Sentiment Analysis Equivalence in Political Campaign Ads

Kasra Lekan
kl5sq@virginia.edu

Sidhardh Burre
ssb3vk@virginia.edu

Matt Whelan
mw3shc@virginia.edu

Adarsha Poudel
xuk2ez@virginia.edu

1 Introduction

In the 2000 presidential election, George W. Bush won the state of Florida with 507 votes (15) leaving Florida with not just the reputation as a swing state but as a bellwether state (a state that predicts the final result of the election). More recently, Florida has become significantly more red, i.e. Republican-dominated. In the most recent governor election in 2022, the GOP candidate won by a 10% margin, the highest in recent history (13). Further, the GOP gained a two-thirds super-majority in the state's legislatures. Giving the GOP party nearly absolute control over what was a swing state just 20 years ago. This raises questions about the mechanisms that can convert a neutral-state into a political stronghold, of one party or another.

If we inspect the largest and most left-leaning county in Florida, Miami-Dade County, we find that over time, Miami-Dade has began to vote more and more right-leaning (8). In 2016, Hilary Clinton won the county by 30% points while in 2022, a Republican governor won the county by 10% points, a significant shift. By inspecting the demographic makeup of Miami-Dade County, more information can be wrought (4). We can see that Miami-Dade County experienced a 4% increase in its Hispanic population at the cost of Miami-Dade's white and black populations. An effect that is consistent among all five of Florida's most populous counties between the years 2010 and 2020.

Further, despite Latino voters leaning left-wards nationwide, Florida has seen a reversal of this trend with Latino voters voting more right-ward in recent elections (6). Although some of this can be attributed to fear-mongering in regard to the "socialist" policies of the Democratic party, a touchy topic for Venezuelan and Cuban immigrants, a portion of this shift can be attributed to a difference in ad spending between the Democratic and Republican parties (14). In Florida alone, the Republican party spent roughly 10 million dollars more on Spanish language advertising than the Democratic party as seen in Figure 1. And this spending is merely on a governor's race within a single state.

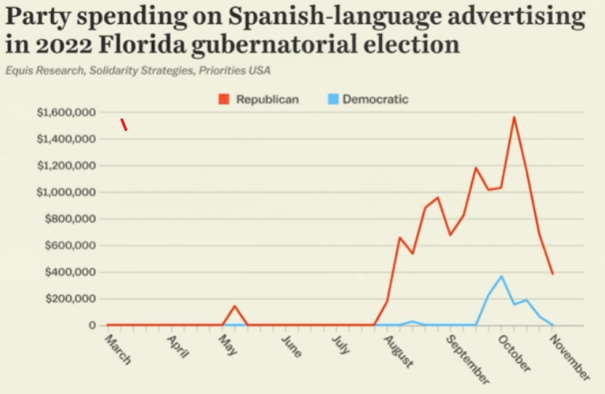

Figure 1: Party-based Spending on Spanish-language Ads

With many states possessing a growing Spanish-speaking minority, particularly swing states, more and more resources are being allocated by politicians to curry favor with this demographic. Figure 2 displays states where there is a significant hispanic population. Of these states with a large hispanic population, Texas, Nevada, New Mexico, Florida, and Colorado are known swing states. And the prominent hispanic population in these states justifies the increase in Spanish-language ad spending that both parties are dedicating to these states. In various state senate states, parties are willing to dedicate between over $3 million dollars only towards Spanish language advertising (11). Therefore, it is prudent to review the messages that these crucial voters are receiving. To do so, advertisements from various recent elections will be inspected in order to detect a class-based difference in sentiment. Specifically, the sentiment contained in a variety of official campaign advertisements

posted to the Meta (then Facebook) platform will be analyzed.

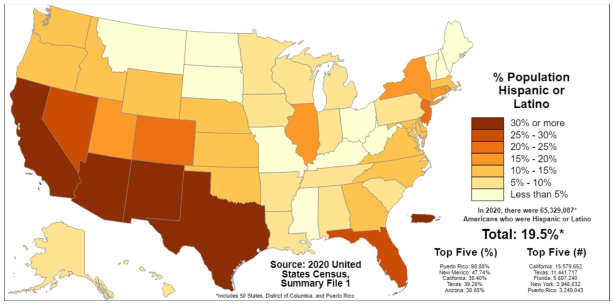

Figure 2: Hispanic/Latinx population in US states

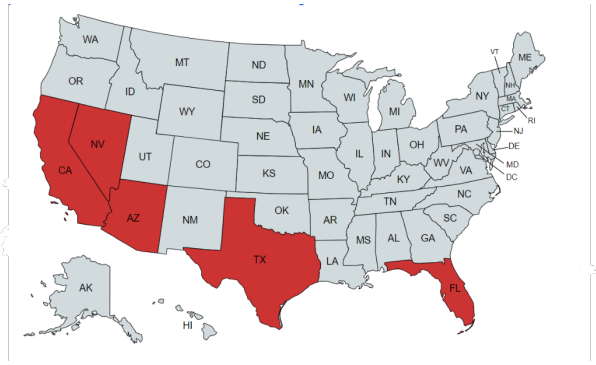

Figure 3: States Included within Dataset

2 Methods

2.1 Data

To gather data, we began by identifying elections of interest that occured in the past decade. Some characteristics that caused an election to be relevant were:

- the election occuring in a swing state
- the state contained a large spanish minority

Therefore, we set out by identifying a list of relevant elections to our cause. We found the following elections, enumerated in Table 1 to of great interest and therefore crafted queries to scrape the Meta ads API for those selected elections.

We collected Facebook Ads in English and Spanish using the Meta Ad Library API from the 2020 US Presidential Campaign and relevant 2022 US State Governor Races. We supplemented our dataset with matched human translations of a 2020 US US Presidential Debate. We also used a baseline dataset to test our models containing matched English-Spanish pairs from the MultiUN (n=1000), Multilingual UN Parallel Text 2000—2009 dataset (https://www.statmt.org/wmt13/translation-task.htm).

The final dataset spans 25 elections and 5 states for a total of 320 Spanish-English pairs and over 700 other unpaired samples. The states from which we have results from are displayed in Figure 3

2.2 Experiments

Similar to metamorphic testing in software testing research as first discussed in (5), we use the underlying assumption that for a text in English $\mathcal{E}$ and the corresponding translation in Spanish $\mathcal{S}$ theoretically should have the same sentiment score. In practice, there may be some baseline δ.

To assess whether or not Spanish language ads carry the same sentiment as English language ads, we first used the T5-small model for translation paired with the BERT model for sentiment classification (HuggingFace nlptown/bert-base-multilingual-uncased-sentiment) to define a baseline sentiment shift for translation from English to Spanish (δ_{e2s}) and vice-versa (δ_{s2e}). The BERT model's sentiment shift baseline was established by comparing equivalent human-translated texts within the UN dataset and measuring the difference in sentiment. The T5-translated text's language usage was also tested to determine a difference between using English versus T5-English. We expect these numbers to be near 0 if our models are adequate for our experiments. Subsequently, the sentiment shift between individual translated pairs ($\mathcal{E}$ and $\mathcal{S}$) of political ads (δ_P) from the same campaign will be assessed.

With the sentiment shift baseline and the sentiment shift seen in political campaign ads, we will then be able to determine the difference in sentiment shift within political ads versus standard machine translation. We will compare the baseline sentiment shift (δ_{e2s} and δ_{s2e}) with the political ad sentiment shift (δ_P) using a two-sided T-test with a significance threshold of $\alpha = 0.05$.

Our experiments compare the sentiment as well as the sentiment drift for political campaigns from Republican and Democratic Party candidates. Our evaluation will be based on statistical significance determined by the T-test (=0.05).

3 Implementation Details for Reproducibility

As stated in the proposal, we collected our data via the Meta Ad Library API. We first split the data by political party and combined all of our data sources into one dataframe. We also extracted and separated English and Spanish texts for each record and adjusted language labels to clearly denote the language of each text.

We used the `transformers` library for BERT's BertTokenizer and BertForSequenceClassification. We used these to make a SentimentClassifier class to perform sentiment analysis on the text, based the pretrained bert-base-multilingual-uncased-sentiment from nlptown. We used this model's tokenizer, which had a vocab size of 105879.

We also used the T5ForConditionalGeneration and T5Tokenizer from the `transformer` library to translate texts between English and Spanish. The T5-small model's tokenizer had a vocab size of 32128.

We used `torch` for basic functions like softmax and dot product to convert the model's output logits to a sentiment rating between 1 and 5. We used `pandas` for basic data preprocessing using dataframes. Other than these packages, we created the code by ourselves and did not use other people's code.

The GitHub repository is available at this link.

4 Experimental Results

Despite our initial hypothesis that Republicans spent more on Spanish-language advertisements (and therefore had more Spanish language ads), based on Florida's ad-spending trends, we found that this trend did not hold within our dataset. Instead, our dataset indicates that Democrats had *more* Spanish-language ads than Republicans. Florida may be an outlier within our dataset, a state that Democrats have seceded to Republicans due to the entrenched position that the Republican party may have in the state.

Looking closer at the data, we found that there were very few aggregate differences between ads that Republicans and Democrats employed. Both the broad strategies and sentiments employed by the parties were generally similar, causing similar shaped sentiment distributions for Republicans and Democrats alike.

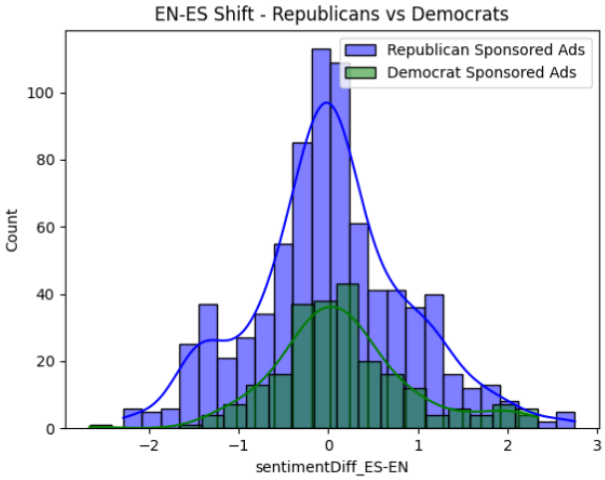

Figure 4: English to Spanish Translation Sentiment Shifts

Figure 4 displays the sentiment shift that occurs when ads in English are translated to Spanish. A p-value of 0.8315 (> 0.05) indicates that there is no significant difference in sentiment shift when English ads from either party are translated to Spanish. The peaked, nearly symmetrical distribution indicates that shifts in sentiment are randomly and evenly distributed. This proffers very promising results for the rest of the experiment. While zero sentiment shift would be preferred, the fact that the sentiment shift overall means towards zero indicates that at least in aggregate the BERT sentiment classifier preserves the validity of the source language.

Now that the validity of our sentiment classifiers is established, we will analyze language-specific sentiments as opposed to sentiment shifts. By establishing that translators apply neutral sentiment bias in aggregate, it is possible to assume that the sentiment classifiers maintain the sentiment contained in the source language. The following figure, Figure 5 displays the sentiment contained within ads written in English.

The following figure displays similar paired bimodal distributions for both Republicans and Democrats in regards to the sentiment encoded within their political advertisements. With a p-value of 0.0885 (> 0.05), we can conclude that there is no significant difference between the sentiment of English language ads used by the two parties. But, these bimodal peaks indicates that each distribution encompasses two separate and distinct populations. Two separate and distinct sentiment populations are to be expected in regards

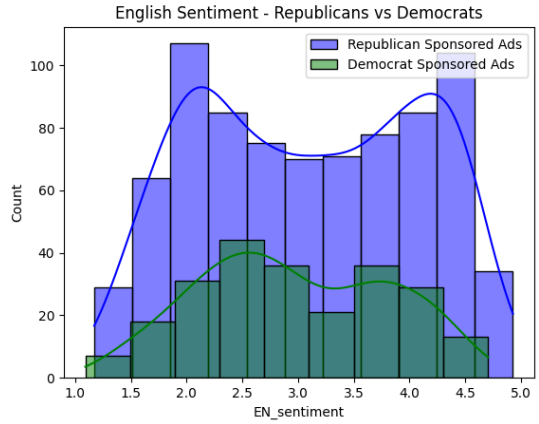

Figure 5: English Sentiment based on Party

to political ads. It has been established that voters are more likely to respond to provocative advertisements; advertisements that are either highly negative or positive (10). As such, the presence of two distinct and separate distributions, one containing highly negative ads and the other highly positive is a symptom of the party's mastery of human emotion within their advertisements. To put it another way, politicians have mastered manipulating emotions to sway voters in their favor.

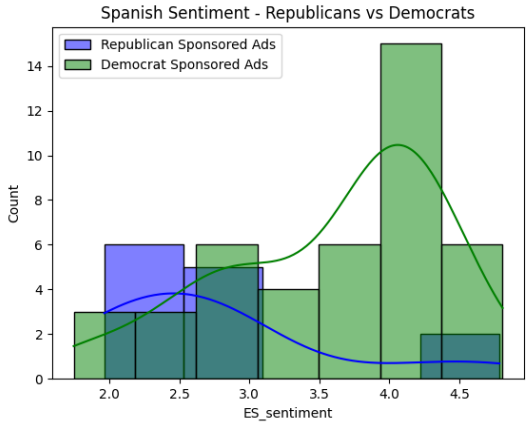

Figure 6: Spanish Sentiment based on Party

One major difference that we observed, however, were the specific tactics used by each party for Spanish language advertising, as shown in Figure 6. We found that, with a median sentiment score of around 3.75, Democrats placed greater emphasis on positive advertising for Spanish voters whereas the Republicans had an emphasis on more negative advertising, with a median sentiment score of around 2.75. This is a statistically significant difference with a p-value of 0.002 (< 0.05). This indicates that there *is* a difference in advertising when it is targeted toward Spanish-speaking voters between parties. Although there were less ads in Spanish for Republicans overall, this echoes the trend that we found historically; negative advertising has become a main tool for the Republican party over the last few elections (9). But, in the English advertising, we found a generally bimodal distribution for both parties with peaks generally around 2 and 4. This indicates that advertisements are provocative and can incorporate elements of positive and negative advertising, due to the rise in "contrast" advertisements (12).

The departure from the strategies used in English language advertising is highly interesting. Political parties have a completely divergent approach when it comes to currying favor with Spanish speaking voters. Instead of deploying ads that are both positive and negative, both parties select to mostly deploy ads of one sentiment over the other. What makes this strategy even more interesting is that both parties have *different* selected sentiments. Democrats employ significant positive sentiment and politicians of the Republican party instead employ highly negative advertisements. And while these results may not indicate the blatant misinformation claimed in (3) (2) (7) (1), it does indicate a distinctly different, and potentially harmful, difference in strategies employed by both parties and justifies further research.

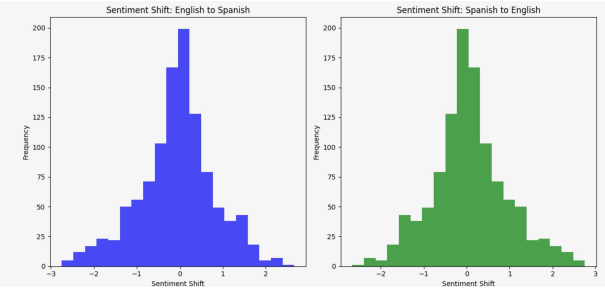

Figure 7: Sentiment Shift Histograms

In Figure 7's left plot, showing the English to Spanish sentiment shift histogram, suggests that translations from Spanish to English have minor sentiment changes, which is consistent across both parties' ads. The sentiment shifts are symmetric and cluster around zero, implying a close sentiment match between the two languages after translation. The right plot of 7, depicting the Spanish to English shift, indicates a similar sentiment alignment spread post-translation. Considering there

are fewer Spanish advertisements, the data might imply that English ads undergo careful translation to maintain sentiment when adapted to Spanish, which is crucial given their relatively smaller number overall.

5 Conclusion

The results observed within our data set meet the original hypothesis and introduce areas for further research. The validity of the sentiment shift applied by machine translators was verified. Despite imperfect sentiment shift, the knowledge that the machine translators are mostly able to preserve the sentiment of the source is of great importance. Not only are translators able to convert semantic meaning, but they are also able to preserve some of the implied sentiment. Further, when inspecting the sentiment used within languages by the different parties, it was possible to ascertain the different strategies employed by both parties. Further, it was possible to establish statistical significance in the difference between the sentiments contained within Spanish language ads.

Despite this, we found that we still had some areas for improvement. Firstly, our dataset was lacking in the areas of Spanish advertisements. Further, it was difficult to find specific pairings of English and Spanish. To rectify this potentially more digital ad-sources could be scraped. Given sufficient resources, video based ads could be a suitable format for this type of analysis in the future.

References

[1] Why Spanish-language mis- and disinformation is a huge issue in 2022.

[2] 2021. Tech companies blasted for inaction on Spanish-language disinformation.

[3] 2022. Latino groups target Spanish-language misinformation ahead of midterm elections. Section: Government.

[4] US Census Bureau. Racial and Ethnic Diversity in the United States: 2010 Census and 2020 Census. Section: Government.

[5] T Y Chen. 1998. Metamorphtic Testing: A New Approach for Generating Next Test Cases y.

[6] CNN. Exit polls for Midterm Election Results 2022 | CNN Politics.

[7] Robert Downen. 2022. U.S. Rep. Joaquin Castro, Latino groups warn about rampant misinformation targeting Spanish speakers ahead of the elections.

[8] FDOS. Florida Department of State - Election Results.

[9] Brett R. Gordon, Mitchell J. Lovett, Bowen Luo, and James Reeder. 2021. How much do campaign ads matter?

[10] David J. Grüning and Thomas W. Schubert. 2022. Emotional Campaigning in Politics: Being Moved and Anger in Political Ads Motivate to Support Candidate and Party. *Frontiers in Psychology*, 12.

[11] Ben Kamisar. 2022. Here are the congressional races drawing the most Spanish ad spending.

[12] Denise-Marie Ordway and John Wihbey. 2016. Negative political ads and their effect on voters: Updated collection of research.

[13] Politico. Governor Election Results 2022: Live Map | Midterm Races by State.

[14] Equis Research. 2022 Post Mortem.

[15] UCSB. 2000 | The American Presidency Project.

Election Type	Election Location	Election Name	Election Year
Presidential	Nation-wide	Trump vs Biden	2020
Gubernatorial	California	Newsom vs Dahle	2022
Gubernatorial	California	Newsom vs Cox	2018
Gubernatorial	Nevada	Lombardo vs Sisolak	2022
Gubernatorial	Nevada	Laxalt vs Sisolak	2018
Gubernatorial	Arizona	Hobbs vs Lake	2022
Gubernatorial	Arizona	Ducey vs Garcia	2018
Gubernatorial	New Mexico	Grisham vs Ronchetti	2022
Gubernatorial	New Mexico	Grisham vs Pearce	2018
Gubernatorial	Texas	Abbott vs O'Rourke	2022
Gubernatorial	Texas	Abbott vs Valdez	2018
Gubernatorial	Florida	DeSantis vs Crist	2022
Gubernatorial	Florida	DeSantis vs Gillum	2018
HoR	California, District 24	Allen vs Carbajal	2022
HoR	California, District 34	Gomez vs Kim	2022
HoR	California, District 40	Kim vs Mahmood	2022
HoR	California, District 46	Correa vs Gonzales	2022
HoR	Florida, District 26	Diaz-Balart vs Olivo	2022
HoR	Florida, District 27	Salazar vs Taddeo	2022
HoR	Florida, District 9	Soto vs Moore	2022
HoR	Florida, District 7	Mills vs Green	2022
HoR	Texas, District 15	Cruz vs Vallejo	2022
HoR	Texas, District 23	Gonzales vs Lira	2022
HoR	Texas, District 29	Garcia vs Schafranek	2022
HoR	Texas, District 20	Castro vs Sinclair	2022
HoR	Arizona, District 3	Chaplik vs Kolodin	2022
HoR	Arizona, District 7	Cook vs Marshall	2022
HoR	Arizona, District 29	Montenegro vs Smith	2022

Table 1: Elections of Interest