

ScholarWorks@GSU

Violent Political Rhetoric in Congressional E-Newsletters

Authors	Massey, Miles
Citation	Massey, Miles. "Violent Political Rhetoric in Congressional E-Newsletters." Thesis, Georgia State University, 2024. https://doi.org/10.57709/hwc8-ec93
DOI	https://doi.org/10.57709/hwc8-ec93
Download date	2026-03-08 21:28:37
Link to Item	https://hdl.handle.net/20.500.14694/12704

Violent Political Rhetoric in Congressional E-Newsletters

by

Miles Massey

Under the Direction of Michael P. Fix, PhD

A Thesis Submitted in Partial Fulfillment of the Requirements for the Degree of

Master of Arts

in the College of Arts and Sciences

Georgia State University

2024

ABSTRACT

Recently, violent rhetoric and its use by political leaders in the United States has become a primary focus of elites, media, and researchers alike. This research unveils whether members of Congress use violent political rhetoric strategically in online communications, and if the rates at which they use it is consistent with intuitive and theoretical understanding. Using a dataset of Congressional e-newsletters from 2011 to 2020, I use a random forest classifier to predict the rates at which violent political rhetoric is used and analyze them over time and along party lines. I find that violent political rhetoric is prevalent among newsletters and has increased overtime in partial support of my theory, but that Democrats use it at greater rates than Republicans in contrast to what was postulated. These findings implicate our understanding of the role of violent rhetoric in politics as a driver of negative partisanship, polarization, and political violence.

INDEX WORDS: Violent rhetoric, Congressional communication, E-newsletters, American politics, Political communication

Copyright by
Miles McKinley Massey
2024

Violent Political Rhetoric in Congressional E-Newsletters

by

Miles Massey

Committee Chair: Michael Fix

Committee: Jeffrey Lazarus

Louis-Alexandre Berg

Electronic Version Approved:

Office of Graduate Services

College of Arts and Sciences

Georgia State University

August 2024

DEDICATION

To my mother, Davina Taylor, and father, Jerome Massey, for supporting me thus far.

ACKNOWLEDGEMENTS

I would like to acknowledge my committee members, Dr. Michael P. Fix, Dr. Jeffrey Lazarus, and Dr. Louis-Alexander Berg, for their help in formulating this thesis. I would also like to acknowledge Dr. Ryan Carlin and Dr. Will Horne for the opportunities and guidance they offered during my time as an MA student.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS	V
LIST OF TABLES	VII
LIST OF FIGURES	VIII
1 INTRODUCTION.....	1
2 LITERATURE REVIEW	3
2.1 Online Political Communication.....	3
2.2 Aggressive Political Communication.....	6
3 THEORY	9
4 DATA	12
5 METHODS	17
5.1 Coding Rules.....	18
5.2 Model Performance.....	20
5.3 Statistical Methods.....	21
6 RESULTS	22
7 DISCUSSION AND CONCLUSION	29
REFERENCES.....	31

LIST OF TABLES

Table 1: Descriptive Statistics of Variables	15
Table 2: Logistic Regression of Violent Political Rhetoric.....	26

LIST OF FIGURES

Figure 1: Count of e-Newsletters in Each Year 2011-2020	14
Figure 2: Confusion Matrix of Random Forest Model Test Set Performance.....	21
Figure 3: Bar Graph of Proportion of Violent Rhetoric by Party	22
Figure 4: Line Graph of Proportion of Violent Rhetoric Over Time by Party.....	24
Figure 5: Bar Graph of Proportion of Violent Rhetoric by Election Year Status	25
Figure 6: Predicted Probability of Violent Rhetoric by Party	27
Figure 7: Predicted Probability of Violent Rhetoric by Election Year Status	28

1 INTRODUCTION

Today, more than any other time in recent history, political elites, media outlets, and researchers are paying close attention to the words politicians use to promote political activity among their base or would-be supporters. Trump's incendiary language took a major spotlight during his second impeachment trial as prosecutors argued the January 6th insurrection at the Capitol was in large part due to his rhetoric. Given the global rise of populism, growing concerns over democratic erosion, and recent incidents of political violence, it has become more important than ever to understand the role violent rhetoric plays in political movement, including the rate at which politicians use it and why.

In terms of the long arc of history, political violence has seen a significant global decline since the latter half of the 20th century (Pinker 2012). Over the past decade, however, political violence in the United States has been trending upward, and its form has taken a different shape (Kleinfeld 2021). Political violence was more commonly conducted by the fringe left in the 1960s and 70s, and it was mostly in the form of the destruction of property. In the late 1970s, violence began to shift rightward with the rise of white supremacist, anti-abortion, and militia groups, and, though the number of incidents declined, the target moved from property to people, such as minorities, abortion providers, and federal agents (Kleinfeld 2021). Today, political violence in the United States is rising on both the left and right, but overwhelmingly more so on the right (Pape 2021).

As political violence has evolved over the past half-century, so too has political communication, with no evolution more drastic than the advent of electronic communication. Internet availability and capacity has expanded globally over the past several decades, and politicians have innovated their campaigns, outreach, and networking to take advantage of the

new technology (Karpf 2013). The emergence of online communication is certainly a momentous stride in the expansion of political communication, but its impact is not strictly positive. As internet platforms have grown, aggressive political speech by platform users, including hate speech, violent speech, and violent political rhetoric, have remained pervasive (Matthew, et al. 2019). Violent political rhetoric is particularly harmful and dangerous as research supports that it leads to increased support for political violence, negative partisanship, and polarization (Kalmoe, Gubler and Wood 2018). Moreover, politician's use of violent language signals the in-group to mobilize, can reduce the taboo of violence against targeted groups, and pave the way for future violence (Yanagizawa-Drott 2014; Kaufman 2001).

Significant scholarly attention has been paid to this incendiary and deleterious political speech on social media platforms by the general public (Siegel 2020; Kim 2022), and recent research analyzes use and effects of violent language by politicians generally (Fisher et. al. 2019; Nugent et al 2022). However, scant attention has been paid to whether, how, and under what circumstances the political elite use aggressive political speech on internet platforms, an increasingly important stream of communication from politicians to constituents, supporters, and the general public. This research seeks to remedy this scholarly dereliction by assessing the frequency and partisan breakdown of violent political rhetoric by U.S. politicians in email newsletters from 2011 to 2020 using supervised machine learning. The primary theoretical contention of this research, that politicians have increasingly used violent political rhetoric strategically to boost support and spur engagement and that this strategy is used by Republicans more than Democrats, is supported in part by the findings that show increased use of violent rhetoric over the time period of the study and in election years, but not by the findings that show greater use by Democrats than Republicans.

2 LITERATURE REVIEW

Political science research has long looked at political communication, the interplay between politicians and constituents, and the effect of political speech on constituent behavior, policy making, and enforcement of the law. Increasingly, recent research has investigated social media as a tool for political communication by the political elite, including judges, members of Congress, and, on an international scale, world leaders (Lassen and Brown 2011; Barberá and Zeitzoff 2017; Curry and Fix 2019). This study seeks to contribute to the literature in three ways. First, it intends to expand our understanding of the internet and social media as a tool for political communication by the political elite for campaigning, communicating directly with constituents, and building a political brand. Second, it adds to the growing understanding of aggressive political communication online by considering violent political rhetoric by political elites in broadly-disseminated mass media communication. Third, it aims to identify which U.S. political party is more likely to perpetrate violent political rhetoric in online communications (specifically e-newsletters in this case) and how the frequency of violent political rhetoric by politicians has trended over time.

2.1 Online Political Communication

Political communication has evolved greatly over the past century as advancements in technology continue to progress; radio, television, and the Internet led to a rapid change in how politicians campaign and interact with the electorate. The innovation of utilizing the Internet in modern political campaigns has led to changes, such as increased small-dollar donor activity and a new focus on data analytics, while other features remain the same, such as fundamental participatory inequalities (Karpf 2013). Following the first election of former president Barack Obama, which many regard as the first successful social media driven campaign, researchers

found stark differences between liberal and conservative Internet use in party networking, information diffusion, and communication. Liberal websites and blogs tended to offer more community engagement architecture than conservative sites (Shaw and Benkler 2012).

The adoption of social media by politicians was relatively rapid following Obama's successful social media campaign in 2008. By 2010, senators, House leadership, and members with larger constituencies were likely to have adopted Twitter (Lassen and Brown 2011). Nearly all members of Congress had adopted Twitter within six years. Research finds that most politicians use Twitter primarily for broadcasting information, but those that also use it to engage with users receive more approval from constituents (Lee and Shin 2012; Lyons and Veenstra 2016). Politicians also engage in "self-personalization" on Twitter, which is exhibited by tweets that focus on individual statements, personal photos, and non-campaign related comments (Evans, Cordova, and Sipole 2014; McGregor, Lawrence, and Cardona 2017). These personalization campaigns vary by candidate, research shows, with some distinct gendered differences in how candidates make personal appeals to their audience (McGregor, et. al. 2017; Meeks 2016). Indeed, research on Internet-based political communication outlines two dimensions of online communication by politicians, both aimed at building a brand and encouraging constituent participation and support: a top-down, information dissemination (or agenda-setting) platform and an interactive, community-building platform (Dahlgren 2006; Lilleker 2015).

E-newsletters, on the other hand, represent a distinct form of political communication; even when one considers only forms of online communication, e-newsletters are still discernably different from things like tweets and other social media posts. Bloom, Cormack, and Shoub (2022) show that frequency of communication by members of Congress vary across mediums

and that first-term members and members representing more populated districts have an increased likelihood of using e-newsletters while Democrats and more ideologically extreme members are less likely to use e-newsletters. Moreover, Congress members differ in their partisan positioning and the way they portray their ideological disposition across different mediums; speech in e-newsletters portrays more partisan polarization than traditional venues (e.g., press releases) but less than social media posts (Green, Shoub, and Cormack 2024). In terms of audience, e-newsletter subscribers are on average older, wealthier, and more partisan members of their electorate (Cormack 2017). The assumed audience for e-newsletters is comprised of primarily constituents of the author; aware of this, politician's change their ideological portrayal depending on the base-to-swing voter ratio that composes their district (Cormack 2016). Members of Congress also take cues from their party leader and other party members in a "follow the leader" style of politics, such as distributing similar post-card style holiday e-newsletters or touting some of the same beliefs as the president regarding Covid-19 strategies. However, they may break from their party leader's messaging regarding important, factual information for their constituents, such as statements regarding the importance of wearing masks during the Covid pandemic (Cormack and Meidlinger 2021).

Further research from the U.K. shows that e-newsletters share some similarities with other forms of online communication, such as being direct communication from the politician to the constituent and being cheap, easy to produce, and able to bypass mass media outlets, yet they are a mostly one-sided form of communication with a narrower audience of mostly supporting constituents (Jackson 2005; Jackson 2006). In this way, the content of political e-newsletters generally focuses on communicating with party members rather than attracting new supporters (Jackson 2005). Research on Canadian lawmakers shows that e-newsletters serve as a tool for

self-presentation, but that lawmakers present themselves differently in e-newsletters than other platforms designated for a more general audience, such as their websites (Koop and Marland 2013). Overall, e-newsletters serve as a direct line of communication from politicians to their political base for self-personalization, promoting constituency services, and relaying unmediated targeted messages (Jackson 2008; Jackson and Lilleker 2007).

2.2 Aggressive Political Communication

In the broadest sense, there are two categories of aggressive political communication: hate speech and violent rhetoric. Hate speech takes on many definitions and conceptualizations within the literature but can be generally defined as communication that disparages a person or a group on the basis of characteristics such as race, color, ethnicity, gender, sexual orientation, nationality, religion, or political affiliation (Castaño-Pulgarín, et al. 2021). Violent political rhetoric, which appears less prominently as its own concept in the literature, is defined as expressing the intention of physical harm against political opponents and includes both direct expressions of intent to inflict physical harm and violent metaphors that may not necessarily carry the intent to commit violence (Kim 2022). Aggressive political communication occurs offline and on social media platforms, of which scholarship has looked at the effects of both. Similarly, these forms of toxic speech are produced and propagated by both political elites and non-elites.

Research into the effects of aggressive political communication yield results with causal arrows pointing both ways. Studies show that hate speech by political figures boost domestic terrorism (e.g., Piazza 2020) while, at the same time, incidents of extremist violence may lead to increased incidents of hate speech online (e.g., Olteanu, et. al. 2018). Many recent incidents of political violence in the U.S. are rhetorically linked to speech made by right-wing political

leaders and commentators, such as former president Donald Trump or former Fox News personality Tucker Carlson, by news outlets such as the New York Times. On the other hand, political violence conducted by extremist groups composed of minority populations, such as attacks by Muslim groups in Western countries, contributes to an increase in online hate groups. Rather than a chicken vs. egg situation, the literature seems to show a self-perpetuating cycle of aggressive political communication fanning the flames for acts of political violence that in turn lead to the self-proclaimed affirmation and proliferation of hateful and violent speech (Wahlström et. al. 2021; Fisher et. al. 2018).

Though it is not often that political elites use hate speech or directly express violent intent, it is common in all forms of political communication through every medium to use violent metaphors (Flusberg, Matlock and Thibodeau 2018). War metaphors are perhaps the most frequent exhibitions of violent metaphors; examples include the war on drugs, war on poverty, and so on. The literature offers two main reasons as to why metaphors like this are so prominent in U.S. politics. The first is that they allow rather complex or abstract topics to be thought of and talked about as something simpler and more concrete (Thibodeau and Boroditsky 2011). When discussing a topic that is a complex sociopolitical issue with an array of causes and consequences, such as the Opioid Epidemic, it is made much simpler by referring to it as a ‘disease’ and a crisis that we have declared ‘war’ on to solve. The second is that these metaphors evoke notable emotional responses (Flusberg, Matlock and Thibodeau 2018). For example, referring to drug addiction as a disease may invoke feelings of compassion or disgust depending on the context, while referring to a political program as a war on something may invoke fear or a sense of threat (Elwood 1995).

Violent metaphors in politics go beyond just war metaphors; they include references to ‘fight for’ or ‘against’ something, portrayals of supporters as combatants, and depictions of political objects as weapons (e.g., “attack ad”) (Kalmoe, Gubler and Wood 2018). These violent metaphors are absorbed by citizens as aggressive cues that prime aggressive responses, which in turn produce higher-aggression partisanship and increased polarization (Kalmoe 2014). This is even more true for individuals who have higher trait aggression – aggressive cues prime aggressive responses to the extent that they resonate with trait aggression in audiences (Anderson and Bushman 2002). More than just increasing partisanship, Kalmoe, Gubler, and Wood (2018) show that violent metaphors increase support for political violence among respondents of a nationally representative survey. Like hate speech and direct expressions of violent intent, violent metaphors exacerbate polarization, foster violent attitudes, and potentially lead to incidents of political violence.

If this is the case, one may reasonably ask why politicians continue to use these forms of aggressive political speech. This is because they are politically incentivized to use these forms of speech to motivate their supporters; the literature supports that aggressive political speech has a serious impact on many political outcomes, including participation, information seeking, inter-group evaluations, and deliberative attitudes (Henson, et. al 2013; Sydnor 2019; Gervais 2019). The rise of negative partisanship in U.S. politics has created a political environment in which politicians are incentivized to depict their political opponents as enemies, which is conducive to various forms of aggressive political communication (Abramowitz and Webster 2016, 2018). With respect to this growing issue, there is also significant literature on detecting and discouraging hateful speech and violent political rhetoric online and elsewhere (Waseem and Hovy; Davidson et al., 2017; Zimmerman, Kruschwitz and Fox 2018; Siegel 2020).

3 THEORY

The literature illustrates that violent political rhetoric takes the form of both direct expressions of violence and violent metaphors or other forms of figurative language. This volatile language, when spoken by political elites, can spur audiences into political action and make complex topics more digestible, but also promote aggression, polarization, and potentially pave the way for incidents of political violence. Why, then, do politicians use violent rhetoric despite the negative effects it can have on peace and democracy? Here, I argue that politicians, namely members of Congress in this study, use violent political rhetoric strategically to promote engagement and support from their political base in spite of the negative effects it may have. In communications where there is a broad general audience, politicians are disincentivized from using the more vitriolic forms of violent rhetoric as the general public typically dislikes uncivil political discourse (Frimer and Skitka 2018). E-newsletters, however, have a primary audience of supporters, reducing the probability of losing floating voters due to poor public relations and thus the cost of using violent rhetoric. Moreover, e-newsletters are top-down, one-way communication from politician to constituent or supporter, reducing the probability of direct pushback or backlash for using polarizing or incendiary speech one might get on other more interactive mediums such as Twitter and Facebook. As a result, e-newsletters are a prime medium for political communication containing strategic violent rhetoric.

While this study does not compare usage of violent political rhetoric across platforms, it does test whether its use, if strategic, is consistent with intuitive and scholarly understanding that its use will have increased as politics have become more contentious, it will be utilized more as the stakes increase (as elections approach), and that Republicans will find its use less costly than Democrats. Here, I formulate three tests to determine whether use of violent political rhetoric is

consistent with the theoretical framework of its strategic application by members of Congress. The first proposes that violent rhetoric usage will be greater among Republicans than Democrats. This is primarily attributed to the rise of right-wing populism, negative partisanship, and grievance-based politics associated with the tea-party, Donald Trump, and the new style of hardline Republican politics (Tesler 2016; Kreis 2017; Abramowitz and McCoy 2018; Hodwitz and Massingale 2020). This style of politics is more conducive for violent rhetoric than a more liberal, base-broadening approach that is expected of Democratic Congress members. Additionally, some research suggests that violent metaphors may be more common among Republicans as popular rhetorical phrases such as the ‘war on terror’ are largely attributed to the Reagan, Bush I, and Bush II brand of Republican politics that emphasizes the use of political and cultural narratives to frame policy issues (Fisher et. al. 2018). Though ‘war on poverty’ was coined by Democratic president Lyndon B. Johnson, predating Reagan, the use of war metaphors since has been primarily a Republican strategy. This all informs the first hypothesis:

H1: Democratic members of Congress will use violent rhetoric in email newsletters less than Republican members.

The second theoretical argument for strategic use of violent rhetoric is that it has increased over time alongside the rise of negative partisanship and grievance-based politics. The grounding for this is essentially the same as H1, but it takes into account the use of violent rhetoric by both parties. Although Trumpian Republicanism is more conducive to violent political rhetoric, negative partisanship has increased on both sides as of the aisle, promoting affective polarization and combative rhetoric regarding the other party (Abramowitz and Webster 2016). As such, there is no counteracting force against the speculated increased use of

violent rhetoric on the right; instead, overall use of violent rhetoric will increase overtime as Democrats continue to use violent rhetoric, only at rates less than Republicans.

H2: Use of violent political rhetoric in email newsletters will increase between 2011 and 2020.

Lastly, if usage of violent political rhetoric is indeed strategic, it should appear in e-newsletters at a greater rate around elections than any other time. One of the primary benefits from using violent rhetoric in political speech is that it increases political activity, elevates in-group evaluations while attitudes towards out-groups suffer, and diminishes deliberative attitudes (Henson et. al. 2013; Gervais 2019). However, this is all at the risk of losing supporters due to general attitudes against incivility in politics. As such, a Congress member's cost-benefit analysis for using violent political rhetoric is more likely to encourage its use when the pressure to promote voter turnout and support for oneself or party is high, i.e., as elections approach. In turn, members may be increasingly discouraged from using violent political rhetoric the further away elections are from the e-newsletters publication as they avoid the risk of losing supporters.

H3: Use of violent political rhetoric in email newsletters will be higher in election years than non-election years

4 DATA

To assess the prevalence of violent rhetoric in communications made by U.S. politicians, I use a dataset of over 123,000 email newsletters from U.S. members of Congress between 2011 and 2020. This dataset comes from DCinbox, an academic project conducted by students and faculty at New York University and Stevens Institute of Technology. The data contains the subject and textual content of the newsletters, the author (or name of the politician whose office distributed the newsletter), date and timestamps, and other important information, including partisan affiliation of the author. Due to missingness in key variables, several observations from the original dataset were dropped from the analyses, resulting in a final dataset of 122,176 observations.

As discussed in previous sections, e-newsletters are certainly not representative political communication broadly. Some key differences include e-newsletters much narrower audience of constituents and their focus being on policy and news as opposed to brand-building, tent-widening, or self-personalization (Cormack 2016; Cormack 2017; Jackson 2005). Moreover, the communications through these venues have already been shown to differ in terms of partisan positioning, use of polarizing language, and overall frequency of use by party, ideology, gender, and other variables (Blum, Cormack, and Shoub 2022). As such, the prevalence of violent rhetoric in political speech may differ based on the target audience or purpose of the speech, but that is beyond the scope of this study. Nonetheless, e-newsletters remain an important part of a politician or political party's arsenal for communicating and disseminating information and thus represent an important form of political speech. In this way, developing a greater understanding of the use of violent rhetoric in this form of political speech aids in our broader understanding of violent rhetoric in U.S. politics.

The dependent variable in this study, *violent rhetoric*, is a dichotomous variable indicating whether a newsletter contains violent rhetoric, indicated as one, or does not, indicated as zero. This label does not exist in the original dataset; instead, this is a predicted label generated by a machine learning model trained on a prelabelled subset of the data and used to predict the occurrence of violent rhetoric throughout the full dataset. Further detail regarding this machine learning process, including model specificity, training, and coding rules, will be outlined in the methods section proceeding this one. When it is all said and done, just over a quarter of all e-newsletters in the dataset, or 30,563, are labelled as containing violent rhetoric; this indicates that violent rhetoric is a relatively frequent occurrence in this form of political communication, but it still is absent from e-newsletters more often than not.

There are three independent variables in this analysis. The first is partisan affiliation, measured as a dummy variable. In the analysis, this variable is referred to as *Democrat* where an instance of zero indicates that the politician who authored the newsletter is a Republican or Independent while an instance of one indicates he or she is a Democrat. In the data, there are 511 Republicans and Independents and 398 Democrats, resulting in a total of 78,618 e-newsletters from Republican or Independent Congress members and 43,961 from Democratic members. There is a sizeable representation of both Democrats and non-Democrats in the data, so there is no expectation of skewed or biased results. The second independent variable is *year*, which refers to the year in which the e-newsletter was distributed. Figure 1 displays the count of e-newsletters in each year, showing a relatively consistent number of e-newsletters each year until a large spike in 2020, which likely is a result of the COVID-19 pandemic. *Election year* is the third independent variable and indicates whether the e-newsletter was released during an election year or not. This is not a perfect measure of closeness to elections as there are certainly smaller

divisions of time that may enable more granulated analysis, and some members of Congress are not always up for re-election in a given election year. However, this measure implies that Congress members campaign not only for themselves but also for their party during election cycles and are thus more inclined to use speech that promotes political activity, which violent rhetoric is theorized to do, even when they are not up for re-election. Moreover, election campaigns do not begin and end within a few weeks of elections, but often unfold over the course of the election year as primary elections in most places take place in the spring or early summer and candidates enter and withdraw from the race throughout the year. As such, using election years as a measure of proximity to elections is appropriate as it takes into account the long life of an election campaign and a member of Congress' prerogative to promote their party's policy positions and accomplishments regardless of whether they are individually up for re-election.




Figure 1: Count of e-Newsletters in Each Year 2011-2020

There are two additional variables included in the analysis to account for variance in the modelled relationship that is not explained by the theory: *seniority* and *district presidential vote*. *Seniority* indicates how senior a Congress member is in their given chamber, measured numerically by the number of terms they have been elected to the chamber. *District presidential vote*, which is the share of the vote won in a Congress members district by the Democratic presidential candidate in the most recent election, serves as a measure of constituent ideology for that member of Congress. It refers to the district-level share of the vote for members of the House and statewide vote-share for members of the Senate, which has been shown to affect the behavior of Congress members (Waldman 1967; Cohen et al 2000; Parker and Parker 1985). This variable at face value seems strongly colinear to partisan affiliation, but it is fundamentally different as districts and states can often split the ballot in relation to presidential and Congressional elections, separating a Congress member's own party identification from the partisan affiliation of their district in presidential elections. Including this variable brings this study in line with standard practices when analyzing the U.S. Congress. These variables are not theoretically linked to the presence of violent rhetoric in a Congress member's e-newsletter and thus have no hypothesized correlation, but may nonetheless impact how Congress members communicate and are included in the analysis as controls.

Table 1: Descriptive Statistics of Variables

Variable	Mean	Standard Deviation	Minimum	Maximum
Democrat	0.36	-	0	1
Year	2015.85	2.93	2011	2020

Variable	Mean	Standard Deviation	Minimum	Maximum
Election Year	0.51	-	0	1
Seniority	4.98	4.09	1	30
District Presidential Vote	48.57	14.88	17.46	96.99

5 METHODS

The first step in this analysis is to generate the predicted labels that serve as the dependent variable. To do this, I train a random forest classifier model on a random subset of 375 pre-labelled e-newsletters. Using Breiman's (2001) random forest algorithm, the model constructs an ensemble of decision trees during training and outputs the mode of the classes of the individual trees, enhancing predictive accuracy and controlling overfitting by averaging multiple decision trees trained on random subsets of the data and features. Essentially, it generates a large number of decision trees and each tree votes on the predicted class, allowing the model to produce a predicted label and a measure of how confident it is in each label prediction.

I choose to implement a random forest model as opposed to other, more cutting-edge classifier models such as support vector machines or neural networks primarily due to feasibility. Due to having to personally hand-label the training set, there is a relatively low-N subset of pre-labelled data, which shrinks the benefits of have a more refined machine learning model. Meanwhile, the costs of such a model remains high in terms of time to train, monetary costs to run, and taxation on the memory usage and processing power of my personal machine. Additionally, I do not choose to use other low-cost models such as Naïve Bayes or k-Nearest Neighbor because the random forest model performs just as well or better, thus serving as a middle ground choice between simplistic classifier models to the more refined, high-cost forms of machine learning.

At its core, there are three steps to supervised machine learning: train, validate, and test. First, the machine learning model is trained using random or nonrandom subset of data, e-newsletters in this case, with pre-labelled classes hand-coded according to a set of coding rules.

This pre-labelled data is next split into a training set and a test set, the first of which the model uses to learn the important features that predict the text's classification. The model is validated to determine how well it is expected to perform on unseen data and prevent overfitting it to the training data. Then, the model is run on the test set to validate how well the model predicts the texts' classification compared to the hand-coded labels. Subsequently, the model hyperparameters are tuned to maximize its precision, accuracy, or other relevant measures of performance while balancing the variance-bias tradeoff and re-tested on the pre-labelled data. Finally, once the model's performance has reached an acceptable level, it is run on the full dataset to provide predicted labels that provide a class for each observation.

5.1 Coding Rules

For this research, there are two categories by which an e-newsletter can be labelled. Class 1 refers to e-newsletters that contain violent rhetoric, which includes those in which the author directly expresses an intention of physical harm (e.g., "we will attack our enemies with full force"), and those in which the author discusses essentially non-violent politics using a violent metaphor or other forms of figurative speech but lacking in violent intent (e.g., "I will continue to fight for you in Congress" or "we must limit job-killing regulation"). Class 2 refers to e-newsletters that do not contain either direct expressions of an intent to act violently or violent figurative speech, including those that may contain violent language that are statements of (perceived) fact, such as messages referencing violent events or normalized violence the death penalty (e.g., the death penalty or the armed forces). It is important that violent speech that is factual is differentiated and excluded from violent political rhetoric because members of Congress often expected to talk about both policies related to violence, such military affairs and the death penalty, or incidents of violence, such as mass shooting or terrorism. This speech is

fundamentally different violent political rhetoric because its primary purpose is not to invoke or appeal to an emotion, but rather to disseminate information about real-world events.

The decision rules are comprised of three sets of sequential questions and answers:

Question 1: Does the author express a personal or externalized intention of physical harm that can be interpreted literally?

Answer: If yes, Class 1. If no, go to Question 2.

Question 2: Does the author use violent figurative speech AND do you think that the figurative speech is used to express an intention of physical harm?

Answer: If yes, Class 1. If no, go to Question 3.

Q3: Does the author use violent figurative speech AND do you think that the figurative speech is used to refer to essentially non-violent political events, issues, or phenomena?

Answer: If yes, Class 1. If no, Class 2

In the case of this study, direct expressions of violent intent and violent figurative speech are combined into one general category of violent rhetoric, but this may not necessarily be the case in analyses of other mediums of political speech. Direct expressions of physical harm are too rare among e-newsletters to be included as its own category distinct from figurative language; this may be because of a relatively thorough vetting process prior to distribution or some other mechanism that future research may discover. Distinguishing between different forms of violent rhetoric certainly aids in explaining its use by political elites and impact on political behavior but is not practical in this case due to the nature of the data. In this way, e-newsletters are expected to reach Class 1 primarily through Question 3, and those that do so by way of Questions 1 and 2 are not expected to occur at a substantively meaningful rate. All cases

in which the answer is no to all three questions result in categorization of Class 2, allowing for speech that discusses violence in a non-rhetorical manner to be appropriately classified.

5.2 Model Performance

In this study, the total N of the pre-labelled dataset amounts to 375 observations. I divide this dataset into a 75-25 split between the training and test sets respectively. I use 10-fold cross validation to provide a robust estimate of model performance on the training set. Figure 2 displays a confusion matrix of the model's final performance on the test set in which it predicted labels with approximately 75% accuracy, 64% precision, and 55% recall. In plain English, the model correctly predicts the class three out of four times overall, correctly identifies newsletters containing violent rhetoric a little less than two-thirds of the time when it makes a positive prediction, and identifies a little over half of the actual cases of violent rhetoric, meaning it misses nearly half of the true positives. These model performance results are not perfect, but can be largely attributed to the size of the training data, which is a shortcoming not easily remedied. Beyond that, however, there is expected to be a class imbalance in the data – there are fewer e-newsletters containing violent rhetoric than those that do not, making it more challenging to achieve high precision and recall. For the purposes of this study, analyzing over half of all potential cases in which there is violent rhetoric and, in doing so, achieving a sufficiently high N is satisfactory.




Figure 2: Confusion Matrix of Random Forest Model Test Set Performance

5.3 Statistical Methods

After generating predicted labels, I perform visual and statistical tests for each of the hypotheses. The visual tests are detailed in the results section below. To statistically test for the effect of partisan identification and election proximity on the likelihood of the presence of violent rhetoric in a newsletter, I estimate a logit model to analyze the effect of *Democrat*, *election year*, *seniority*, and *district presidential vote* on the predicted label for *violent rhetoric*. I exclude year as a variable from this model to avoid problematic autocorrelation and rely solely on a visual test to support the related hypothesis.

6 RESULTS

Figure 3 below illustrates the difference in the proportion of violent rhetoric use between the two political parties as a visual test for the first hypothesis. I use proportions as opposed to counts to account for difference in the total number of e-newsletters from each party that exist in the data. These findings run directly counter to the contention that Democrats use violent rhetoric in e-newsletters less than their Republican counterparts. Instead, the graph shows that approximately 27% of Democrat's e-newsletters contain violent rhetoric while the same is true for only around 24% of Republican e-newsletters. While a 3% difference between the parties is relatively small in magnitude, the fact that Democrats use violent rhetoric at a greater rate is notable in and of itself.

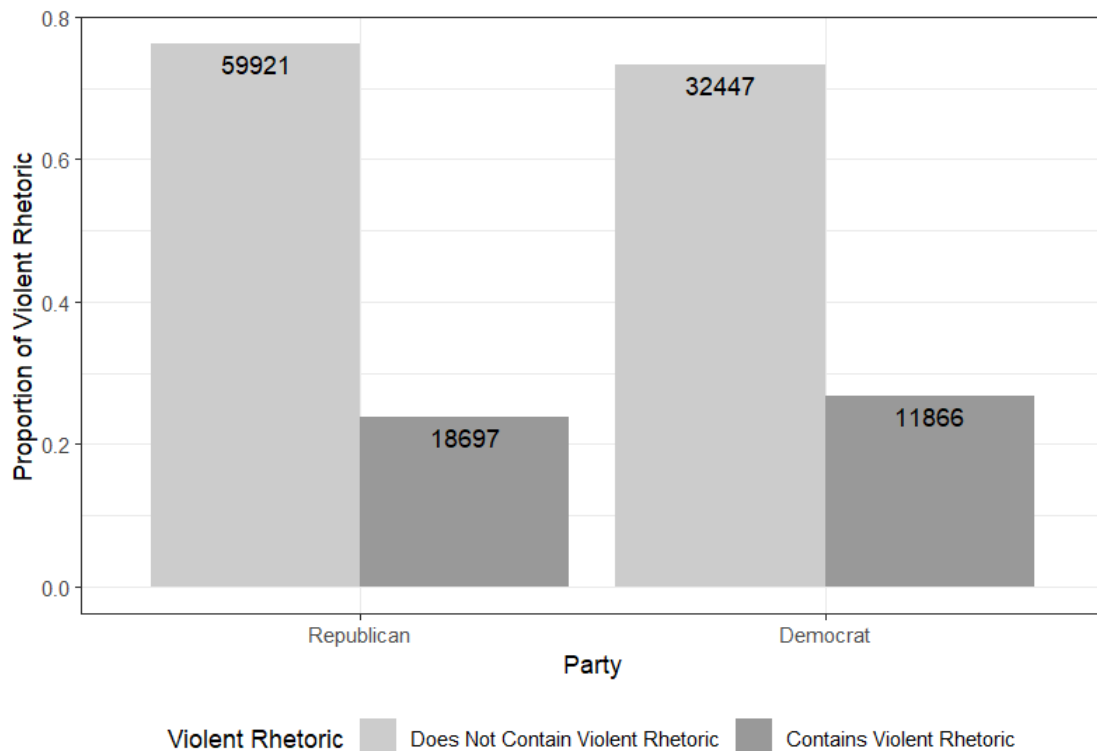


Figure 3: Bar Graph of Proportion of Violent Rhetoric by Party

The second hypothesis contends that violent rhetoric in e-newsletters has increased over time. Figure 4 displays this overall trend in the proportion of e-newsletters containing violent rhetoric, as well as how the use of violent rhetoric has trended by party. I use proportions as opposed to counts to account for both differences in total number of e-newsletters produced each year and the greater number of Republican-authored e-newsletters than Democrat-authored ones that exists in the data. The graph shows a sizeable general increase in the proportion of all e-newsletters that labelled as containing violent rhetoric by the random forest model. By 2020, violent rhetoric occurred in about 26% of e-newsletters as compared to around 16% in 2011. The trend is not linear, however; use of violent rhetoric sharply rose from 2013 to 2016, drastically dropped in 2017, then rose back to 2016-level use where it mostly plateaus. Notably, the changes in the overall trend are largely driven by changes in usage by Republican members of Congress, likely due to their larger representation in the data. Democrats, counterintuitively, show higher proportions of violent rhetoric use than Republicans in almost every year, and their usage peaks in 2017, the same year that Republicans exhibit their largest year-to-year decrease. This may indicate that big national election losses incentivized Democrats to use more violent rhetoric to evoke emotional responses that would increase policy salience or political activity among constituents while big wins for Republicans allowed them to curb their use of violent rhetoric. However, this visual test can only sufficiently show that use of violent rhetoric in proportion to the total use of e-newsletters has indeed increased over time, in support of the second hypothesis, and that Democrats on average use more violent rhetoric in a year than Republicans, providing further evidence against the first hypothesis.

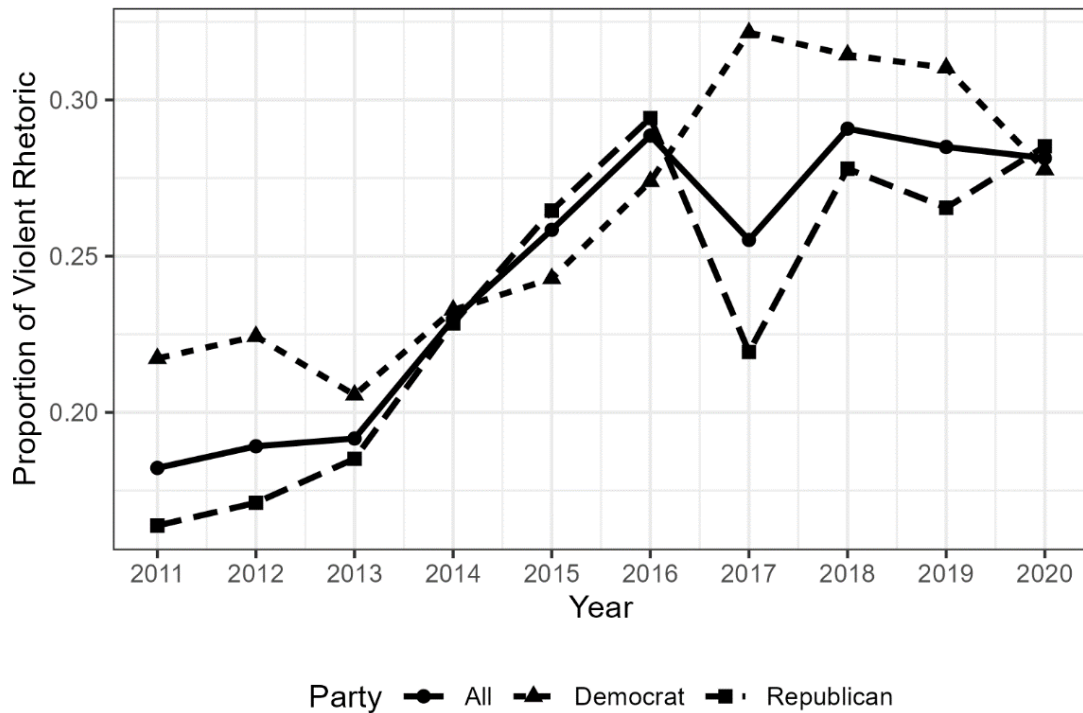


Figure 4: Line Graph of Proportion of Violent Rhetoric Over Time by Party

Figure 5 compares the proportion of e-newsletters that contain violent rhetoric in years with a national election and non-election years. According to the graph, election years see an approximately 2.3% greater proportion of e-newsletters containing violent rhetoric than non-election years with a total proportion of about 26% and 23.7% respectively containing violent rhetoric. This finding supports the third hypothesis, that violent rhetoric is greater in election years than in non-election years. Again, however, the substantive magnitude of the difference is relatively small.

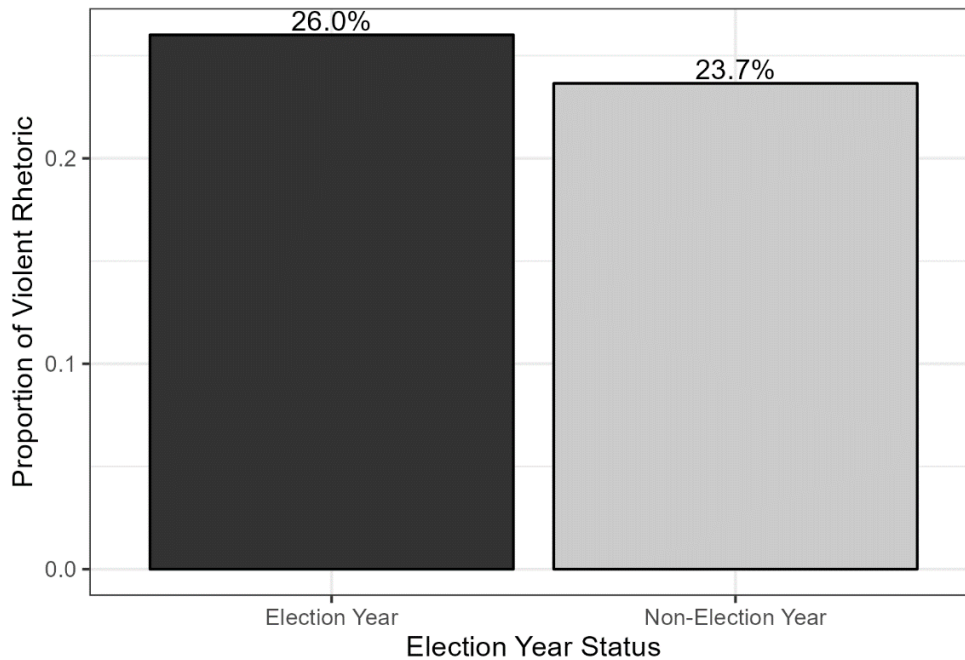


Figure 5: Bar Graph of Proportion of Violent Rhetoric by Election Year Status

The results of nested and non-nested logistic regression models are displayed in Table 2 below. Nested models are displayed to address potential model overspecification; as the variables maintain their statistical significance in both nested and non-nested models, the full model is likely not over-specified. The model results show statistical significance for the two theoretical variables at the 0.001 level. I use a more stringent cut-off for statistical significance of 0.001 as opposed to the more conventional cut-off of 0.05 due to the high number of observations in the analysis that make it much easier to achieve statistical significance at the conventional cut-off. Nonetheless, both *democrat* and *election year* are statistically significant and positive, indicating that e-newsletters from Democrats and e-newsletters in election years are more likely to be labelled as containing violent rhetoric than those from Republicans or those from non-election years. This is again in contention with the first hypothesis but supportive of the second. To note, the control variable *district presidential vote* is also statistically significant

and positive, indicating there may be a relationship between increased right-wing partisan preference among constituents and increased use of violent rhetoric.

Table 2: Logistic Regression of Violent Political Rhetoric

	<i>Violent Rhetoric</i>		
	Logit 1	Logit 2	Logit 3
Democrat	0.159*		0.346*
	(0.014)		(0.020)
Election Year		0.127*	0.123*
		(0.013)	(0.013)
District Presidential Vote			-0.009*
			(0.001)
Seniority			0.001
			(0.002)
Constant	-1.165*	-1.172*	-0.885*
	(0.008)	(0.010)	(0.030)
Observations	122,930	122,930	122,930
Wald χ^2	135.971 *	871.440 *	91.681 *
Log Likelihood	-68,873.440	-68,895.120	-68,741.920
Akaike Inf. Crit.	137,750.900	137,794.200	137,493.800

Notes

* $p < 0.001$. Robust standard errors in parentheses.

For increased interpretability of these model results, Figures 6 and 7 display the predicted probability of violent rhetoric by party and election year status respectively. These graphs allow us to visualize the magnitude of the change in the likelihood of an e-newsletter containing violent rhetoric based on the two variables. Figure 6 shows a relatively small but still substantively meaningful increase in the likelihood of violent rhetoric from Republicans to Democrats; there is about a 21.5% predicted probability of a Republican-authored e-newsletter containing violent rhetoric compared to just under 28% predicted probability for Democrat-authored ones. An increase of almost 6.5% is relatively sizeable, especially given the direction of the change. Figure 7 shows a substantively insignificant change in predicted probabilities for violent rhetoric between election years and non-election years. E-newsletters in election years have just over a 23.5% predicted probability of containing violent rhetoric compared to just under 21.5% in non-election years. This change of around 2%, especially given the intuitive directionality of the change, is not functionally meaningful.

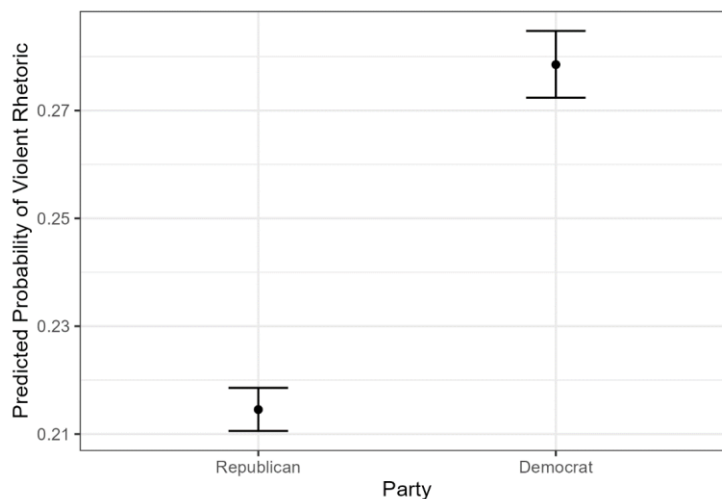


Figure 6: Predicted Probability of Violent Rhetoric by Party

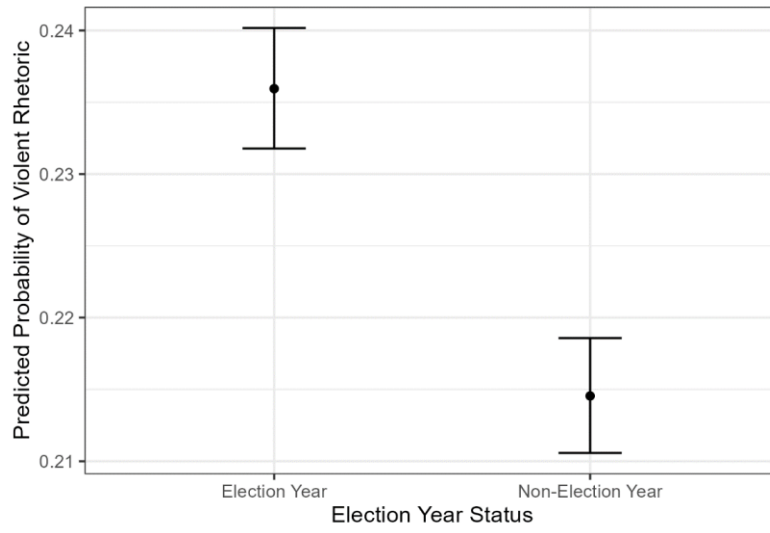


Figure 7: Predicted Probability of Violent Rhetoric by Election Year Status

7 DISCUSSION AND CONCLUSION

This study finds that violent political rhetoric is pervasive among email newsletters distributed by U.S. members of Congress, and its use increased over the course of the past decade. The prevalence of violent rhetoric in these e-newsletters changes in relation to the partisan affiliation of the author – Democrats are more likely to use violent rhetoric in e-newsletters than Republicans. The reason for this is not immediately clear; it may be the case that Democrats compensate for their lack of violent political rhetoric in other mediums in e-newsletters, as they are less likely to lose would-be supporters in e-newsletters as compared to other mediums of communication. As such, Democrats may use violent rhetoric more than Republicans in e-newsletters but not in other forums for communication, though this is highly speculative. Violent political rhetoric usage also increases based on proximity to elections at a statistically significant rate.

While e-newsletters are not representative of all forms of political speech, they are an important part of politician-constituent interaction; unveiling the use of violent rhetoric in these e-newsletters illustrates how politicians use this form of speech to elevate their political platform. These findings implicate our understanding of violent rhetoric in U.S. political speech, of how U.S. members of Congress use aggressive political speech broadly, and the role of researchers in identifying and understanding relationships between violent speech, political-constituent relationships, and acts of political violence. Much of the literature links Donald Trump, right-wing populism, and other aspects of the right-wing political movement in the rise of vitriolic politics and aggressive political communication, but the findings of this study show that Democrats also play a key role in this.

All political speech is strategic speech used to advance a politician's political agenda, disseminate information about a policy or campaign, or to dictate the narrative surrounding political issues and events. In this way, the use of violent rhetoric in political speech is a political tactic that calls upon the emotional sensibilities of the audience to incite political activity or drum up support for a political issue. However, such speech may have negative side effects of increasing partisanship and support for political violence as shown in the literature. In today's political landscape, the role of incendiary language on political violence and democratic erosion is highlighted more than ever by media outlets, politicians, and researchers. It is important to continue to advance our understanding of the use of violent rhetoric in political speech, especially in relation to incidents of political violence.

Future research related to violent political rhetoric may expand upon this research design to include speech from political elites beyond members of Congress, such as presidents, candidates, and media outlets, as well as speech from other mediums of communication, both online and in-person. For example, including social media posts, interviews, and speeches made by U.S. politicians may yield interesting results related to how usage of violent rhetoric varies across mode of speech and how its relationship with key variables changes depending on the intended target audience. Key differences may also exist between speech that is carefully curated, like e-newsletters, and speech that is typically more improvised, such as impromptu speeches or responses to interview and debate questions. Finally, future research should continue to analyze the relationship between violent rhetoric and incidents of violence in order to curb the role political speech in inciting political violence.

REFERENCES

- Abramowitz, A. I., & Webster, S. (2016). The rise of negative partisanship and the nationalization of U.S. elections in the 21st century. *Electoral Studies*, *41*, 12–22. <https://doi.org/10.1016/j.electstud.2015.11.001>
- Abramowitz, A. I., & Webster, S. W. (2018). Negative Partisanship: Why Americans Dislike Parties But Behave Like Rabid Partisans. *Political Psychology*, *39*(S1), 119–135. <https://doi.org/10.1111/pops.12479>
- Abramowitz, A., & McCoy, J. (2019). United States: Racial Resentment, Negative Partisanship, and Polarization in Trump's America. *The ANNALS of the American Academy of Political and Social Science*, *681*(1), 137–156. <https://doi.org/10.1177/0002716218811309>
- Anderson, C. A., & Bushman, B. J. (2002). Human aggression. *Annual Review of Psychology*, *53*(1), 27–51. <https://doi.org/10.1146/annurev.psych.53.100901.135231>
- Barberá, P., & Zeitzoff, T. (2018). The New Public Address System Why Do World Leaders Adopt Social Media. *International Studies Quarterly*, *62*(1), 121–130. <https://doi.org/10.1093/isq/sqx047>
- Blum, R., Cormack, L., & Shoub, K. (2023). Conditional Congressional communication: how elite speech varies across medium. *Political Science Research and Methods*, *11*(2), 394–401. doi:10.1017/psrm.2022.28
- Breiman, L. (2001). Random Forests. *Machine Learning*, *45*(1), 5–32. <https://doi.org/10.1023/A:1010933404324>
- Castaño-Pulgarín, S. A., Suárez-Betancur, N., Vega, L. M. T., & López, H. M. H. (2021). Internet, social media and online hate speech. Systematic review. *Aggression and Violent Behavior*, *58*, 101608. <https://doi.org/10.1016/j.avb.2021.101608>

- Cohen, J. E., Bond, J. R., Fleisher, R., & Hamman, J. A. (2000). State-Level Presidential Approval and Senatorial Support. *Legislative Studies Quarterly*, 25(4), 577–590.
<https://doi.org/10.2307/440435>
- Cormack, L. (2016). Extremity in Congress: Communications versus Votes. *Legislative Studies Quarterly*, 41: 575-603. <https://doi.org/10.1111/lsg.12126>
- Cormack, L., & Meidlinger, K. (2021). Congressional Communication in a Pandemic: “Follow the Leader” Politics and Responsive Representation. *Congress & the Presidency*, 49(3), 273–298. <https://doi.org/10.1080/07343469.2021.1955037>
- Curry, T. A., & Fix, M. P. (2019). May it please the twitterverse: The use of Twitter by state high court judges. *Journal of Information Technology & Politics*, 16(4), 379–393.
<https://doi.org/10.1080/19331681.2019.1657048>
- Davidson, T., Warmsley, D., Macy, M.W., & Weber, I. (2017). Automated Hate Speech Detection and the Problem of Offensive Language. *International Conference on Web and Social Media*. <https://doi.org/10.1609/icwsm.v11i1.14955>
- Elwood, W. (1995). Declaring War on the Home Front: Metaphor, Presidents, and the War on Drugs. *Metaphor and Symbol*, 10, 93–114. https://doi.org/10.1207/s15327868ms1002_2
- Evans, H. K., Cordova, V., & Sipole, S. (2014). Twitter Style: An Analysis of How House Candidates Used Twitter in Their 2012 Campaigns. *PS: Political Science & Politics*, 47(2), 454–462. <https://doi.org/10.1017/S1049096514000389>
- Flusberg, S., Matlock, T., & Thibodeau, P. (2018). War metaphors in public discourse. *Metaphor and Symbol*, 33, 1–18. <https://doi.org/10.1080/10926488.2018.1407992>

- Fisher, D. G., Dugan, L., & Chenoweth, E. (2019). Does US presidential rhetoric affect asymmetric political violence? *Critical Studies on Terrorism*, 12(1), 132–150.
<https://doi.org/10.1080/17539153.2018.1494120>
- Frimer, J. A., & Skitka, L. J. (2018). The Montagu Principle: Incivility decreases politicians' public approval, even with their political base. *Journal of Personality and Social Psychology*, 115(5), 845–866. <https://doi.org/10.1037/pspi0000140>
- Gervais, B.T. (2019), Rousing the Partisan Combatant: Elite Incivility, Anger, and Antideliberative Attitudes. *Political Psychology*, 40: 637-655. <https://doi.org/10.1111/pops.12532>
- Green, J., Shoub, K., Blum, R., & Cormack, L. (2024). Cross-Platform Partisan Positioning in Congressional Speech. *Political Research Quarterly*, 77(3), 653-668.
<https://doi.org/10.1177/10659129241236685>
- Henson, B., Reynolds, B. W., & Fisher, B. S. (2013). Fear of Crime Online? Examining the Effect of Risk, Previous Victimization, and Exposure on Fear of Online Interpersonal Victimization. *Journal of Contemporary Criminal Justice*, 29(4), 475-497.
<https://doi.org/10.1177/1043986213507403>
- Hodwitz, O., & Massingale, K. (2023). Rhetoric and hate crimes: Examining the public response to the Trump narrative. *Behavioral Sciences of Terrorism and Political Aggression*, 15(3), 303–320. <https://doi.org/10.1080/19434472.2021.1936121>
- Jackson, N. (2005a). Party E-Newsletters in the UK. *Journal of E-Government*, 1(4), 39–62.
https://doi.org/10.1300/J399v01n04_04
- Jackson, N. (2005b). Party E-Newsletters in the UK. *Journal of E-Government*, 1(4), 39–62.
https://doi.org/10.1300/J399v01n04_04

- Jackson, N. (2006). An MP's role in the internet Era – The impact of e-newsletters. *The Journal of Legislative Studies*, 12(2), 223–242. <https://doi.org/10.1080/13572330600739561>
- Jackson, N. (2008). MPs and their e-Newsletters: Winning Votes by Promoting Constituency Service. *The Journal of Legislative Studies*, 14(4), 488–499. <https://doi.org/10.1080/13572330802442923>
- Jackson, N. A., & Lilleker, D. G. (2007). SEEKING UNMEDIATED POLITICAL INFORMATION IN A MEDIATED ENVIRONMENT: The uses and gratifications of political parties' e-newsletters. *Information, Communication & Society*, 10(2), 242–264. <https://doi.org/10.1080/13691180701307495>
- Kalmoe, N. P. (2014). Fueling the Fire: Violent Metaphors, Trait Aggression, and Support for Political Violence. *Political Communication*, 31(4), 545–563. <https://doi.org/10.1080/10584609.2013.852642>
- Kalmoe, N. P., Gubler, J. R., & Wood, D. A. (2018). Toward Conflict or Compromise? How Violent Metaphors Polarize Partisan Issue Attitudes. *Political Communication*, 35(3), 333–352. <https://doi.org/10.1080/10584609.2017.1341965>
- Karpf, D. (2013). The Internet and American Political Campaigns. *The Forum*, 11. <https://doi.org/10.1515/for-2013-0051>
- Kaufman, S. J. (2001). *Modern Hatreds: The Symbolic Politics of Ethnic War*. Cornell University Press. <http://www.jstor.org/stable/10.7591/j.ctt1tm7gbw>
- Kim, T. (2023). Violent political rhetoric on Twitter. *Political Science Research and Methods*, 11(4), 673–695. <https://doi.org/10.1017/psrm.2022.12>
- Kleinfeld, R. (2021). The Rise of Political Violence in the United States. *Journal of Democracy*, 32, 160–176. <https://doi.org/10.1353/jod.2021.0059>

- Koop, R., & Marland, A. (2012). Insiders and Outsiders: Presentation of Self on Canadian Parliamentary Websites and Newsletters. *Policy & Internet*, 4(3–4), 112–135.
<https://doi.org/10.1002/poi3.13>
- Kreis, R. (2017). The “Tweet Politics” of President Trump. In *Journal of Language and Politics* (Vol. 16, Issue 4, pp. 607–618). John Benjamins.
- Lassen, D. S., & Brown, A. R. (2011). Twitter: The Electoral Connection? *Social Science Computer Review*, 29(4), 419-436. <https://doi.org/10.1177/0894439310382749>
- Mathew, B., Dutt, R., Goyal, P., & Mukherjee, A. (2019). Spread of Hate Speech in Online Social Media. *Proceedings of the 10th ACM Conference on Web Science*, 173–182.
<https://doi.org/10.1145/3292522.3326034>
- McGregor, S. C., Lawrence, R. G., & Cardona, A. (2017). Personalization, gender, and social media: Gubernatorial candidates’ social media strategies. *Information, Communication & Society*, 20(2), 264–283. <https://doi.org/10.1080/1369118X.2016.1167228>
- Meeks, L. (2016). Gendered styles, gendered differences: Candidates’ use of personalization and interactivity on Twitter. *Journal of Information Technology & Politics*, 13(4), 295-310.
<https://doi.org/10.1080/19331681.2016.1160268>
- Nugent, W. R., Abrams, T. E., & Joseph, A. A. (2022). The Relationship between Violent Political Rhetoric and Mass Shootings. *Journal of Social Service Research*, 48(2), 246–258.
<https://doi.org/10.1080/01488376.2021.2018089>
- Oddo, J. (2014). Variation and continuity in intertextual rhetoric: From the “War on Terror” to the “Struggle against Violent Extremism.” In *Journal of Language and Politics* (Vol. 13, Issue 3, pp. 512–537). John Benjamins.

- Olteanu, A., Castillo, C., Boy, J., & Varshney, K. (2018). The Effect of Extremist Violence on Hateful Speech Online. *Proceedings of the International AAAI Conference on Web and Social Media*, 12(1). <https://doi.org/10.1609/icwsm.v12i1.15040>
- Pape, R., & Chicago Project on Security and Threats. (2021). "Understanding American Domestic Terrorism: Mobilization Potential and Risk Factors of a New Threat Trajectory." *The University of Chicago*.
- Parker, G. R., & Parker, S. L. (1985). Correlates and Effects of Attention to District by U. S. House Members. *Legislative Studies Quarterly*, 10(2), 223–242.
<https://doi.org/10.2307/439725>
- Piazza, J. A. (2020). Politician hate speech and domestic terrorism. *International Interactions*, 46(3), 431–453. <https://doi.org/10.1080/03050629.2020.1739033>
- Pinker, Stephen. (2012). *The Better Angels of Our Nature: Why Violence Has Declined*. New York: Penguin Publishing Group.
- Shaw, A., & Benkler, Y. (2012). A Tale of Two Blogospheres: Discursive Practices on the Left and Right. *American Behavioral Scientist*, 56(4), 459-487.
<https://doi.org/10.1177/0002764211433793>
- Siegel, A. A. (2020). Online Hate Speech. In N. Persily & J. A. Tucker (Eds.), *Social Media and Democracy* (pp. 56–88). chapter, Cambridge: Cambridge University Press.
- Sydnor, E. (2019). *Disrespectful democracy: The psychology of political incivility*. New York: Columbia University Press.
- Tesler, M. (2016). *Post-racial or most-racial?: Race and politics in the Obama era*. University of Chicago Press.

- Thibodeau, P. H., & Boroditsky, L. (2011). Metaphors we think with: The role of metaphor in reasoning. *PLoS ONE*, 6(2), Article e16782. <https://doi.org/10.1371/journal.pone.0016782>
- Wahlström, M., Törnberg, A., & Ekbrand, H. (2021). Dynamics of violent and dehumanizing rhetoric in far-right social media. *New Media & Society*, 23(11), 3290–3311. <https://doi.org/10.1177/1461444820952795>
- Waldman, L. K. (1967). Liberalism of Congressman and the Presidential Vote in their Districts. *Midwest Journal of Political Science*, 11(1), 73–85. <https://doi.org/10.2307/2109327>
- Waseem, Z., & Hovy, D. (2016). *Hateful Symbols or Hateful People? Predictive Features for Hate Speech Detection on Twitter* (p. 88-93). <https://doi.org/10.18653/v1/N16-2013>
- Webster, S. W., & Abramowitz, A. I. (2017). The Ideological Foundations of Affective Polarization in the U.S. Electorate. *American Politics Research*, 45(4), 621–647. <https://doi.org/10.1177/1532673X17703132>
- Yanagizawa-Drott, D. (2014). PROPAGANDA AND CONFLICT: EVIDENCE FROM THE RWANDAN GENOCIDE. *The Quarterly Journal of Economics*, 129(4), 1947–1994. <https://doi.org/10.1093/qje/qju020>
- Zimmerman, S., Kruschwitz, U., & Fox, C. (2018, May). Improving hate speech detection with deep learning ensembles. In *Proceedings of the eleventh international conference on language resources and evaluation (LREC 2018)*.

