Is the News Always Negative? Using Deep Learning to Track News Sentiment During the Covid-19 Pandemic

By

Gabriel Nicholson

The University of Chicago

May 1, 2023

Advisor(s):
Bryon Aragam, University of Chicago Statistics and Econometrics Department
Bruce Sacerdote, Dartmouth College Economics Department

A paper submitted in partial fulfillment of the requirements for the Master of Arts degree in Computational Social Science

Abstract

The news is notorious for its tendency to focus on negative events, giving rise to the adage, "If it bleeds, it leads." In this study, I examine whether news coverage during the Covid-19 pandemic mainly focused on negative events while downplaying positive developments such as decreasing Covid-19 cases. Utilizing a state-of-the-art fine-tuned language model, I analyzed the sentiment of over 900,000 Covid-19 related news articles from March 2020 to April 2022 across the United States, Canada, and the United Kingdom. The results indicate that the news is far more negative than positive—even when Covid-19 cases and hospitalizations are decreasing. This negativity is most pronounced in Op-Ed articles, front-page news articles, and articles published by large news organizations (e.g., New York Times, BBC, Fox News). However, non-Op-Ed news articles do become more positive as Covid-19 cases decrease, contradicting previous research findings. These discrepancies can be attributed, in part, to differences in model accuracy, as the model I trained is approximately 20% more accurate than other models used in the literature. Further, when dividing U.S news by the publisher's political ideology, clear differences emerge: both left-wing and right-wing sources are much more negative than centrist news sources. Surprisingly, these differences in sentiment are about as large as the difference between regular news and Covid-19 news sentiment, indicating substantial differences in news reporting across political lines. These findings provide new insights into news reporting patterns during the pandemic and carry important implications for public health messaging and news reporting practices.

1 Introduction

Anyone who kept up with the news during the Covid-19 pandemic would have likely concluded it was almost always negative. But no human has the time to analyze hundreds of news articles on a daily basis, so our subjective assessment of the news's negativity is limited to extremely small sample sizes. It is exactly these situations where machine learning models thrive by scaling the sentiment analysis process beyond human reading capabilities. This is especially important for analyzing the news media environment since only with a large sample of news articles can any serious claims be made about how the news environment, as a whole, responded to the Covid-19 pandemic.

To address this question, I trained a neural network language model on 1150 news article samples related to Covid-19. This model was then used to predict the sentiment of news articles originating from the United States, Canada, and the United Kingdom. Using these predictions, I examined the relationship between the news's daily sentiment, the Covid-19 positivie rate, and the number of Covid-19 hospital admissions. To the best of my knowledge, this article presents the first assessment of how news sentiment changed in relation with health metrics throughout the entire pandemic: testing the hypothesis that the news has a bias toward negativity, even when health statistics are improving.

The pandemic provides a unique opportunity for tracking news sentiment as it was characterized by a series of both positive and negative events. The bad events are all too obvious with over a million lives lost and numerous individuals still affected by long Covid. On the other hand, there were remarkable advancements in vaccine development, as well as successful efforts to reduce Covid-19 cases through vaccination, social distancing, and natural immunity. Understanding how news outlets reported on such a diverse array of events is not initially obvious, highlighting the need for a more in-depth examination.

An important question to consider is what fair news coverage would look like in terms of the news's sentiment. I argue that equitable news coverage should display a sentiment trend that correlates, at least to some extent, with objective health metrics, and that the strength of this correlation should be consistent during both positive and negative events. For example, we would expect news articles, *on average*, to adopt a more positive tone when Covid-19 cases are declining and a more negative tone when cases are on the rise.

To give an example of where this can go wrong, assume during the sharp decline in cases after the Omicron wave, news articles reported almost exclusively on the few cities with increasing Covid cases, while neglecting to mention the overall decline in cases across the country. In such a scenario, the articles would not only fail to correlate with health metrics, but also present an inaccurate representation of the current state of the pandemic to their audience.

The primary objective of my research question is to examine whether such reporting patterns occurred during the pandemic.

Some may argue that overly negative news reporting during the pandemic is not a bad thing, and in fact, may be warranted under a "better safe than sorry" philosophy. However, there is strong evidence from cognitive psychology that suggests being exposed to lots of negative news is detrimental to both our well-being and having an accurate understand of the world (McLaughlin, Gotlieb, and Mills 2022). For instance, negative news has been shown to increase stress, lower trust in news sources, and lead to unrealistic beliefs about the world (Association 2017; Johnston and Davey 1997; Slovic Paul 1987). It may not be a coincidence that the sharp increase in negative news stories starting around 1990, coincides with a general increase in news consumption and a decrease in U.S happiness that is still continuing to this day (Twenge 2019). The CDC themsleves recommended to "Take breaks from news stories," during the Covid-19 pandemic to help relieve stress (*Coping with Stress* 2023). Moreover, trust in U.S. media sources reached an all-time low in 2021, partly attributed to high negative news consumption (Brenan 2021). Buneviciene et al. (2021) demonstrated that individuals exposed to the most negative news during the pandemic became less interested in Covid-19 news and less likely to consult healthcare professionals for Covid-19 information.

Given the importance of trust in news sources and being well-informed during a pandemic, analyzing the relationship between news sentiment and health statistics can provide valuable insights into the strengths and weaknesses of news reporting patterns. If news outlets persistently highlight negative events and overlook positive developments when they arise, audiences may become desensitized and less likely to respond appropriately to future viral outbreaks—making news sentiment a potentially life-altering factor that warrants careful consideration.

2 Literature Review

2.1 News Negativity

If you are a regular news reader, you probably noticed doom and gloom dominating headlines even before the Covid-19 pandemic. This tendency is not an artifact of the internet or of contemporary society; it can be traced back to as early as 1965 when a team of researchers noticed that negativity was a core component of what made a story "newsworthy"(Galtung and Ruge 1965). On a broader scale, the news has been growing increasingly gloomy since around 1945, with a marked decline in positive news stories after 1990 (Leetaru 2019). This is a surprising finding consdering that during the same time that the news has become more depressing, societies across the globe have been consistently making progress across metrics such as crime, health, income, relative inequality, human rights, and technological innovation

(Pinker 2018). Hence, it is well established that the news has a proclivity toward focusing on negative news stories.

The exact causes of the increase in negative news within the last few decades are not well-established, whether it stems from increased consumer demand, shifting incentives for news reporters, or a combination of the two (S. N. Soroka 2006). Not everything is unknown, though; recent research from Cognitive Psychology helps explain the demand for negative news. For example, a highly replicable finding is that there is an innate tendency to focus more on negative information—a predisposition that likely proved advantageous for our ancestors' survival (Pinker 2002; Shoemaker 1996; Tversky and Kahneman 1973). The reason is simple: negative information often carries life-or-death consequences, while positive information does not. Disregarding a threat of violence could prove fatal, while ignoring uplifting stories about acts of kindness would not have the same dire consequences. Our preference for negative information is also reflected in how we engage with news stories. S. Soroka and McAdams (2015) showed that people are more focused when reading negative news stories compared to positive ones and could better retain their content. Other research has also shown that even when participants said they wanted to read more positive news, they still gravitated toward reading the negative stories (Trussler and S. Soroka 2014).

So it should not be surprising that media publishers have catered toward their readers demand for negativity: S. Soroka and McAdams (2015) also found that magazine sales increased by around 30% when the front-cover was about something negative. A Russian based newspaper "The City Reporter" decided to run an experiment where they only published positive news stories for a day, as a result, their daily viewership dropped by around 67% (News BBC 2014). In summary, there are strong reasons to believe that negative news sells more than positive news.

There have been some recent studies that analyzed news sentiment during the pandemic for insights into health measure responses (Krawczyk et al. 2021; S. Liu and J. Liu 2021); economic outcomes (Federal Reserve Bank of San Francisco et al. 2020); political polarization (Asthana 2021); and country-level differences in Covid news reporting (Ahmad et al. 2022; Sacerdote, Sehgal, and Cook 2020; Singh, Jakhar, and Pandey 2021). The prevailing consensus is that news reporting during the pandemic has been more negative than positive—although to varying degrees depending on the study. Yet none of these studies explicitly analyze the relationship between news sentiment and objective health outcomes. It remains unclear whether the news was predominantly negative during periods of decreasing Covid-19 cases, and whether this varies by political ideology, country, or the type of news organization.

2.2 Sentiment Analysis

Sentiment analysis algorithms have been around since at least the 1990s, which until 2018, have mostly been dictionary-based approaches along with supervised machine learning algorithms such as support vector machines (SVM) (Yadav and Vishwakarma 2020). Dictionary based approaches work by assigning a fixed sentiment score to specific words. When given an input, it simply counts the sentiment scores associated with each word and then takes a weighted average to make its sentiment prediction. The advantage with this approach is interpretability and computational efficiency. It is highly interpretable because the researcher can see exactly how and why the model makes a prediction. The computationally efficient part is because all it needs to do is iterate over words and calculate a weighted sum. The downside with this approach is that it does not take into account the context surrounding a word and completely fails with negation. For example, the statement "I got Covid, but I'm asymptomatic" would likely be classified as negative. Human readers can clearly tell this would be a mistake. The word "but" clarifies that the person is okay, and as a result, we would subsequently discount the fact that they got Covid and interpet the statement as being positive or neutral.

Supervised learning algorithms such as SVM use more complicated classification rules by taking advantage of non-linear patterns and higher order interactions in the text data (Hartmann et al. 2019). These models have higher accuracy than dictionary models on average, but as a cost, they are less interpretable. Similar to how linear regression models become less interpretable once you start adding in quadratic and interaction terms.

With the advances in natural language processing over the last five years, there are deep learning models that go far beyond previous approaches. These models typically go by the name of "transformers" and were proposed by Vaswani et al. (2017). The key innovation of the transformer model is its self-attention mechanism, which allows the model to weigh the importance of different parts of the input sequence when generating the output sequence. This self-attention mechanism enables the model to take advantage of long-range dependencies and complex relationships between elements in the sequences, making it highly effective for various natural language processing tasks.

Transformer models have millions of parameters and require millions of training samples before becoming effective. But once these models are trained, they can be adapted to a number of different specialized tasks simply by replacing the output layer and retraining them on a few hundred annotations (Heitmann et al. 2020; Lin et al. 2022). This is why the model used in this paper is effective at predicting hundreds of thousands of news articles despite only being trained on 1150 annotations.[1]

RoBERTa is one such transformer model that was developed by researchers at Facebook and Washington University in 2019 (Y. Liu et al. 2019). It has shown strong results on standardized benchmarks such as the famous SS2 and GLUE datasets, and outperforms many other transformer models for sentiment classification. Transformer models have been shown to have around 20% higher accuracy on average compared to lexicon-based approaches (Heitmann et al. 2020). But transformer models are not a panacea—they achieve greater accuracy at the cost of a complete loss in interpretability. The reasons why a specific piece of text is classified as positive or negative is not provided, the model simply makes the prediction.

Further, the majority of previous studies that used sentiment analysis on Covid-19 related news topics use either dictionary-based approaches or transformer models that were not fine-tuned on the specific task at hand. Both approaches are not as powerful compared to a fine-tuned transformer model. When analyzing hundreds of thousands of articles, an accurate sentiment classifier is critical for having reliable results, especially given how ambiguous and noisy sentiment classification tends to be (Heitmann et al. 2020).

3 Methods

3.1 Covid-SiEBERT Neural Network Model

The neural network language model used to classify article sentiment is based on "SiEBERT," a fine-tuned checkpoint from the original RoBERTa-Large model (Heitmann et al. 2020). SiEBERT has 24 layers and over 500 million parameters. The advantage of using SiEBERT as a foundation for my own model is because it was trained on a variety of news sources, whereas BERT and many other popular large language models were trained on less relevant text sources such as social media data or online reviews. News articles tend to be more direct than social media text and possess other distinct semantic properties. However, since SiEBERT was not explicitly trained on Covid-19 news articles, it does not inherently understand the sentiment of Covid-19 related topics. For instance, when asked to predict the sentiment of "Covid cases are increasing across the country," it fails to classify the statement as negative. Covid-19 did not exist when the model was initially trained so it never showed up in the training data. To address this issue, I further fine-tuned the SiEBERT model by adding an additional linear layer and classification layer, and then retrained the model on 1150 hand-picked annotated article statements related to Covid-19. After retraining, the new model could accurately classify

[1] Part of the reason why they are called "transformer" models is because of their ability to generalize to a wide range of sequence-based tasks.

Covid related sentiment.

Throughout this paper I refer to this fine-tuned version of SieBERT as "Covid-SiEBERT." Now the statement, "*Thirty-seven states are seeing sustained reductions in Covid cases*" is classified as positive, and "*The decision comes as a warning sign to other Western countries, including the U.S., that had hoped to put the pandemic behind them thanks to successful vaccination campaigns*" is classified as negative.

All transformer language models have a limit on the input size they can process, in the case of Covid-SiEBERT, that limit is 512 characters. Since most news articles exceed 512 characters, each article was truncated to 18 sentences in length. These sentences were then further divided into five groups each containing three sentences. Once the model reads in these three sentences as input, it outputs either a 1 if it predicts the sentences are positive (-1 if negative) and a confidence score in its prediction (between 0.5 and 1). When the model encounters a neutral statement or unfamiliar input, it is usually biased toward predicting them as positive, and assigning a confidence score below 0.85. In total, less than 2% of predictions had a confidence score below 0.85, and removing them did not alter any of the conclusions.

Using the entire article is computationally expensive, and previous work has found that using only a small subset of the article is sufficient to accurately classify the general sentiment of an article (Heitmann et al. 2020). This is not surprising as writers typically present their main argument in the introduction to capture their reader's attention. An interesting finding that corroborates this claim—as well as the hypothesis that negative news sells more—is that the average negativity of an article is highest at the beginning and monotonically decreases within each block of three sentences (Table 2).

3.2 Validating Model Performance

Table 1: Classification Report for Covid-SiEBERT and Other Models

	F1 Score	Accuracy Score	Test Set Sample Size
Covid-SiEBERT	0.961	0.956	245
Base SiEBERT	0.712	0.722	245
VADER	0.642	0.644	245
Distil-BERT	0.621	0.633	245
Covid-RoBERTa (3 class)	0.853	0.892	148

The performance of Covid-SiEBERT, along with five other sentiment classifiers commonly used in the literature was assessed using an out-of-sample test set (N=245) created by randomly

Table 2: Average Sentiment Score for Each Sentence Group

Sentence Group	US	Canada	UK	US Regular News	Canada Regular News	UK Regular News
1-3	-0.260	-0.258	-0.308	-0.054	-0.039	-0.042
4-7	-0.228	-0.193	-0.244	-0.020	-0.010	-0.019
8-10	-0.214	-0.208	-0.200	-0.007	0.014	0.011
11-14	-0.193	-0.188	-0.185	0.004	0.061	0.030
15-18	-0.174	-0.148	-0.183	0.033	0.065	0.074

sampling articles from the main dataset. The labels were manually annotated and double-checked with an additional reviewer for accuracy. The Covid-SiEBERT model outperforms the other binary classification models by around 20 percentage points, boasting an impressive 96.1% accuracy and 95.6% F1 score (Table 1). K-fold cross-validation was also used during the initial training of Covid-SiEBERT ($k = 5, N = 230$). The test set results are comparable to the accuracy and F1 scores obtained through cross-validation, indicating that the model was not overfit.

To further validate the Covid-SieBERT model, a different model with the ability to classify positive, negative and neutral statements was also fine-tuned using 300 different training samples. This model was built from the origianl RoBERTa-base model which is a smaller model than SiEBERT (230 million fewer parameters) and not as accurate.[2] The drop in accuracy is to be expected since binary classification tasks are strictly easier than three-class prediction problems. Comparing the sentiment trends of these two models can help verify that neutral statements are not biasing the predictions from Covid-SiEBERT.

The stark difference in classification accuracy shown in Table 1 suggests that the predictions made using Covid-SiEBERT are more reliable than previous models used in the literature. To illustrate the impact of a 20% difference in accuracy, the sentiment predictions from both VADER and Distil-BERT are also computed on the same data and compared to Covid-SieBERT (Figure 4).

3.3 Sentiment Score Calculation

Predicting the sentiment score for a single article is done by taking an average over all five predictions (one for each of the five sentence blocks). For example, if the model predicted the first 3 sentence blocks as negative and the last two blocks as positive, then that article's sentiment would be equal to $-\frac{1}{5}$. From here there are two ways to calculate the news's sentiment in a given day using these individual article predictions.

[2] SiEBERT was only trained on binary labels, so it would be impossible to change the model to accommodate three-class predictions without requiring many more training samples.

The first method is to take an average over all the predictions within an article, and then taking an average over each of these predictions (weighting each article equally). The downside of this approach is that the daily sentiment score is hard to interpret. For example, a sentiment score of -0.1 does not necessarily mean there were more negative articles than positive articles that day, it could be that the majority of the articles are slightly positive but the negative articles are very negative, resulting in a negative sentiment score for that day.

The second approach solves this problem by classifying each article as being negative, neutral, or positive, using a hard cutoff rule. If the mean sentiment of an article is greater than 0.1 it is classified as positive, if it is below -0.1 then it is negative, otherwise it is classified as neutral:

$$\gamma_i(\bar{x_i}) = \begin{cases} 1 & \text{if } \bar{x_i} > \frac{1}{6} \\ -1 & \text{if } \bar{x_i} < -\frac{1}{6} \\ 0 & \text{otherwise} \end{cases} \tag{1}$$

Where γ_i is the i-th article sentiment classification and $\bar{x_i}$ is the original sentiment score for article i. In the example above where the article had an average sentiment of $-\frac{1}{5}$, it would be classified as neutral and given a value of 0. Calculating the news sentiment on a single day can be done by simply taking an average over all γ_i.

With this approach, a daily sentiment score equal to -0.10 does mean that there were more negative articles than positive articles that day.[3] One of the potential downsides of this approach is if many of the articles have a sentiment score of $|\frac{2}{5}|$, then that would skew the overall trend since $\bar{x_i} = |\frac{2}{5}|$ and $\bar{x_j} = |\frac{5}{5}|$ are weighted the same. Thankfully, from Figure 10, we can see this is not the case for any of the countries. Also, the correlation between the two sentiment score methods is greater than 0.98 for each country, so the trends themselves do not differ by much. Throughout the analysis, the first method is typically used, while the second method is employed whenever a claim about the number of positive or negative articles being published is made.

3.4 Time-Series Comparisons

When comparing time-series, it is important to ensure both trends are stationary before analyzing the data to avoid spurious correlations. In my analysis, I only compute correlations after transforming the series to be stationary, which is done by differencing the series. Both

[3] It could be possible that there are more neutral articles than either positive or negative articles, but within this analysis, neutral article classifications are always a minority, so the edge case where there are more neutral than positive and negative articles never occurs.

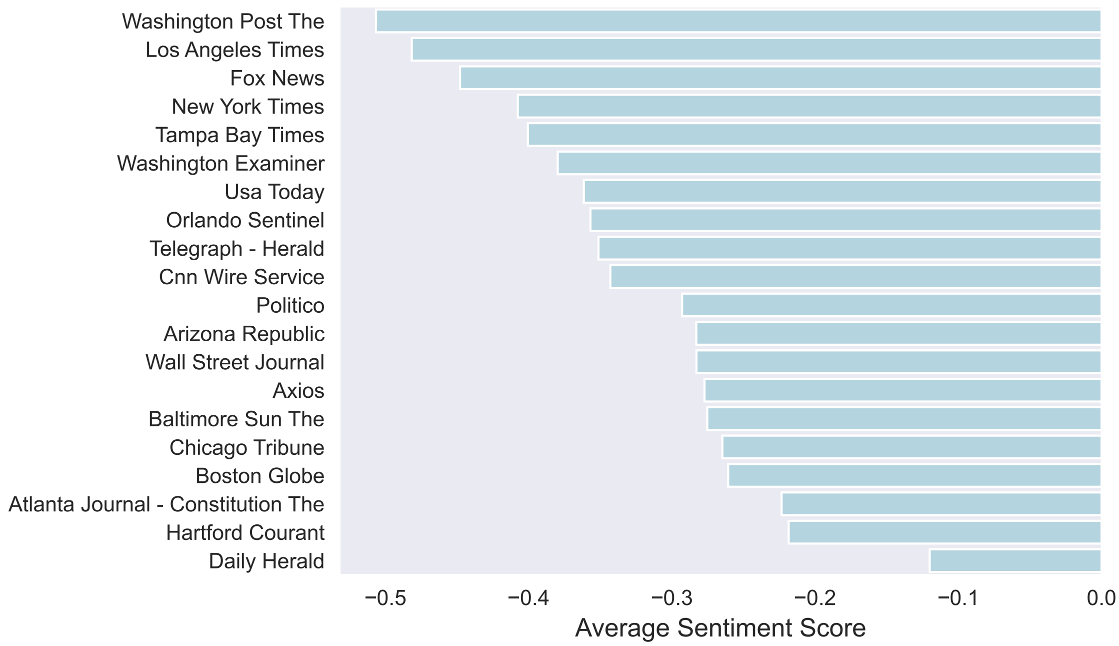

Figure 1: Top 20 U.S News Sources

the Kwiatkowski-Phillips-Schmidt-Shin (KPSS) and Augmented Dickey Fuller (ADF) tests are used to verify that the trends are stationary. Finally, Spearman's rank correlation (denoted as r_s), is used to correlate the series. This method is favored over the Pearson correlation since it can accommodate non-linear trends and outliers while also not assuming the data are normally distributed.

4 Data

4.1 News Article Data

The news article data is from ProQuest TDM Studio which owns the copyrights to thousands of different news publications. Figure 1 shows the 20 highest sample size U.S sources from the ProQuest dataset and their associated average sentiment score across the entire pandemic timeline. In the United States, there are around 2.2 million articles between March 2020 and April 2022. After filtering these articles using Covid-19 keywords, 325,102 articles remained. Table 3 shows a summary of the number of articles associated with each country and their average sentiment score. The filtering technique is a keyword matcher that takes into account the number of keywords an article has, and the length of the article (longer articles required more keywords to be kept). Some of the keywords are “covid”, “omicron”, “delta”, “pandemic”, and “vaccine”. The complete preprocessing steps can be found in the GitHub repository (link is at the bottom of the paper).

News articles unrelated to Covid-19 were also collected and are used as a reference point for comparing the sentiment of Covid-19 articles. Throughout my analysis, I refer to these as "regular news" articles which typically cover subjects such as the economy, politics, sports, and local news stories.

The filtering process was validated by randomly sampling approximately 200 articles and checking if they were relevant. The limitation of this approach is that the false positive rate is almost certainly not zero and so some irrelevant articles are bound to make it through. However, this should not ruin any conclusions, and would likely only bias the sentiment score upwards (Covid articles tend to be the most negative articles) along with adding extra noise. This was inadvertendly validated during the analysis process after discovering 7500 spanish articles that went past the filtering process; after removing these articles, no results changed, suggesting that the trends are robust and have a high tolerance for false positive articles.[4]

Finally, the 1150 annotations used to train the Covid-SieBERT came from a variety of hand-picked news articles and can be found on the GitHub page. They were reviewed by two additional annotators to ensure reliability. There was over 94% overlap in reviewer annotations, making them a reliable source for ground truth sentiment classifications. The test set was created using the same process.

4.2 Publisher Classification

News publishers were categorized as either "national" or "regional" depending on how popular they are within their country. A publisher is considered national if they rank among the top 20 most viewed news networks, the remaining publishers are classified as regional. These terms should not be taken literally, as some publishers focus more on global news than others. The terms are really meant to contrast large news organizations with medium to small ones.

Publishers were also classified as being left-wing, centrist, or right-wing. The majority of these classifications were done by referencing AllSides.com. Unfortunately proquest does not contain many right-wing news sources. To address this, I manually scraped approximately 7,000 Covid-19 related articles from The Epoch Times and Fox News—two prominent right-wing news sources. After this adjustment, around 8% of the sample are from right-wing sources, 34% from centrist sources, and 58% from left-wing sources.

One of the limitations of this analysis is that it is not clear how generalizable these findings are for making claims about right-wing news sources given the disparity in sample size. But a potential argument against having balanced data would be that many popular news publishers

[4]The language model was only trained on English texts so the model tends to give random predictions when it reads text from another language.

Table 3: Number of Samples per Country and Mean Sentiment Score

	Article Sample Size	Mean Sentiment Score	Standard Deviation
United States	325,102	-0.202	0.129
Canada	328,768	-0.197	0.124
United Kingdom	281,580	-0.224	0.126
US Regular News	253,866	-0.045	0.060
Canada Regular News	185,438	-0.014	0.081
UK Regular News	177,531	0.007	0.066

tend to be left-wing, so this asymmetry is inherent in our current news environment. In this case, having anything close to 50/50 representation in the data would be biased if the goal was to have an unbiased representation of the United State's overall news sentiment.

4.3 Covid-19 Health Statistics

The Covid-19 health data is from the Our World in Data Covid-19 database (Ritchie et al. 2020). The primary metric used in this analysis is the *Covid-19 positive rate* which is defined as "the share of Covid-19 tests that are positive, given as a rolling 7-day average." The positive rate is a good indicator of the overall health of a country since it is not affected by the number of tests being conducted. For example, if a country is conducting more tests, then the number of positive cases will increase, but the positive rate will remain the same.

5 Results

The United States's news sentiment between April 2020 to April 2022 is displayed in Figure 2, along with the Covid-19 positive rate and the regular news sentiment trend. The sentiment trends were smoothed using a LOWESS non-parametric model (the prediction interval was also calculated from the LOWESS model with $\sigma = 2$). From this Figure it is clear that Covid-19 news is much more negative than regular news during almost the entire timeline with a mean absolute difference of 0.16 between their sentiment scores. The only exceptions to this pattern are around May 2021—when Covid cases reached an all-time low and the vaccine rollout peaked—and March 2022, which marked the end of the Omicron wave. Even during these periods of decline, the news was more negative than positive, with average sentiment scores of -0.05 and -0.14, respectively.

Figure 2 also reveals a negative correlation between the Covid-19 positive rate and news sentiment trend ($r_s = -0.39$). As Covid cases decrease, the news becomes more positive,

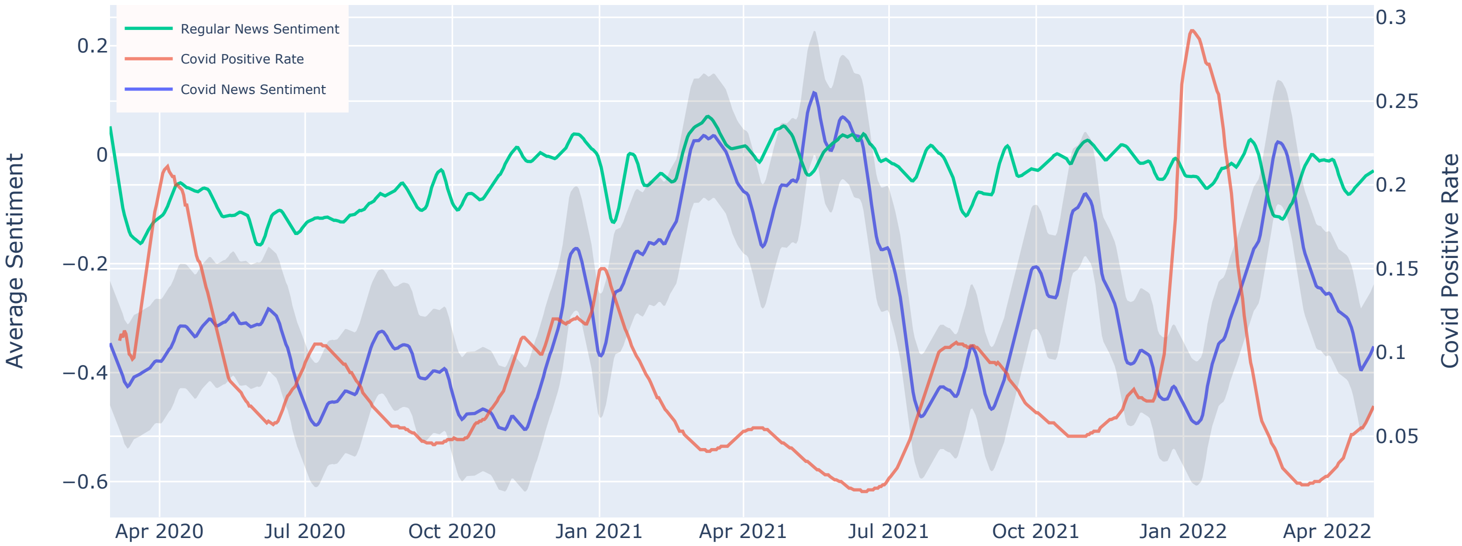

Figure 2: United States Sentiment Score Trend

and conversely, as cases increase, the news becomes more negative. Nonetheless, even during periods of sharp declines in Covid-19 cases, the news continues to publish more negative than positive articles, and is usually more negative than regular news articles.

5.1 Op-ed Article Comparisons

Figure 3 separates the previous sentiment trend line into Op-Ed articles and non-Op-Ed articles. Op-Ed articles are much more negative than non-Op-Ed articles with an average sentiment score of -0.48 compared to -0.26, and have fewer fluctuations. They also have no discernible correlation with the trend in Covid-19 cases ($r_s = -0.02$). In contrast, non-Op-Ed articles demonstrate a much stronger correlation ($r_s = -0.42$).

To test if the lack of correlation was due to the small sample size of Op-Ed articles, non-Op-Ed articles ($N = 307,242$) were resampled to match the sample size of Op-Ed articles ($N = 17,860$), and the correlation was recalculated; this process was repeated $10,000$ times. The result is that the strength and direction of the correlation remain robust despite the reduced sample size ($r_s = -0.38 \pm 0.05$).

There are also interesting visual trends between the two types of news articles. Prior to February 2021, Op-Ed and non-Op-Ed articles moved in tandem, with only a 0.10 difference in their average sentiment. Only when Covid cases begin to decrease after February 2021 do the two trends completely diverge and their mean difference more than doubles to 0.21. Unlike non-Op-Ed articles, Op-Ed articles did not become more positive during the decline in Covid cases over the 2021 summer or at the end of the Delta and Omicron waves. However, the two

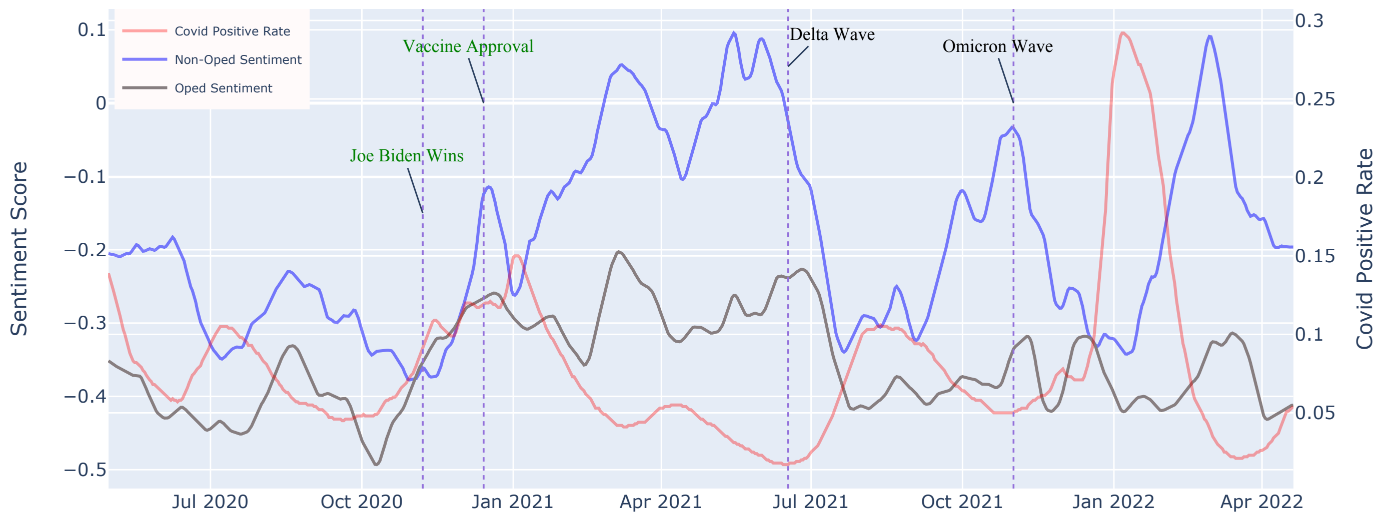

Figure 3: Sentiment Trend for Op-Ed and Non Op-Ed Articles

trends do realign themselves when Covid cases are high such as at the start of the Omicron and Delta waves, suggesting that they publish similar stories only during the worst of times.

5.2 Comparing Different Models

As discussed in the literature review, most studies using sentiment analysis for Covid-19 related topics used dictionary or transformer models that were not fine-tuned for their specific tasks. Figure 4 displays the sentiment trends from the two most popular models, VADER and Distil-BERT, alongside my own model, Covid-SiEBERT. There are clear differences among the three models: Distil-BERT predicts the news to be very negative with an average sentiment score of -0.52 while VADER leans more neutral, predicting an average sentiment score of -0.05. Both VADER and Distil-BERT exhibit less variability than the fine-tuned model, as evidenced by their relatively flat slopes between April 2020 and April 2022. This difference is also reflected in their correlations with the Covid-19 positive rate: VADER ($r_s = -0.04$), Distil-BERT ($r_s = -0.11$), and Covid-SiEBERT ($r_s = -0.39$).

Not only are the three model's sentiment trends different, but the empirical conclusions drawn from each model are also markedly distinct. The Distil-BERT model predicts that news sentiment has a small negative correlation with the Covid-19 positive rate and that the news is extremely negative. While VADER's news sentiment trend shows no correlation with the Covid-19 positive rate and predicts the news to be mostly neutral while becoming more negative over time. Both of these conclusions are incompatible with Covid-SiEBERT's predictions: the news is mostly negative and has a moderately strong negative correlation with the Covid-19 positive rate.

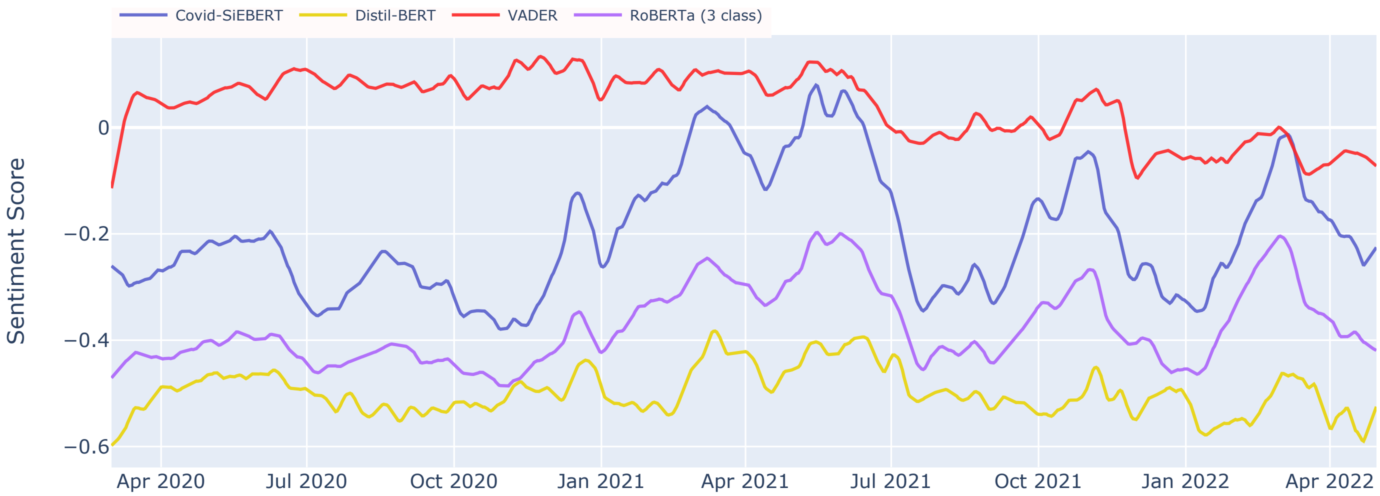

Figure 4: Comparing News Sentiment Trends Across Models

These results underscore the importance of fine-tuning a model for a specific task, since it is very likely that the Covid-SiEBERT model is capturing a more accurate representation of the news's true sentiment trend, given its superior accuracy.

Additionally, Covid-RoBERTa's predicted sentiment trend is also shown. Both models exhibit almost identical trends, with the only difference being a level discrepancy in their sentiment scores. This similarity is reassuring, as the Covid-RoBERTa model was used to test whether Covid-SiEBERT was systematically misclassifying neutral statements. If this was the case, then we would expect the two predicted trends to be different.[5]

5.3 What Makes a Story "Newsworthy"

The articles a news publisher features on the front page betray what they perceive to be the most important story—or at least the story that will quickly grab the readers attention. It turns out that front-page news articles are the most negative, corroborating the previous literature that negative news sells. Table 4 displays the average sentiment score associated with each page number for both Covid and regular news articles. For all three countries, there is a clear trend where front-page articles are the most negative. Figure 5 quantifies the difference between the average sentiment of front-page articles and all other pages. The distribution was created by bootstrapping the samples and recomputing the mean each time. The p-values were computed using a permutation test.

Furthermore, when narrowing the analysis to times when Covid cases are decreasing, such

[5] It is important to note that the level difference is not relevant since the sentiment score trends cannot be interpreted in the same way. For instance, in the three-class case, a sentiment score of -0.2 could imply more positive articles than negative ones, depending on how many articles are classified as neutral.

Table 4: Sentiment Score by Page Number and Country

Page Number	US	Canada	UK	US Regular News	Canada Regular News	UK Regular News
1	-0.236	-0.246	-0.300	-0.139	-0.159	-0.152
2	-0.179	-0.114	-0.205	-0.010	-0.008	-0.032
3	-0.183	-0.151	-0.147	-0.020	-0.061	0.083
4	-0.212	-0.199	-0.232	-0.108	-0.134	-0.081
5	-0.194	-0.197	-0.207	-0.056	-0.130	-0.116
6	-0.210	-0.182	-0.235	-0.119	-0.115	-0.072

as the summer (January 2021 to June 2021) and during the decline in cases after Omicron (January 2022 to March 2022), the front-page articles still have a negative sentiment score (-0.08 and -0.07, respectively). This is true even though the other pages have an average sentiment score that is positive (0.04 and 0.01, respectively)—indicating that the news is publishing slightly more positive than negative articles on these other pages. This demonstrates that the news does indeed publish positive articles on the subject of Covid-19; however, they tend to not make it on the front page. Whether this difference in sentiment is because the front-page article is writing about the decline in Covid cases with a negative tone or writing about a completely different topic, is unclear. Given the large number of articles in the sample, both situations are likely occuring to varying degrees.

Front page coverage is important because it is the first thing people read, and in many cases, the only thing people read as they pass by a newsstand or scroll through their phones. But the total number of articles published in a given day is also important. The number of daily articles published are shown in Figure 6 alongside the Covid positive rate. Clearly more articles are being published when Covid cases are at their highest and the least number of articles are published when cases are at their lowest. A striking example of this phenomenon is during the Omicron wave when the number of articles published hit around 550 a day. But once cases begin dropping off—about as fast as they increased in the first place—the number of Covid articles published decreases to about 150 a day (-72% decrease). The same phenomenon occurs during the Delta wave (-56% decrease). In general there are around twice as many articles published during the peak of a wave, compared to two weeks after once Covid cases begin decreasing.

On the one hand, this may seem obvious: increasing Covid cases is "news," and decreasing Covid cases is not. However, one hypothesis about why negative news sells more than positive news is that negative events happen quickly, while positive developments tend to unfold slowly, sometimes taking years to accumulate.

The prediction here is that the news would report positive events if they were to happen

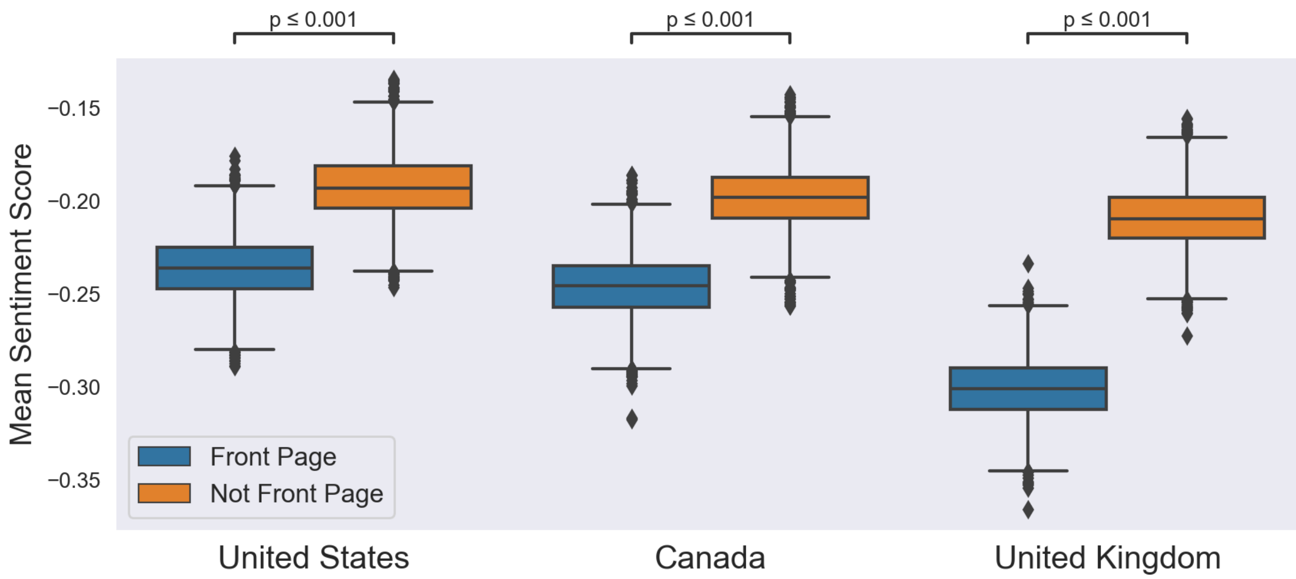

Figure 5: Front-Page Covid News Articles Compared to All Other Pages

fast enough to become emotionally salient—which is already an inherent feature of negative stories. Yet, the Omicron wave example from Figure 6 challenges this hypothesis: Covid cases unexpectedly drop about as fast as they initially increased, but this event did not receive nearly the same amount of coverage. This example underscores that being "news-worthy" is not solely due to being a salient event, but rather that it is a negatively salient event.

These two findings—the number of articles published and the sentiment of front-page articles—support the hypothesis that the news is drawn to report on negative events rather than just emotionally salient events. It also appears that negative Covid stories sell more, as news publishers would not feature these types of stories on the front page if they did not think they were the most eye-catching for their viewers. Since news publishers place more negative stories on the front page during periods of declining Covid cases, this phenomenon cannot be explained by readers' risk perception (i.e., readers demanding more negative news because they perceive their own health risks to be higher and want to be well-informed).

5.4 Differences Across U.S Political Ideologies

The United States is notorious for its toxic divide between political parties. These differences are also apparent from Figure 8, as there are notable differences in their news sentiment trends. Right-wing sources are slightly more negative (-0.27) than left-wing sources (-0.26), and both are much more negative compared to centrist sources (-0.14). Prior to the Delta wave, the difference between left-wing and centrist sources reached as high as 0.19. To put this difference into perspective, regular U.S news and Covid-19 news differ only by 0.16 points (Table 3).

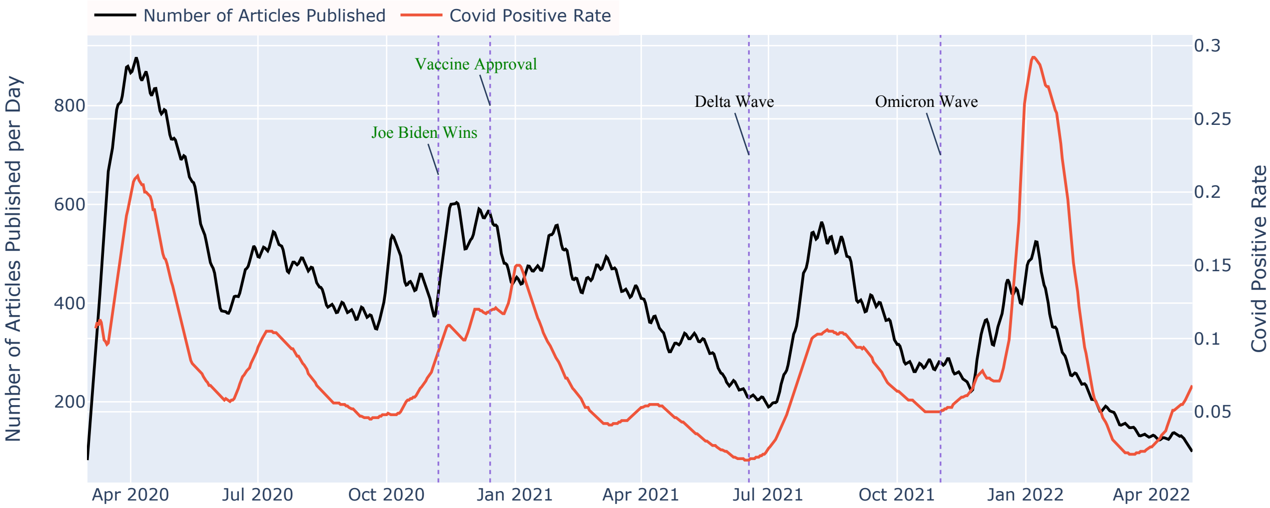

Figure 6: Number of Covid-19 Articles Published per Day

Furthermore, the finding that centrist sources are the least negative does not appear to be unique to Covid-19 news topics, since these results replicate exactly what Rozado, Hughes, and Halberstadt (2022) found when analyzing news headlines between 2000 and 2019. It is disconcerting that there are large differences across political ideologies when the topic being reported on should ideally be apolitical.

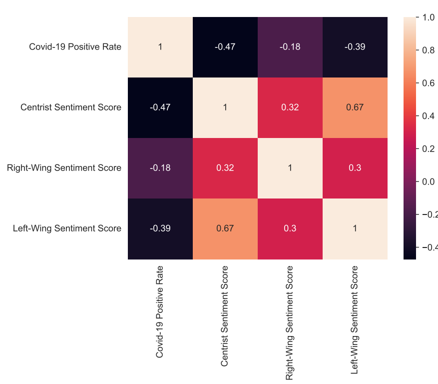

Figure 7: Correlation Between U.S Political Sentiment Trends

Figure 7 shows the correlation between each of the sentiment trends and with the Covid-19 positive rate. Some notable findings are that Centrist sources are the least negative and have the strongest correlation with the Covid-19 positive rate (-0.47). Left-wing sources also have a fairly strong correlation (-0.39) while right-wing publishers having the lowest correlation (-0.18).

Right-wing sources start off very similar to centrist news sources, but then completely diverge once vaccines become approved by the CDC. After that point, they become very negative and fail to correlate with changes in the Covid-19 positive rate. In general, right-wing news sources tend to be more negative than both left-wing and centrist news sources. So it is interesting that they were more positive than left-wing sources during the first 8 months of the pandemic.

Figure 9 highlights the differences in the distribution of sentiment score predictons through-

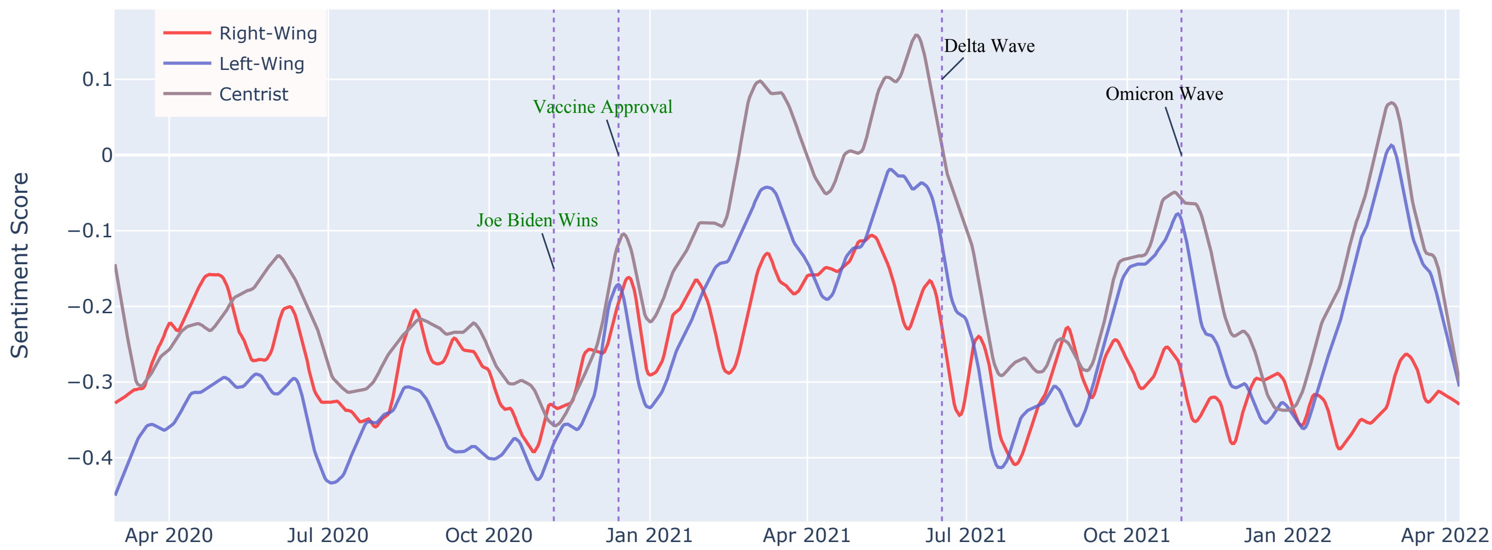

Figure 8: Sentiment Score by Publisher Political Ideology

out the entire pandemic. A striking finding is is that left-wing and right-wing sources did not have a single day where they published more positive than negative articles. In contrast, centrist sources had 90 days where they published more positive than negative articles.

Finally, there are also some notable differences in the number of articles published. During the decline in cases after the Delta and Omicron waves, both right-wing and left-wing news sources had a decrease in the number of articles published, but left-wing sources had around a 20% sharper decrease. Similarly, when the Omicron and Delta waves began, left-wing sources were publishing around 30% more articles compared to right-wing sources. Offering support for a view that left-wing news sources were more focused on negative events during the pandemic compared to right-wing and centrist sources. Whether this difference is caused by viewer demand, or the publishers themselves, is unclear.

5.5 Is the United States's News Enviornment an Outlier Compared to Other Western Countries?

Previous work from Sacerdote, Sehgal, and Cook (2020) used a dictionary model for sentiment classification and found that United States media was much more negative compared to other countries during Covid-19. However, my own analysis using different data and a different sentiment classifier, arrives at a different conclusion (Figure 10). Canada, United States, and the United Kingdom are all very negative relative to regular news, but they have around the same average sentiment score, with the United Kingdom being slightly more negative than the United States (Table 3). It also turns out that U.S news sentiment has a larger correlation with

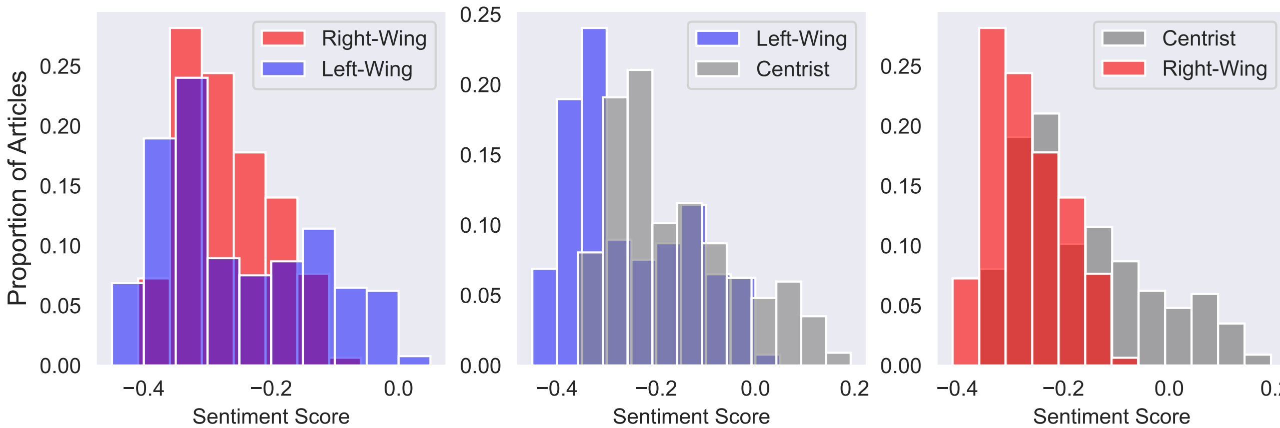

Figure 9: Sentiment Score Distribution by Political Ideology

the Covid-19 positive rate ($r_s = -0.39$) compared to Canada ($r_s = -0.14$) and the United Kingdom ($r_s = -0.25$). So if anything, U.S media is a slight outlier for having a stronger correlation between their article sentiment and changing Covid cases.

Additionally, the correlation between Op-Ed articles from Canada and the United Kingdom and the Covid-19 positive rate is very small. ($r_s = -0.03$ and $r_s = -0.01$, respectively). Canada's Op-Ed articles have a far more negative average sentiment score compared to non-Op-Ed articles (-0.57 vs -0.25). While the United Kingdom's Op-Ed articles have an average sentiment score of -0.61 compared to -0.29 for non-Op-Ed articles. Once again these are practically significant differences given that there is a smaller difference between their regular news and Covid-19 news articles (Table 3). These comparisons offer evidence that U.S Op-Ed articles are not unique for being very negative and not correlating with objective health measures. Nor is the U.S very different for their front-page coverage being the most negative (Table 4).

5.6 National vs Regional News Sources

Assuming there are strong supply and demand pressures for producing negative news, we would expect that national news organizations are more negative than smaller regional sources because of a heightend demand from their larger customer base along with more subject matter to draw from since they are not focused on a single juridisction.

Indeed, this is exactly what we see in the data across all three countries. Both sources tend to have similar sentiment trends, but national sources are consistently more negative during the same time periods. This is clearly visualized by the national sentiment trend line always being below the regional sentiment trend line (Figure 12). In fact, for all three countries, national news sources never publish more positive than negative articles. In contrast, local

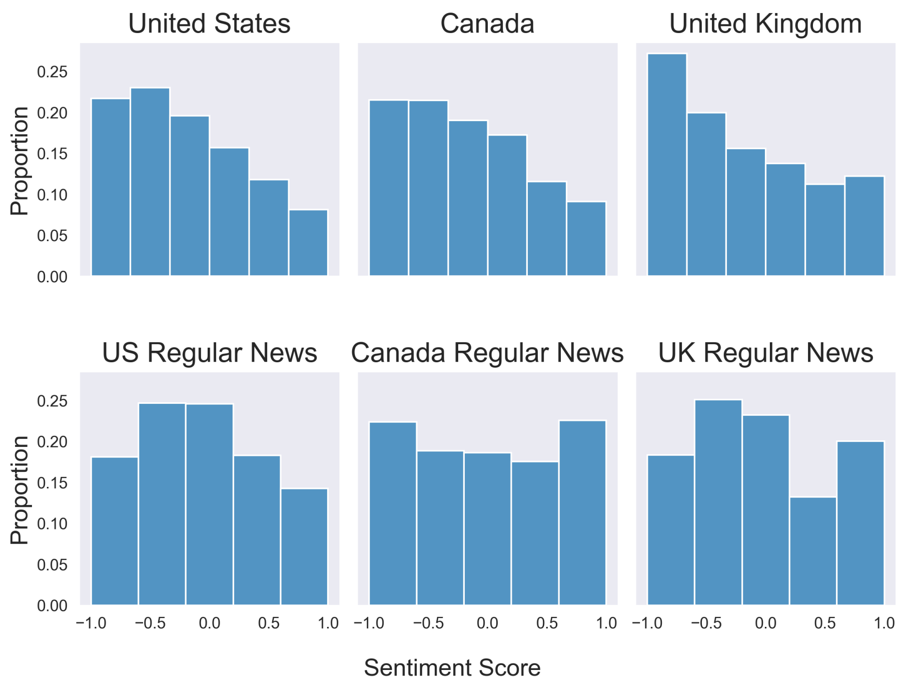

Figure 10: Sentiment Score Distribution by Country

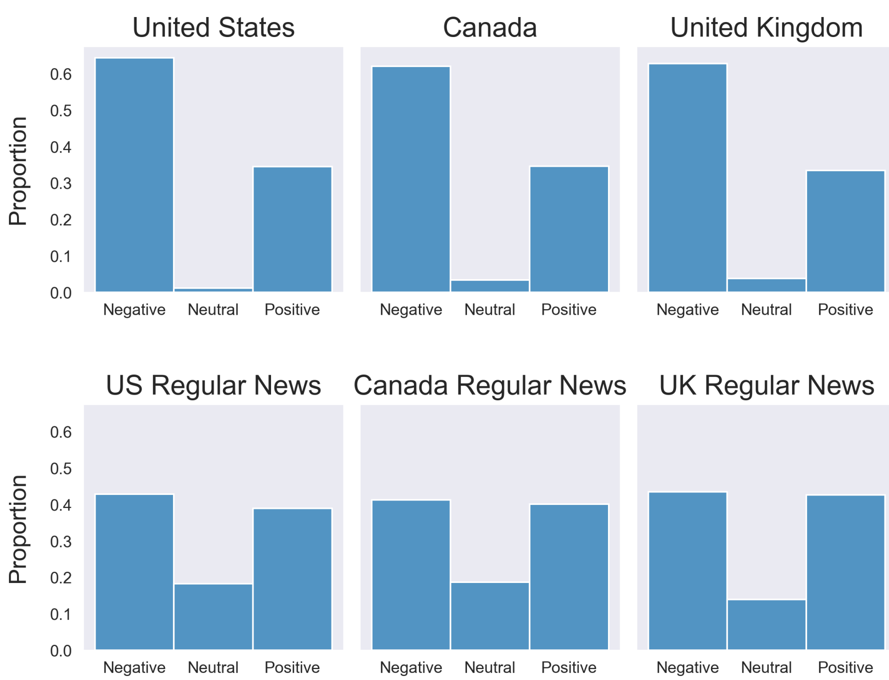

Figure 11: Article Classification Distribution by Country

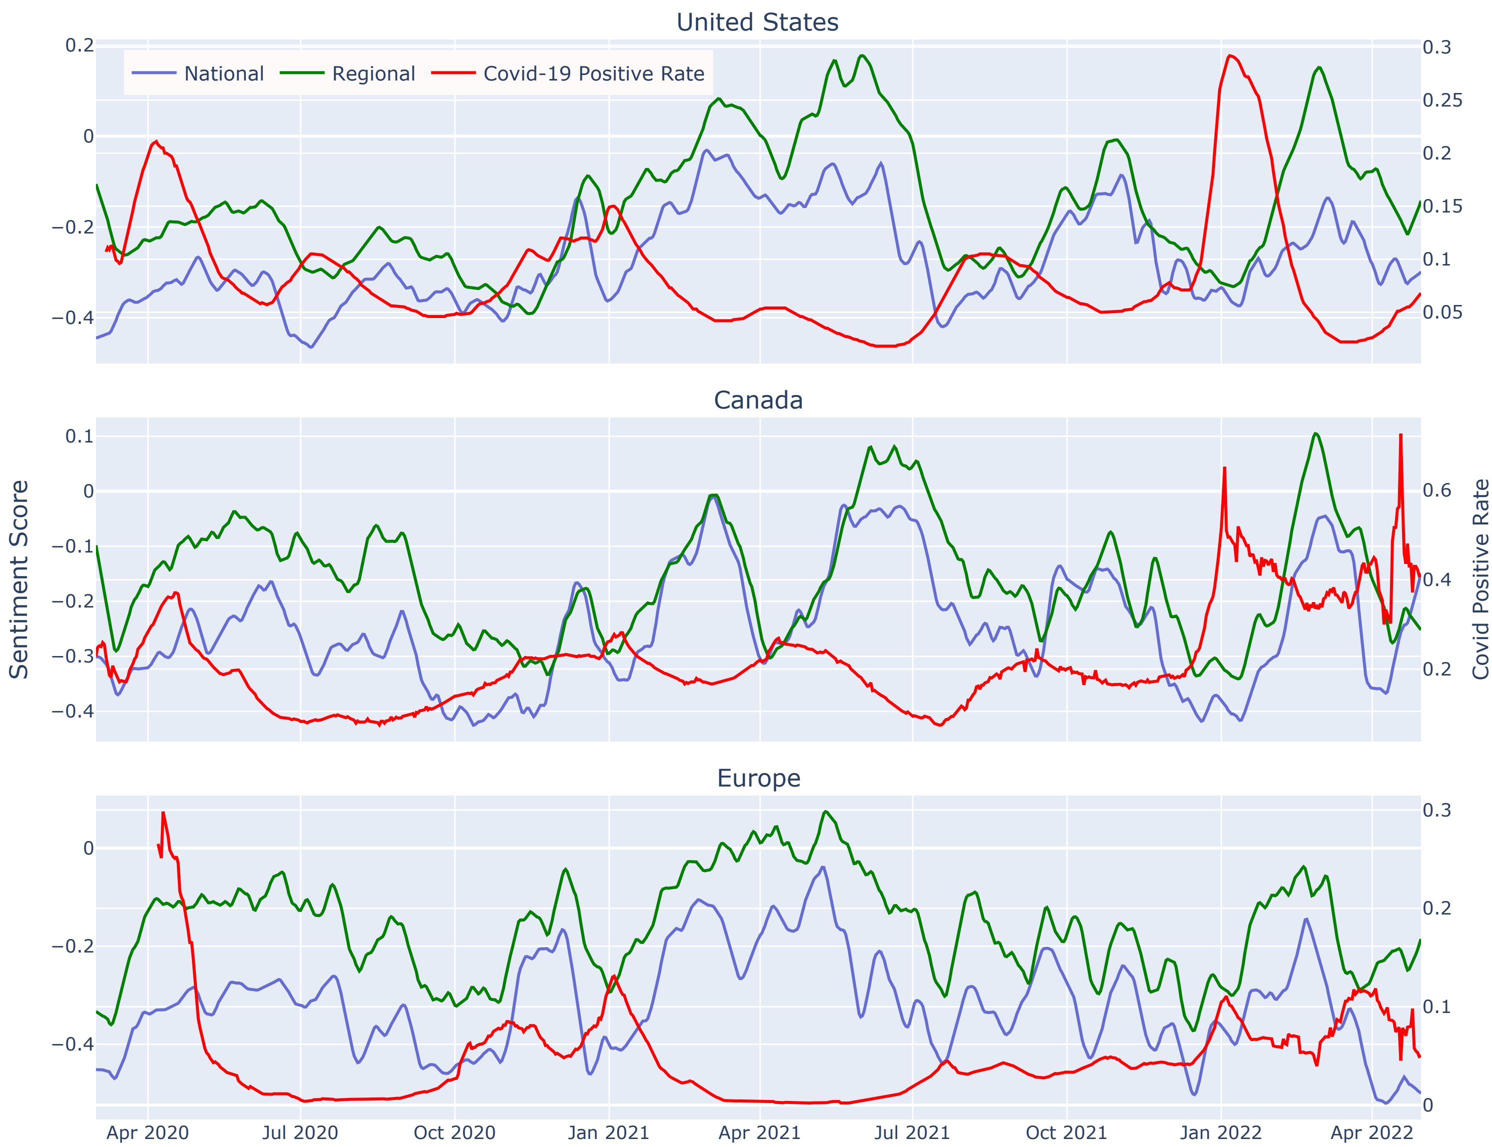

Figure 12: Sentiment Score of National News vs Regional News

Table 5: Summary of National vs Regional News

	United States		Canada		United Kingdom	
	Average Sentiment	Correlation	Average Sentiment	Correlation	Average Sentiment	Correlation
Regional	-0.17	-0.28	-0.16	-0.25	-0.15	-0.31
National	-0.27	-0.41	-0.13	-0.13	-0.26	-0.19

sources have a number of days where their sentiment score rises above zero, typically during times when Covid cases are decreasing.

Table 5 outlines the mean sentiment and correlation with the Covid-19 positive rate across these groups. For both the U.S and the U.K, their local news sources have a stronger correlation with The Covid-19 positive rate and are much more positive compared to national sources.

Canada breaks this trend since both sources have about the same average sentiment and also have a similar correlation. This is also clearly seen from the graph where their trend lines closely overlap. It is not clear why this is the case, but one strong possibility is that Canada does not have a "large" national news organization the same way that the U.S and the U.K do (e.g., BBC, The Guardian, Fox News, New York Times, etc.). If this is true, then Canada might not be a directly comparable case. Suggesting that this sentiment difference between news sources is only prominent for very large news organizations. As far as I can find, there has been no prior research empirically examining this difference, potentially making this a novel finding and a potential avenue for future research.

6 Discussion

The initial research question was if Covid-19 news sentiment correlated with objective health metrics over the course of the pandemic, and if this differed depending on if Covid-19 cases were increasing or decreasing. Based on the results above, the answer is mixed. On the one hand, there is a moderate negative correlation between news sentiment and the Covid-19 positive rate (at least for the United States). But this correlation fails to hold for Op-Ed type articles which are also much more negative than non-Op-Ed news articles. This makes sense since non-Op-Ed news articles have more constraints on what they can report on. Op-Ed articles have no such boundary and can be focused on any topic. It appears that this freedom has motivated publishers to select stories that are relatively more negative, likely caused by viewer demand.

A common theme throughout this analysis is that the kinds of news reporting practices that lead to the most negative news also have the smallest correlation with objective health metrics. We saw this in the case of Op-Ed articles, front-page news articles, centrist vs non-centrist publishers, and national news organizations. Many of these patterns were shared across countries, indicating that this is not an isolated phenomenon, but rather a general feature of news reporting. It is probably not a coincidence that the most negative aspects of the news come from Op-Ed articles and front-page articles, which are also the sections over which news organizations have the most control.

Similarly, when positive events are occuring, there are fewer articles on Covid-19 being published—a pattern that still applies even when Covid-19 cases drop unexpectedly. It also turns out that front-page articles are always more negative than positive, even during times where Covid-19 cases are decreasing. It is not the case that the news does not report on these positive developments, they likely do, but these stories are relegated to less prominent pages and get less viewership. These findings suggest that readers who focus on front-page articles or Op-Ed articles, will likely come away with an overly negative outlook of the current state of the pandemic at times when Covid cases are decreasing. How many people fall into this category is unknown, but it is likely a non-trivial amount. The American Press Institute ran a survey in 2014 and found that 60% of U.S citizens admitted that their weekly consumption of news is mostly from reading headlines.

These results may partly explain the finding that news readers exposed to lots of negative news tend to have less accurate beliefs about the world (Slovic Paul 1987). It would be hard to argue that this same phenomenon could not apply toward Covid-19 news consumers. Worse still, there is some evidence that after listening to lots of negative news about Covid-19, individuals report having "Covid fatigue," a phenomenon where readers are less likely to trust healthcare professionals and tune out Covid-19 news entirely (Buneviciene et al. 2021; McLaughlin, Gotlieb, and Mills 2022). The risk of Covid fatigue makes studying news sentiment important, not just for a stronger understanding of news reporting behavior, but because of its potential to illuminate future risks when media reporting becomes too negative.

Usually the United States is an outlier compared to other western countries when it comes to health outcomes, political division, and economic prosperity. However the evidence outlined in section 5.5 suggests that the United States's news sentiment is not unusually more negative compared to Canada and the United Kingdom. Also, the United States's news sentiment has a noticeably higher correlation with their Covid-19 positive rate compared to Canada and the United Kingdom.

These findings in particular clash with previous research that analyzed news articles and broadcast transcripts between Apirl 2020 and December 2020 using a dictionary model (Sacerdote, Sehgal, and Cook 2020). The divergent findings could be potentially explained by different models, different data or a combination of factors. One explanation is that broadcast transcripts are qualitatively similar to Op-Ed articles which do not correlate with health outcomes. Another explanation can be because of the differences in model accuracy: from Figure 4 we can see that the dictionary model VADER provides a very different interpretation than Covid-SieBERT—hinting that the difference is at least partly due to a difference at the level of model accuracy.

One of the main contributions of this paper is showcasing the importance of having an accurate model for sentiment analysis. Before looking at any data, one could assume that a less accurate model would result in the same trend as a more accurate model except with higher variance. From Figure 4 we see that this is not the case. Instead, having lower accuracy results in a completely different predicted trend with a different empirical conclusion.

6.1 Limitations

The Covid-SiEBERT model is very accurate, but it does not give a nuanced prediction of *how* negative a news article is. For example, the sentences "I am sick with Covid," and "Everyone in my family died from Covid," are both classified as negative, but clearly the latter is much more negative than the former. The model is not able to quantify how negative an article is and this is likely an important factor that is not available for inspection with the current transformer model. Another downside with the model is that you can not interpret what combination of words are causing the prediction. For example, it isn't clear how much vaccinations are contributing to the overall sentiment score trend or how articles talking about different health aspects such as hospitalizations or Covid cases may be contributing to the sentiment score. For example, during the Omicron wave, increasing Covid cases become decoupled with hospitalizations and the correlation between them drops from 0.97 to 0.68. There could be some articles talking about decreasing hospitalizations while some others talking about increasing Covid cases—in these situations, the sentiment score for that day would be split and appear somewhat neutral. It seems that there is something qualitatively different between two neutral articles vs one positive and one negative article which is not captured by the current methodology.

One misleadingly easy fix to this problem is to train a transformer model that can classify sentiment on a 5 or 10 point scale. The issue here is the drop in accuracy associated with increasing the output space. And we know from Figure 4 that this does not come without cost.

As mentioned before, the data from ProQuest may not be a random sample of news publishers and is biased towards having more left-wing and centrist sources. There are also issues with classifying publishers as national or regional. There are over 400 publishers in the U.S data alone, and many of these were classified as regional if they were not in the top 100 news sources in the sample. It is possible that there are important differences between regional and local news sources. But because of limited resources, this was not done, so both kinds of news publishers are lumped together.

6.2 Future Research Opportunities

There are several future research opportunities that can be derived from this paper. First, there is no data on differences in news reporting depending on whether the publisher is privately owned or publicly funded. It is possible that publicly funded news organizations do not have the same reporting incentives, therefore they will tend to be more positive. This was roughly observed when comparing national and regional sources, although that is not a direct comparison. One possible experiment would be to test how a news publisher's sentiment changes after being suddenly acquired by a private or public company.

In my analysis, I made assumptions about the effects of Covid news by extending the findings from controlled psychological studies on news viewership (section 2.2 reviews this literature). However, none of these studies used Covid articles. It is possible that people react differently to Covid articles since these stories are directly relevant to the viewers who are living through the pandemic. Future research could examine the causal effect of Covid news and see if this replicates the negative psychological findings found when studying other kinds of negative news stories. Furthermore, the potential impact of negative news on "Covid-fatigue," is still a debated topic which deserves future exploration given the findings presented here.

Finally, the causal effects of negative news on public health policy prescriptions remains unknown. A possible explanation for the unusually negative Op-Ed articles from Canada and the United Kingdom is that they both had stricter lockdowns compared to the United States, which could be partly influenced by the types of articles in the news. It could also be a case of reverse causation if the negativity in the news was driven by stricter restrictions.

7 Conclusion

In conclusion, it does appear that the news does have a bias toward negativity and this bias is most pronounced among both left-wing and right-wing sources, front-page articles, Op-Ed articles, and national news organizations. In each of these four examples, the estimated sentiment score trend is below zero throughout the entire timeline, meaning that they never published more positive articles than negative articles. And this pattern exists even when Covid-19 cases are decreasing.

Whether this negativity is warranted is a different question and beyond the scope of this descriptive analysis. Instead, this analysis highlights surprising differences that emerge from the data which deserve further consideration. For example, the fact that the average sentiment differences between political news sources are about as large as the difference between regular news and Covid-19 news sentiment. And that this difference also occurs, to a lesser extent, between national and regional news sources. It is also concerning that articles from the most negative news sources have the a smaller correlation with health metrics compared to their more positive counterparts.

It is also true that these patterns are not a universal fact about the news as a whole: centrist news sources, non-front-page articles, and non-Op-Ed articles have days where there are more positive than negative articles, and these instances always happen when Covid-19 cases are at their lowest or are decreasing. This finding contradicts previous studies and presents a more nuanced interpretation of the news's negativity. Furthermore, these news reporting patterns are not unique to a single country, they are found within the United States, United Kingdom, and Canada—suggesting that this is an inherent feature of news reporting.

Many of these findings support existing psychological literature on the news's negativity bias, such as the observation that front-page articles are the most negative and the considerable discrepancy between the number of articles published during good and bad times. This paper also offers new insights, such as the high negativity of Op-Ed articles that don't correlate with health measures, and the fact that national sources are more negative than regional sources. Additionally, the analysis of disparities across political ideologies contributes valuable information to the growing understanding of the contemporary U.S. political divide.

Overall, these findings provide a new perspective on how news sentiment changed in response to health metrics. This was achieved using a very high sample size of news articles and a state-of-the-art language model. However, several questions remain unexplored: whether private and public news organizations differ in the same way as national and regional organizations, if the political differences found in this analysis can be extended to countries other than the United States, or if these findings can be generalized to news reporting practices in non-western countries and other mediums such as TV and podcasts.

The code for the current study is available in the Covid-Thesis GitHub repository: `https://github.com/GabeNicholson/covid-thesis`. There are also interactive graphs available on my **website.**

References

Ahmad, Waseem et al. (Nov. 2022). "Enhanced Sentiment Analysis Regarding COVID-19 News from Global Channels". In: *Journal of Computational Social Science*. ISSN: 2432-2725. DOI: 10.1007/s42001-022-00189-1. (Visited on 04/22/2023).

Association, American Psychological (Nov. 2017). "APA Stress in America™ Survey: US at 'Lowest Point We Can Remember;' Future of Nation Most Commonly Reported Source of Stress". In: *Press Release*.

Asthana, Prakul (2021). "An Application of Sentiment Analysis with Transformer Models on Online News Articles Covering the Covid-19 Pandemic". PhD thesis. UCLA. (Visited on 03/19/2023).

Brenan, Megan (Oct. 2021). *Americans' Trust in Media Dips to Second Lowest on Record*. https://news.gallup.com/poll/355526/americans-trust-media-dips-second-lowest-record.aspx. (Visited on 05/08/2022).

Buneviciene, I. et al. (July 2021). "COVID-19 Media Fatigue: Predictors of Decreasing Interest and Avoidance of COVID-19–Related News". In: *Public Health* 196, pp. 124–128. ISSN: 0033-3506. DOI: 10.1016/j.puhe.2021.05.024. (Visited on 03/19/2023).

Coping with Stress (Jan. 2023). https://www.cdc.gov/mentalhealth/stress-coping/cope-with-stress/index.html. (Visited on 03/22/2023).

Federal Reserve Bank of San Francisco et al. (Mar. 2020). "Measuring News Sentiment". In: *Federal Reserve Bank of San Francisco, Working Paper Series*, pp. 01–49. DOI: 10.24148/wp2017-01. (Visited on 04/09/2022).

Galtung, Johan and Mari Holmboe Ruge (Mar. 1965). "The Structure of Foreign News: The Presentation of the Congo, Cuba and Cyprus Crises in Four Norwegian Newspapers". In: *Journal of Peace Research* 2.1, pp. 64–90. ISSN: 0022-3433. DOI: 10.1177/002234336500200104. (Visited on 04/07/2022).

Hartmann, Jochen et al. (Mar. 2019). "Comparing Automated Text Classification Methods". In: *International Journal of Research in Marketing* 36.1, pp. 20–38. ISSN: 0167-8116. DOI: 10.1016/j.ijresmar.2018.09.009. (Visited on 04/22/2023).

Heitmann, Mark et al. (2020). "More than a Feeling: Benchmarks for Sentiment Analysis Accuracy". In: *Communication & Computational Methods eJournal*.

Johnston, Wendy M. and Graham C. L. Davey (1997). "The Psychological Impact of Negative TV News Bulletins: The Catastrophizing of Personal Worries". In: *British Journal of Psychology* 88.1, pp. 85–91. ISSN: 2044-8295. DOI: 10.1111/j.2044-8295.1997.tb02622.x. (Visited on 04/09/2022).

Krawczyk, Konrad et al. (June 2021). "Quantifying Online News Media Coverage of the COVID-19 Pandemic: Text Mining Study and Resource". In: *J Med Internet Res* 23.6, e28253. ISSN: 1438-8871. DOI: 10.2196/28253.

Leetaru, Kalev (May 2019). "Sentiment Mining 500 Years Of History: Is The World Really Darkening?" In: *Forbes*.

Lin, Tianyang et al. (Jan. 2022). "A Survey of Transformers". In: *AI Open* 3, pp. 111–132. ISSN: 2666-6510. DOI: 10.1016/j.aiopen.2022.10.001. (Visited on 04/22/2023).

Liu, Siru and Jialin Liu (Sept. 2021). "Public Attitudes toward COVID-19 Vaccines on English-language Twitter: A Sentiment Analysis". In: *Vaccine* 39.39, pp. 5499–5505. ISSN: 1873-2518. DOI: 10.1016/j.vaccine.2021.08.058.

Liu, Yinhan et al. (July 2019). "RoBERTa: A Robustly Optimized BERT Pretraining Approach". In: *arXiv:1907.11692 [cs]*. arXiv: 1907.11692 [cs]. (Visited on 05/08/2022).

McLaughlin, Bryan, Melissa R. Gotlieb, and Devin J. Mills (Aug. 2022). "Caught in a Dangerous World: Problematic News Consumption and Its Relationship to Mental and Physical Ill-Being". In: *Health Communication* 0.0, pp. 1–11. ISSN: 1041-0236. DOI: 10.1080/10410236.2022.2106086. (Visited on 03/19/2023).

News BBC (Dec. 2014). "Russia: 'Good News Day' Decimates Website's Readership". In: *BBC News*. (Visited on 05/01/2023).

Pinker, Steven (2002). *The Blank Slate: The Modern Denial of Human Nature*. Viking.

— (2018). *Enlightenment Now: The Case for Reason, Science, Humanism and Progress*. New York: Viking.

Ritchie, Hannah et al. (Mar. 2020). "Coronavirus Pandemic (COVID-19)". In: *Our World in Data*. (Visited on 05/23/2022).

Rozado, David, Ruth Hughes, and Jamin Halberstadt (Oct. 2022). "Longitudinal Analysis of Sentiment and Emotion in News Media Headlines Using Automated Labelling with Transformer Language Models". In: *PLOS ONE* 17.10. Ed. by Sergio Consoli, e0276367. ISSN: 1932-6203. DOI: 10.1371/journal.pone.0276367. (Visited on 11/13/2022).

Sacerdote, Bruce, Ranjan Sehgal, and Molly Cook (2020). *Why Is All COVID-19 News Bad News?* NBER Working Paper Series no. w28110. Cambridge, Mass: National Bureau of Economic Research.

Shoemaker, Pamela J. (1996). "Hardwired for News: Using Biological and Cultural Evolution to Explain the Surveillance Function". In: *Journal of Communication* 46.3, pp. 32–47. ISSN: 1460-2466. DOI: 10.1111/j.1460-2466.1996.tb01487.x. (Visited on 04/10/2022).

Singh, Mrityunjay, Amit Kumar Jakhar, and Shivam Pandey (Mar. 2021). "Sentiment Analysis on the Impact of Coronavirus in Social Life Using the BERT Model". In: *Social Network*

Analysis and Mining 11.1, p. 33. ISSN: 1869-5469. DOI: 10.1007/s13278-021-00737-z. (Visited on 04/22/2023).

Slovic Paul (Apr. 1987). "Perception of Risk". In: *Science* 236.4799, pp. 280–285. DOI: 10.1126/science.3563507. (Visited on 04/09/2022).

Soroka, Stuart and Stephen McAdams (Jan. 2015). "News, Politics, and Negativity". In: *Political Communication* 32.1, pp. 1–22. ISSN: 1058-4609. DOI: 10.1080/10584609.2014.881942.

Soroka, Stuart N. (May 2006). "Good News and Bad News: Asymmetric Responses to Economic Information". In: *The Journal of Politics* 68.2, pp. 372–385. ISSN: 0022-3816. DOI: 10.1111/j.1468-2508.2006.00413.x. (Visited on 04/07/2022).

Trussler, Marc and Stuart Soroka (July 2014). "Consumer Demand for Cynical and Negative News Frames". In: *The International Journal of Press/Politics* 19.3, pp. 360–379. ISSN: 1940-1612. DOI: 10.1177/1940161214524832. (Visited on 04/06/2022).

Tversky, Amos and Daniel Kahneman (Sept. 1973). "Availability: A Heuristic for Judging Frequency and Probability". In: *Cognitive Psychology* 5.2, pp. 207–232. ISSN: 0010-0285. DOI: 10.1016/0010-0285(73)90033-9.

Twenge, Jean (Mar. 2019). *The Sad State of Happiness in the United States and the Role of Digital Media*.

Vaswani, Ashish et al. (2017). "Attention Is All You Need". In: *Advances in Neural Information Processing Systems*. Vol. 30. Curran Associates, Inc. (Visited on 04/22/2023).

Yadav, Ashima and Dinesh Kumar Vishwakarma (Aug. 2020). "Sentiment Analysis Using Deep Learning Architectures: A Review". In: *Artificial Intelligence Review* 53.6, pp. 4335–4385. ISSN: 1573-7462. DOI: 10.1007/s10462-019-09794-5.