Tiago do Nascimento Silva

# Evolution of distrust in electronic voting: a sentiment analysis

Brasília

2023

Tiago do Nascimento Silva

# Evolution of distrust in electronic voting: a sentiment analysis

Monografia apresentada ao Curso de Graduação em Economia, Universidade de Brasília, como requisito parcial à obtenção do grau de Bacharel em Ciências Econômicas.

Universidade de Brasília - UnB

Faculdade de Administração Contabilidade e Economia - FACE

Departamento de Economia - ECO

Orientador: Prof. Daniel Oliveira Cajueiro, Dr.

Brasília

2023

Tiago do Nascimento Silva
Evolution of distrust in electronic voting: a sentiment analysis/ Tiago do Nascimento Silva. – Brasília, 2023-
46p. : il. (algumas color.) ; 30 cm.

Orientador: Prof. Daniel Oliveira Cajueiro, Dr.

Monografia (Graduação) – Universidade de Brasília - UnB
Faculdade de Administração Contabilidade e Economia - FACE
Departamento de Economia - ECO , 2023.

1. Urnas eletrônicas. 2. Análise de sentimento. 3. Processamento de Linguagem Natural. II. Universidade de Brasília. III. Faculdade de Administração, Contabilidade e Economia - FACE. IV. Departamento de Economia IV. Evolution of distrust in electronic voting: a sentiment analysis

Tiago do Nascimento Silva

# Evolution of distrust in electronic voting: a sentiment analysis

Monografia apresentada ao Curso de Graduação em Economia, Universidade de Brasília, como requisito parcial à obtenção do grau de Bacharel em Ciências Econômicas.

Trabalho aprovado. Brasília, 20 de julho de 2023:

**Prof. Daniel Oliveira Cajueiro, Dr.**
Orientador

**Professor**
Convidado 1

**Professor**
Convidado 2

Brasília
2023

# Acknowledgements

Agradeço à Deus por ter me permitido realizar esse curso e ter me abençoado durante toda a trajetória.

Aos meus pais por cada ensinamento e por terem dado o máximo de si para que eu tivesse uma boa educação. Agradeço ao meu irmão pelo companheirismo.

Agradeço ao Professor Daniel Cajueiro por ter me orientado na realização desse trabalho de forma solícita.

À todos os colegas que me acompanharam durante o período de graduação.

# Resumo

Neste trabalho, nós avaliamos a percepção sobre as urnas eletrônicas no Brasil. Nós submetemos mais de 57 mil sentenças de notícias, de 1996 a 2023, à quatro modelos de análise de sentimento pré-treinados. Depois, construímos uma série temporal do índice sentimento, que fornece uma medida quantitativa para a percepção sobre as urnas. Nossos resultados indicam que recentemente o sentimento sobre as urnas eletrônicas tem se tornado cada vez mais mais negativo e que as eleições presidenciais de 2014 foram um ponto chave na série temporal.

**Palavras-chave**: Urnas Eletrônicas, Análise de Sentimento, Notícias, Índice de Sentimentos, Confiança

# Abstract

In this work we evaluate the perception towards electronic voting machines in Brazil using sentiment analysis. We submit over 57K news sentences, from 1996 to 2023, to four pre-trained sentiment analysis models in order to evaluate the polarity of these sentences. Then, we construct a sentiment index series, which provides a quantitative measure for the sentiment over time. Our results suggest a increase in negative sentiment towards electronic voting and that the 2014 Brazilian presidential elections were a key point in the sentiment index series.

**Keywords**: Electronic Voting Machines, Sentiment Analysis, News, Sentiment Index, Trust

# List of Figures

# List of Tables

# Contents

# 1 Introduction

The understanding of public perception towards the electoral process is crucial to comprehend the electoral participation and the confidence in institutions in any democratic system. A greater trust in electoral process may lead to a more active popular participation and to a more stable democracy. This ensures a more representative democracy and fosters a favorable environment for development.

Since the adoption of electronic voting machines in Brazil, politicians and population discuss the fairness of the elections during every electoral year. Specially in recent years, this topic has become one of the most important debates in Brazilian politics. In this work, we propose a quantitative analysis of the sentiment towards electronic voting in Brazil using natural language processing. We collect over 57K news sentences from 1996 to 2023 and evaluate their polarity using four different pre-trained sentiment analysis models. Then, we define a sentiment index and construct a series to track the evolution of public perception towards electronic voting over time. Our results indicate a growing negativity in sentiment, and that the 2014 Brazilian presidential elections was a key point in this trend.

Brazil adopted the electronic voting technology in the 1996 municipal elections, and since the 2000 municipal elections, all votes have been computed electronically. The introduction of electronic voting aimed to keep human interference away from the electoral process and prevent electoral frauds. The adoption of electronic voting also accelerated the vote counting process. Fujiwara (2015) finds that electronic voting reduced residual voting (votes not assigned to a candidate and discarded from the tallying of results) in state legislature elections. And this reduction in residual voting enfranchised millions of voters who would not have their votes counted when using a paper ballot.

Despite continuous testing and improvement of the security of electronic voting machines, it is clear that in recent years, a part of the population has questioned the trust in electronic voting machines. The main accusations are vulnerabilities in the security system and impossibility of auditing the machines[1]. Ruediger and Grassi (2018) show that between October and September 2018 a supposed fraud in electronic voting machines was the most cited fake news on Twitter, with more than one million tweets. This content, in addition to being widely reproduced, has a practical interference in electoral process

[1] The Brazilian Superior Electoral Court have publicized many materials to deny these accusations. On the vulnerability of electronic voting system: <https://www.tse.jus.br/o-tse/escola-judiciaria-eleitoral/publicacoes/revistas-da-eje/artigos/revista-eletronica-eje-n.-6-ano-4/por-que-a-urna-eletronica-e-segura> (accessed July 14, 2023). On the impossibility of auditing the voting machines: <https://www.tre-sp.jus.br/comunicacao/noticias/2021/Julho/o-voto-eletronico-brasileiro-e-auditavel> (accessed July 14, 2023).

(DOURADO, 2020).

The use of natural language processing allows us to evaluate a large dataset and provide a quantitative measure for the sentiment. We submit our data to four different pre-trained models from Hugging Face in order to verify if our results are consistent. Despite our sentiment index series have different magnitudes in each model, they have similar paths.

This work is organized as follows. In the next section we discuss the relevant literature related to our work. In section 3 we describe the methodology used to construct our sentiment index series. In 4 we present our data. In section 5 we present and discuss the achieved results. Finally, in section 6 we summarize and conclude our work.

# 2 Literature review

Natural language processing can capture latent concepts, such as ideology. Rheault and Cochrane (2019) train a word embedding model on parliamentary corpora from Britain, Canada, and the United States to estimate the ideological placement of political parties. They compare their findings with several indicators, such as ideology scores, surveys and data from the Comparative Manifestos Project, and achieve good results in capturing the ideology of the parties.

In the context of Brazilian elections, Sakiyama et al. (2019) develop a breaking news event detector employing sentiment analysis. In their work, they use a convolutional neural network to classify tweets related to presidential candidates in the 2018 Brazilian elections as positive, negative or neutral. They then construct a time-series and use an unsupervised time-series anomaly detector to monitor it. Their model has good results in classifying sentiment in tweets and in detecting breaking news events.

On the trust in electronic voting machines, Ruediger and Grassi (2020) collect posts from Facebook and Twitter between 2014 and 2020 to investigate the circulation of content related to the idea of fraud in electronic voting in Brazil. They find that the circulation of such content is intensifying, not only during election periods, but also between elections.

DOURADO (2020) analyses a sample of fake news collected during the 2018 Brazilian presidential elections. Her work appoints that "fraud in electronic voting" was the principal thematic set found and the most shared, indicating the strength of this idea. She also argues that the sharing of fake news by the professional press contributed to strengthen the distrust in electronic voting in Brazil. In her work, she affirms that this dissemination of fake news has a practical effect on the electoral process.

Rodrigues et al. (2022) conduct a neo-institutional analysis to examine the transition from trust to distrust in electronic voting in Brazil. The neo-institutional analysis considers the development and interaction of formal and informal institutions. They state that the electronic voting machines emerged as a formal institution to combat fraud in electoral process. Initially, the collective confidence in electoral institutions supported the acceptance of electoral results. However, this collective confidence did not exempt the confidence in electoral process from threats. They discuss how declarations of Leonel Brizola[1] raised doubts to the security of electronic voting in the early years of electronic voting adoption. They understand that the declarations of Brizola are the genesis of a contemporary series of exogenous irritations in the electoral system confidence. They also affirm that the Aécio

[1] Leonel Brizola (1922-2004) was a prominent Brazilian politician. He ran for the presidency in the 1989 presidential elections and is the only politician elected governor in two Brazilian states. Brizola founded the PDT (Democratic Labour Party) and was known as a charismatic leader.

Neves' questions about the electoral results marked the moment when these exogenous irritations began to seriously threaten confidence in the electoral system, and that this debate was amplified during 2018 general elections.

# 3 Methodology

We collect sentences from news that contain the words *electronic voting machine*, *electronic voting machines*, *ballot box*, *ballot boxes*, *electronic vote*, *electronic votes* or *electronic voting*[1] in order to measure the sentiment towards electronic voting machines. We evaluate the polarity of these sentences using four pre-trained sentiment analysis models from Hugging Face. After doing so, we define our measure of sentiment for each year and model. In the following subsections, we will elaborate on sentiment analysis, the selected models, and our sentiment index.

## 3.1 Sentiment Analysis

Sentiment analysis is an application of natural language processing where the aim is to classify the polarity of a given text. The polarity of a text refers to the sentiment that is expressed, whether it is positive, negative or neutral. There are several approaches to build a sentiment analysis model. We can first compute a statistic, such as the tf-idf, or train a language model, as Word2Vec or BERT, and then use it to train a classification model using supervised learning.

### 3.1.1 Tf-idf

Tf-idf (*term frequency-inverse document frequency*) is a statistical measure used to assess the importance of a given word in a document. Tf-idf aims to give more importance to words that are most common in one document, but appear in a few documents in the corpus. In this way, words that are common in a text but are present in a lot of documents will have little importance. This is achieved by computing for each word in a document the relative frequency of the word in the document and weighting by the logarithm of the inverse relative frequency of documents that the word appears. Equation 3.1 describes tf-idf, for term $t$ in document $d$:

$$tf\text{-}idf_{t,d} = tf_{t,d} \times log(\frac{N}{df_t}) \tag{3.1}$$

where $tf_{t,d}$ is the frequency of term $t$ in document $d$, $N$ is the total number of documents and $df_t$ is the number of documents with term $t$.

[1] These are a literal translation from Portuguese. The original words are respectively *urna eletrônica*, *urnas eletrônicas*, *urna*, *urnas*, *voto eletrônico*, *votos eletrônicos* and *votação eletrônica*.

Once we have the statistic for each word in each document, we can represent each document as a vector of these statistics and then utilize supervised learning techniques to classify the sentiment.

### 3.1.2 Word2Vec

Word2Vec is a language model used to generate distributed vector representations of words. For each word in the corpus vocabulary, the model generates a dense vector of determined dimension. The main advantage of Word2Vec is that the generated vectors preserve syntactic and semantic relations.

The model, proposed by Mikolov et al. (2013a), exists in two versions: CBOW and Skip-gram. CBOW (Continuous Bag of Words) takes words from context on both left and right windows and attempts to predict the target word. The Skip-gram takes the target word and tries to predict contexts words from both the left and right windows.

The learned vectors reproduce syntactic and semantic relations from the words. For example, words with similar meaning will have vectors with greater similarity. Mikolov et al. (2013b) show that simple algebraic operations with the vectors result in vectors with high similarity to the vector of the expected word.

After obtaining the word embeddings, we can proceed to classify the sentiment of each document. We can do this by taking the average of the vectors of all words in the document and then training a supervised learning model.

### 3.1.3 Transformer

Transformer is a deep learning architecture based on a self-attention mechanism used to generate sequences from an input sequence. The self-attention mechanism captures dependencies of different positions in a sequence in order to learn the most important parts and then generate a representation of each token with context information (Vaswani et al., 2017). Unlike recurrent neural networks (RNNs), the self-attention mechanism does not require the previous hidden state to compute the current hidden state. This allows more parallelization, reducing the computational cost, and improves the model performance by attending to dependencies regardless of distances within the sequence.

The self-attention mechanism works from three different representations of tokens in a sequence: queries, keys and values. As we explain further, the embeddings layers transform each token in the sequence in an embedding and add it to a positional encoding vector. These embeddings together form a matrix of dimension $L \times d_{model}$, where $L$ is the number of tokens and $d_{model}$ is the chosen model dimension. Then, we project this matrix into three separate representations using different linear transformations, resulting in the queries and keys matrices, both of dimensions $L \times d_k$, and the values matrix, with dimension $L \times d_v$. There are many functions to calculate attention. Transformer architecture uses the Scaled Dot-Product Attention, described in equation 3.2:

$$Attention(Q, K, V) = softmax(\frac{QK^T}{\sqrt{d_k}})V \tag{3.2}$$

where $Q$ is the queries matrix, $K$ is the keys matrix and $V$ is the values matrix.

We multiply the queries matrix by the transpose of the keys matrix to obtain a matrix of dimension $L \times L$. Then, to improve the results of the softmax function, we scale each element of this matrix by $\frac{1}{\sqrt{d_k}}$ to reduce the magnitude of the matrix elements. After doing so, we apply the softmax function for each row in the matrix and multiply by the values matrix. This results in a matrix of dimension $L \times d_v$ with context information for each token. Figure 1 illustrates the self attention mechanism.

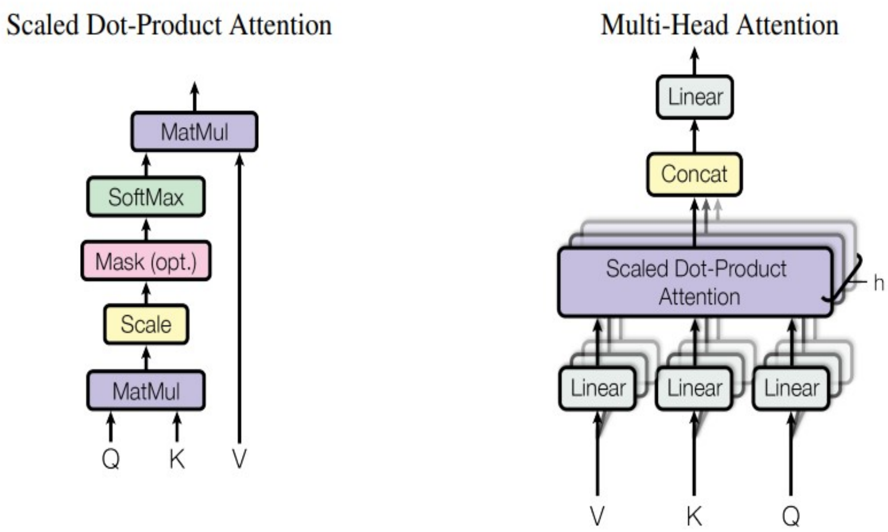


Figure 1 – Scaled Dot-Product Attention on left and Multi-Head Attention on right. Source: Vaswani et al. (2017)

Vaswani et al. (2017) find it beneficial to project the queries, keys and values matrices into $h$ representations using different linear transformations and compute the attention function in parallel for each of these representations, and then concatenate the attention matrices. This technique, known as Multi-Head Attention, enables the attention mechanism to learn dependencies from different representations at different positions in the sequence. Figure 1 illustrates the Multi-Head Attention and equation 3.3 describes it:

$$MultiHead(Q, K, V) = Concat(head_1, ..., head_h)W^O \tag{3.3}$$

where $head_i = Attention(QW_i^Q, KW_i^K, VW_i^V)$, with $i = 1, ..., h$, $W_i^Q \in \mathbb{R}^{d_{model} \times d_k}$, $W_i^K \in \mathbb{R}^{d_{model} \times d_k}$ and $W_i^V \in \mathbb{R}^{d_{model} \times d_v}$ are respectively the queries, keys and values projections matrices for head $i$ and $W^O \in \mathbb{R}^{hd_v \times d_{model}}$ is the projection matrix for the multi-head attention mechanism.

The transformer architecture is composed by an encoder-decoder structure. Figure 2 shows the transformer architecture.

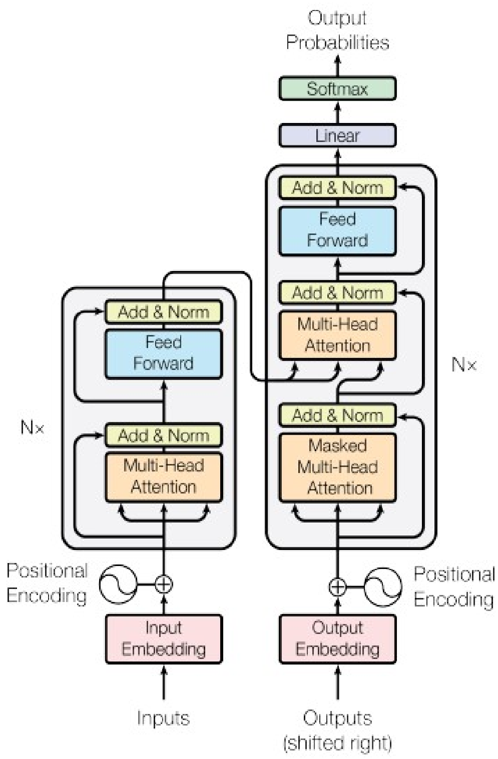


Figure 2 – Transformer architecture. Source: Vaswani et al. (2017)

For each token in the input sequence, we generate an embedding of dimension $d_{model}$. This embedding is the same for a determined word, regardless of its occurrence in different positions or contexts. In order to provide information about the word position, Vaswani et al. (2017) place a positional encoding vector for each token, also of dimension $d_{model}$. The positional encoding vector is defined using sine and cosine functions with different frequencies, described by the functions 3.4:

$$\begin{aligned} PE(pos, 2i) &= sin(\frac{pos}{10000^{2i/d_{model}}}) \\ PE(pos, 2i+1) &= cos(\frac{pos}{10000^{2i/d_{model}}}) \end{aligned} \tag{3.4}$$

where $pos$ is the word position, and $i$ is the dimension index. For even indexes the value of positional encoding comes from the sine function and for odd indexes the value comes from the cosine function.

We add the positional encoding vector to each token embedding, so the embedding contains information about the word and its position. Then, we pass the embedding into the encoder.

The encoder has six layers. Each layer is composed by a multi-head attention sub-layer and a feed-forward sub-layer. We pass the matrix with word and position information into the multi-head attention sub-layer that generates a matrix with context information for each token, as explained above. We add the generated matrix to the matrix with word and position information and then normalize it. So after the multi-head attention sub-layer, we obtain a matrix with word, position and context information. We pass this matrix into the feed-forward sub-layer. The feed-forward sub-layer perform two linear transformations on the input matrix, and apply a ReLu activation function between them, as shown in equation 3.5.

$$FFN(x) = max(0, xW_1 + b_1)W_2 + b_2 \tag{3.5}$$

where $x$ is the input, $W_1$ and $W_2$ are linear transformations and $b_1$ and $b_2$ are biases.

As we do in the multi-head attention sub-layer, we add the output matrix from the feed-forward sub-layer to the input matrix and normalize. Then, we pass the output to the next encoder layer, and this process continues until we reach the final layer

The decoder has six layers. Each layer is composed by three sub-layers: a masked multi-head attention, a multi-head attention and a feed-forward sub-layer. Similar to the encoder, after each sub-layer we add the sub-layer output to its input and normalize. The decoder initializes with the output sequence, which is transformed into a matrix with word and position information of each token. Then, we pass this matrix into the masked multi-head attention sub-layer. The masked multi-head attention sub-layer works similar to the multi-head attention, but before we compute the softmax function at each token, we mask the scores of the tokens at the right of the current token, setting them to $-\infty$. By doing this, when we compute the softmax, the subsequent tokens will have a value of zero, retaining information only from the previous and current tokens. This is necessary to prevent the model from having any information about the next tokens while trying to predict them. The next sub-layer is the multi-head attention. In this sub-layer we generate the queries from the output of the previous layer, and the keys and values are generated from the encoder output and we perform the attention computation. After doing so, we pass the output into a feed-forward sub-layer, similar to what we do in the encoder. This process is repeated in each decoder layer. Finally, we apply a linear transformation in the feed-forward sub-layer output of the last decoder layer and perform a softmax to predict the next token.

Several domains use the transformer architecture, such as language models or computer vision. Specially in language models, the transformer-based models achieves state-of-the-art in many tasks. One essential language model that utilizes the transformer architecture is BERT (Bidirectional Encoder Representations from Transformers).

BERT is a language model trained on the transformer architecture proposed by Devlin et al. (2019). Unlike other language models, BERT is a bidirectional model, which means that it learns token representations using both left and right context. BERT is pre-trained using two unsupervised tasks. The first task is randomly masking some token from the input using a special token and them attempting to predict the masked token. The second task is trying to predict whether a sentence is the next sentence in a sequence. After pre-training the model, it can be fine-tuned to a specific task.

BERT is trained on the BooksCorpus and English Wikipedia datasets. The model was trained in two sizes. The $\text{BERT}_{\text{BASE}}$ has 12 transformer blocks and 110M parameters, and $\text{BERT}_{\text{LARGE}}$ has 24 transformer blocks and 340M parameters. Devlin et al. (2019) show that BERT achieves state-of-the-art in many tests for different tasks. In this way, BERT is widely used for fine-tuning other models and for developing new versions of the model, such as the models that we will present later.

## 3.2 Hugging Face

Hugging Face[2] is a platform for sharing datasets and pre-trained machine learning models. The majority of models available on Hugging Face are transformer-based and they have several applications, such as text classification, image classification or translation. We can download and use the available pre-trained models on Hugging Face using the transformer library for Python. The transformer API simplifies the use of pre-trained models by providing pipelines that automatically pre-process the input and use the models.

We select four text classification pre-trained models from Hugging Face to evaluate the polarity of each sentence in our dataset. The selected pre-trained models are described below:

- **cardiffnlp/twitter-xlm-roberta-base-sentiment[3]:** a multilingual XLM-roBERTa-base model proposed by Barbieri et al. (2022). The model is trained on a dataset of more than 190 million tweets and fine-tuned for sentiment analysis. We will refer to this model as *cardiff*;
- **lxyuan/distilbert-base-multilingual-cased-sentiments-student[4]:** a DistilBERT model trained on a multilingual sentiments dataset, which contains texts from product reviews and tweets, for example. We will call it *distilbert*;
- **lucas-leme/FinBERT-PT-BR[5]:** a fine-tuned version of BERTimbau (Souza et al.,

[2] <https://huggingface.co/>
[3] Available on: <https://huggingface.co/cardiffnlp/twitter-xlm-roberta-base-sentiment>
[4] Available on: <https://huggingface.co/lxyuan/distilbert-base-multilingual-cased-sentiments-student>
[5] Available on <https://huggingface.co/lucas-leme/FinBERT-PT-BR>

2020) for sentiment analysis with financial news from Brazilian newspapers proposed by Santos et al. (2023). We will call it *finbert*; and

- **citizenlab/twitter-xlm-roberta-base-sentiment-finetunned[6]:** a fine-tuned version from the Cardiff NLP Group sentiment classification model, trained on a collection of Wikipedia comments. We will refer to this model as *cardiff-FN*.

## 3.3 Sentiment Index

For each model, we define our sentiment index as defined by Hiew et al. (2022):

$$SI_{m,t} = \frac{Pos_{m,t} - Neg_{m,t}}{Pos_{m,t} + Neu_{m,t} + Neg_{m,t}} \tag{3.6}$$

where $m$ is the model, $t$ is the year and $Pos_{m,t}$, $Neu_{m,t}$, and $Neg_{m,t}$ are respectively the number of sentences labeled as positive, neutral and negative in model $m$ in year $t$.

So, the closer the sentiment index is to 1, the more positive the sentiment is towards the sentences. On the other hand, the closer the sentiment index is to -1, the more negative the sentiment is towards the sentences.

[6] Available on: <https://huggingface.co/citizenlab/twitter-xlm-roberta-base-sentiment-finetunned>

# 4 Data

Our data[1] contains 57.531 sentences extracted from over 29K news, from 1996 to 2023[2]. We collect over 500K news from the politics section of four different Brazilian newspapers: Folha de São Paulo, Valor Econômico, Gazeta do Povo and Correio Braziliense[3]. We select the news that contain the words: *electronic voting machine*, *electronic voting machines*, *ballot box*, *ballot boxes*, *electronic vote*, *electronic votes* or *electronic voting*. Then, for each news we extract the sentences that contain at least one of the mentioned words.

The data from Folha de São Paulo ranges from 1996 to 2023. For Valor Econômico the data is from 2011 to 2023. The data from Gazeta do Povo ranges from 2005 to 2023. And for Correio Braziliense the data is from 2008 to 2021.

Since our data comes only from news, it may reflect the perception of the journalists. Despite news articles can influence the sentiment and public opinion, we need to be careful when generalizing our results.

Figure 3 shows the number of sentences for each year and figure 4 presents the number of sentences per newspaper.

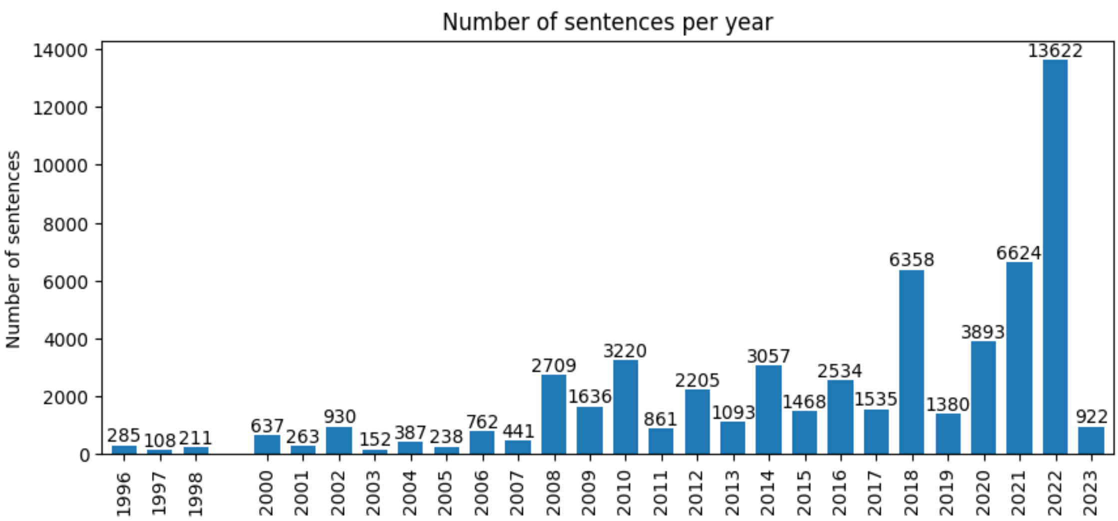


Figure 3 – Number of sentences per year in our dataset.

1 Our data and code are available on: <https://github.com/tiagonsilva/Monografia/tree/main>
2 Our data for 2023 is only up to May.
3 Folha de São Paulo: <https://www.folha.uol.com.br/>
Valor Econômico: <https://valor.globo.com/>
Gazeta do Povo: <https://www.gazetadopovo.com.br/>
Correio Braziliense: <https://www.correiobraziliense.com.br/>

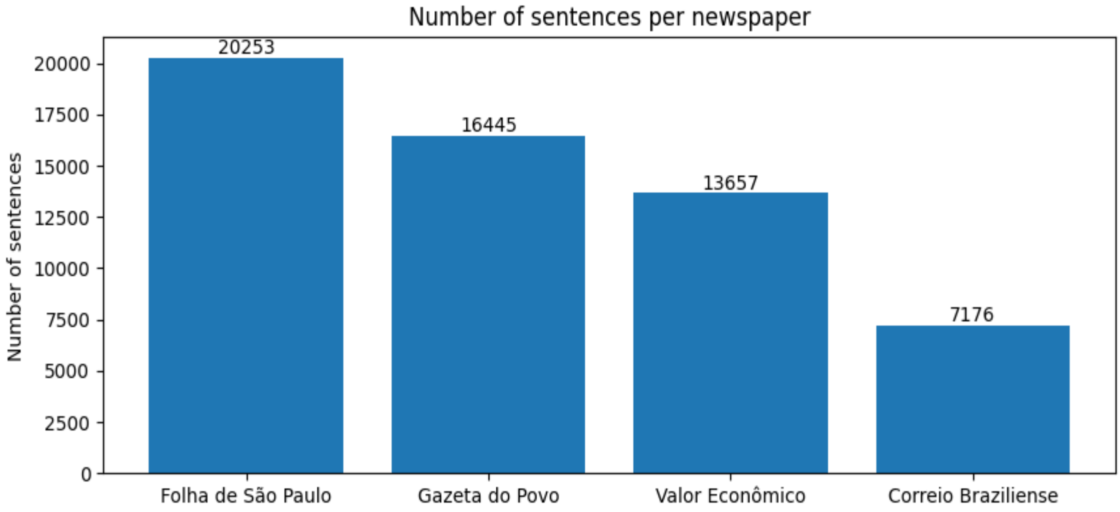


Figure 4 – Number of sentences per newspaper in our dataset.

# 5 Results

Figure 5 presents the sentiment index series for each model. We provide the tables with the results for each model in Appendix A.

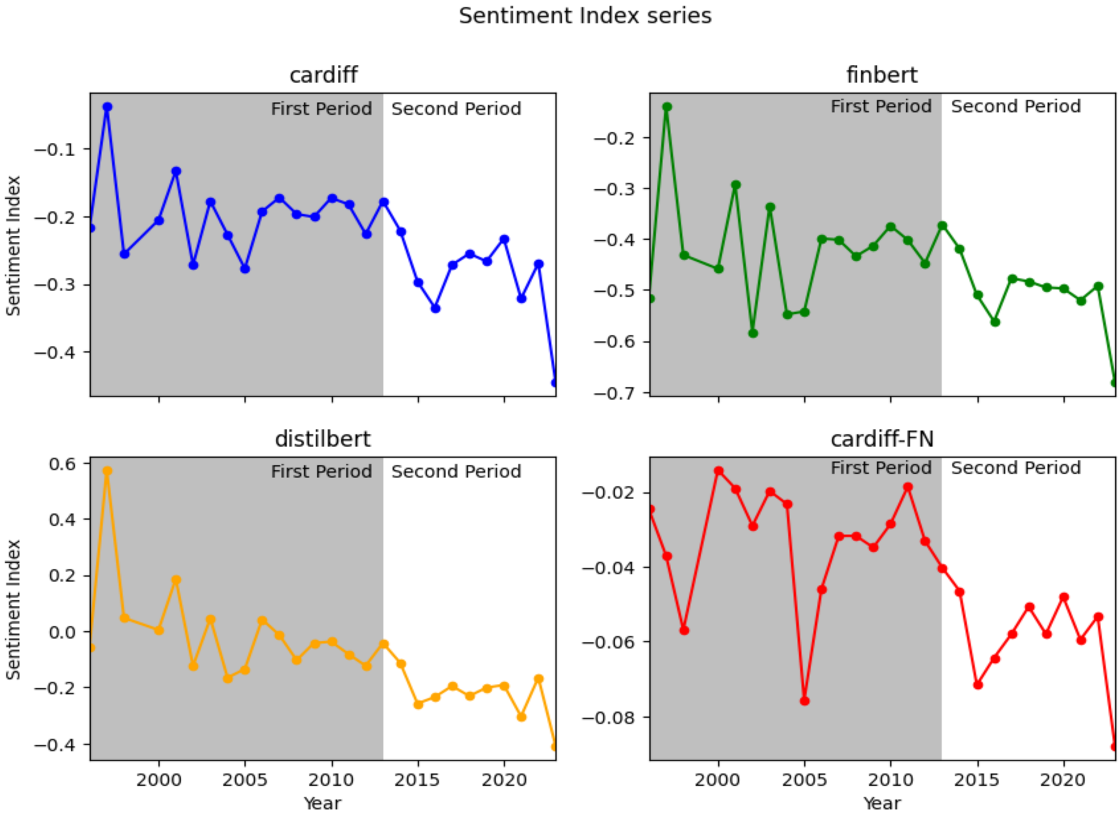


Figure 5 – Sentiment Index series for each model. First period is the period between 1996 and 2013 and the second period is the period between 2014 and 2023.

Despite the sentiment index series has different magnitudes for each model, it has similar time paths in both models. In the *cardiff*, *finbert* and *cardiff-FN* models, all values of sentiment index series are negative. The *cardiff-FN* model labels a large number of sentences as neutral, so the sentiment index magnitudes are smaller. In the beginning of *distilbert* sentiment index series, there are some positive values, indicating that the model captures the sentiment in those year as more positive. However, as the series progresses, the model follows a similar pattern to the other models.

In all models, we observe the same behavior. In general, the values of the sentiment index series are higher before 2014 than after 2014. Except for the years of 1998, 2002, 2004 and 2005, we can state this behavior. However, in these years the effects on sentiment index are one-off events, since the series returns to the previous level in the next periods. On the other hand, the series maintain or decrease their levels after the level decrease observed in 2014. Thus, the models suggest that the sentiment towards the electronic voting machines has become more negative since 2014.

In the following sections we divide our series in two periods: from 1996 to 2013 and from 2014 to 2023, to further explore this behavior. We also discuss the polarization of this sentiment.

## 5.1 First period: 1996-2013

The electronic voting machines were created as a solution to keep human interference away from the electoral process and then prevent electoral frauds. However, it did not take long for the electronic voting machines become the target of suspicion. In 2002, there were several accusations made against the machines, specially accusations from Leonel Brizola about the possibility of frauds. However, the affirmations about the vulnerabilities of electronic voting machine did not have a generalized impact on the trust of electoral system. Our sentiment index series reflect this, given that in the first period the sentiment index remains stable, except for the years of 1998, 2002, 2004 and 2005. We inspected the sentences for these years to understand why there are valleys in the series for these years.

In 1998, the electronic voting machines were used for the first time in Brazilian general elections. Two thirds of the population voted in electronic voting machines. Previously, the machines were used only in the 1996 municipal elections, by around one third of the population. In 2005 was performed the Brazilian firearms and ammunition referendum, which asked whether the sale of firearms and ammunition should be banned. The referendum was performed totally using electronic voting machines. For both years we provide the same explanation for the valleys in sentiment index series. The majority of sentences labeled as negative in these years were related to issues such as delays in voting, difficulties in using the electronic voting machines, defects in the machines or canvassing at poll places. Unlike occurs in the other valleys, the sentences were not related to the security or fairness of the voting.

In 2002, the Brazilian presidential elections were performed totally using electronic voting machines for the first time. The majority of sentences labeled as negative for this year were in fact related to doubts and concerns about the security of electoral process. We can report the affirmations about the fear of potential fraud in the machines from Leonel Brizola[1] and from Ciro Gomes[2]. Some specialists related problems in the compilation of the programs, constraints in accessing the machines source codes, difficulties in auditing

[1] Available on: <https://www1.folha.uol.com.br/folha/brasil/ult96u33089.shtml> (accessed July 14, 2023).

[2] Available on: <https://www1.folha.uol.com.br/folha/brasil/ult96u29760.shtml> (accessed July 14, 2023).

the machines and the possibility of identifying the voters[3]. In this election, around 5% of electronic voting machines were equipped with printers to allow for election recount, and it was the expected that all machines would have a print by 2004[4].

In 2004, around 30% of votes were computed electronically in the American presidential elections. The use of electronic voting machines in United States was widely debated and had a lot of polemics about the security of the machines. This had a large repercussion and affected our sentiment index series, what may explain the valley in our series for that year.

However, these years did not have a systematic impact in the trust on electronic voting machines. Actually, the sentiment index series remain in the same levels and look stable in the period from 2006 to 2013. Despite the controversies and discussions surrounding the use of electronic voting machines in these years, the overall sentiment towards electronic voting machines did not show significant long-term shifts during the subsequent years.

## 5.2 Second period: 2014-2023

In the year of 2014 there is a level decrease in the sentiment index series for both models. This is due to the PSDB (Brazilian Social Democracy Party) audit request after the defeat in 2014 Brazilian presidential elections for Dilma Roussef, from PT (Workers' Party). After the tight result in the presidential election, PSDB requested an election audit for the Superior Electoral Court in order to verify the fairness of the elections. The party argued that the society was questioning the trust in the electronic voting machines and the audit was necessary for avoid this sentiment. Initially, the Superior Electoral Court denied providing the data for all parties as requested by PSDB, but provided the data only for PSDB. The PSDB report did not identify any fraud on the electoral process, but stated that the electoral system could not be effectively audited.

Unlike the previous shocks, the one occurred in 2014 perpetuated through the next periods in the sentiment index series. This can be explained by the organization of the Brazilian right since 2013. In 2013, Brazil witnessed a great social movement that was started with protests against the increase in bus ticket prices but quickly evolved into demonstrations against the government, the parties, specially PT, and the corruption. Many right groups were created and strengthen during this movement, what allowed the emergence of a more organized right in Brazil. With a more arranged right, the discourse of fraud in the electoral process could be systematically reproduced. Thus, the organization

[3] Avalilable on: <https://www1.folha.uol.com.br/fsp/informat/fr2108200201.htm> (accessed July 14, 2023).

[4] Available on: <https://www1.folha.uol.com.br/fsp/informat/fr2108200203.htm> (accessed July 14, 2023).

of Brazilian right in 2013 allowed PSDB affirmations after the 2014 presidential elections perpetuate through the next years, as captured in our sentiment index series.

Jair Bolsonaro is a prominent figure in the discussions and questions surrounding electronic voting in the period from 2018 to 2023. Both during election years and between elections, while he was in the presidency, Bolsonaro made numerous accusations against the fairness of the electoral process. After the results of the first round in 2018 presidential election, Bolsonaro affirmed that he was not elected in the first round, because the electronic voting machines were rigged[5]. He repeated this statement several times in following years. During 2022 presidential elections, Bolsonaro once again made various accusations against the machines, including suspicions of fraud during a meeting with ambassadors[6] and requests to invalidate votes in machines with alleged problems[7]. In the period of Bolsonaro government, our sentiment index series is slightly decreasing.

The acts on January 8th were the most serious expression of the distrust in electronic voting machines. In that day, a group of radicals that did not trust and accept the results of electoral process assaulted the Palácio do Planalto, the National Congress and the Supreme Federal Court. This event can be related to the valley observed at the end of the sentiment index series.

## 5.3 Polarization

In addition to assessing the sentiment towards electronic voting machines, we also investigate whether this sentiment has become more polarized over time. When sentiment is polarized, individuals tend to express extreme opinions rather than being neutral. Our measure for polarization is the proportion of sentences labeled as neutral. If the proportion of neutral sentences is decreasing, it suggests an increasing polarization in sentiment. Figure 6 shows the polarization for each model.

We observe that the polarization patterns generally align with the sentiment index patterns for each model, except for the *distilbert* model, which classifies a relatively small number of sentences as neutral. The models indicate that the polarization in sentiment towards electronic voting machines is increasing over time. Similar to the sentiment index, our polarization measure also indicates a decrease in 2014, which continues to influence the polarization in the following years.

In the Brazilian politics, Fuks and Marques (2022) state an affective polarization

[5] Available on: <https://www1.folha.uol.com.br/poder/2018/10/bolsonaro-diz-que-foi-alvo-de-fraude-e-pede-mobilizacao-a-eleitores.shtml> (accessed July 14, 2023).

[6] Available on: <https://www1.folha.uol.com.br/poder/2022/07/bolsonaro-repete-teorias-da-conspiracao-e-ataca-urnas-stf-e-tse-a-embaixadores.shtml> (accessed July 14, 2023).

[7] Available on: <https://www1.folha.uol.com.br/poder/2022/11/pl-endossa-golpismo-de-bolsonaro-e-usa-relatorio-sem-provas-para-pedir-invalidacao-de-votos.shtml> (accessed July 14, 2023).

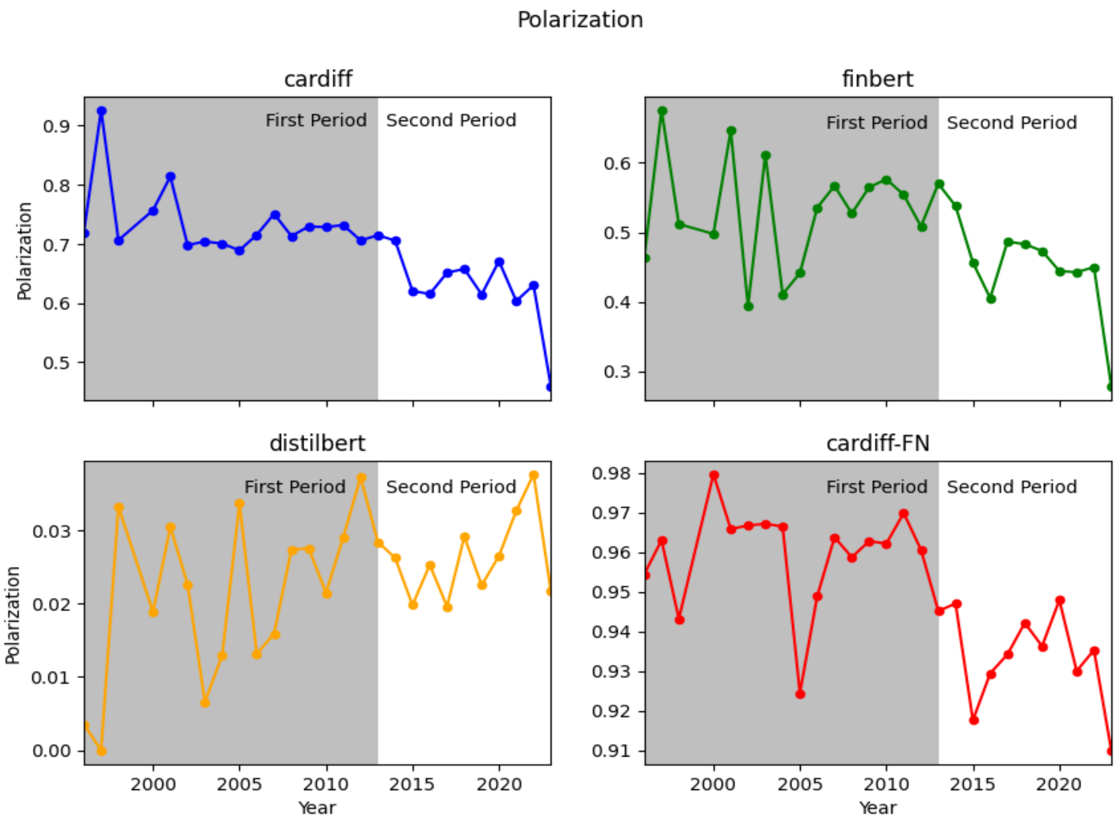


Figure 6 – Polarization for each model. First period is the period between 1996 and 2013 and the second period is the period between 2014 and 2023.

started in 2014, related to the strengthen of right. By 2018, this affective polarization became even more visible. Our polarization measure reveals a similar polarization in the sentiment towards electronic voting machines. This suggests that the sentiment polarization towards electronic voting can be a reflection of the affective polarization in the country, as extreme discourses adopted by politicians can influence public opinion. Therefore, in addition to becoming more negative, the sentiment polarization is increasing.

# 6 Summary and conclusion

In this work we use a sentiment analysis approach to evaluate the perception about electronic voting machines over time in Brazil. We submit over 57K sentences extracted from news to four pre-trained sentiment analysis models to evaluate the polarity of the sentences. We define a sentiment index and construct a sentiment index series for each model. The series in both models have similar paths, suggesting a growing negative sentiment towards electronic voting since 2014. Our results show that the 2014 Brazilian presidential elections were a key point in the perception about electronic voting machines, since our sentiment index series show a level decrease in this year that persists in next years. We divide our analysis in two periods: before 2014 and after 2014. In the first period, we state that despite there were suspicions about the machines, these suspicions did not have a systematic impact in the trust on the machines. However, in the second period, the affirmations made by PSDB perpetuated in the next years due to the strengthen of Brazilian right in 2013. Our study can provide a measure for a quantitative analysis of the perception about electronic voting machines.

# Bibliography

Barbieri, F.; Anke, L. E.; Camacho-Collados, J. *XLM-T: Multilingual Language Models in Twitter for Sentiment Analysis and Beyond*. 2022. Cited on page 26.

Devlin, J. et al. *BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding*. 2019. Cited on page 26.

DOURADO, T. *Fake news na eleição presidencial de 2018 no Brasil*. Tese (Doutorado) — Programa de Pós-Graduação em Comunicação e Culturas Contemporâneas, Universidade Federal da Bahia, Salvador, 2020. Cited 2 times on pages 18 and 19.

Fujiwara, T. Voting technology, political responsiveness, and infant health: Evidence from brazil. *Econometrica*, v. 83, n. 2, p. 423–464, 2015. Cited on page 17.

Fuks, M.; Marques, P. H. Polarização e contexto: Medindo e explicando a polarização política no brasil. *Opinião Pública*, v. 28, n. 3, p. 560–593, Dec 2022. Cited on page 34.

Hiew, J. Z. G. et al. *BERT-based Financial Sentiment Index and LSTM-based Stock Return Predictability*. 2022. Cited on page 27.

Mikolov, T. et al. *Efficient Estimation of Word Representations in Vector Space*. arXiv, 2013. Available on: <https://arxiv.org/abs/1301.3781>. Cited on page 22.

Mikolov, T. et al. *Distributed Representations of Words and Phrases and their Compositionality*. arXiv, 2013. Available on: <https://arxiv.org/abs/1310.4546>. Cited on page 22.

Rheault, L.; Cochrane, C. Word embeddings for the analysis of ideological placement in parliamentary corpora. *Political Analysis*, v. 28, n. 1, p. 112–133, 2019. Cited on page 19.

Rodrigues, A. L.; Silva, L. G.; Nóbrega, F. F. A desconfiança no sistema eleitoral como fator de comprometimento do capital social eleitoral: A descredibilidade que retrocede nos avanços democráticos. *Revista Jurídica da UFERSA*, v. 5, n. 10, p. 85–106, 2022. Cited on page 19.

Ruediger, M. A.; Grassi, A. Desinformação na era digital: amplificações e panorama das eleições 2018. *Policy Paper*, FGV DAPP, 2018. Cited on page 17.

Ruediger, M. A.; Grassi, A. Desinformação on-line e eleições no brasil: a circulação de links sobre desconfiança no sistema eleitoral brasileiro no facebook e no youtube (2014-2020). *Policy Paper*, FGV DAPP, 2020. Cited on page 19.

Sakiyama, K. M.; Silva, A. Q. B.; Matsubara, E. T. Twitter breaking news detector in the 2018 brazilian presidential election using word embeddings and convolutional neural networks. In: *2019 International Joint Conference on Neural Networks (IJCNN)*. 2019. p. 1–8. Cited on page 19.

Santos, L.; Bianchi, R.; Costa, A. Finbert-pt-br: Análise de sentimentos de textos em português do mercado financeiro. In: *Anais do II Brazilian Workshop on Artificial Intelligence in Finance*. Porto Alegre, RS, Brasil: SBC, 2023. p. 144–155. ISSN 0000-0000.

Available on: <https://sol.sbc.org.br/index.php/bwaif/article/view/24960>. Cited on page 27.

Souza, F.; Nogueira, R.; Lotufo, R. Bertimbau: Pretrained bert models for brazilian portuguese. In: ______. 2020. p. 403–417. ISBN 978-3-030-61376-1. Cited on page 27.

Vaswani, A. et al. *Attention Is All You Need*. 2017. Cited 4 times on pages 11, 22, 23, and 24.

# Appendix

# APPENDIX A – Results tables

In this appendix we provide the tables with the results for each model. The column *Negative* is the number of sentences labeled as negative, *Neutral* is the number of sentences labeled as neutral, *Positive* is the number of sentences labeled as positive and *Sentiment Index* is the value of our sentiment index.

Table 1 we present the results for *cardiff* model, in table 2 we present the results for *finbert* model, in table 3 we present the results for *distilbert* model and in table 4 we present the results for *cardiff-FN* model.

| Year | Negative | Neutral | Positive | Sentiment Index |
|---|---|---|---|---|
| 1996 | 71 | 205 | 9 | -0.22 |
| 1997 | 6 | 100 | 2 | -0.04 |
| 1998 | 58 | 149 | 4 | -0.26 |
| 2000 | 143 | 482 | 12 | -0.21 |
| 2001 | 42 | 214 | 7 | -0.13 |
| 2002 | 267 | 649 | 14 | -0.27 |
| 2003 | 36 | 107 | 9 | -0.18 |
| 2004 | 102 | 271 | 14 | -0.23 |
| 2005 | 70 | 164 | 4 | -0.28 |
| 2006 | 182 | 545 | 35 | -0.19 |
| 2007 | 93 | 331 | 17 | -0.17 |
| 2008 | 655 | 1932 | 122 | -0.20 |
| 2009 | 386 | 1193 | 57 | -0.20 |
| 2010 | 716 | 2345 | 159 | -0.17 |
| 2011 | 194 | 630 | 37 | -0.18 |
| 2012 | 574 | 1556 | 75 | -0.23 |
| 2013 | 253 | 781 | 59 | -0.18 |
| 2014 | 788 | 2155 | 112 | -0.22 |
| 2015 | 496 | 910 | 61 | -0.30 |
| 2016 | 912 | 1559 | 63 | -0.34 |
| 2017 | 476 | 999 | 59 | -0.27 |
| 2018 | 1899 | 4179 | 277 | -0.26 |
| 2019 | 450 | 848 | 82 | -0.27 |
| 2020 | 1094 | 2611 | 187 | -0.23 |
| 2021 | 2379 | 3997 | 248 | -0.32 |
| 2022 | 4355 | 8584 | 683 | -0.27 |
| 2023 | 454 | 424 | 44 | -0.44 |

Table 1 – *cardiff* model results.

| Year | Negative | Neutral | Positive | Sentiment Index |
|---|---|---|---|---|
| 1996 | 150 | 132 | 3 | -0.52 |
| 1997 | 25 | 73 | 10 | -0.14 |
| 1998 | 97 | 108 | 6 | -0.43 |
| 2000 | 306 | 317 | 14 | -0.46 |
| 2001 | 85 | 170 | 8 | -0.29 |
| 2002 | 553 | 366 | 11 | -0.58 |
| 2003 | 55 | 93 | 4 | -0.34 |
| 2004 | 220 | 159 | 8 | -0.55 |
| 2005 | 131 | 105 | 2 | -0.54 |
| 2006 | 329 | 408 | 25 | -0.40 |
| 2007 | 184 | 250 | 7 | -0.40 |
| 2008 | 1228 | 1428 | 53 | -0.43 |
| 2009 | 694 | 924 | 18 | -0.41 |
| 2010 | 1284 | 1856 | 80 | -0.37 |
| 2011 | 365 | 477 | 19 | -0.40 |
| 2012 | 1035 | 1121 | 49 | -0.45 |
| 2013 | 438 | 623 | 32 | -0.37 |
| 2014 | 1347 | 1644 | 66 | -0.42 |
| 2015 | 773 | 670 | 25 | -0.51 |
| 2016 | 1465 | 1027 | 42 | -0.56 |
| 2017 | 760 | 747 | 28 | -0.48 |
| 2018 | 3180 | 3070 | 108 | -0.48 |
| 2019 | 705 | 653 | 22 | -0.49 |
| 2020 | 2049 | 1729 | 115 | -0.50 |
| 2021 | 3569 | 2931 | 124 | -0.52 |
| 2022 | 7098 | 6127 | 397 | -0.49 |
| 2023 | 646 | 257 | 19 | -0.68 |

Table 2 – *finbert* model results.

| Year | Negative | Neutral | Positive | Sentiment Index |
|---|---|---|---|---|
| 1996 | 150.0 | 1.0 | 134.0 | -0.06 |
| 1997 | 23.0 | 0.0 | 85.0 | 0.57 |
| 1998 | 97.0 | 7.0 | 107.0 | 0.05 |
| 2000 | 311.0 | 12.0 | 314.0 | 0.00 |
| 2001 | 103.0 | 8.0 | 152.0 | 0.19 |
| 2002 | 512.0 | 21.0 | 397.0 | -0.12 |
| 2003 | 72.0 | 1.0 | 79.0 | 0.05 |
| 2004 | 223.0 | 5.0 | 159.0 | -0.17 |
| 2005 | 131.0 | 8.0 | 99.0 | -0.13 |
| 2006 | 360.0 | 10.0 | 392.0 | 0.04 |
| 2007 | 220.0 | 7.0 | 214.0 | -0.01 |
| 2008 | 1457.0 | 74.0 | 1178.0 | -0.10 |
| 2009 | 830.0 | 45.0 | 761.0 | -0.04 |
| 2010 | 1635.0 | 69.0 | 1516.0 | -0.04 |
| 2011 | 453.0 | 25.0 | 383.0 | -0.08 |
| 2012 | 1197.0 | 82.0 | 926.0 | -0.12 |
| 2013 | 554.0 | 31.0 | 508.0 | -0.04 |
| 2014 | 1659.0 | 80.0 | 1315.0 | -0.11 |
| 2015 | 908.0 | 29.0 | 530.0 | -0.26 |
| 2016 | 1532.0 | 64.0 | 938.0 | -0.23 |
| 2017 | 901.0 | 30.0 | 602.0 | -0.20 |
| 2018 | 3818.0 | 185.0 | 2352.0 | -0.23 |
| 2019 | 813.0 | 31.0 | 536.0 | -0.20 |
| 2020 | 2267.0 | 103.0 | 1522.0 | -0.19 |
| 2021 | 4208.0 | 216.0 | 2200.0 | -0.30 |
| 2022 | 7679.0 | 512.0 | 5431.0 | -0.17 |
| 2023 | 639.0 | 20.0 | 263.0 | -0.41 |

Table 3 – *distilbert* model results.

| Year | Negative | Neutral | Positive | Sentiment Index |
|---|---|---|---|---|
| 1996 | 10.0 | 272.0 | 3.0 | -0.02 |
| 1997 | 4.0 | 104.0 | 0.0 | -0.04 |
| 1998 | 12.0 | 199.0 | 0.0 | -0.06 |
| 2000 | 11.0 | 624.0 | 2.0 | -0.01 |
| 2001 | 7.0 | 254.0 | 2.0 | -0.02 |
| 2002 | 29.0 | 899.0 | 2.0 | -0.03 |
| 2003 | 4.0 | 147.0 | 1.0 | -0.02 |
| 2004 | 11.0 | 374.0 | 2.0 | -0.02 |
| 2005 | 18.0 | 220.0 | 0.0 | -0.08 |
| 2006 | 37.0 | 723.0 | 2.0 | -0.05 |
| 2007 | 15.0 | 425.0 | 1.0 | -0.03 |
| 2008 | 99.0 | 2597.0 | 13.0 | -0.03 |
| 2009 | 59.0 | 1575.0 | 2.0 | -0.03 |
| 2010 | 107.0 | 3098.0 | 15.0 | -0.03 |
| 2011 | 21.0 | 835.0 | 5.0 | -0.02 |
| 2012 | 80.0 | 2118.0 | 7.0 | -0.03 |
| 2013 | 52.0 | 1033.0 | 8.0 | -0.04 |
| 2014 | 152.0 | 2893.0 | 10.0 | -0.05 |
| 2015 | 113.0 | 1346.0 | 8.0 | -0.07 |
| 2016 | 171.0 | 2355.0 | 8.0 | -0.06 |
| 2017 | 95.0 | 1433.0 | 6.0 | -0.06 |
| 2018 | 345.0 | 5987.0 | 23.0 | -0.05 |
| 2019 | 84.0 | 1292.0 | 4.0 | -0.06 |
| 2020 | 195.0 | 3689.0 | 8.0 | -0.05 |
| 2021 | 429.0 | 6160.0 | 35.0 | -0.06 |
| 2022 | 804.0 | 12740.0 | 78.0 | -0.05 |
| 2023 | 82.0 | 839.0 | 1.0 | -0.09 |

Table 4 – *cardiff-FN* model results.