checks. Section 5 presents the main conclusions and key insights. Annexes are provided at the end of the paper.

2. LITERATURE REVIEW

Sentiment analysis has emerged as a powerful tool to capture real-time information beyond what traditional macroeconomic indicators and surveys can provide, transforming textual data from news, social media, and financial reports into quantifiable indicators of economic expectations. A growing body of literature demonstrates that sentiment measures can improve forecasts of GDP, consumption, unemployment, and financial market variables, often outperforming conventional models during periods of uncertainty or crisis. However, much of this evidence comes from developed economies such as the United States and Europe while applications in emerging economies remain scarce.

This section reviews key international contributions for both developed countries and Latin American contexts that apply sentiment analysis to macroeconomic forecasting, alternative methodologies to the implementation of pre-trained language models for sentiment analysis, the usefulness of LASSO regularization for condensing high-frequency information as well as the

adequacy of ARDL models in this type of analysis and the implementations and mechanisms for language processing using VADER, TextBlob and SpaCy.

2.1. Sentiment Analysis in Macroeconomic Forecasting

Ellingsen et al. (2022) analyzed a corpus of 22,5 million news articles from Dow Jones Newswires to forecast U. S. macroeconomic activity, comparing text-derived predictors with the standard FRED-MD database of economic indicators. Using advanced machine learning models such as LASSO and random forest, they show that non-linear methods extract more predictive value from text than linear models. Their results indicated that news-based sentiment features significantly enhanced forecast accuracy by capturing information absent from conventional indicators. The most notable improvements appeared in consumer spending forecasts: models incorporating news sentiment consistently outperformed FRED-MD benchmarks for consumption growth, suggesting that media sentiment provides timely signals about household behavior that traditional economic series fail to capture. The study demonstrated essentially that textual sentiment can substantially improve macroeconomic forecasts, particularly for consumer-related indicators and during periods of economic stress.

Building on this line of research, Rambaccusing and Kwiatkowski (2020) apply text classification techniques (support vector machines) to approximately 395.000 articles from 12 major UK newspapers covering the period 1990-2018. Their analysis revealed a mixed impact of news sentiment on forecast accuracy. For inflation, incorporating sentiment measures offered little benefit, indicating that price dynamics are already well captured by quantitative indicators or that news content reacts too slowly to inflationary pressures. In contrast, forecasts of real activity and labor market outcomes improved markedly. Tailored sentiment indices—such as keyword counts related to "unemployment" or "growth"—and even single-newspaper sentiment series produced significant gains in predictive accuracy. Importantly, when media sentiment was incorporated, GDP forecasts improved beyond the Bank of England's own projections for up to seven quarters ahead (Rambaccusing & Kwiatkowski, 2020). These results underscored that news sentiment can complement official forecasts, particularly by foreshadowing changes in output and employment more effectively than it does for prices.

Lu et al. (2023) used topic modeling to extract a set of prevalent economic news topics and compute their daily intensity (frequency and sentiment) through a specialized Chinese sentiment dictionary, yielding a daily News Prosperity Index for China. Their findings revealed that Chinese news data indeed contained valuable signals about economic fluctuations helping to track China's GDP growth more closely as well as improving the detection of turning points in the business cycle. Moreover, the authors even succeeded

to identify the six most influential news topics driving China's economic prosperity, which included themes like reform and innovation, infrastructure construction, COVID-19 impacts, economic development (industry and trade), international communication, and corporate innovationN.

Ashwin et al. (2021) depicted how the addition of newspaper sentiment yields significant gains in nowcasting accuracy, as text-based indicators showed substantial improvement over both the European Central Bank's official projections and PMI-based models in predicting current-quarter GDP. The benefit of news sentiment was found to be especially large during the Great Recession (2008-09) and the COVID-19 pandemic, periods when sudden shifts in the outlook were captured in news media before showing up in slower-moving indicators. Furthermore, the authors emphasized that the choice of sentiment dictionary influences performance in different crises. A finance-focused dictionary (or lexicon) demonstrated to be very effective at detecting the downturn in 2008-09 but failed to capture the COVID-19 shock, which involved public health and lockdown terms outside typical financial vocabulary. In contrast, a more general-purpose sentiment dictionary proved more robust across both crises, picking up negative signals in pandemic-related news as well. This suggests that for practical use, sentiment analysis must be tailored to the nature of the shock.

2.2. Sentiment Analysis in Macroeconomics Forecasting for Emerging Economies

Pitta de Jesus and da Nóbrega Besarria (2025) examined how textual sentiment derived from central bank reports can improve macroeconomic forecasts in Brazil. Their study showed that incorporating a time-varying sentiment index significantly enhanced GDP growth predictions. Specifically, models augmented with machine-learning-based sentiment scores consistently reduced forecast errors for real-time GDP growth compared to benchmarks without sentiment. These enriched models also produced more accurate nowcasts and one-quarter-ahead forecasts. Moreover, the authors reported that part of the variation in GDP forecasts from market surveys could also be explained by shifts in the sentiment index, suggesting that central bank communications provide additional real-time information not fully captured by official statistics. These findings have important implications for Latin American forecasting, as they highlight that policymakers and forecasters can systematically use government narratives to refine economic outlooks. By formally incorporating this technique, analysts can improve short-term forecasts and better anticipate turning points in the business cycle.

Delice and Pinto (2025) investigated how sentiment extracted from newspaper articles in Mexico relates to key macroeconomic indicators. Their research demonstrated that a lexicon-based sentiment index closely tracks

inflation dynamics. For example, when consumer price inflation rises, the lexicon-based index tends to move in parallel, outperforming alternatives to sentiment measures. In classification tests, this method achieved 80-88 % accuracy substantially higher than results from more standard rule-based models. Furthermore, Granger-causality tests revealed that the lexicon sentiment index contained predictive information for the consumer price index, reinforcing its value as a timely economic indicator (Delice & Pinto, 2025). Overall, the study highlighted that lexicon-based sentiment measures are both coherent with known inflationary episodes and capable of capturing relevant short-term signals, making them a satisfactory tool for inflation monitoring and forecasting in emerging markets.

Adebiyi et al. (2022) take a different approach by relying on unstructured text from social media rather than traditional news or policy documents, reflecting public opinion in the digital space. They constructed a sentiment index for inflation using Twitter data, classifying tweets into positive, negative, or neutral categories and aggregating them into a monthly index of inflation sentiment. Their findings are striking: incorporating this sentiment index consistently improves inflation forecasts, yielding lower errors across all components, including food, core and headline inflation (Adebiyi et al., 2022). In practical terms, the predicted inflation trajectory by their model aligned more closely with actual outcomes when sentiment is included. The authors also noted that the sentiment-augmented forecasts captured subtle month-to-month adjustments in expectations, while statistical tests confirmed a significant reduction in mean squared forecast errors. These results provide strong evidence that social-media-derived sentiment contributes incremental and valuable information, tightening inflation forecasts for Nigeria and demonstrating its potential as a real-time policy tool in developing countries.

2.3. ARDL Models

Autoregressive Distributed Lag (ARDL) models are a class of econometric models designed to capture both the short-run dynamics and long-run equilibrium relationships between variables. They are flexible as they incorporate lags of both the dependent variable and explanatory variables, allowing the careful examination of how current outcomes relate not only to current predictors but also to past values of both the outcome and predictor variables. ARDL models are particularly useful when studying both short-term and long-term trends because they separate out immediate adjustments from deeper, more persistent relationships. In the short run, transitory shocks may cause variables to deviate from their long-term path; ARDL models can pinpoint how quickly and through which channels the system returns to equilibrium. Over the long term, ARDL models help identify stable, long-run coefficients that indicate the underlying relationship between variables, ensuring that analysts

understand the fundamental linkages that persist over time, even after short-term fluctuations wash out.

The proposed ARDL framework in this paper offered distinct advantages over alternative methods commonly used in the academic literature, particularly for the case of Ecuador's economic forecasting. Traditional methods, like vector autoregressive (VAR) models or vector error correction models (VECM), are popular for capturing relationships between macroeconomic variables but are often ill-suited to handle mixed-frequency data and high-dimensional datasets like those involving sentiment indicators. VAR and VECM also require all variables to share the same integration order, which limits their flexibility when incorporating sentiment data that might behave differently from traditional economic variables (Pennington et al., 2014). Additionally, the rigidity of these models in the presence of structural breaks, such as those caused by the 2016 earthquake or the COVID-19 pandemic in Ecuador, often leads to reduced forecast accuracy (Barbaglia et al., 2024). In contrast, machine learning models, such as neural networks or ensemble methods, are increasingly applied in forecasting tasks due to their ability to capture complex, non-linear relationships.

However, these models often lack interpretability and require large volumes of data to perform well, making them less practical for countries like Ecuador, where historical datasets may be limited in scope and frequency. Unlike black-box machine learning methods, the ARDL models retain the interpretability of traditional econometric techniques while leveraging the benefits of feature selection to enhance forecasting accuracy. For instance, neural networks tend to overfit when used with small or medium-sized datasets, while ARDL models-enhanced by LASSO regularization in contrast ensures that only the most relevant predictors are included, reducing noise and improving generalization (Medeiros & Mendes, 2017).

2.4. LASSO Regularization

Medeiros and Mendes (2017) demonstrated the value of this regularization by estimating stationary ARDL models with generalized autoregressive conditional heteroskedasticity (GARCH) errors through LASSO penalization. Their study showed that adaptive LASSO effectively identified relevant variables with a probability converging to one, thereby achieving oracle efficiency—that is, the estimator's distribution aligns with that of an oracle-assisted least squares estimator, which assumes prior knowledge of the relevant variables (Medeiros & Mendes, 2017). Ultimately, they illustrated that the lasso estimator can serve as a basis for constructing initial weights in the adaptive LASSO procedure, thereby improving performance in finite samples.

Moreover, Jiang et al. (2022) displayed the usefulness of lasso regularization when addressing the challenge of modeling and forecasting economic

variables in a globalized context, where interdependencies among regions and high-dimensional datasets complicate traditional econometric approaches. In their study, the researchers proposed a time-varying parameter global vector autoregressive (TVP-GVAR) model integrated with machine learning techniques, including the aforementioned LASSO-type regularization. Their results showed how the LASSO-enhanced TVP-GVAR outperformed traditional GVAR and other econometric models in forecasting accuracy (Jiang et al., 2022). The time-varying parameters captured the dynamic relationships between variables more effectively than static models. Furthermore, by selecting a subset of significant predictors, the model offered clearer insights into the economic drivers and their evolving impacts over time. Ultimately, the model's ability to identify key economic interdependencies provided valuable insights for policymakers, particularly in understanding how global shocks propagate across regions.

2.5. Natural Language Processing Algorithms

Natural language processing (NLP) is an interdisciplinary field that combines linguistics and computer science to apply mathematical and computational methods to natural language. Its applications are diverse, encompassing text-to-speech conversion, automatic translation, text correction, information retrieval, and many other areas (Lukauskas et al., 2022). NLP is widely used by various companies for these tasks and others, including sentiment analysis.

Traditionally, sentiment analysis has relied on statistical methods that calculate indices based on the frequency of certain words or phrases within a text corpus (Buckman et al., 2020). These frequency-based approaches, often referred to as bag-of-words models, consider each word independently without accounting for context or linguistic nuances. While effective to some extent, these methods can miss subtleties such as sarcasm, negations, or the influence of surrounding words, potentially leading to less accurate sentiment assessments.

Lexicon-based sentiment analysis algorithms, on the other hand, offer a holistic view by incorporating dictionaries of sentiment-laden words along with rules that consider context and grammatical structures. Tools like VADER (Valence Aware Dictionary and Sentiment Reasoner) or TextBlob aim to combine heuristics to evaluate the intensity and polarity of sentiments expressed in text extracts. This approach allows for integrating higher variation when performing statistical analysis on social phenomena, as it accounts for modifiers, intensifiers, and negations that affect the sentiment conveyed (Hutto & Gilbert, 2014). Moreover, lexicon-based methods are transparent and interpretable since they rely on predefined sentiment scores assigned to words and phrases. They are particularly effective in domains where labeled training data is scarce or when the goal is to understand specific language patterns contributing to

overall sentiment. However, they require regular updates to the lexicon to stay current with evolving language use, slang, and domain-specific terminology.

The results from lexicon-based sentiment analysis can enrich and sharpen economic analysis, such as in forecasting economic and financial variables (Einav & Levin, 2014). Sentiment derived from news is especially useful when predicting macroeconomic variables, as it allows the state of the economy to be monitored in real-time—unlike official releases of macroeconomic data that are seldom available and often significantly delayed (Einav & Levin, 2014). Consoli et al. (2022) suggest that understanding human sentiments can provide better and clearer insights into market dynamics in economic analysis. Consequently, improving forecasting models performance (Malandri et al., 2018) and providing more accurate results that serve as the basis for more informed economic and financial decisions (Chang et al., 2016).

2.6. Valence Aware Dictionary and Sentiment Reasoner

Developed by C.J. Hutto and Eric Gilbert in 2014, VADER (Valence Aware Dictionary and Sentiment Reasoner) is a lexicon and rule-based sentiment analysis tool specifically designed for text from online platforms such as tweets, comments, and web page reviews. It combines a comprehensive sentiment lexicon with heuristic rules to evaluate the intensity of emotions conveyed in text. Since VADER is lexicon-based, it does not require training on labeled datasets, saving time and resources. Additionally, VADER's lexicon is customizable, allowing users to add or modify words and their associated sentiment scores to tailor it to specific domains or applications. Although primarily designed for English, VADER's methodology can be adapted for other languages by integrating it with external statistical models trained in the desired language.

At the heart of VADER lies its sentiment lexicon, a carefully curated list of words, phrases, and symbols annotated with sentiment scores. Each word or phrase in the lexicon is assigned a valence score that reflects its emotional intensity. Key aspects of the sentiment lexicon include:

- **Annotated scores**: Words like "happy" and "excellent" are assigned positive scores (e.g., +2,1, +3,0), while words like "sad" and "terrible" receive negative scores (e.g., -2,0 -3,5). Neutral words have a score close to 0.
- **Coverage of informal language**: The lexicon includes slang, emoticons, abbreviations, and colloquial expressions such as "lol," "meh," or ":)", ensuring compatibility with informal text like tweets and digital media posts.

- **Intensifiers and modifiers**: Words like "very" and "extremely" amplify the sentiment intensity of adjacent words, while negations like "not" or "barely" reduce or invert sentiment.

After processing the text using its lexicon and rules, VADER computes sentiment scores. These scores are presented in four categories:

- **Positive**: The proportion of the text that conveys positive sentiment.
- **Negative**: The proportion of the text that conveys negative sentiment.
- **Neutral**: The proportion of the text that is neutral or lacks strong sentiment.
- **Compound score**: A normalized, weighted composite score that combines positive, negative, and neutral sentiments into a single value ranging from -1 (most negative) to +1 (most positive).

Due to its adaptability and versatility, VADER has been increasingly utilized to analyze sentiment across different domains. For instance, Soni et al. (2023) demonstrated the immediate impact of news headlines on a company's stock performance using NLP. By comparing traditional machine learning algorithms with VADER sentiment analysis, they found that the VADER-based approach provided superior accuracy in predicting stock market trends, offering a more reliable tool for investors and businesses navigating the highly volatile financial landscape (Soni & Mathur, 2023).

Similarly, Rutkowska and Szyszko (2024) explored sentiment analysis in central bank communications by assessing sentiments using four different lexicon dictionaries. Through a series of analyses—including lexicon content comparison, performance tests for highly positive and negative messages, and statistical tests of dictionary alignment and correlation—they concluded that the choice of dictionary significantly impacts the detection of central bank intentions (Rutkowska & Szyszko, 2024). Their findings suggested that applying dictionary methods, such as those used in VADER, to assess policy release sentiments in small open economies can effectively capture and analyze whether the intended messages by central authorities are properly conveyed to the public.

In politics, VADER has been extensively used to study public opinion, campaign strategies, and policymaker communication. Its ability to process short, informal texts makes it an effective tool for analyzing sentiment in political discourse. Researchers have leveraged VADER to evaluate the impact of politicians' social media activity on public sentiment and electoral outcomes. For instance, Ali et al. (2022) conducted a large-scale sentiment analysis of 7,6 million tweets pertaining to the 2020 U. S. Presidential Election using VADER. Their novel approach included identifying tweets and user accounts

that were later deleted or suspended, allowing them to observe sentiments across accessible, deleted, and suspended tweets and accounts. They discovered that deleted tweets posted after Election Day were more favorable toward Joe Biden, while those leading up to Election Day were more positive about Donald Trump. Additionally, older Twitter accounts tended to post more positive tweets about Joe Biden (Ali et al., 2022). This study underscored the importance of conducting sentiment analysis on all posts captured in real time, including those now inaccessible, to determine the true sentiments surrounding significant political events.

Collectively, these applications of VADER at the intersection of economics and politics sentiment highlight its innovative and powerful utility. By exploring sentiments in central bank communications and analyzing tweets and addresses by government officials to assess public sentiment surrounding fiscal policies, VADER provides actionable insights into the interplay between political rhetoric and economic decision-making. These applications reveal how sentiment extracted via VADER can inform researchers and policymakers about public reactions to fiscal and monetary policies, offering valuable tools for shaping effective communication strategies and making informed decisions in both economic and political spheres.

2.7. TextBlob

TextBlob is a Python library for processing textual data for sentiment analysis by leveraging pre-trained models and a lexicon-based approach. It processes text at the sentence level, breaking the input into sentences and analyzing each one independently. This allows it to handle variations in sentiment across a text document, providing a more granular analysis (Loria, 2018). For instance, if one sentence is positive and another is negative, the tool calculates separate scores for each and then averages them to produce an overall sentiment for the entire document. This methodology helps to capture mixed sentiments in texts like reviews or news articles.

A key advantage of TextBlob is its simplicity and integration into natural language processing workflows. Unlike more advanced deep learning models, TextBlob requires no training or additional setup. Regarding score calculations, TextBlob assigns a sentiment polarity score to text through the polarity ranges from -1 (negative sentiment) to +1 (positive sentiment), and a subjectivity score ranging from 0 (objective) to 1 (subjective). This dual scoring system allows TextBlob to capture not only the overall sentiment but also the degree of personal opinion versus factual content in the text.

These scores are calculated by evaluating individual words and phrases in the text against a predefined lexicon of sentiment-bearing terms. The underlying sentiment lexicon in TextBlob contains words with predefined polarity and subjectivity values (Loria, 2018). For example, words like “great”

or "terrible" have strong positive or negative polarity values, respectively. TextBlob aggregates these individual word scores across the text, adjusting them based on linguistic features such as negations ("not good" flips polarity) or modifiers ("very good" amplifies positivity). This approach ensures that sentiment analysis accounts for contextual nuances rather than relying solely on traditional word-level scores. Besides, TextBlob works directly in the users setup making it accessible for developers and analysts who need quick and reliable sentiment analysis. While its lexicon-based approach may lack the sophistication of neural network-based models, it remains effective for many use cases, particularly for straightforward text data.

TextBlob's utility extends to multilingual sentiment analysis, including Spanish, thanks to its support for language translation and processing through the integration of external libraries like Google Translate. This capability is particularly useful for Spanish sentiment analysis in cases where no native Spanish lexicon is available, as TextBlob can translate Spanish text into English before performing sentiment analysis using its English lexicon (Unipython, 2019). While this introduces some dependency on translation accuracy, it provides a direct solution for handling Spanish texts, enabling sentiment analysis even when no specialized tools for Spanish are available. Furthermore, TextBlob's simplicity and adaptability make it a practical choice for preprocessing Spanish text before applying more advanced or domain-specific models. For Spanish sentiment analysis, it can also be combined with custom lexicons or rule-based enhancements to address language-specific nuances through other language processing models. This flexibility, coupled with its ease of use, makes TextBlob a versatile tool for researchers and developers working with Spanish-language sentiment analysis. Overall, TextBlob is a valuable tool for standard sentiment analysis tasks where interpretability and simplicity are priorities. Despite its strengths, TextBlob has limitations as it relies heavily on its predefined lexicon, making it less effective for domain-specific texts or slang-heavy language that might not be well-represented in the lexicon (Unipython, 2019). Additionally, TextBlob struggles with complex linguistic constructs like sarcasm or context-dependent sentiment(s).[1]

2.8. SpaCy Large News Model

The SpaCy large news model is a robust and pre-trained natural language processing (NLP) model designed to process and analyze text, particularly within the domain of news and general information. Built on a combination of machine learning techniques and linguistic rules, this model enables users to extract meaningful insights from large text datasets. Moreover, its architecture,

[1] To account for deeper understanding and adaptability, the approach was integrated with VADER interface to obtain a more suitable solution for the context-specific scenario of news articles.

pre-training, and processing pipeline make it a powerful tool for handling the complexities of natural language. At its core, the SpaCy large news model leverages a deep learning framework based on transformers or convolutional neural networks (CNNs) to encode semantic and syntactic features of text (Honnibal & Montani, 2020). This pre-training allows it to generalize effectively across various news-related tasks, such as named entity recognition (NER), dependency parsing, and part-of-speech (POS) tagging, which are foundational to understanding text structure and meaning.

Additionally, the model's pipeline consists of several key components, each contributing to the processing and analysis of text. Tokenization is the first step, where the text is broken down into individual tokens, such as words or punctuation marks, while preserving language-specific rules, such as contractions or compound words to ensure accurate downstream analysis. After tokenization, POS tagging assigns grammatical roles (e.g., nouns, verbs) to each token while dependency parsing identifies syntactic relationships between tokens, such as subject-verb-object structures (Pennington et al., 2014). These tools enable the model to map grammatical and logical relationships within a sentence.

Another key aspect of the model is its use of vector embeddings to capture semantic meaning. SpaCy employs word embeddings, such as GloVe or transformers-based embeddings, to represent words as dense numerical vectors in high-dimensional space. These embeddings encode the relationships between words based on their context in the training data. For instance, the model understands that "bank" in the context of "river bank" differs from "bank" in "financial institution." This semantic understanding enhances its ability to disambiguate words and phrases in diverse contexts, making it a versatile tool for analyzing nuanced language.

The SpaCy large news model is also highly customizable and efficient, which makes it suitable for a wide range of NLP tasks. Users can fine-tune the model on domain-specific corpora, such as financial or scientific articles, to improve its performance in specialized contexts. Additionally, the model is optimized for speed and scalability, enabling it to process large text datasets quickly. This efficiency makes it particularly valuable in news analytics, where real-time or near-real-time processing of vast quantities of data is often required. The combination of pre-trained knowledge, linguistic capabilities, and adaptability ensures that the SpaCy large news model remains a go-to solution for comprehensive text analysis in the news domain (Vaswani et al., 2017). In addition, SpaCy contains specialized text cleaning and normalization functions required for pre-processing the data before running sentiment analysis on the text extracts.

3. MATERIALS AND METHODS

3.1. Data

The data were gathered from publicly accessible sources, including the websites of the National Institute of Statistics and Censuses (INEC, by its initials in Spanish), the Central Bank of Ecuador (CBE), and the Trading Economics economic data provider. The dependent variable of this study is Ecuador's annual GDP growth rate, sourced from the CBE's quarterly reports. To provide a comprehensive analysis of economic activity, monthly data on national exports and remittances sent to Ecuador from abroad were also included. To account for international influences, the research incorporated the U. S. Federal Reserve's interest rate and West Texas Intermediate (WTI) crude oil prices, both obtained from the ALFRED database. For structural information about economic expectations, data from the Economic Expectations Index (IEE) published by the CBE was also included.

These survey indicators represent the expectations from medium-high and high-income companies across the four key sectors of the Ecuadorian economy published by CBE through an analytical bulletin detailing the results of the IEE from 2001-2024. This index is derived from the Monthly Business Opinion Survey (EMOE, by its initials in Spanish), which targets companies with the highest sales income in their economic sectors. The IEE captures the opinions of company directors in construction, commerce, manufacturing, and services regarding the current economic situation and future prospects of their companies.

Furthermore, the dataset also included the Consumer Confidence Index (ICC) from CBE, which is a weighted average measuring households' perceptions of their economic situation, consumption patterns, and the country's economic condition over the previous month and the next three months. The Ecuador Business Confidence Index (BUSS_CONFIDENCE) records the ease of doing business and expectations about the company's situation and the economy as a whole, serving as an indicator of the overall business climate. The dataset also included sentiment indicators for the health, economy, politics, security, and society sectors, built through the identification strategy described in the sentiment identification subsection.

To account for pivotal exogenous macroeconomic events, the dataset compiled key recessionary periods and shocks in the last 14 years as dummy variables. These periods were identified based on news data from IndexMundi news repository and CBE´s report of economic cycles for 2022 and 2023. By explicitly modeling these periods, the analysis ensured that the effects of these external shocks were appropriately isolated allowing the models to better capture the underlying relationships between survey and sentiment indicators on the target variable in their respective frameworks. The periods included:

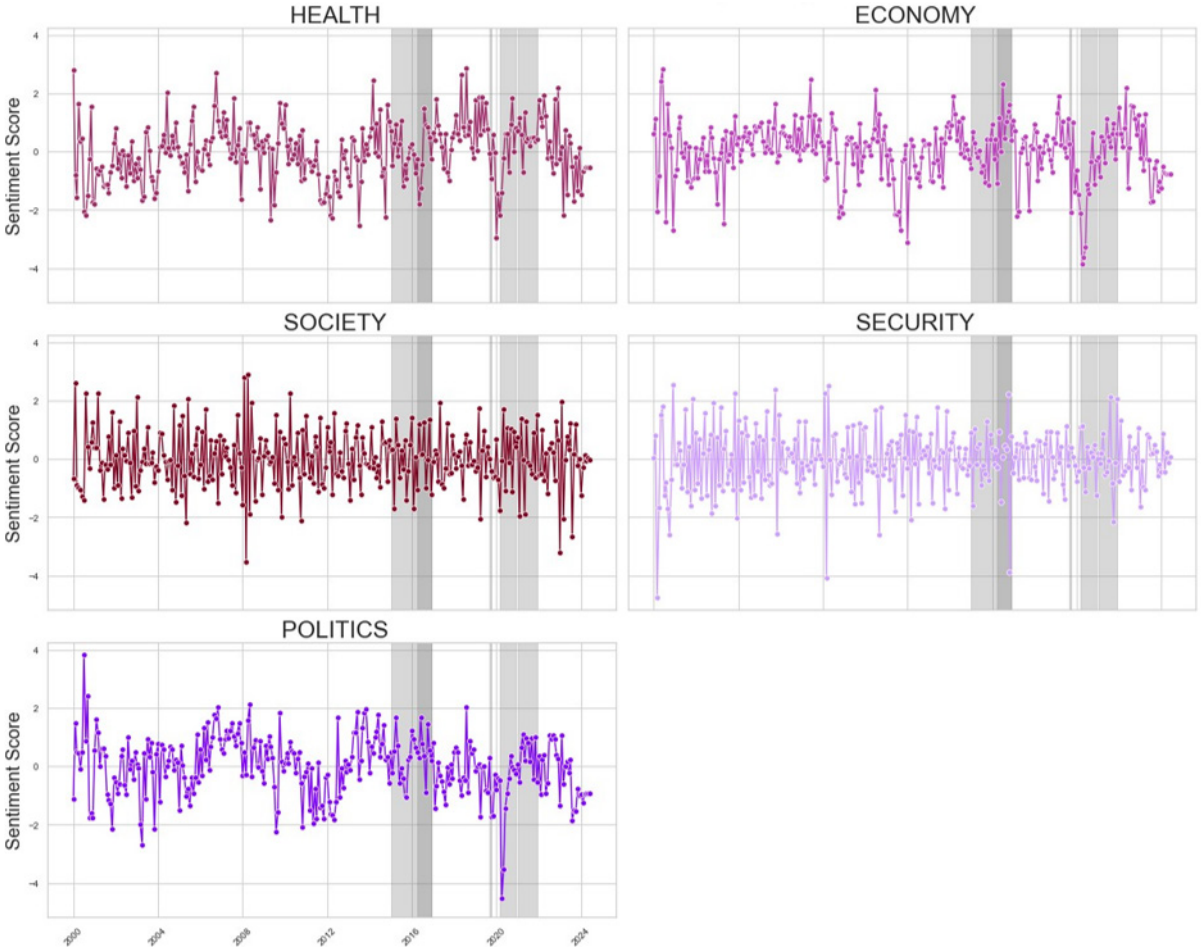
- Recession of 2015-2016 (“2015-01-01” to “2016-12-01”).
- 2016 Earthquake (“2016-04-01” to “2016-12-01”).
- 2019 Fuel Subsidy Removal (“2019-10-01” to “2019-11-01”).
- COVID-19 Pandemic (“2020-03-01” to “2020-12-01”).
- Post-pandemic recession (“2021-01-01” to “2021-12-01”).

For the sentiment dataset, USFQ DataHub provided digitized newspaper titles, extracts, and full articles, totaling 1,5 million records. After preprocessing and analysis, a stratified random sample was drawn—2.000 articles per section—resulting in approximately 10.000 articles per year and 240.000 in total.

3.1.1. Descriptive Statistics

The analysis of sentiment across different sectors in Ecuador depicted in Figure 1 shows the sentiment indicators of economic activity together with recessionary periods. The information was aggregated from daily frequency to a monthly frequency by averaging the values within each month and standardizing the measures to a normal distribution with mean 0 and variance 1, in order to make them comparable across indicators. The results revealed significant volatility, particularly during recessionary periods, and brought out the intricate interplay between various domains of public perception.

Figure 1. Time Series of the Standardized News-Based Sentiment Indicators for Different Sectors



Note: The sentiment is averaged within each month and sampled at a daily frequency. The shaded areas represent the recession established from IndexMundi information.

The *health sector* exhibited pronounced fluctuations, with significant peaks and troughs in sentiment, especially during times of economic downturn. The COVID-19 pandemic marked a drastic drop in health-related sentiment, reflecting heightened public concern over the healthcare system's capacity to handle the crisis. Recovery in this sector appears sluggish, with the effects of the pandemic extending into the post-pandemic recession of 2021. This delayed recovery indicated deep-seated public anxiety over systemic issues and demonstrated the sector's high reactivity to external shocks. The *economic sentiment* trends displayed consistent oscillations but were sharply impacted by major downturns such as the recession of 2015-2016, the lifting of fuel subsidies in 2019, and the pandemic-related disruptions in 2020. During the 2015-2016 recession, economic sentiment suffered a prolonged decline, showcasing a slow and gradual recovery that underscores the depth of economic instability during that period. The 2019 fuel subsidy removal further

exacerbated economic concerns, leading to one of the steepest declines in sentiment. Even as economic sentiment began to stabilize post-pandemic, it did not fully recover to pre-crisis levels, reflecting lingering structural challenges and the lasting impact of recessionary periods on public perception.

The *societal sentiment* is characterized by frequent fluctuations closely aligned with crises that impact social structures. The 2016 earthquake caused a steep decline in societal sentiment as the immediate social and infrastructural repercussions dominated public discourse. Similarly, the pandemic brought sustained negativity, reflecting extensive social disruptions caused by lockdowns, unemployment, and strained community networks. Interestingly, societal sentiment demonstrated quicker rebounds compared to the health and economy sectors, possibly due to an inherent adaptability to changes in social narratives and resilience in the face of social challenges. The *security sentiment* maintained relative stability over time, with significant declines occurring primarily during acute crises. Interestingly, the lifting of fuel subsidies in 2019 caused a minimal drop in security sentiment, likely to public uncertainty of the consequences of the strike on social security and governmental capacity to manage the ensuing protests. The pandemic also saw a notable decline, pointing to heightened perceptions of insecurity during periods of societal strain. However, recovery in this sector tended to be faster, as security-related concerns often normalize once immediate crises subside.

The *political sentiment* displayed notable relations with other indicators, acting seemingly both a driver of and response to shifts. Periods of economic and social turmoil corresponded with sharp declines in political sentiment, highlighting public dissatisfaction with governmental policies and their perceived inadequacy in addressing crises (Starke, 2024; Center, 2019; Méda, 2024). Political sentiment showed delayed improvement after the 2015-2016 recession but stabilized more quickly following the pandemic, possibly due to government efforts in managing recovery measures. This mixed recovery pattern underscores the critical role of effective governance in shaping public perception during and after crises.

The interconnectedness of these sectors becomes particularly evident during recessionary periods and crises. On general, the recession of 2015-2016 had a profound impact on sentiment across all sectors. Economic sentiment saw a sustained decline with slow recovery, indicating the deep and lasting nature of this downturn. Health and societal sentiments also declined, albeit less severely, pointing to the interconnected effects of economic instability on public health and social cohesion. Security and political sentiments demonstrated more resilience during this period, reflecting a degree of stability in these domains despite economic struggles.

Similarly, the 2016 earthquake primarily affected societal and security sentiments, with marked dips in public perception tied to the disaster's immediate social and safety consequences. While economic and political sentiments also declined during this period, the impact was less drastic, as the earthquake's

effect on broader economic and policy narratives was more contained. Recovery from this event was relatively rapid compared to other crises, reflecting the resilience of societal narratives and effective disaster response measures.

The fuel subsidy policy change in 2019 caused one of the sharpest declines in sentiment across sectors. Economic, societal, and health sentiments were particularly affected, reflecting widespread public dissatisfaction with the policy's immediate impact on living costs and social stability. Political sentiment also suffered as public perceptions of governance and accountability seems to have been deteriorated. Recovery from this period was uneven, with societal and security sentiments rebounding faster, while economic and political perceptions remained subdued, highlighting the enduring impact of economic policies on public sentiment.

The pandemic emerged as the most synchronized crisis, with all sectors showing sharp declines in sentiment. Health and societal sentiments were the most heavily impacted, reflecting concerns over healthcare system strain, lockdowns, and social isolation. Economic sentiment also suffered a steep decline, mirroring the global economic fallout of the pandemic. Political sentiment dipped sharply, indicating public dissatisfaction with governmental responses. Recovery across sectors was slow, particularly for health and the economy, highlighting the pandemic's deep and multifaceted impact on public perceptions and the prolonged effects of recessionary periods. During the post-pandemic recession, economic sentiment remained subdued, revealing lingering structural challenges. Other sectors, including health and society, began to show signs of stabilization, with gradual recovery in public perceptions. Political sentiment showed some improvement, reflecting renewed public confidence in recovery measures. Security sentiment, which had normalized after the pandemic, remained stable during this period, indicating reduced societal volatility.

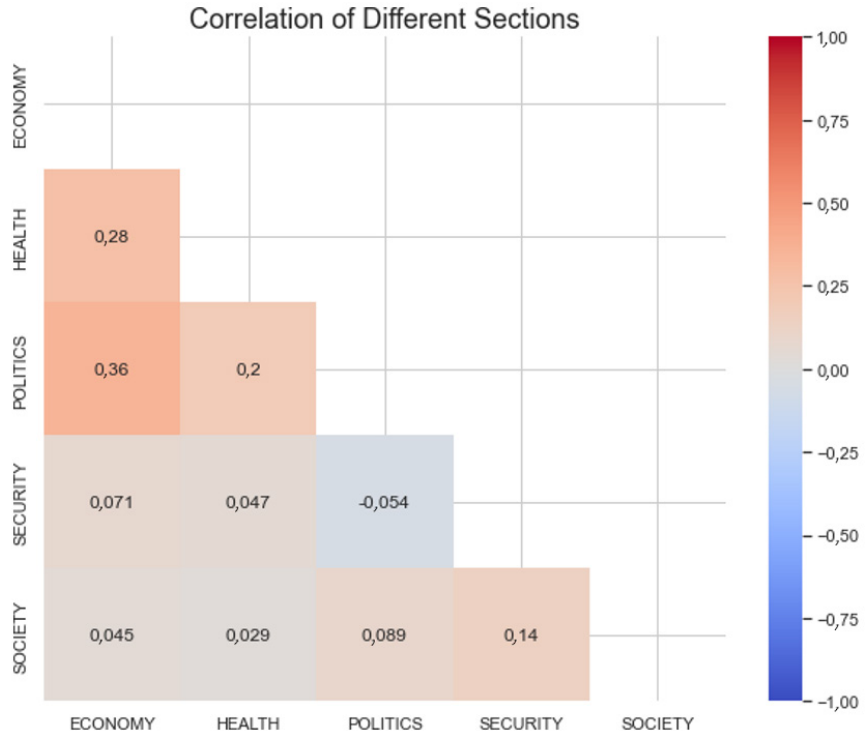
Volatility across sectors varies significantly, with health, economy, and politics showing the highest reactivity to external shocks. These sectors are deeply intertwined, as economic downturns often cascade into health crises and political instability (Forum, 2023). Security and societal sentiments are relatively stable, with more rapid recovery patterns, suggesting a certain resilience to acute shocks. The interconnectedness of these sectors is particularly evident during crises, where economic disruptions consistently affect societal perceptions, and health crises amplify security and political concerns (Fathi, 2022). Political sentiment, in turn, seems to shape and is shaped by economic and social dynamics, demonstrating its central role in public discourse during periods of instability. In addition, recovery patterns highlight differences in resilience across sectors. Security and societal sentiments tend to stabilize quickly, reflecting the public's ability to adapt to changes in social and security narratives. In contrast, health and economic sentiments exhibit prolonged recovery periods, indicating the lasting impacts of structural challenges and the complexities involved in restoring public confidence in these

areas. Political sentiment recovery is mixed, as it often depends on the effectiveness of governance and policy responses to crises.

The time series results from the sentiment indicators suggest a series of co-movement in sentiment across different sections. To analyze this commonality, Figure 2 displays correlation estimates between different sentiment indicators calculated at a monthly frequency. The most significant correlation is observed between Economy and Politics, indicating a moderately strong positive relationship. This reflects how economic conditions influence public perceptions of political governance and policy effectiveness. Furthermore, the correlation between Economy and Health highlights a positive relationship, albeit less pronounced than that between Economy and Politics. Economic perceptions seem to have a direct effect on public health sentiment, as recessions can strain healthcare systems, reduce access to medical services, and exacerbate public concerns about health.

The correlation between Health and Politics reflects a weaker yet still positive relationship displaying how public sentiment toward healthcare systems is still influenced by political decision-making, especially during health crises. While the sentiment related to Security exhibits weak correlations with most other sectors, with the exception of a slightly negative correlation with Politics. The lack of strong associations suggests that security sentiment is relatively insulated from broader public perceptions of economic, health, or political conditions. Besides, the sentiment for Society shows weak positive correlations with other sectors, such as Politics and Economy. These weak correlations suggest that societal sentiment, which reflects social cohesion and community resilience, is only marginally influenced by economic or political conditions. Instead, societal sentiment may be driven by broader cultural and social factors, such as community solidarity or generosity.

Figure 2. Correlation Estimates on the Sentiment Indicators Across Sections

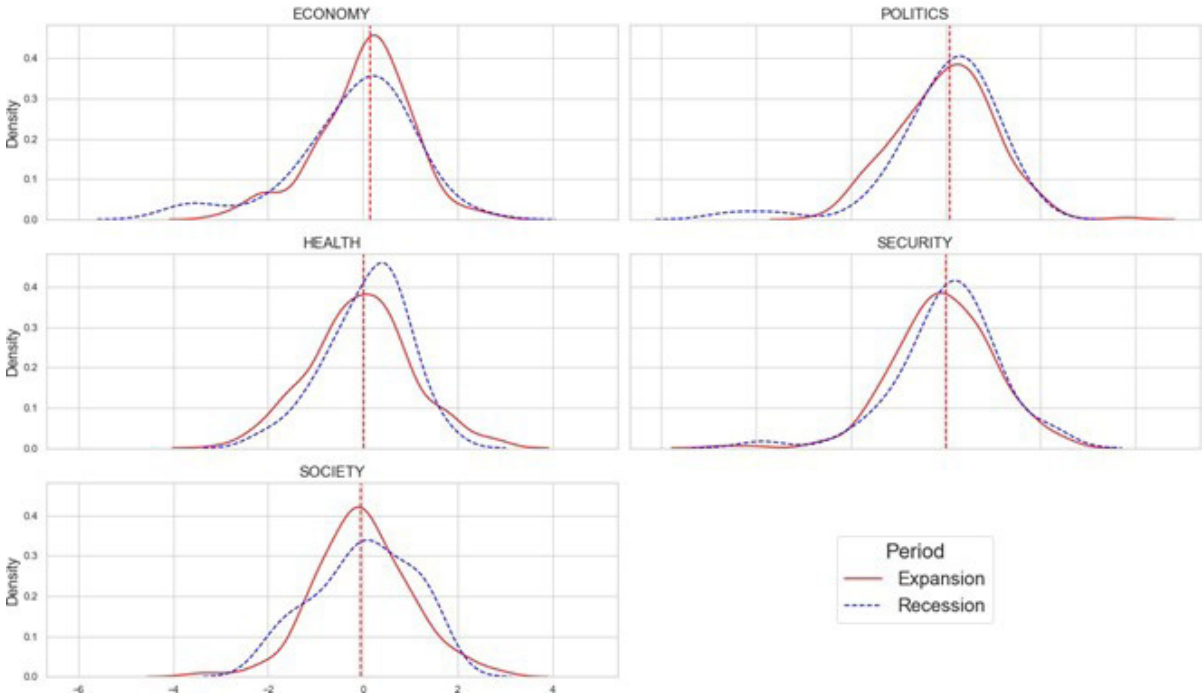


Note: Redder colors indicate a large correlation in absolute value.

The dependence of the scaled sentiment scores on the business cycle is apparent in Figure 3, which displays the behavior of the estimated indicators during expansionary and recessionary periods. By examining sentiment distributions across key sectors during these periods, the research aimed to gain critical insights into how economic cycles influence public perceptions through kernel density estimation (KDE) plots. These plots reveal how sentiments vary in magnitude, spread, and recovery patterns during times of stability and crisis across different sectors.

Economic sentiment exhibited the most pronounced decline during recessions, with the recessionary KDE curve shifting significantly leftward. This shift depicted the severity of negative perceptions during economic downturns, while the broader spread reflected heightened variability and uncertainty among the public. A significant proportion of sentiment scores fell below the mean value depicted by the dotted red line, emphasizing the profound impact of economic instability on public confidence. The sector's slow recovery further highlighted the lasting scars that economic crises leave on public perception, making it one of the most sensitive domains during downturns.

Figure 3. Kernel Density Estimates of the Standardized Sentiment Indicators During Expansions and Recessions by Section



In contrast, Politics sentiment showed a more modest leftward shift during recessions compared to economic sentiment. The overlap between expansionary and recessionary distributions indicated that while recessions negatively impacted perceptions of governance, these effects were less severe and more contained. The symmetry of the curve suggested that Politics sentiment stabilized more quickly, as the public's focus shifted toward governance solutions during crises. Nonetheless, sharp dips in sentiment during significant economic and social crises highlighted the importance of effective government responses in maintaining public confidence during challenging times.

The health sector displayed a marked decline in sentiment during recessions, with a significant leftward shift and a notable concentration in the negative range. Health-related crises, such as the COVID-19 pandemic, amplified public anxieties, particularly when healthcare systems faced strain. The wider spread of the recessionary curve suggested sustained public concerns and heightened variability in perceptions.

The slow recovery of health sentiment indicated systemic challenges in restoring public confidence, underscoring the importance of addressing public health concerns comprehensively during and after economic crises.

Security sentiment, unlike other sectors, exhibited minimal leftward shift during recessions, with significant overlap between expansionary and recessionary distributions. This stability reflected the public's relatively consistent perception of security, even amid economic and social disruptions.

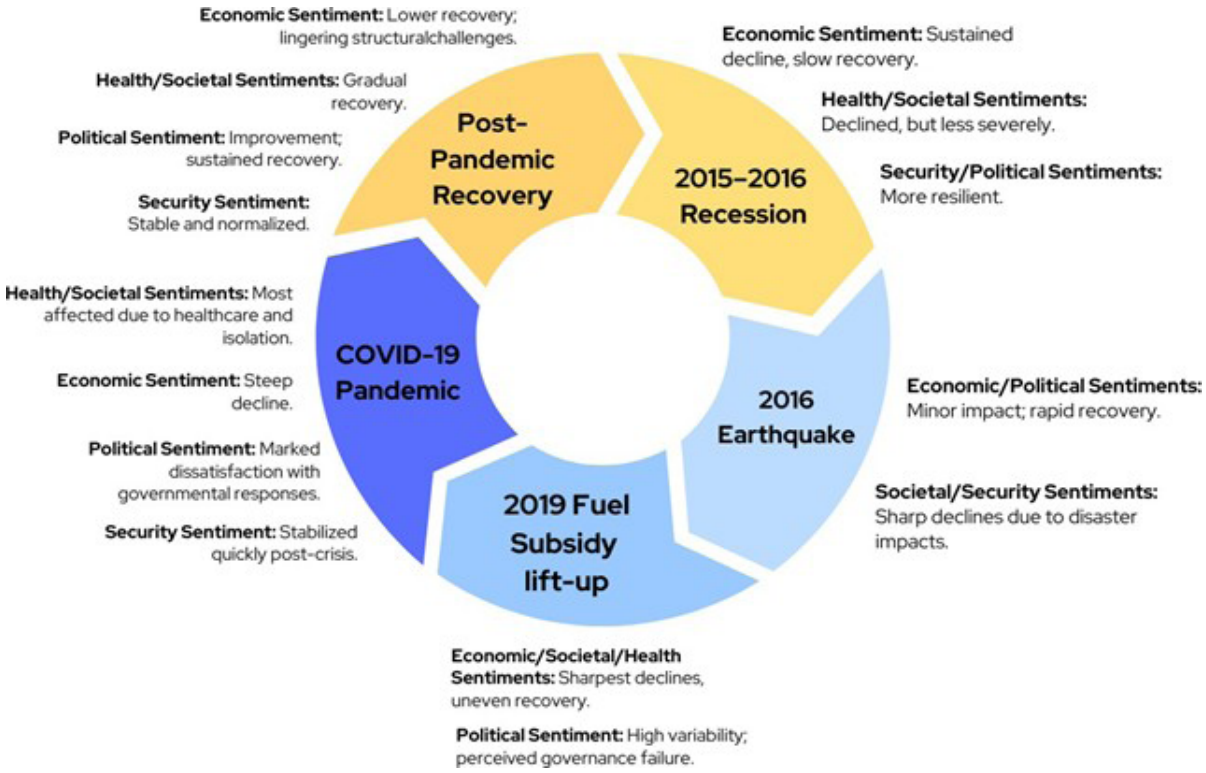
While security sentiment experienced occasional declines during acute crises like protests or pandemics, it normalized quickly once stability returned. The narrower spread during recessions further highlighted its resilience, making it one of the least volatile sectors.

Societal sentiment displayed a more interesting multimodal structure during recessions, mainly due to a complex and heterogeneous public discourse. This suggests that societal sentiment does not respond uniformly to economic downturns; instead, it reflects a blend of narratives shaped by various factors such as media coverage, cultural resilience, and differing socio-economic experiences. A more nuanced analysis from a VAR framework[2] suggested a bidirectional relationship between macroeconomic survey indicators and societal sentiment, creating a dynamic interplay that could drive or stabilize public perceptions. Declining consumer and business confidence likely contributed to the dominant negative narrative—the main bump in the leftward shift of the KDE. The coexistence of multiple sentiment clusters, as evidenced by the secondary bumps in the KDE, could reflect polarized public discourse, where societal concerns are both influenced by and further exacerbate declining economic expectations.

Finally, recovery patterns further illustrated differences in resilience across sectors. Security and societal sentiments stabilized quickly, reflecting the public's ability to adapt to changes in social and security narratives. In contrast, health and economic sentiments exhibited prolonged recovery periods, underscoring the structural challenges in these areas. Politics sentiment recovery varied depending on governance effectiveness, highlighting the critical role of policymaking in restoring public confidence. Economic and health sentiments seemed to be highly sensitive and slow to recover, emphasizing the profound and lasting impact of economic crises on public perceptions in these domains. Politics sentiment's stability depended on effective governance during crises, while security and societal sentiments demonstrated resilience by stabilizing more rapidly due to the public's adaptability. The fact of recognizing these dynamics is crucial for policymakers aiming to address public concerns and foster recovery during and after economic downturns.

[2] See annexes section for the specifications of the model construction and further analysis on Granger-causality and instantaneous-causality.

Figure 4. Effects of Major Crises on Different Sentiments and Their Recovery Patterns



3.2. Methodology

This section explains the methodology for sentiment analysis and the econometric modeling application. To perform sentiment analysis on Spanish-language news articles, an NLP algorithm was customized by integrating SpaCy, VADER and TextBlob libraries from Python. Specifically, SpaCy's large news model was employed for robust preprocessing due to the languages' linguistic complexities capabilities such as handling of gendered nouns, verb conjugations, and data cleaning protocols, while VADER and TextBlob advanced features supported the tasks of sentence segmentation, tokenization, lemmatization, and named entity recognition. This allowed the analysis to focus on meaningful linguistic units, improving accuracy and contextual relevance by processing texts in 5.000-character chunks to ensure scalability and efficient handling of articles.

3.2.1. Sentiment Identification Strategy

To structure and analyze sentiment indicators, articles from the main newspapers companies like *El Comercio, La Hora, Diario El Universo, Expreso, Metro, El Mercurio*, among others from year 2001 to 2024 were digitized

to identify five key sectors of interest that have been historically prone to systemic instability in the Ecuadorian economy: ECONOMY, POLITICS, HEALTH, SOCIETY, SECURITY. To categorize the newspaper extracts into these sectors, a customized function was applied that used a predefined dictionary to perform the classification task; extracts that couldn't be fitted into any of these categories were labeled as "OTHER."

The function takes a text input and first checks whether the text is null or an empty string after trimming any leading or trailing spaces. If the text is empty, the function returns the label "OTHER"; otherwise, it converts the text to lowercase to ensure consistency in processing. The SpaCy's functions for language processing analyze each classified phrase for unrecognized characters (e.g @#$) and replace them with a word based on its most direct semantical root[3]. After that, the function initializes a dictionary called *match_counts* with keys representing the defined sectors and values set to zero. For each sector, the function counts the number of occurrences of predefined keywords associated with that sector within the text, then the counts are updated in the dictionary accordingly. These keywords are specific terms indicative of the sector's content. After processing all sectors, the function identifies the category with the highest keyword match count. If the highest count is greater than zero, it returns that category as the label for the text; otherwise, it returns "OTHER."

Next a loop executes the code for each newspaper from 2000 to 2024, dynamically generating the identification labels. For each generated label, the function stored the processed identifier of the news extract variable in a new column called *section*. This process effectively categorized the text data across multiple yearly data frames. Finally, the identified sections were extracted from all the processed data frames and compiled into a new dataset. This dataset included the time index (dates), serial code of the news from the source dataset of the news extract, the assigned sector (section), and the news extract itself. This organized dataset served as the foundation for the subsequent sentiment analysis.

3.2.1.1. Data Selection and Sampling Strategy

The full USFQ DataHub corpus contained over 1,5 million raw news articles. To make the analysis the dataset was pre-filtered by topic and date and then randomly sampled, following approaches commonly used in the literature[4]. Moreover, prior work indicates that relatively modest samples are sufficient

[3] For further details, see the subsection on the SpaCy large news model in the Materials and Methods section.

[4] For instance, Rathje et al. (2024) analyzed a random sample of 2.000 tweets for sentiment and Huang et al. (2023) constructed news classification datasets by randomly selecting 2.000 articles per category.

for robust results. Chaturvedi et al. (2023) show that sample sizes above 1.000 typically achieve stable classification metrics in NLP tasks, suggesting that a few thousand documents are enough to capture sentiment distributions accurately while keeping computation manageable.

This strategy of random sampling substantially reduced computational demands while still supporting valid inference. To implement it, Python's random module was applied within each stratum, giving every article equal probability of selection. This procedure mirrors standard practices in large-scale sentiment studies, which balance efficiency and representativeness through random selection (Xiong, 2023). On this basis, a sample of 2.000 articles per section (10.000 per year) was deemed large enough to yield reliable sentiment estimates, with an expected margin of error of about 2% for broad categorical proportions.

3.2.1.2. Digitized Newspapers Collection

Digitized archives of historical newspapers provide a consistent source of economic and social information that cannot be replicated by digital-born media alone. In Ecuador, leading newspapers such as *El Comercio* (Quito, founded in 1906), *El Universo* (Guayaquil, 1921), *Expreso*, and *La Hora* have published daily editions for decades, shaping public opinion and policy debates across generations. Their extensive numbers, amounting to thousands of articles per year long before the rise of the internet, make them indispensable for constructing a long-term record of national sentiment. Since fully developed online news outlets of Ecuadorian news media became fully matured around 2010, this research relied on the collaboration with USFQ-DataHub digitization and OCR of printed editions to ensure both continuity and consistency of the text data.

Digitization transforms these archives into a continuous text corpus that spans years when digital media were still experimental. This approach not only enables the construction of extended sentiment indicators but also lays the groundwork for integrating archival projects into broader research and policy applications. Ultimately, the strength of print archives lies in their ability to bridge historical gaps, offering a stable and comprehensive baseline for sentiment analysis. While digitization and OCR of newspapers are resource-intensive, the outcome is a uniquely long, coherent, and reliable sentiment series—something digital-born sources alone cannot provide in Ecuador.

For macroeconomic forecasting, such continuity is critical. Van Binsbergen et al. (2024) demonstrate that analyzing text from 200 million pages of 13.000 U. S. local newspapers allowed the construction of a 170-year sentiment index that vastly outperforms shorter series in predictive power. By contrast, restricting analysis to digital-native sources such as blogs, digital news portals, or social media for the case of Ecuador would only limit coverage to the last

decade, excluding the historical impact of significant episodes like Ecuador's dollarization in 2000. Archival newspapers therefore provide the temporal depth and historical richness necessary to capture sentiment across multiple economic volatile contexts, crises, and recoveries—elements essential for robust long-term forecasting.

3.2.2. Text Cleaning Protocol

Text cleaning and normalization are essential processes in natural language processing (NLP) to handle irregular or non-standard inputs. Words like *enc@m$ass*, which include unusual characters, present challenges for models and must be transformed into a more understandable and structured form. The cleaning process typically involves identifying and removing unwanted elements, ensuring uniformity, and approximating meaningful content. This transformation relies on regular expressions (*regex* function) and other techniques to clean and normalize text systematically. The four key steps involved in this process are described below.

The first step involved identifying and removing non-alphanumeric characters. Consider the case of the word *enc@m$ass* that contains symbols like @ and $, which are neither letters nor numbers and provide no semantic value in most contexts. These characters can be detected using a regex pattern such as $[a-zA-Z0-9]$, which matches anything that is not a letter or number. Replacing these characters with an empty string or a space effectively removes them, leaving only the core alphanumeric components. In the case of *enc@m$ass*, this step results in the cleaned string *encmass*. By eliminating extraneous symbols, this step ensures that the word is stripped of noise while retaining its semantic essence.

Once non-alphanumeric characters are removed, the second step involves ensuring case uniformity. Natural language often includes a mix of uppercase and lowercase letters, but for most NLP tasks, treating words in a case-sensitive manner is unnecessary and can introduce inconsistencies. Converting all letters to lowercase guarantees that words like *Enc@m$ass* and *enc@m$ass* are processed identically, yielding *encmass* after cleaning. This normalization step is straightforward yet crucial, as it helps standardize text for downstream processing, reducing the number of unique tokens in the data.

The third step addressed the presence of multiple consecutive non-alphanumeric characters, often found in informal text from social media or user-generated content. If such characters are detected, they are either consolidated into a single space or removed entirely to produce a more readable and structured result. Once cleaned, the algorithm approximates the word to a valid dictionary term or a more recognizable form. For instance, while the string *encmass* is free of noise, it may not correspond to a valid word. Using techniques like Levenshtein distance or phonetic algorithms, the cleaned word

is compared to a lexicon to identify the closest match, such as *encompass*. These methods assess string similarity or phonetic alignment, helping to recover the intended word and correcting typographical errors or unconventional spellings.

3.2.3. Denton Method for Temporal Disaggregation of Time Series

The Denton method is a widely used technique for temporal disaggregation, designed to convert low-frequency data into high-frequency data while preserving the original series' consistency. For this study, the goal was to disaggregate quarterly GDP growth rates, remittances, and exports into monthly series using the monthly Consumer Confidence Indicator (CCI), denoted as CONS_CONFIDENCE in the dataset, as the guiding variable. The process estimates smooth and accurate monthly series for these variables, consistent with observed quarterly growth rates. Essentially, the method minimizes distortions in the proportional relationship between the high-frequency indicator and the target series to ensure the smoothness of the series (Sax & Steiner, 2013). As a result, the disaggregated monthly series maintained consistency with quarterly data while incorporating meaningful high-frequency variations suggested by the CCI.

When integrated with ARDL framework, it underscores its suitability by harmonizing mixed-frequency data—a frequent challenge in economic forecasting for developing countries. Unlike dynamic factor models, which are often used for similar purposes, the Denton method maintains a smooth and consistent relationship between high and low frequency data without imposing rigid assumptions about the underlying processes (Denton, 1971). This adaptability is critical for Ecuador, where macroeconomic indicators like GDP and remittances are typically reported quarterly but sentiment information is available at much higher frequencies. In this context, the quarterly GDP annual growth rate, exports and remittances were disaggregated into a monthly series using the monthly indicator of Consumer Confidence Index as the high-frequency movement guide for the disaggregation process. By combining the Denton method's data harmonization capabilities with the robust regularization of LASSO, the proposed approach not only captures short and long term dynamics but also ensures stability in the presence of high-frequency text-derived sentiment data, addressing the unique forecasting challenges posed by Ecuador's volatile economic environment.

3.2.4. Forecasting Models

The goal of the model is to understand both short-term and long-term relationships between the annual percentual growth rate of GDP and multiple explanatory variables. The baseline specification considered is as follows:

$$y_t = \beta_t X_t + \eta_t S_{t-j} + \delta_t D_t + \theta_t Z'_{t-j} + \varepsilon_t \quad \textbf{(1)}$$

Where:

- y_t represents GDP annual growth rate at time *t*.
- X_t indicates a vector of different control variables and survey indicators.
- Z'_{t-j} is a vector of lagged control variables, survey indicators, and the GDP annual growth rate.
- S_{t-j} is a vector of sentiment indicators from digitalized newspapers from 2000-2024, with their respective lags. These sentiments aim to track the flow of information available to forecasters up to time *t*.
- D_t a set of dummy variables that account for structural changes, such as recessions, natural disasters, and pandemics.
- β_t, φ_t, γ_t, δ_t, θ_j are the parameters to be estimated.
- ε_t is the forecast error for the monthly release of GDP growth based on the information available at time *t*.

The model is dynamically constructed in the statistical software R to integrate all the key components described above. It is implemented in two formats: one that focuses solely on the effects of survey variables as a benchmark (ARDL) and another that incorporates both survey and sentiment variables (ARDLS). Using the dynlm package in R, the ARDL benchmark model seamlessly includes autoregressive terms, contemporaneous and lagged explanatory variables, real-time sentiment indicators, and structural breaks through dummy variables. This structure provides a robust framework for estimating short-term dynamics while capturing long-term relationships.

Due to the high variability inherent in sentiment information, it was necessary to implement a regularization process to estimate the parameters of the ARDLS model. To address this challenge, LASSO regularization was applied within the ARDLS framework, as it enhances model selection and estimation accuracy, particularly in the presence of numerous predictors. Compared to other regularization methods such as ridge regression, which shrinks coefficients without variable elimination, LASSO is especially useful because it simultaneously performs shrinkage and variable selection, yielding more parsimonious and interpretable models.

The coefficients of the ARDLS model were estimated by penalizing the absolute magnitude of regression coefficients through LASSO to effectively select the most significant regressors while shrinking irrelevant coefficients to zero. In return, this allowed the model to improve its robustness and generalizability by simplifying the model and reducing noise from the high-dimensional sentiment data through the regularization parameter (λ). This parameter was

optimized using cross-validation to ensure a balance between complexity and predictive accuracy. The annexes section contains a graphical comparison for the model fit before and after LASSO regularization for better visual reference.

3.2.5. Limitations

Despite numerous advantages, the integration of news-based sentiment analysis into GDP forecasting also has several limitations that warrant careful consideration. A primary concern lies in the inherent subjectivity of sentiment analysis, where news content is shaped by editorial policies, biases, and agendas, which can influence the tone and framing of articles. As a result, sentiment indicators derived from such data may not fully reflect a fully objective state of public opinion or economic conditions. For example, during politically charged events, media outlets with differing political leanings might portray economic developments in contrasting lights, introducing noise or bias into the sentiment analysis process.

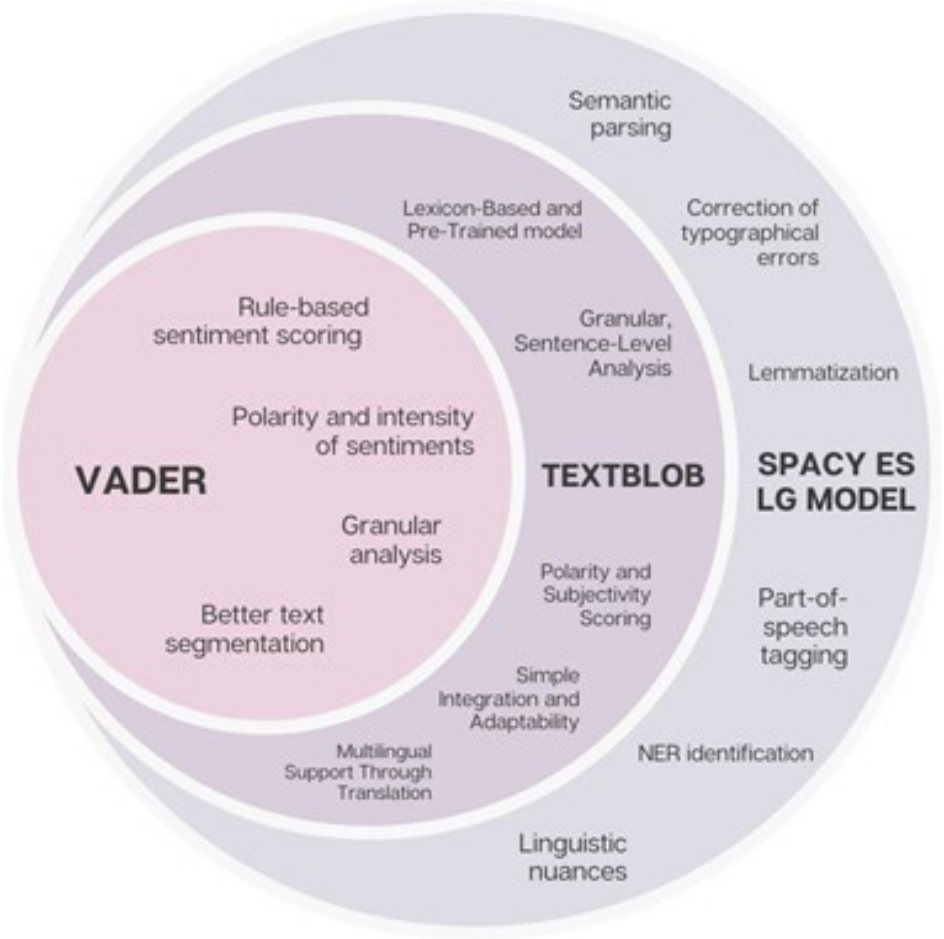
Additionally, NLP algorithms, while powerful, are not immune to inner biases. Pre-trained models such as VADER or TextBlob rely on lexicons and rules that may not adequately account for highly specific regional linguistic nuances or evolving language patterns such as idiomatic expressions and colloquialisms. As a result, they may misinterpret the message of the analyzed texts, potentially skewing sentiment scores by unknown degrees depending on the intensity of the sentiment. Even advanced models, while more context-aware, require significant fine-tuning to accurately process the subtleties of local language and discourse. These limitations highlight the need for continuous refinement of NLP tools to ensure that they capture sentiments accurately and meaningfully in context-specific scenarios.

In addition, the overreliance on high-frequency textual data, which while providing timely insights, can also amplify short-term noise or anomalies, is another factor to be considered. Events that generate intense but fleeting public reactions, like controversial political statements or short-lived market fluctuations, might disproportionately influence sentiment indicators. This can lead to an overestimation of their impact on GDP forecasts, especially in highly dynamic periods. Without robust mechanisms to filter or contextualize such outliers, there is a risk of producing forecasts that are overly reactive to changes in transient sentiment.

Finally, the proposed LASSO-ARDL framework, while effective for short-term forecasting, demonstrated diminishing returns for longer horizons. This limitation comes from the fact that the influence of sentiment-based indicators weakens over time as structural economic trends and traditional macroeconomic variables become more dominant. Although this does not diminish the utility of the approach for short-term applications, it underscores the need for complementary methods to enhance long-term predictive power.

Addressing these limitations thus requires further methodological refinements, such as incorporating additional data sources, fine-tuning NLP models for local contexts, and developing hybrid forecasting frameworks that balance short-term responsiveness with long-term stability.

Figure 5. Integration of Multiple Tools and Models for Sentiment Analysis: VADER, TextBlob and SpaCy ES LG Model for Comprehensive Textual Data Analysis



Note: Each layer builds on the previous one, progressing from basic sentiment scoring to more context-aware analysis and aspect-based aggregation. This multi-layered approach aims to produce the most accurate and cost effective, and nuanced sentiment extraction, as well as effective trend summarization across domains.

4. RESULTS AND LIMITATIONS

4.1. In-Sample Analysis

This section discusses a comparative analysis between the benchmark model ARDL model and the LASSO-regularized ARDL (LASSO-ARDLS) model with sentiment indicators for an in-sample scenario. Results reveal significant differences in their forecasting performance and stability over time, based on several statistical metrics including the Iterated root mean square error (RMSE), Diebold-Mariano (DM) test results, and stability tests such as the cumulative sum (CUSUM) and cumulative sum of squares (CUSUMSQ).

In essence, results displayed a clear advantage of the LASSO-ARDL model in short-term GDP forecasting, particularly in terms of root mean square error (RMSE), a critical metric for assessing prediction accuracy. For instance, the model's RMSE for 1-month forecasts is considerably lower than that of the benchmark ARDL model, indicating that it can provide more precise predictions of Ecuador's GDP fluctuations within this short horizon. To put this into perspective, if policymakers rely on GDP forecasts to plan fiscal adjustments, a lower RMSE translates to fewer deviations between expected and actual outcomes reducing the likelihood of overestimating or underestimating economic performance. This is particularly important for Ecuador, where accurate forecasts are essential for managing debt obligations and social spending amidst volatile commodity markets.

Figure 4 showcases iterated RMSE values for both models showing noticeable spikes that likely correspond to periods of economic shocks or structural changes in the underlying dataset. In terms of general performance, the ARDL model exhibits significant volatility in RMSE values, with higher pronounced spikes indicating periods where the model struggles to capture sudden changes in the data. The average RMSE for the ARDL model is 1,3831, suggesting a relatively higher prediction error overall. In contrast, the LASSO-ARDLS model demonstrates a more stable RMSE trajectory, with fewer extreme variations. Its average RMSE is 0,9122, notably lower than that of the ARDL model. This lower RMSE indicates that the LASSO-ARDLS model provides more accurate forecasts on average. This stability can be attributed to the LASSO regularization technique which seems to effectively mitigate overfitting by penalizing the absolute size of the regression coefficients, thus simplifying the model and enhancing its ability to generalize to new data.

Table 1. Diebold-Mariano Test Results for LASSO-ARDLS Model at Different Forecasting Horizons (In-Sample)

Forecasting Horizon	DM Coefficient	p-value
1 month	1,3874	0,0027
3 months	1,0928	0,0250
6 months	0,8215	0,0699
12 months	0,6423	0,1078
24 months	0,5246	0,1688

Note: DM statistic is set as errors ARDL and errors LASSO.

The practical significance of these improvements becomes evident when considering real-world economic scenarios. For example, during the COVID-19 pandemic, rapid changes in economic activity required timely and accurate forecasts to allocate emergency funds effectively. A sentiment-enhanced LASSO-ARDL model, with its lower RMSE and ability to incorporate real-time public perceptions from news, could have provided early indications of economic downturns or recovery trends. This timeliness showcases a useful tool for policymakers to preemptively address challenges such as unemployment spikes or disruptions in public services, which traditional models might detect only after substantial delays.

Additionally, the LASSO-ARDL model's lower RMSE indicates its robustness in capturing the nuanced effects of high-frequency data like sentiment indicators. During the 2019 fuel subsidy removal in Ecuador, widespread public dissent was reflected in media sentiment, which traditional economic variables might not immediately capture. A model with a lower RMSE ensures that such immediate public reactions are effectively translated into actionable insights. This responsiveness is particularly valuable in Ecuador's context, where political and economic instability often amplify the importance of short-term decision-making.

Table 1 presents the Diebold-Mariano (DM) test results for forecasting horizons of 1, 3, 6, 12, and 24 months, comparing the models predictive accuracy in these horizons. The DM test yields statistically significant p-values of 0,0027 and 0,0250 for the 1-month and 3-month horizons, respectively. These low p-values indicate that the differences in forecasting accuracy between the two models are statistically significant at 5 % level. Nonetheless, as the forecasting horizon extends the statistical significance diminishes as shown by the 6-month, 12-month, and 24-month horizons p-values, where ultimately there is no statistically significant difference between the forecasting ability of the models. This trend suggests that while the LASSO-ARDLS model has a clear advantage in the short term, its superiority diminishes over longer forecasting

horizons as both models' display a similar ability to capture long-term trends, making their performances converge over time.

Figure 6. Comparison of Iterated RMSE Over Time for ARDL and LASSO-ARDLS Models

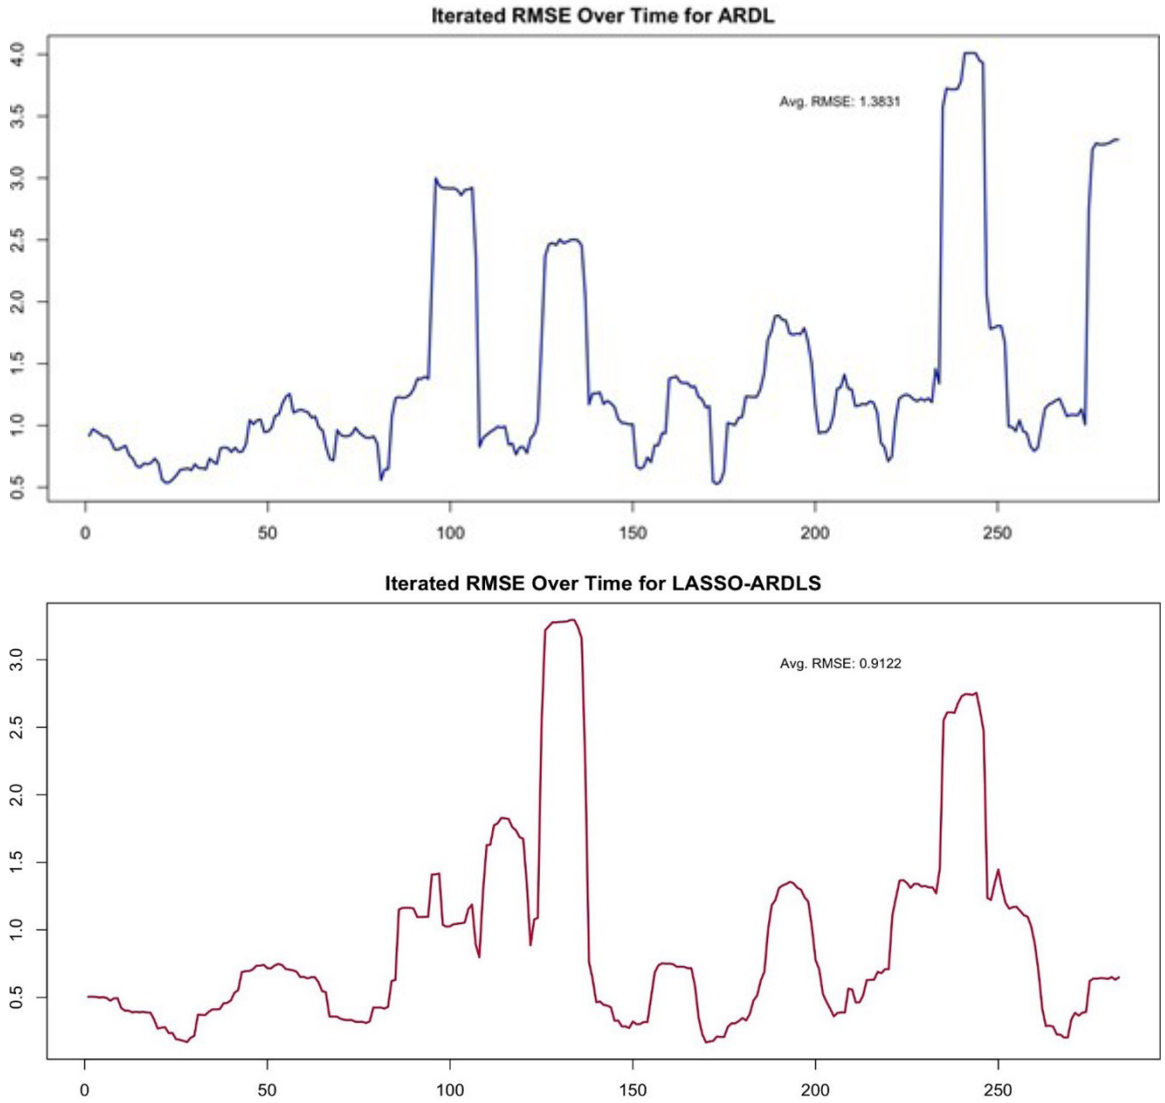
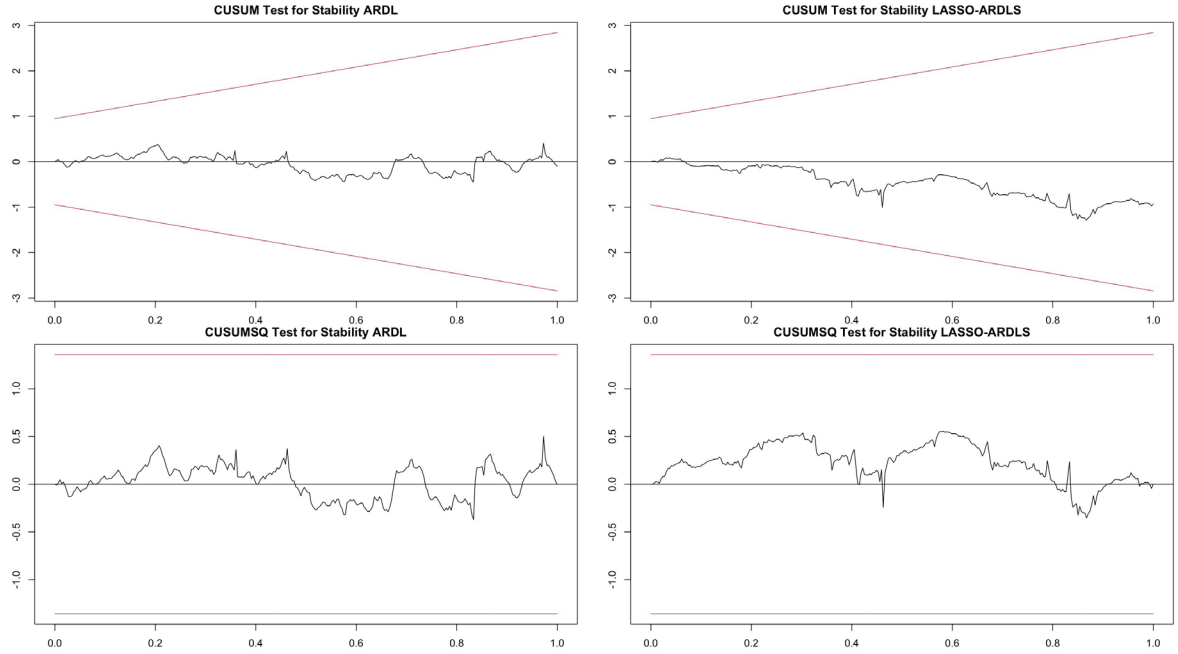


Table 2. Average RMSE for ARDL and LASSO-ARDLS models

Model	Average RMSE
In-Sample ARDL	1,3831
In-Sample LASSO-ARDLS	0,9122
Out-of-Sample ARDL	11,4027
Out-of-Sample LASSO-ARDLS	0,0368

In terms of stability and robustness, the CUSUM and CUSUMSQ tests in Figure 7 shows that the cumulative sum of residuals remains within the critical bounds, indicating structural stability in each model specifications. However, a closer inspection reveals that the LASSO-ARDLS model displays smaller and more stable deviations compared to the ARDL model, highlighting its robustness to shocks or parameter shifts. The CUSUMSQ test results reinforce this observation, as it depicts greater variability in the ARDL model, particularly during periods of volatility. In contrast, the LASSO-ARDLS model appears less sensitive to variance changes, making it a more robust choice under dynamic conditions.

Figure 7. Comparison of CUSUM and CUSUMSQ Tests for Stability of ARDL and LASSO-ARDLS Models (In-Sample)



Altogether, the statistical evidence shows that the LASSO-ARDLS model is better suited for short-term forecasting, where it demonstrates superior accuracy and stability based on its lower RMSE, significant DM test results at shorter horizons and robust performance under the CUSUM and CUSUMSQ tests. Moreover, the LASSO-ARDLS model's robustness to structural shifts suggests it may still offer an edge under dynamic scenarios. These findings not only highlight the statistical robustness of the sentiment integrated model but also underscore its practical value for addressing real-world challenges in economic forecasting. By providing more accurate short-term forecasts, this approach presents a scalable framework for policymakers and businesses stakeholders to make informed decisions. Thus, the integration of sentiment analysis with traditional econometric methods offers a forward-looking solution to

the limitations of conventional models, especially in rapidly changing and data-limited environments like Ecuador. Nonetheless, for longer horizons the differences between the two models become less pronounced suggesting that the information of sentiment indicators diminishes over time.

4.2. Out-Sample Analysis

Table 3. Diebold-Mariano Test for LASSO-ARDLS Model at Different Forecasting Horizons (Out-Sample)

Forecasting Horizon	DM Coefficient	p-value
1 month	3,0868	0,0029
3 months	2,2811	0,027
6 months	1,835	0,0599
12 months	1,6247	0,1078
24 months	1,3874	0,1788

Note: DM statistic is set as errors ARDL and errors LASSO.

This section evaluates the results of the performance analysis for the LASSO-ARDLS model in comparison with the benchmark model ARDL. The analysis was conducted using various robustness checks such as bootstrap RMSE, average rolling RMSE, cross-validation with varying lambda values, and the Diebold-Mariano (DM) test at different forecasting horizons on the validation portion of the dataset.

Firstly, the model achieved a bootstrap RMSE mean of 1,512247 with a standard deviation of 0. This zero standard deviation across bootstrap samples indicates exceptional stability, suggesting that the model performs consistently on different subsets of test data. Such consistency is a strong indicator of the model's robustness, as it implies that its predictive accuracy does not vary when trained and tested on different random samples. When using a 60-month rolling window, the model's average rolling RMSE reported a value of 0,874, significantly lower than the bootstrap RMSE. This reduction suggests that the model adapts well to temporal shifts, maintaining high accuracy across different time segments. The ability to adjust to changes over time is crucial in time series forecasting, as it demonstrates the model's capacity to handle evolving patterns in the data.

A cross-validation analysis with rolling forecasting origin resampling was performed to evaluate the effect of varying lambda values, ranging from 0,001 to 0,096, on the model's performance. It was observed that as lambda increased, the RMSE and mean absolute error (MAE) decreased, while the

R-squared value improved. The optimal lambda was identified at 0,096, where the model achieved the lowest RMSE of 1,751 and an R-squared of 0,248. This means that approximately 24,8 % of the variance in the target variable is explained by the model at this lambda value. The improvement in performance metrics with increasing lambda highlights the effectiveness of LASSO regularization in enhancing the model's predictive power by penalizing less informative predictors and preventing overfitting.

Coefficients in Table 3 show the results from the Diebold-Mariano test further supporting the model's performance. At the 1-month forecasting horizon, the DM coefficient was 3,0868 with a p-value of 0,0029, indicating a statistically significant improvement over the ARDL model at the 1 % level. For the 3-month horizon, the DM coefficient was 2,2811 with a p-value of 0,0270, showing significance at the 5 % level. At the 6-month horizon, results reflected marginal significance at the 10 % level. However, for the 12-month and 24-month horizons, the improvements were not statistically significant.

These results highlight that the LASSO-ARDLS model significantly outperforms the ARDL benchmark for short-term forecasting horizons of 1 to 3 months. As the horizon extends beyond 6 months, the performance advantage diminishes and becomes statistically insignificant. This suggests the model is particularly well-suited for short-term forecasting, while its advantage in long-term forecasts may be limited. Overall, the LASSO-ARDLS model demonstrates strong out-of-sample accuracy and robustness, as shown by the consistent bootstrap RMSE and low rolling RMSE indicating stable performance across diverse data segments, supporting its reliability for forecasting. The low RMSE values across cross-validation and rolling windows confirm the model's capacity to generalize well to new data, which is essential for making accurate predictions on unseen datasets. Furthermore, the observed improvements in R-squared and reductions in RMSE with increasing lambda underscore the benefits of regularization in model tuning. By effectively balancing predictive accuracy and generalizability, the integration of sentiment indicators model is well-suited for stable and accurate forecasting in time series applications.

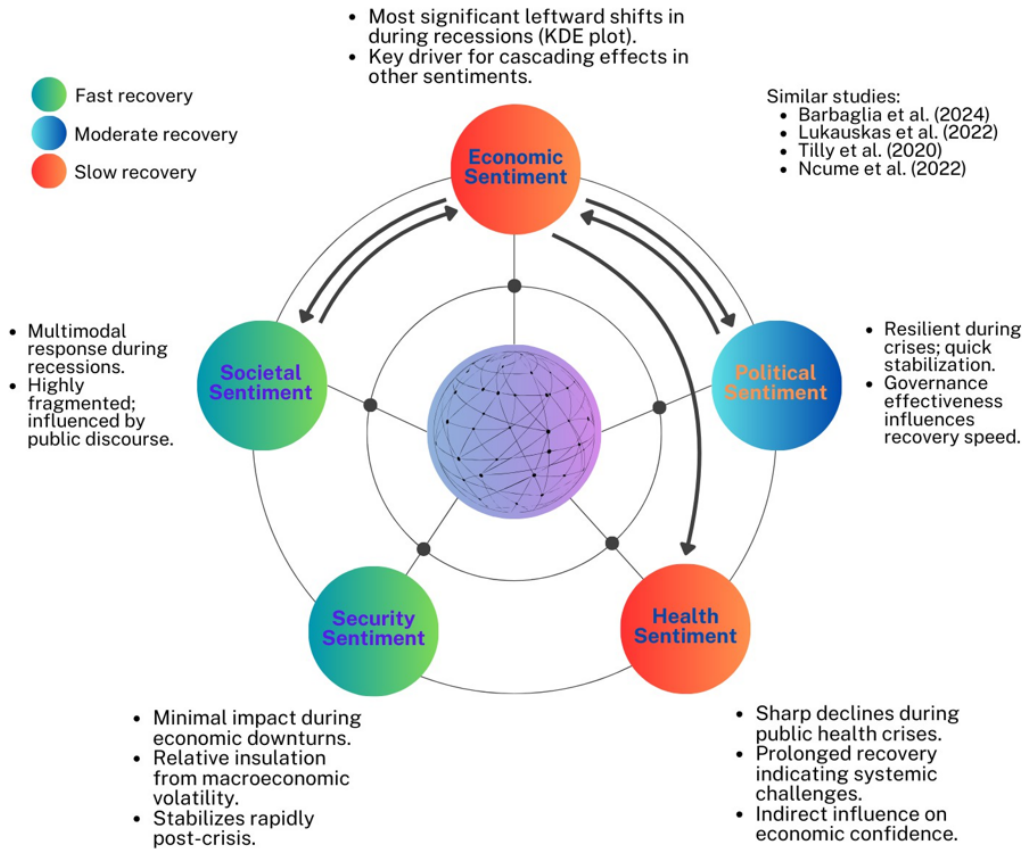
The results of this study highlight the critical role of sentiment analysis in understanding how economic cycles influence public perceptions across various domains. The normalized sentiment scores show a strong dependence on the business cycle, as evidenced by sharp declines during recessions and varied recoveries across sectors. For example, economic sentiment exhibited the most significant leftward shift in recessionary periods in the KDE plots, indicating intensified negative perceptions and heightened public uncertainty during downturns. This finding aligns with Barbaglia et al. (2024), who demonstrated that economic sentiment, measured through news-based indicators, is pro-cyclical and highly sensitive to macroeconomic conditions.

The results also underscore the role of Politics sentiment during periods of instability as it showed more resilience compared to Economics sentiment,

stabilizing more quickly during crises. Similarly, Barbaglia et al. (2024) found a bidirectional interaction between political and economic sentiment, with news coverage on monetary policy and governance decisions significantly predicting GDP growth in countries like Spain and Italy. This relationship highlights the central role of politics in shaping economic expectations. Lukauskas et al. (2022) observed a similar pattern during the outbreak of war and the COVID-19 pandemic in Lithuania in 2019, showing that negative political sentiment during crises strongly influences consumer satisfaction and economic forecasts.

The health sector experienced a significant decline in sentiment during economic downturns, with slow recovery patterns indicating systemic challenges. The COVID-19 pandemic exacerbated this effect, highlighting public anxiety about the adequacy of healthcare systems during crises. This mirrors the results of Barbaglia et al. (2024), who noted that while health-related sentiment was not directly incorporated into their GDP forecasting models, it indirectly influenced overall public confidence and macroeconomic sentiment. Furthermore, the findings of Lukauskas et al. (2022) reinforce this perspective by showing that alternative data sources, such as healthcare-related sentiment, can improve forecasts of economic indicators like unemployment and consumer satisfaction during periods of public health stress.

Figure 8. Cascading Effects of Sentiment Dynamics Across the Five Sentiment Indicators



Note: Arrows indicate directional influences between sections. Recovery speeds are color-coded.

Sectoral resilience varied, with security and societal sentiments stabilizing more quickly than economic and health sentiments. Results displayed minimal leftward shifts in security sentiment during crises, reflecting its relative insulation from broader economic volatility. However, Barbaglia et al. (2024) highlight that financial sector sentiment, a component of security sentiment in their framework, showed marked declines during double-dip recessions but significantly contributed to short-term GDP predictions in countries like Germany and the United Kingdom. This suggests that while security sentiment may appear stable overall, specific components like financial security can reveal acute vulnerabilities during crises.

One key distinction of this research is the explicit identification of a feedback loop between societal sentiment and macroeconomic indicators. While Barbaglia et al. (2024) focused on sentiment's predictive power for GDP, particularly through financial and monetary policy indicators, the results of this research suggest that societal sentiment is equally important for economic forecasting especially in understanding public confidence and cohesion. During

recessions, declining societal sentiment magnifies the impact of poor business expectations and consumer confidence, contributing to the broader economic downturn. This contrasts with Barbaglia's emphasis on the direct influence of financial sentiment on GDP forecasts as my results indicate that for the case of Ecuador societal sentiment operates on a more diffuse but equally critical level influencing and being influenced by broader economic sentiment.

In light of these results, it is safe to affirm that news media and public discourse play a significant role during economic cycles. As societal sentiment shifts during recessions, public discourse often focuses on unemployment, inequality, and instability, thereby magnifying the negative feedback loop between societal sentiment and economic expectations. Conversely, in expansionary periods, public discourse pivots to themes of progress and opportunity, reinforcing positive consumer and business confidence. This interaction between societal narratives and economic sentiment, observed in both this study and that of Lukauskas et al. (2022), underscores the importance of real-time monitoring of societal sentiment to capture nuances in public confidence during economic transitions.

Although sentiment-based forecasting has attracted growing interest globally, it remains quite underexplored in Ecuador. In Latin America, however, several notable studies have demonstrated the value of media-derived sentiment as mentioned earlier in the literature review section. By contrast, prior research in Ecuador has not yet applied news-driven sentiment indices to systematic macroeconomic forecasting, making this work a pioneering study for its context; as such, this work is original on several fronts.

It leverages Ecuador's newly digitized press archives accessible through USFQ DataHub for research access, analyzing hundreds of thousands of newspaper articles to produce high-frequency sentiment indicators. As a result, this research extends the Latin American literature on textual forecasting by filling the Ecuadorian gap and by introducing a ready-to-implement tool and replicable framework for upcoming research in the field. Results show that sentiment information effectively complements traditional business surveys with real-time news sentiment, aligning with regional efforts to exploit soft information in policy analysis. This outcome is both a new empirical finding and a practical forecasting system for decision-makers in Ecuador: policymakers and private-sector analysts may use the sentiment-augmented model as a scalable, real-time dashboard for macroeconomic perspectives.

For business owners, these findings highlight opportunities to manage sentiments during volatile stages of the economic environment. As such firms and investors may benefit from the superior short-term accuracy of the sentiment enhanced model by how news-based signals capture shifts in public confidence and consumption intent that traditional economic data do not immediately reflect. In fact, Ellingsen et al. (2022) mention how media text data contain information not captured by the hard economic indicators which are especially useful for forecasting domestic demand. Delice and Pinto

(2025) similarly demonstrate that Spanish language news sentiment is a reliable source for economic analysis, offering considerable insights for market analysis and making informed investment decisions. In practice, businesses can use the information to anticipate market turns and adjust their strategies in terms of production, staff recruitment, inventory, or asset allocation ahead of official announcements. Moreover, the inclusion of regularization techniques like LASSO ensures that only the most relevant predictors are used, which sharpens the signal for decision- makers. These findings reinforce prior evidence that news sentiment improves short-term output forecasts as shown by Rambaccussing and Kwiatkowski for example, suggesting that Ecuador's private sector could harness sentiment-based GDP forecasts as an early-warning tool to manage demand risk and investment timing.

For the Central Bank of Ecuador (CBE), the enhanced forecasting framework offers a means to augment traditional models with timely soft information. Monetary authorities rely on GDP forecasts to set interest rates, guide reserve policy, and communicate outlooks to the public. By incorporating high-frequency sentiments it can sharpen these forecasts, particularly around turning points. This aligns with findings by Ashwin et al. (2021), who report that newspaper sentiment has materially improved nowcasts of real GDP growth when higher-frequency indicators are scarce. During the first half of a quarter when conventional data like official GDP reports or official surveys are not yet available, news-based sentiment can give the CBE an earlier read on growth momentum. Central banks in other emerging economies have seen similar gains, as mentioned by Adebiyi et al. (2022), who found that adding a computed sentiment index to Nigeria's forecasting model consistently lowers forecast errors for inflation, enhancing decision-making for the monetary policy committee. By analogy, the CBE could use sentiment signals to improve its short-term growth outlook, allowing quicker policy studies and suggestions to shocks such as commodity price swings or political events. Furthermore, the results from the stability tests show that the sentiment-enhanced forecasts remain robust under structural shifts, implying that the CBE could rely confidently on these forecasts even during volatile periods.

Lastly but not least, the information provided by accurate short-run GDP forecasts is crucial for fiscal planning that the Ministry of Finance can leverage to better manage budgets and fiscal policy. A more precise 1-3 month outlook enables the government to adjust revenue projections and spending plans when economic conditions change unexpectedly due to external shocks or policy shifts. For instance, if sentiment indicators signal an imminent slowdown, the fiscal authorities could preemptively tighten spending or accelerate budget reallocations. Delice and Pinto (2025) argued in their research that sentiment analysis provides a bridge between public discourse and policy-making, offering insights for crafting economic policies in real time. Thus, the combination of sentiment-driven GDP forecasts with fiscal models may prove beneficial for the Ministry of Finance to anticipate revenue shortfalls or surpluses more

promptly than with hard and delayed economic indicators alone. In addition, the regularization technique allows the usage of the most informative news indicators so policymakers may have the confidence that the signals reflect genuine economic trends rather than noise. In consequence, Ecuador's fiscal authorities could use sentiment-enhanced forecasts to fine-tune fiscal policy to manage debt or social spending more efficiently and in touch with social needs, especially when facing scenarios of considerable recession and uncertainty. Ultimately, the results of this research underscore that sentiment-augmented forecasting can support more agile and informed fiscal decision-making in Ecuador's dynamic economic environment.

5. CONCLUSIONS AND DISCUSSION

The integration of sentiment analysis into GDP forecasting models demonstrated significant potential for improving short-term economic predictions, especially in volatile changing environments like Ecuador. Leveraging news-based sentiment indicators alongside traditional survey data highlighted the usefulness of capturing high-frequency economic signals to provide actionable insights as shown by LASSO-ARDLS model superior performance for short-term horizons. This approach not only complements traditional macroeconomic indicators but also addresses their limitations, offering a scalable and adaptive framework for economic forecasting and underscoring the value of incorporating real-time data.

Moreover, the sectoral analysis of sentiment indicators revealed important insights into the interconnectedness of public perceptions across different domains of the economy. Recessions were characterized by sharp declines in sentiment across all sectors, with recovery patterns varying significantly and reflecting the complex interplay between structural challenges and public confidence. While traditional economic indicators often capture the aftermath of structural shocks, sentiment-based indicators provide immediate signals, thereby enhancing the responsiveness of forecasting models. For policymakers and private stakeholders, these results emphasize the importance of timely interventions and strategic communication to manage public perceptions and adjust commercial or operational strategies to strengthen resilience. Overall, this study highlights the potential of sentiment indicators to improve economic planning, resource allocation, and policy design by offering a practical tool to quantify qualitative dimensions of the economy. By bridging the gap between conventional forecasting instruments and alternative data sources, the research contributes a robust framework for navigating uncertainty and supporting evidence-based decision-making in Ecuador's fragile economic context. These findings also pave the way for more dynamic and adaptive forecasting tools tailored to the country's volatile conditions.

Looking forward, future studies may explore the application of advanced machine learning techniques, such as BERT (Bidirectional Encoder Representations from Transformers) models to enhance the extraction of sentiment indicators from textual data. BERT's ability to understand context and nuances in language at higher level would likely improve the accuracy of sentiment scores, particularly for a context with such complex narratives like Ecuador. Moreover, developing NLP models tailored specifically for Ecuadorian Spanish would significantly enhance the analysis of local news and public discourse as Ecuadorian Spanish includes unique regional idioms, expressions, and slang, which pre-trained models for standard Spanish may not adequately capture. Customizing models to account for these linguistic nuances would greatly ensure more precise sentiment analysis and deeper insights into local economic conditions.

In addition, further research endeavors could expand the scope of sentiment analysis to explore its applications in other domains, such as sectoral forecasting or social policy evaluation. For instance, sentiment indicators could be applied to predict sector-specific outcomes such as trends in tourism, agriculture, or remittances, which are critical to Ecuador's economy. Exploring the interplay between sentiment indicators and long-term economic trends in such areas could provide a more comprehensive understanding of their impact. These efforts would not only refine the methodological framework presented here but also solidify sentiment analysis as a cornerstone of economic forecasting in the Ecuadorian context.

Acknowledgments

The author extends his sincere gratitude to USFQ DataHub for providing the newspaper articles and source materials essential to this research. The author is also deeply grateful to Rodrigo López, PhD, and Pablo Astudillo-Estevez, PhD, for their invaluable guidance and constructive feedback, which greatly contributed to the quality of this work.

AI Disclaimer

The author used an AI-based text processing tool solely for language refinement and clarity of concatenation of ideas between certain paragraphs. The tool was not employed in any circumstance to generate content, conduct analyses, or interpret results. All research design, data analysis, and conclusions are entirely the work of the author.

BIBLIOGRAPHY

Adebiyi, M. A., Adenuga, A. O., Olusegun, T. S., & Mbutor, O. O. (2022). Big data and inflation forecasting in Nigeria: A text mining application. *CBN Economic and Financial Review*, 60(1):1–23.

Ali, R. H., Pinto, G., Lawrie, E., & Linstead, E. (2022). A large-scale sentiment analysis of tweets pertaining to the 2020 us presidential election. *Journal of Big Data*, 9(1):79.

Ashwin, J., Kalamara, E., & Saiz, L. (2021). Nowcasting Euro area GDP with news sentiment: A tale of two crises. *Journal of Applied Econometrics*, 39(2):447–465.

Barbaglia, L., Consoli, S., & Manzan, S. (2024). Forecasting GDP in Europe with textual data. *Journal of Applied Econometrics*, 39(2):338–355.

Buckman, S., Shapiro, A., Sudhof, M., & Wilson, D. (2020). News sentiment in the time of covid-19. *FRBSF Economic Letter*, 8-5.

Center, P. R. (2019). Many across the globe are dissatisfied with how democracy is working. Pew Research Center. Accessed: 2025-09-27.

Chang, C.-Y., Zhang, Y., Teng, Z., Bozanic, Z., & Ke, B. (2016). Measuring the information content of financial news. In *Proceedings of COLING 2016 - 26th International Conference on Computational Linguistics*, 3216–3225.

Chaturvedi, J., Shamsutdinova, D., Zimmer, F., Velupillai, S., Stahl, D., Stewart, R., & Roberts, A. (2023). Sample size in natural language processing within healthcare research.

Consoli, S., Barbaglia, L., & Manzan, S. (2022). Fine-grained, aspect-based sentiment analysis on economic and financial lexicon. *Knowledge-Based Systems*.

Delice, P. A. & Pinto, D. (2025). Decoding economic insights: The analytical power of news content. *Journal of Scientometric Research*, 14(1):365–372.

Denton, F. T. (1971). Adjustment of monthly or quarterly series to annual totals: An approach based on quadratic minimization. *Journal of the American Statistical Association*, 66(333):99–102.

Einav, L. & Levin, J. (2014). Economics in the age of big data. *Science*, 346(6210):715–721.

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

Fathi, K. (2022). *Multi-Resilience - Development - Sustainability Requirements for securing the future of societies in the 21st century*. Springer.

Forum, W. E. (2023). Global risks report 2023. Accessed: 2024-11-18.

Honnibal, M. & Montani, I. (2020). *SpaCy 101: Everything you need to know*. Explosion AI.

Hutto, C. & Gilbert, E. (2014). Vader: A parsimonious rule-based model for sentiment analysis of social media text. *AAAI Press*, 8(1):216–225.

Jiang, Y., Wang, X., Xiong, Z., Yang, H., & Tian, T. (2022). Interpreting and predicting economic flows: A time-varying parameter global vector autoregressive integrated machine learning model.

Loria, S. (2018). Textblob: Simplified text processing. Accessed: 2024-11-22.

Lu, W., Wang, Y., & Zhang, X. (2023). Which news topics drive economic prosperity in china? *PLOS ONE*, 18(10).

Lukauskas, M., Pilinkiene˙, V., Bruneckiene˙, J., Stundžiene˙, A., Grybauskas, A., & Ruzgas, T. (2022). Economic activity forecasting based on the sentiment analysis of news. *Mathematics*, 10(3461).

Malandri, L., Xing, F., Orsenigo, C., Vercellis, C., & Cambria, E. (2018). Public mood-driven asset allocation: The importance of financial sentiment in portfolio management. *Cognitive Computation*, 10(6):1167–1176.

Medeiros, M. C. & Mendes, E. F. (2017). Adaptive lasso estimation for ARDL models with GARCH innovations. *Econometric Reviews*, 36(6-9):622–637.

Méda, D. (2024). Le puissant sentiment d'injustice ressenti par une partie des français explique la puissance de la réaction dans les urnes. *Le Monde*.

Pennington, J., Socher, R., & Manning, C. D. (2014). Glove: Global vectors for word representation. *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 1532–1543.

Pitta de Jesus, D. & da Nóbrega Besarria, C. (2025). Central bank narratives and macroeconomic forecasting: Using textual analysis from machine learning. *IMG Economic Review*.

Rambaccusing, D. & Kwiatkowski, A. (2020). Forecasting with news sentiment: Evidence with UK newspapers. *International Journal of Forecasting*, 36(4):1501–1516.

Rathje, S., Mirea, D.-M., Sucholutsky, I., Marjieh, R., Robertson E.C., Van Bavel, J. J. (2024). GPT is an effective tool for multilingual psychological text analysis. *Proceedings of the National Academy of Sciences of the United States of America*, 121(34), e2308950121. https://doi.org/10.1073/pnas.2308950121

Rutkowska, A. & Szyszko, M. (2024). Dictionary-based sentiment analysis of monetary policy communication: On the applicability of lexicons. *Quality & Quantity*.

Sax, C. & Steiner, P. (2013). Temporal disaggregation of time series. *The R Journal*, *5*(2):1–17.

Soni, J. & Mathur, K. (2023). *Sentiment analysis of news headlines for stock market prediction using VADER*. IEEE, Bengaluru, India.

Starke, P., Kaasch, A. & van Hootegem, F. (2024). Political parties and social policy responses to global economic crises: Constrained partisanship in mature welfare states. *Journal of Social Policy*.

Unipython (2019). Análisis de sentimientos con TextBlob y Vader en Python. Accessed: 2024-11-22.

Van Binsbergen, J. H., Bryzgalova, S., Mukhopadhyay, M., & Sharma, V. (2024). (Almost) 200 years of news-based economic sentiment (Working Paper 32026). National Bureau of Economic Research.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. In *Advances in neural information processing systems*, 5998–6008.

Xiong, X. (2023). Evaluating random sampling bias in sentiment analysis of social media data. *Theoretical and Natural Science*, 25(1):36–42.

APPENDIX

Appendix A. VAR-X Model with Dummy Variables as Exogenous Factors

In order to determine a broader and more precise relationship between the sentiment variables and the annual percentual growth rate of GDP at a monthly frequency, the following section describes a VAR model with exogenous variables (VAR-X). In this case, such role is fulfilled by the dummy variables in D_t. These variables capture specific external conditions regarding key recessionary events in the last 10 years of Ecuadorian history that affect the dependent variable and other endogenous variables in Z_t, but are not influenced by the dynamics of Z_t. The matrix B quantifies the impact of these exogenous factors on the system.

The VAR-X model estimates the parameters A_i and B by maximizing the likelihood of observing Z_t given the lag structure determined by the AIC coefficient and the exogenous effects of the dummy variables.

Thus, the dataset is defined as:

$$Z = [Y, X_{\text{controls}}, X_{\text{survey}}, X_{\text{sentiment}}]$$

where:

- Y is the dependent variable,
- X_{controls} represents the matrix of control variables,
- X_{survey} represents the matrix of survey variables,
- $X_{\text{sentiment}}$ represents the matrix of sentiment variables.

The combined dataset Z is used as the endogenous variable set for the VAR model.

To determine the optimal lag order *p*, the Akaike Information Criterion (AIC) was applied across models with lag orders from 1 to a specified maximum *p*max. The lag *p* is chosen to minimize the AIC, calculated as follows:

$$p = \arg\min_{k \leq p_{\max}} \text{AIC}(k)$$

where

$$\text{AIC}(k) = -2\ln(L) + 2k$$

and *L* is the likelihood of the model at lag *k*.

After selecting the optimal lag order *p*, the final VAR-X model in matrix was specified as:

$$Z_t = A(L)Z_t + BD_t + u_t$$

where:

$$A(L) = I - \sum_{i=1}^{p} A_i L^i$$

and:

- Z_t is the vector of endogenous variables at time *t*;
- $A(L)$ is the lag polynomial of coefficient matrices, with A_i representing the coefficients for each lag *i*;
- D_t is the vector of exogenous variables, specifically the dummy variables in this model;
- B is the matrix of coefficients for the exogenous variables D_t;
- u_t is a vector of residuals at time *t*, assumed to follow a white noise process.

This representation captured the dynamic relationships among endogenous variables Z_t while accounting for the influence of exogenous factors D_t through the matrix B.

Appendix B. Granger-Causality Patterns Analysis

The analysis reveals that survey indicators generally have substantial explanatory power over sentiment variables, as evidenced by the widespread statistical significance (p-value < 0,05) of Granger-causality results for survey-to-sentiment relationships. Specifically, consumer confidence (CONS_CONFIDENCE) exhibits significant Granger-causality for all sentiment variables—including ECONOMY, POLITICS, HEALTH, SECURITY, and SOCIETY—with extremely low p-values (e.g., < 0,00002). This suggests that changes in consumer confidence strongly influence public sentiment across various domains, highlighting the importance of consumer perceptions in shaping societal narratives.

Similarly, manufacturing expectations (MANUF_EXP) demonstrate consistent Granger-causality across all sentiment variables, depicting the influence of sector-specific economic expectations on broader public sentiment. The significant p-values associated with MANUF_EXP indicate that manufacturing sector outlooks are reliable predictors of changes in public sentiment. Services expectations (SERVICES_EXP) also show significant causality for most sentiment variables, indicating the relevance of service-sector perceptions in driving sentiments related to the economy, politics, and society.

In addition to Granger-causality, the instantaneous causality tests (chi-squared tests) highlight strong contemporaneous relationships between surveys and sentiments. For instance, the chi-squared statistic for CONS_CONFIDENCE and various sentiments is highly significant (p-value < 0,00001), suggesting mutual influence within the same time frame. Also, it indicates that surveys capture real-time public perceptions, which closely align with sentiment dynamics.

The strong explanatory power of survey indicators reflects their structured nature. Surveys, by design, aggregate information on consumer, business, and sector-specific expectations, making them reliable predictors of public sentiments. They effectively act as leading indicators, capturing economic and social realities that shape broader sentiment trends. In contrast, the analysis shows fewer significant relationships when sentiments are tested as causes of survey indicators. While some sentiment variables exhibit explanatory power for specific surveys, the overall impact is less pronounced. For example, Economic Sentiment (ECONOMY) shows moderate explanatory power for survey indicators like CONS_CONFIDENCE and BUSS_CONFIDENCE, but the p-values are less compelling compared to the reverse relationship. Political Sentiment (POLITICS) demonstrates Granger-causality for surveys such as MANUF_EXP and SERVICES_EXP; however, the strength and scope of these relationships are limited compared to survey-to-sentiment causality.

While Granger-causality is weaker in this direction, the instantaneous causality results remain highly significant (p-value < 0,05) across most relationships. This suggests that sentiments and surveys are contemporaneously related, even if sentiments alone do not consistently lead changes in survey

indicators. The weaker explanatory power of sentiments may stem from their reactive nature, as sentiments are often shaped by external shocks, social narratives, and underlying economic conditions, which are more systematically captured in survey indicators. Consequently, sentiments seem to follow survey results rather than driving them.

Comparing the explanatory power of survey-to-sentiment and sentiment-to-survey relationships reveals a clear dominance of survey indicators. Surveys such as CONS_CONFIDENCE, MANUF_EXP, and SERVICES_EXP consistently exhibit strong Granger-causality for sentiment variables, providing structured insights into public expectations and making them reliable predictors of sentiment trends. On the other hand, sentiment variables have limited Granger-causality for survey indicators, suggesting that sentiments are more likely to reflect, rather than drive, changes in public perceptions captured by surveys. Both surveys and sentiments exhibit significant instantaneous causality, highlighting their mutual influence within the same time frame and indicating a close interconnection with real-time feedback between the two.

These findings have important implications. From a policy perspective, they underscore the value of survey indicators as tools for policymakers and analysts. By monitoring survey data, stakeholders can anticipate shifts in public sentiment and design interventions to address emerging concerns. Sector-specific surveys, such as those focusing on manufacturing expectations, can provide early warnings of sentiment shifts in related areas like societal or political perceptions.

The results also highlight the interdependence of economic, social, and political factors. Surveys provide a structured representation of these dynamics, while sentiment indicators capture more nuanced and reactive responses. Together, they offer complementary insights into the state of public perceptions. The significant instantaneous causality between surveys and sentiments emphasizes their mutual influence and interdependence.

In essence, the Granger-causality analysis demonstrates that survey indicators generally have greater explanatory power for sentiment variables than vice versa. Surveys seem to act as leading indicators, reflecting structured assessments of public and sector-specific expectations that shape broader sentiment trends. In contrast, sentiments are more reactive and limited in their ability to predict changes in survey results. These findings emphasize the importance of leveraging survey data for forecasting and policy-making while recognizing the reactive and complementary role of sentiments in understanding public perceptions.

Table B1. Granger Causality Test Results for Sentiments-to-GDP Annual Growth Rate (%)

Variable	F-Test	df1, df2	p-value
ECONOMY	1,1438	165, 1616	0,1123
POLITICS	1,3581	165, 1616	0,0026
HEALTH	1,3498	165, 1616	0,0031
SECURITY	1,3818	165, 1616	0,0016
SOCIETY	1,4380	165, 1616	0,0004

The results from the Table B1 reveal that certain sentiment variables significantly predict GDP annual growth rates, highlighting the impact of societal perceptions on economic outcomes. Interestingly, ECONOMY sentiment does not show statistical significance in predicting GDP growth (F-Test: 1,1438, p-value: 0,1123). This may be because traditional economic indicators already capture much of its explanatory power, or due to potential multicollinearity with other variables obscuring its independent effect.

These results highlight the importance of non-economic factors—such as societal, security, political, and health sentiments—in driving GDP growth. Public perceptions of stability, governance, and well-being are intricately connected to economic performance. Monitoring these sentiments provides valuable insights for forecasting economic trends and designing policy interventions.

From a policy perspective, enhancing public confidence in governance, security, and societal well-being can foster an environment conducive to growth, even without significant changes in traditional economic indicators. By quantifying and analyzing people's sentiment policymakers may catch early signals of potential shifts in economic performance. It's important to note that while Granger-causality indicates predictive relationships, it does not imply true causation. The associations identified should be interpreted as indicative rather than definitive causal mechanisms. Further research is suggested to explore the underlying drivers of these relationships, especially for a system as volatile as Ecuador.

Table B2. Granger-Causality Analysis Results for Surveys-to-Sentiments

Cause	Effect	F-Test	P-Value (Granger)	Chi-Square	P-Value (Inst.)
CONS_CONFIDENCE	ECONOMY	1,56787	0,00002	62,80371	0,0000
CONS_CONFIDENCE	POLITICS	1,56787	0,00002	62,80371	0,0000
CONS_CONFIDENCE	HEALTH	1,56787	0,00002	62,80371	0,0000
CONS_CONFIDENCE	SECURITY	1,56787	0,00002	62,80371	0,0000
CONS_CONFIDENCE	SOCIETY	1,56787	0,00002	62,80371	0,0000
BUSS_CONFIDENCE	ECONOMY	1,05127	0,32037	125,24297	0,0000
BUSS_CONFIDENCE	POLITICS	1,05127	0,32037	125,24297	0,0000
BUSS_CONFIDENCE	HEALTH	1,05127	0,32037	125,24297	0,0000
BUSS_CONFIDENCE	SECURITY	1,05127	0,32037	125,24297	0,0000
BUSS_CONFIDENCE	SOCIETY	1,05127	0,32037	125,24297	0,0000
MANUF_EXP	ECONOMY	1,55227	0,00002	134,56653	0,0000
MANUF_EXP	POLITICS	1,55227	0,00002	134,56653	0,0000
MANUF_EXP	HEALTH	1,55227	0,00002	134,56653	0,0000
MANUF_EXP	SECURITY	1,55227	0,00002	134,56653	0,0000
MANUF_EXP	SOCIETY	1,55227	0,00002	134,56653	0,0000
SERVICES_EXP	ECONOMY	1,47435	0,00018	138,34552	0,0000
SERVICES_EXP	POLITICS	1,47435	0,00018	138,34552	0,0000
SERVICES_EXP	HEALTH	1,47435	0,00018	138,34552	0,0000
SERVICES_EXP	SECURITY	1,47435	0,00018	138,34552	0,0000
SERVICES_EXP	SOCIETY	1,47435	0,00018	138,34552	0,0000
CONSTRUC_EXP	ECONOMY	1,30508	0,00778	102,88846	0,0000
CONSTRUC_EXP	POLITICS	1,30508	0,00778	102,88846	0,0000
CONSTRUC_EXP	HEALTH	1,30508	0,00778	102,88846	0,0000
CONSTRUC_EXP	SECURITY	1,30508	0,00778	102,88846	0,0000
CONSTRUC_EXP	SOCIETY	1,30508	0,00778	102,88846	0,0000
ECON_EXP	ECONOMY	1,60664	0,00001	140,24323	0,0000
ECON_EXP	POLITICS	1,60664	0,00001	140,24323	0,0000
ECON_EXP	HEALTH	1,60664	0,00001	140,24323	0,0000
ECON_EXP	SECURITY	1,60664	0,00001	140,24323	0,0000
ECON_EXP	SOCIETY	1,60664	0,00001	140,24323	0,0000

Table B3. Granger-causality analysis results for sentiment-to-surveys

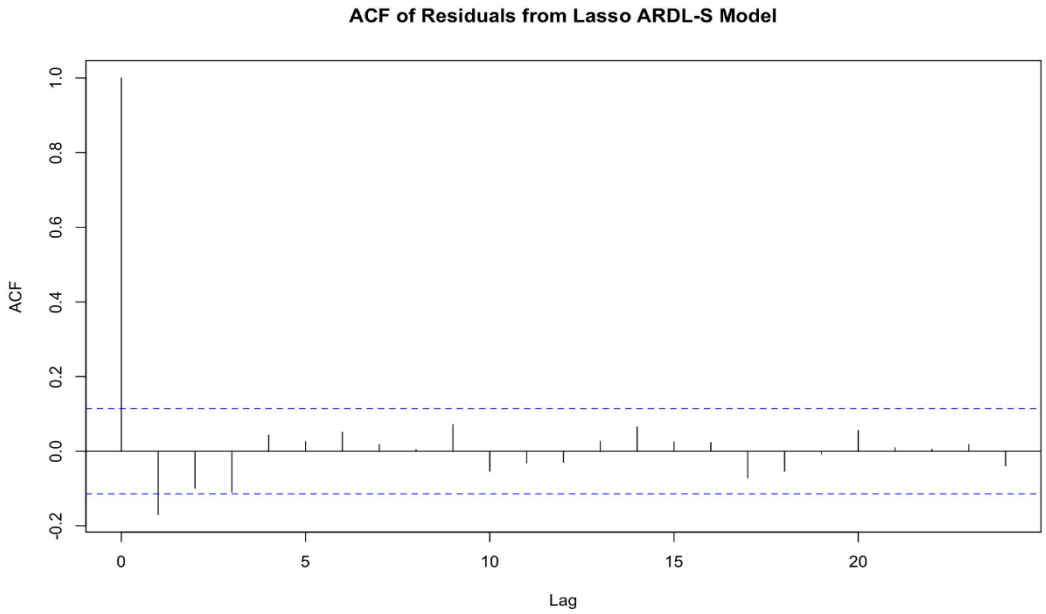
Cause	Effect	F-Test	P-Value (Granger)	Chi-Square	P-Value (Inst.)
ECONOMY	CONS_ CONFIDENCE	1,14376	0,11231	47,47917	0,0000
ECONOMY	BUSS_CONFIDENCE	1,14376	0,11231	47,47917	0,0000
ECONOMY	MANUF_EXP	1,14376	0,11231	47,47917	0,0000
ECONOMY	SERVICES_EXP	1,14376	0,11231	47,47917	0,0000
ECONOMY	CONSTRUC_EXP	1,14376	0,11231	47,47917	0,0000
ECONOMY	ECON_EXP	1,14376	0,11231	47,47917	0,0000
POLITICS	CONS_ CONFIDENCE	1,35815	0,00262	54,56292	0,0000
POLITICS	BUSS_CONFIDENCE	1,35815	0,00262	54,56292	0,0000
POLITICS	MANUF_EXP	1,35815	0,00262	54,56292	0,0000
POLITICS	SERVICES_EXP	1,35815	0,00262	54,56292	0,0000
POLITICS	CONSTRUC_EXP	1,35815	0,00262	54,56292	0,0000
POLITICS	ECON_EXP	1,35815	0,00262	54,56292	0,0000
HEALTH	CONS_ CONFIDENCE	1,34979	0,00313	29,41676	0,0142
HEALTH	BUSS_CONFIDENCE	1,34979	0,00313	29,41676	0,0142
HEALTH	MANUF_EXP	1,34979	0,00313	29,41676	0,0142
HEALTH	SERVICES_EXP	1,34979	0,00313	29,41676	0,0142
HEALTH	CONSTRUC_EXP	1,34979	0,00313	29,41676	0,0142
HEALTH	ECON_EXP	1,34979	0,00313	29,41676	0,0142
SECURITY	CONS_ CONFIDENCE	1,38184	0,00157	51,53013	0,0000
SECURITY	BUSS_CONFIDENCE	1,38184	0,00157	51,53013	0,0000
SECURITY	MANUF_EXP	1,38184	0,00157	51,53013	0,0000
SECURITY	SERVICES_EXP	1,38184	0,00157	51,53013	0,0000
SECURITY	CONSTRUC_EXP	1,38184	0,00157	51,53013	0,0000
SECURITY	ECON_EXP	1,38184	0,00157	51,53013	0,0000
SOCIETY	CONS_ CONFIDENCE	1,43798	0,00043	64,27585	0,0000
SOCIETY	BUSS_CONFIDENCE	1,43798	0,00043	64,27585	0,0000
SOCIETY	MANUF_EXP	1,43798	0,00043	64,27585	0,0000
SOCIETY	SERVICES_EXP	1,43798	0,00043	64,27585	0,0000
SOCIETY	CONSTRUC_EXP	1,43798	0,00043	64,27585	0,0000
SOCIETY	ECON_EXP	1,43798	0,00043	64,27585	0,0000

Appendix C. Non-Zero Coefficients for LASSO-ARDLS With λ = 0,0781

Variable	Coefficient
GDP_AN_GWTH_lag2	-0,0698
GDP_AN_GWTH_lag3	-0,0213
WTI_PRICES	0,0549
EXPORTS	0,2786
REMITTANCES	-0,2649
WTI_PRICES_lag1	-0,1481
WTI_PRICES_lag5	0,058
WTI_PRICES_lag11	0,0037
REMITTANCES_lag1	0,0344
REMITTANCES_lag9	0,0384
CONS_CONFIDENCE_lag1	-1,6114
CONS_CONFIDENCE_lag3	4,6834
CONS_CONFIDENCE_lag4	0,2263
CONS_CONFIDENCE_lag9	0,9159
BUSS_CONFIDENCE_lag10	0,1911
MANUF_EXP_lag4	0,2203
MANUF_EXP_lag5	-0,2391
SERVICES_EXP_lag10	0,0204
CONSTRUC_EXP_lag4	-1,1912
CONSTRUC_EXP_lag8	-0,4232
ECON_EXP_lag3	0,8452
ECON_EXP_lag9	0,3743
POLITICS	-0,0729
ECONOMY_lag6	0,0415
ECONOMY_lag7	-0,0811
ECONOMY_lag8	0,1068
POLITICS_lag1	0,0292
POLITICS_lag5	0,0135
POLITICS_lag7	-0,0005
SECURITY_lag1	0,076
SECURITY_lag3	-0,004
SECURITY_lag5	0,0257
SECURITY_lag7	0,0963
SOCIETY_lag5	-0,0504
SOCIETY_lag6	0,0149
Recession_2015_2016	-0,0408
Pandemic_2020	-0,4773
Post_Pandemic_Rec	0,2036

Appendix D. Graphs

Figure D1. ACF of Residuals from LASSO ARDL-S Model



Note: The autocorrelation function (ACF) results measuring the correlation between data points in a time series and their previous values.

Figure D2. Overall Time Series Sentiment Identification

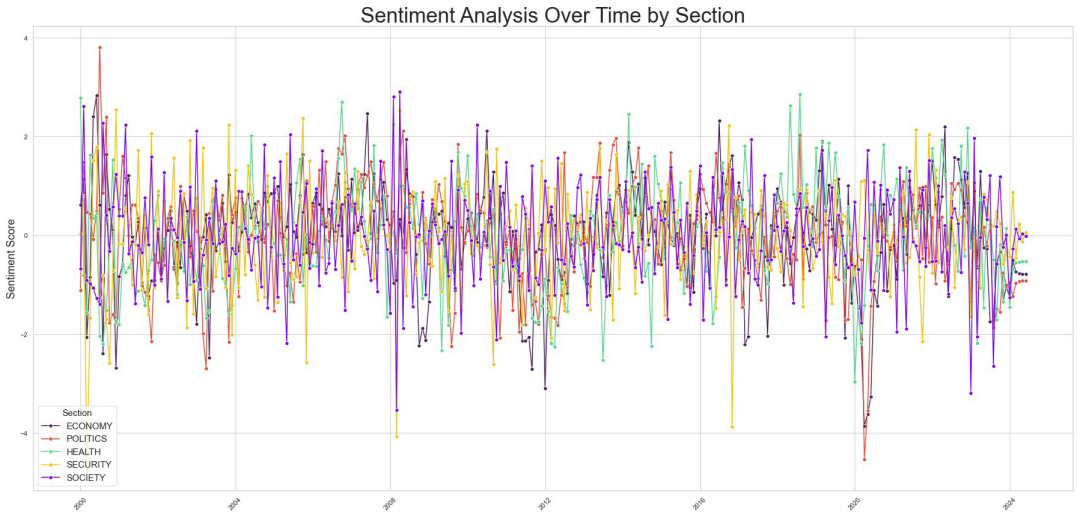


Figure D3. Economy Time Series Sentiment Identification

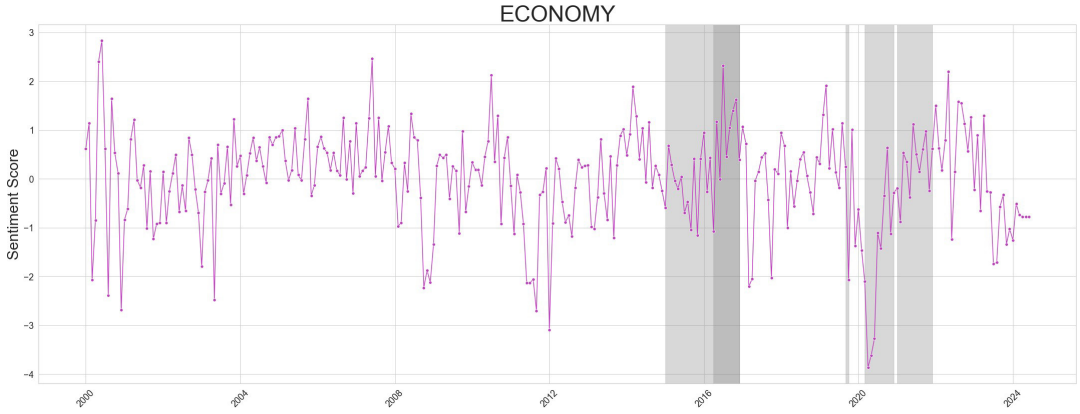


Figure D4. Politics Time Series Sentiment Identification

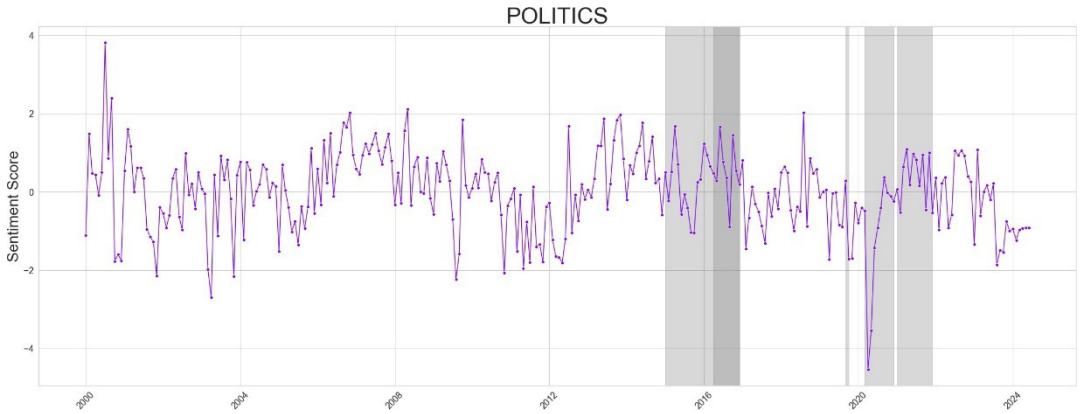


Figure D5. Health Time Series Sentiment Identification

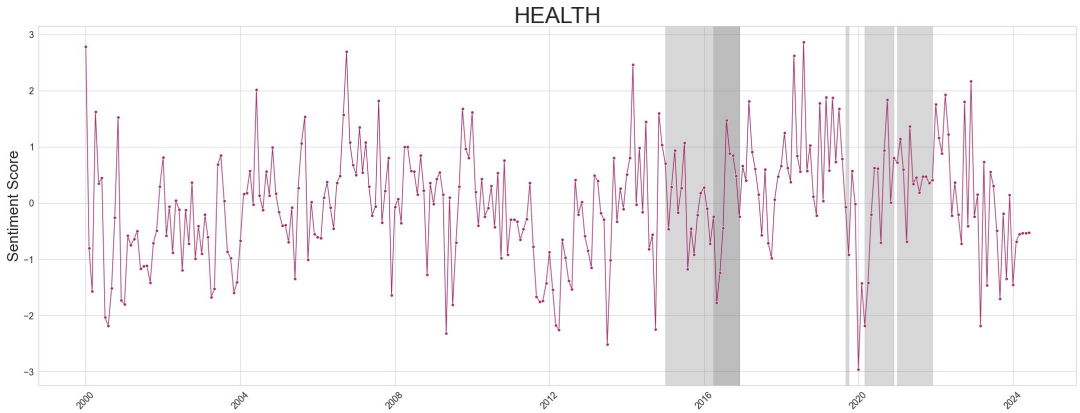


Figure D6. Security Time Series Sentiment Identification

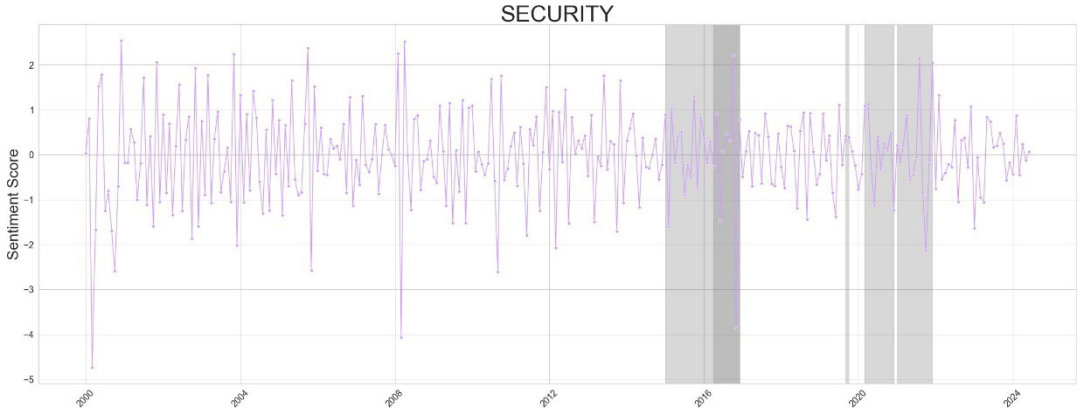


Figure D8. Society Time Series Sentiment Identification

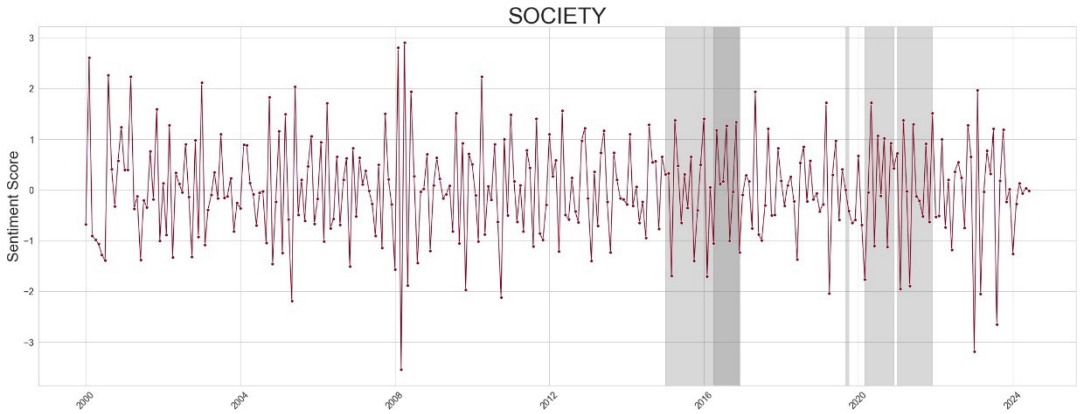


Figure D9. ARDLS Fitted Results Before (Upper Graph) and After LASSO Regularization (Below)

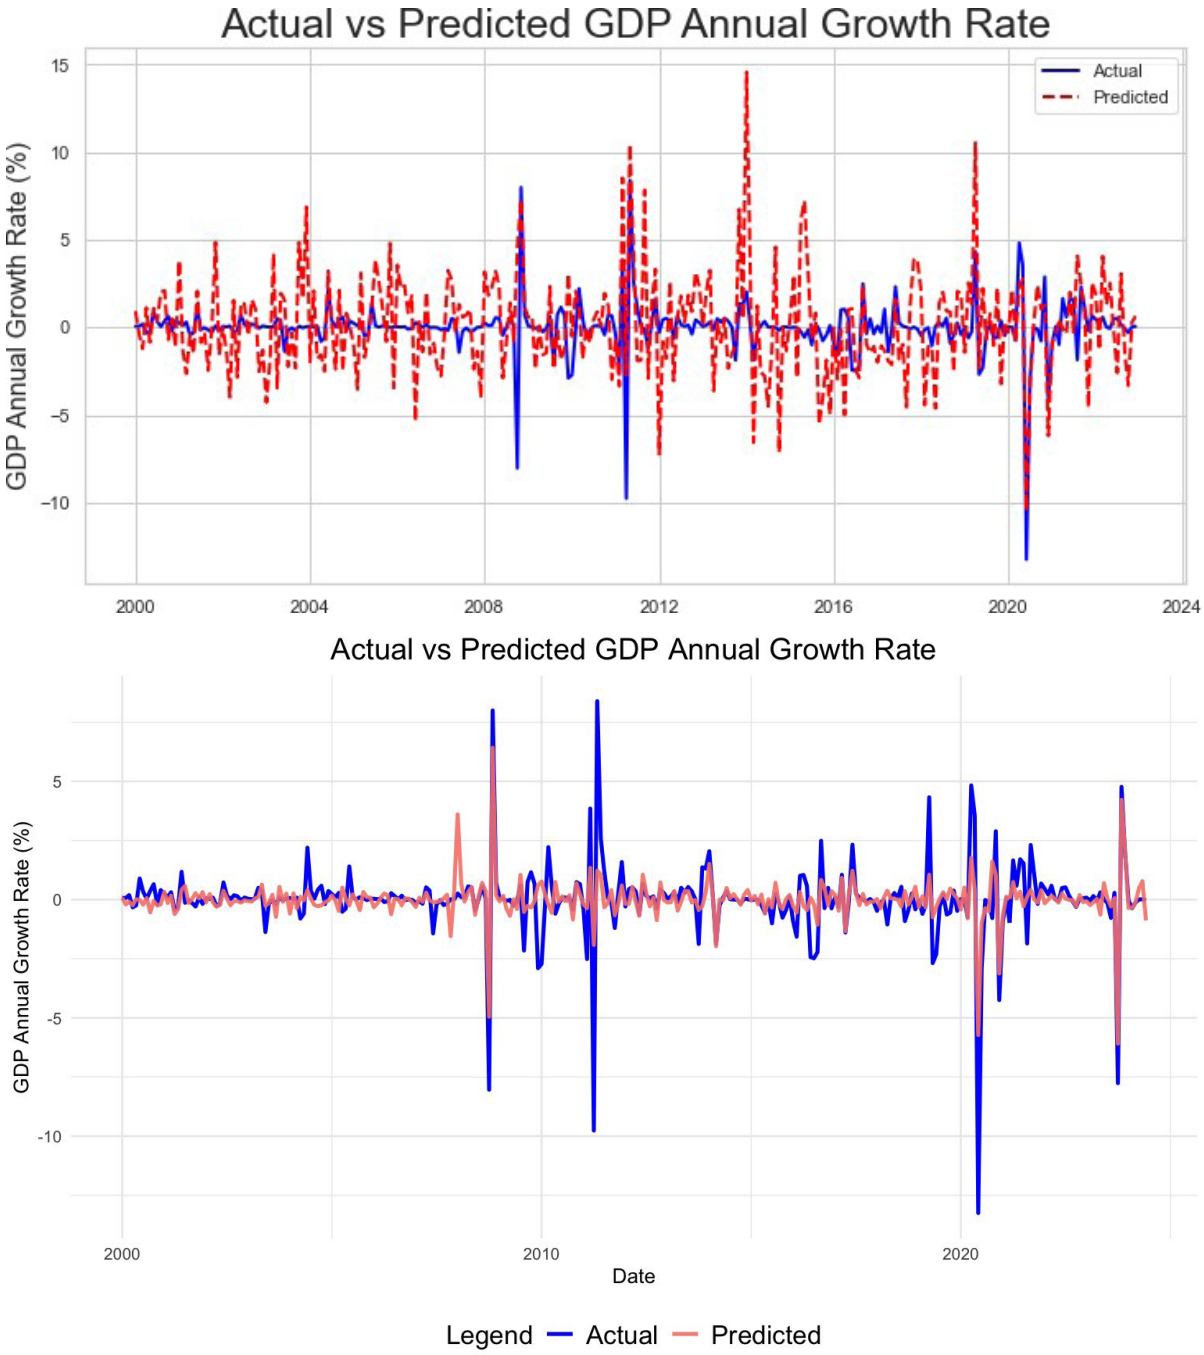
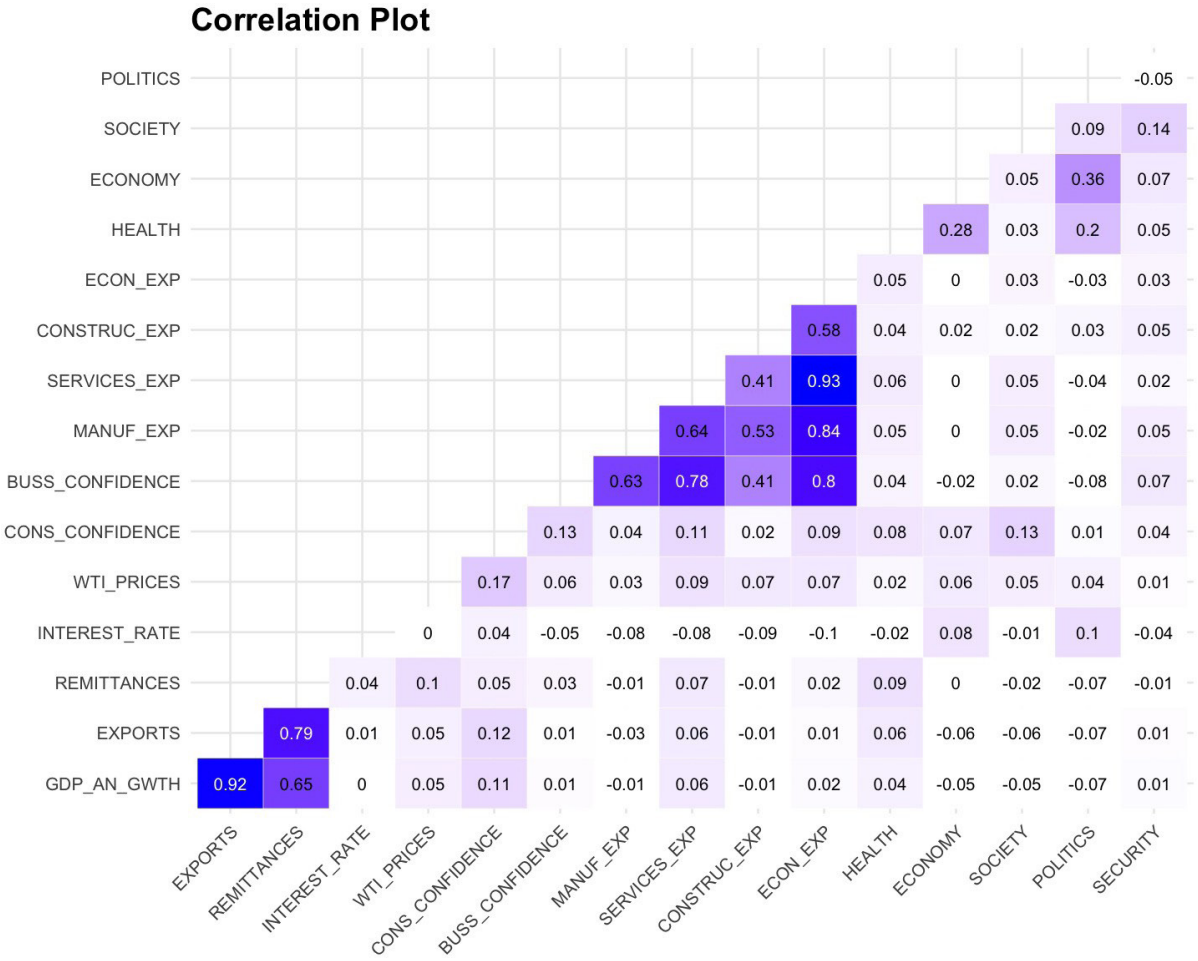


Figure D10. Correlation Plots of all the Different Variables



Appendix E. NLP dictionary (Reduced)

Economy

Economy, trade, business, entrepreneurial, labor, productive sector, company, market, human resources, innovation, expectations, hydrocarbons, trade agreement, stock market, market trends, Ecuador shrimp industry potential, strategic sector, exports, work, surplus, deficit, investment, bank, credit, financial, fund, mercantile, [...], losses, investor, import, export, real estate, currency, foreign exchange, remittances, tariff, customs duty, fiscal, microcredit, money, budget, subsidy, recession, inflation, GDP, tax, income, finances, costs, economic shutdown, industrial, funds, profit, banking, competitiveness, subsidy, circular economy, financing, foreign investment, central banking, investment funds.

Society

Society, citizen security, community, conjuncture, culture, interculturality, education, national assembly, state institution, issues, science, current events, news, events, opinion, national, global panorama, countries, firefighters, service, community, family, citizenship, neighbors, residential, neighborhood, event, social, church, NGO, volunteer, municipality, festival, campaign, foundation, social program, community project, social development, [...], organization, neighbor, environment, daily life, cultural, history, solidarity, party, families, communities, celebration, union, immigration, integration, coexistence, gender equality, human rights, social inclusion, collective.

Health

Health, pandemic, coronavirus, health situation, pandemic effects, health crisis, global fight against the pandemic, public health, diseases, pharmaceuticals, medicines, Ecuadorian Social Security Institute (IESS), vaccination, hospital, clinic, treatment, patient, epidemic, doctor, nurse, illness, contagion, symptoms, medicine, surgery, consultation, cancer, prevention, healthcare, emergencies, [...], therapy, psychology, cardiology, well-being, hygiene, care, rehabilitation, pharmacy, care, mental health, hospitalization, diagnosis, medical services, medical emergencies, biosecurity.

Security

Security, insecurity, crime, violence, prison crisis, street danger, murders, organized crime, criminal gangs, emergency, defense, civil, police, military, murder, robberies, kidnapping, criminal, armed forces, guard, justice, court,

prosecutor's office, prison, penalty, jail, terrorism, armed, conflict, safe, fire, accident, surveillance, [...], crime, operation, escort, national security, self-defense, guard, patrol, weapon, confiscation, pursuit, jurisdiction, civil protection, cybercrime, arms trafficking, public security, security operations, fraud.

Politics

Politics, political conjuncture, government change, elections, election results, political activity, presidential candidates, years, council, presidential elections, government, consensus, parties, politicians, austerity measures, Internal Revenue Service (SRI), congress, regulatory agency, ombudsman, state, minister, president, deputy, mayor, regime, legislative, reform, opposition, [...], democracy, bill, board, administration, council, vice president, campaign, alliance, deputies, senator, governorship, candidate, representative, election, municipality, jurisdiction, senate, mandate, political alliances, coalition, diplomacy, geopolitics, public policies, civil rights, electoral system.