

Righting the Writers: Assessing Bias in Wikipedia’s Political Content — An Event Study and Sentiment Analysis[†]

By GUILLERMO PARRA*

This paper investigates the presence of political bias in Wikipedia through a causal inference framework. Utilizing a dataset of 1,399 politicians from the US, UK, and Canada and 271,400 historical snapshots of their Wikipedia pages, I employ an event study/staggered Difference-in-Differences (DiD) research design combined with a Large Language Model (LLM) for sentiment analysis. The analysis estimates the impact of being affiliated with right-wing versus left-wing parties on the sentiment of politicians’ Wikipedia pages. The findings reveal a statistically significant decrease in the sentiment of these pages following a switch to a more right-wing party, an effect that is not observed with switches to more left-wing parties. These results highlight Wikipedia’s potential ideological biases and continue the discussion on how media platforms influence public perception and discourse.

The idea of political bias and the role it plays in the media has been a long-discussed topic; individuals making these claims are often labelled “conspiracy theorists,” lacking the proper information to back up these ideas. However, in a digital age where information is as accessible as it is influential, Wikipedia has emerged not only as a major source of freely available knowledge but also as a significant influencer of public opinion (Santana, 2010). One notable instance occurred when a YouTube video titled “*Wikipedia is Biased!!*” was discovered in December 2022. The video argued that Wikipedia had a significant political bias and attempted to substantiate this claim with various anecdotes. Despite the video’s intriguing nature, it was challenging to fully accept its conclusions due to the heavy reliance on anecdotal evidence rather

[†]The processed dataset used in this study, along with the code used for data collection and analysis, will be made available in the following GitHub repository. @gillyparra

*Parra; University of British Columbia (UBC), Vancouver School of Economics (VSE) (email: gparra@student.ubc.ca) I first want to thank my paper advisor, Thomas Lemieux. His guidance on research direction and interest in the topic allowed me to select what I believe is a highly successful paper. I also want to thank my friends from back home, Michael Amodio, Ahsan Choudary, Dui Davidsson, Maria Del Corral, and Reid Martel, who took the time to read, critique, and ultimately improve this paper. A special thank you goes to Sein Jone Tao for his Wayback Machine repository; without his help, this paper would not have been possible. Lastly, I would like to thank all of my professors and friends from the Vancouver School of Economics. Through their interactions and teachings, I’ve learned more than I thought possible in a year, and I feel much more prepared for the world of economics.

than a data-based approach. This observation highlighted the need for a more empirical investigation. Over a year and a half later, an opportunity arose to explore this idea, inspired by the issues raised in the video and the desire to approach the topic from a more fact-based and data-driven perspective.

This study looks at a particularly contentious aspect of this influence: the potential political bias in Wikipedia. With Wikipedia’s model of open editing, concerns about biases, whether ideological or accidental, are widespread and perhaps somewhat justified (Ford and Wajcman, 2017). The significance of understanding these biases extends beyond academic interest, as it has the potential to influence public perceptions and democratic discourse.

Prior research has explored biases in digital media but often without the approaches needed for causal inference.¹ This paper extends these studies by conducting an econometric analysis to estimate the effect of changing political affiliation on the sentiment of politicians’ Wikipedia articles, this effect can be argued to be the effect of being right-wing on your Wikipedia page sentiment. Despite the large body of work on media bias, empirical studies examining Wikipedia’s political content are scarce (Dan Bernhardt, 2008). This research aims to fill that gap by analyzing if there is a systematic bias in Wikipedia entries shown in the re-alignments of political figures, utilizing an innovative dataset and a multi-time period Difference-in-Differences approach.

There are three main contributions of this study: it quantifies the extent and nature of political bias on Wikipedia, uses an event study to assess the causal impact of party switches on article sentiment, and discusses the broader implications of such biases on public knowledge and perception.

The paper hypothesizes that Wikipedia exhibits some political biases in its portrayal of political figures, as seen in many previous studies such as Greenstein and Zhu, 2018; Ackerly and Michelitch, 2022; Callahan and Herring, 2011. This bias could become particularly evident and pronounced following party switches. These biases may skew public perception, potentially distorting democratic and political engagement and altering the perception of what constitutes conventional knowledge. This study will begin with a review of the current literature on the topic, followed by a description of the data and the methods used to collect it. The next section presents background information on the notion that Wikipedia harbours biased political content, tracing its origins and evolution over time. It will also portray the main arguments as to why Wikipedia may or may not be biased. This is followed by the presentation of the empirical results and the interpretations of these findings. Finally, the study concludes with a summary and suggestions for avenues of future research.

¹Said research would only consist of surface-level analysis like the graphical representation of sentiment for right-wing vs left-wing politicians over time from Figure 11

I. Existing Literature

There is a long-standing literature regarding bias in media and, more recently, Wikipedia (Hamborg, 2022). It highlights the presence of systematic biases and their potential implications for the public and what is understood to be “knowledge.” This review goes over some relevant studies, focusing on Wikipedia bias, media bias, and media sentiment analysis.

A. Literature on bias in Wikipedia

The first ever recorded example of someone questioning Wikipedia’s neutrality was in 2005 when Joseph M. Reagle looked at the notion of neutrality in his essay *Is the Wikipedia Neutral?*, focusing on Wikipedia’s attempt to achieve neutrality through its Neutral Point of View (NPOV) policy. Despite Wikipedia’s efforts to represent all views fairly without bias, the essay reveals ongoing debates and misunderstandings about neutrality. He also published a blog article called *Can you Trust Wikipedia?* where he stresses that trustworthiness depends heavily on the quality of its sources (Jr., 2005b; Jr., 2005a).

The AER paper *Is Wikipedia biased?* by Greenstein and Zhu analyzes the slant of U.S. political topics on Wikipedia and finds that, initially, Wikipedia’s political entries leaned left. Over time, the bias diminished, although not due to revisions of existing articles but due to the addition of new articles with opposing slants (Greenstein and Zhu, 2012). They also have papers which investigate whether expert-produced or crowd-produced models generate bias, analyzing U.S. political content in Encyclopedia Britannica and Wikipedia. They found that Wikipedia articles are more biased towards left-wing viewpoints compared to Britannica (Greenstein and Zhu, 2018; Greenstein and Zhu, 2016).

Further papers on Wikipedia bias include *Wikipedia and political science: Addressing systematic biases with student initiatives*, which looks at knowledge gaps and biases in political science topics on Wikipedia (Ackerly and Mitchell, 2022). And *Exploring Systematic Bias Through Article Deletions on Wikipedia from a Behavioural Perspective* investigates the potential systematic bias on Wikipedia by examining article deletions and their relation to content of supposed interest to men and women. The study concludes that there is no significant systematic bias against content in terms of article deletions (Worku et al., 2020).

Cultural bias in Wikipedia content on famous persons by Callahan and Herring examines cultural bias in Wikipedia content by comparing articles about famous persons in the English and Polish versions. They find systematic differences in content and perspective, reflecting the distinct cultures, histories, and values of Poland and the United States (Callahan and Herring, 2011). Lastly, the paper *Edit Wars in Wikipedia* introduces a method for detecting

severe conflicts, known as “edit wars,”² In Wikipedia articles across six different languages. The study reveals that only a small fraction of Wikipedia articles are highly controversial (Sumi et al., 2012).

B. Literature using sentiment analysis to detect bias

The paper *Detecting biased statements in Wikipedia* by Hube and Fetahu attempts to detect biased statements in Wikipedia using a supervised classification approach. They propose an automated method for generating a biased word lexicon and demonstrate their model’s effectiveness in identifying biased statements with an accuracy of 74% (Hube and Fetahu, 2018). There is also *Is Wikipedia Politically Biased*, which examines political bias in English-language Wikipedia articles by analyzing the sentiment and emotional tone associated with politically charged terms. It finds that Wikipedia articles tend to associate right-of-center public figures with more negative sentiment and emotions compared to left-of-center public figures (Rozado, 2024a). *Forced transparency: Corporate Image on Wikipedia and What It Means for Public Relations* by DiStaso and Messner looks at the impact of Wikipedia on corporate image and public relations, analyzing Fortune 500 companies from 2006 to 2010. The study finds that Wikipedia articles about these companies have become more negative over time (DiStaso and Messner, 2010).

C. Literature on media bias

Political Polarization and the Electoral Effects of Media Bias by Bernhardt, Krasa, and Polborn develops a model demonstrating how media bias, arising from the profit-maximizing behaviour of media firms catering to partisan audiences, can lead to electoral mistakes. The model shows that even rational voters, aware of the bias, may make suboptimal electoral choices due to the suppression of critical information by biased media outlets (Dan Bernhardt, 2008). In the context of media bias, the Facebook-Cambridge Analytica scandal is a significant example of how data harvesting and targeted advertising can influence public perception and electoral outcomes. The incident involved the unauthorized collection of personal data from millions of Facebook users, which was then used to create psychological profiles and deliver highly personalized political ads (Rehman, 2019)

D. Research Contribution

My study would be the first to isolate the relationship between individuals identifying with left-wing or right-wing political parties on Wikipedia. When writing about politically charged topics, such as a politician’s Wikipedia page, personal bias can alter the tone and style of writing. This research looks to

²Edit Wars are repeated, contentious revisions of content among editors involving back-and-forth changes.

determine whether there is indeed a difference between being perceived as left or right-wing.

II. Background

A. Exploring the Origins of Alleged Bias in Wikipedia

A.1. The Early Years: Foundation and Initial Challenges (2001–2006)

The idea that Wikipedia is biased began to take shape almost from its inception. Founded in 2001, Wikipedia was created on the principle of allowing anyone to edit its content. This openness was to democratize information, but it also introduced the potential for bias. Some of its early criticisms came from concerns about the accuracy and neutrality of its articles. It was acknowledged that the collaborative model enabled a diverse range of contributions, but it also meant that articles could reflect the personal biases of individuals (Jr., 2005b).

One of the earliest documented concerns about bias in Wikipedia was the issue of “POV pushing,” where contributors would insert their personal points of view into articles. This was a problem for Wikipedia’s goal of providing neutral and balanced information. The platform’s foundational policy, the Neutral Point of View (NPOV), was designed to counteract this by requiring that articles be written without bias, representing all significant views fairly (Wikipedia contributors, ndb).

The interpretation and enforcement of NPOV itself became a point of contention. Its critics argued that what constituted a “neutral” perspective was often just the dominant cultural and social views of most contributors. Since, at the time, Wikipedia’s editor base was predominantly composed of young, white males from Western countries, there were concerns about systemic bias in the representation of certain topics. This demographic skew led to an overrepresentation of topics and viewpoints prevalent among this group (Cohen; Sanger, 2011; 2005).

These early observations and criticisms laid the groundwork for ongoing debates about the trustworthiness and bias of Wikipedia. Highlighting the challenges of maintaining neutrality in a platform where content could be influenced by the personal and collective biases of its contributors. As a result, the idea that Wikipedia might be biased persisted (Wikipedia contributors; Lih, nda; 2009).

A.2. Gaining Momentum: Increasing Popularity and Distrust (2006–2010)

Between 2006 and 2010, Wikipedia became one of the most influential online platforms, becoming a ubiquitous source of information. This period saw significant growth in its user base and content, but it also saw a corresponding rise in mistrust regarding its reliability. Wikipedia’s allowing anyone to edit articles became a double-edged sword, creating widespread concerns about the quality of its content.

A notable moment in 2006 was the creation of Conservapedia. Founded by conservative activist Andy Schlafly, it looked to counter what he perceived as Wikipedia’s liberal bias. He argued that Wikipedia’s editors censored conservative viewpoints. He pointed out very specific issues, such as the alleged lack of credit given to Christianity for the Renaissance (Johnson; Zeller, 2007; 2007).

In 2007, the development of WikiScanner by Virgil Griffith further fueled mistrust. WikiScanner exposed instances where individuals and organizations edited Wikipedia entries to serve their interests, revealing how the platform could be manipulated. High-profile cases included edits from computers at Anheuser-Busch, PepsiCo, and Diebold, where content critical of these companies was altered or removed (Hafner, 2007).

All this coincided with criticism from established institutions and public figures. Dale Hoiberg, the editor-in-chief of Encyclopaedia Britannica, scrutinized a study published in *Nature* that claimed Wikipedia was nearly as accurate as Britannica. Journalist John Seigenthaler criticized Wikipedia after defamatory comments about him were published, showing the potential for serious reputational damage from inaccurate entries (Zeller, 2007).

Jay Richards, in his article for AEIdeas, said that Wikipedia’s bias was not random but consistently leaned towards liberal viewpoints. He used anecdotes where Wikipedia’s treatment of controversial subjects, such as climate change and political figures, appeared to favour left-wing perspectives (Richards, 2009).

Despite this, Wikipedia continued to grow. Its high visibility from top positions in search engine results allowed it to remain a primary reference for millions of people (Ford and Wajcman; Santana, 2017; 2010).

A.3. Maturity: Addressing Bias and Expanding Influence (2011–2019)

During the period from 2011 to 2019, Wikipedia continued its ascent and became one of the most mainstream sources of information. However, the concerns regarding bias and reliability did not go away.

In 2015, Zhu explored the comparative bias of Wikipedia and Encyclopedia

Britannica. The research concluded that Wikipedia contained more politically charged language. Despite this, articles that underwent more revisions showed reduced bias, indicating that the crowd-sourced nature of Wikipedia could, over time, balance out initial biases (Zhu, 2015). In 2016, research solidified this idea, as another study found that articles trended towards neutrality as they received more edits. This shift shows the platform’s capacity for self-correction through its collaborative editing process (Bhattacharya, 2016).

In 2018, Poppy Noor’s article in *The Guardian* identified five major biases in Wikipedia: gender, western, language, political, and historical biases. Noor emphasized the male-dominated and western-centric nature of the platform (Noor, 2018). Another study found that nearly half of all edits to location-focused articles on Wikipedia were made by people in France, Germany, Italy, the UK, and the US. This concentration of editorial power from people in high-income countries resulted in a disproportionate representation at the expense of voices from lower-income countries (Temperton, 2015). Martin Körner and Tatiana Sennikova showed that the English version of an article on Russia’s annexation of Crimea had a different balance of references from Ukrainian and Russian sources compared to the German version (Reynolds, 2016).

Kalla and Aronow further investigated the issue of editorial bias in political information. Their experiments on Wikipedia pages of U.S. senators found a bias towards positivity, with negative facts being more likely to be removed than positive ones (Kalla and Aronow, 2015).

A.4. Recent Developments: The Rise of Conspiracy Theories (2020–2024)

The period from 2020 to 2024, marked by the COVID-19 pandemic, heightened fears of misinformation and led to more centralized control over specific articles on the platform. This section explores the recent pushback against the theory of Wikipedia’s political bias.

The COVID-19 pandemic, starting in early 2020, triggered an unprecedented wave of suspected global misinformation. As a result, Wikipedia implemented stricter editorial controls to ensure the accuracy of its content related to the pandemic. These measures, while intended to curb misinformation, were perceived by some as a move towards centralized control, particularly over politically sensitive articles (Barnard, 2020). The conflict between Russia and Ukraine, which began escalating in 2021, intensified these accusations. According to former Wikipedia editor Arseny Natapov, the platform exhibited an increasingly anti-Russian stance. He claimed that participants with pro-Russian views were blocked, and articles highlighting Russian achievements were deleted. Natapov suggested that many admins were Ukrainian or resided in EU countries (Natapov, 2023).

A study from the *Manhattan Institute* in 2024 found evidence that right-of-center public figures in the U.S. were depicted more negatively compared to their left-of-center counterparts. This bias extended to the language used in articles, with terms associated with negative sentiment more frequently linked to conservative figures (Rozado, 2024b). Former editor Jonathan Weiss also noted that the platform showed a clear bias in the selection of sources deemed reliable. Right-leaning news outlets such as Fox News and The Daily Caller were often labelled as unreliable, while left-leaning sources like CNN, MSNBC and even VOX were considered trustworthy (Stossel, 2022).

Critics began labelling Wikipedia as part of a larger conspiracy to promote left-leaning ideologies. Larry Sanger, Wikipedia’s co-founder, was particularly vocal, describing the platform as “propaganda for the left-leaning establishment.” Sanger said that conservative viewpoints were systematically excluded or downplayed while liberal perspectives were prominently featured (Sanger, 2021c).

Furthermore, an investigation by the Telegraph in 2023 highlighted the influence of powerful editors and administrators who controlled the narrative on Wikipedia. These individuals were accused of using their positions to enforce a particular political agenda (de Quetteville, 2023). The reaction to these criticisms was mixed. On the one hand, supporters of Wikipedia’s policies argued that the measures were necessary to combat misinformation, especially during a global health crisis (Barnard, 2020). On the other hand, critics saw these policies as a means to suppress dissenting views and control the narrative.

As the debate continued, the influence of Wikipedia on public perception remained large. Studies showed that Wikipedia articles often influenced the content of AI language models, further amplifying the impact of any potential bias. This period highlighted information control, public trust, and the ongoing battle over perceived political bias on one of the world’s most visited websites (Mastrine, 2024).

B. Perceptions of Bias: Arguments Supporting Wikipedia’s Partisanship

Wikipedia, Founded by Jimmy Wales and Larry Sanger, began with a libertarian vision, emphasizing spontaneous order and crowd-sourced contributions. While this model made Wikipedia one of the most-used websites globally, recent critiques highlight a concerning shift towards political bias, particularly favouring left-leaning perspectives.

An example of this bias can be seen in the treatment of the “Hunter Biden Email Controversy.” Initially, Wikipedia redirected searches for this topic to the “Biden-Ukraine Conspiracy Theory” article, mirroring the narrative pushed by

left-leaning media outlets that dismissed the story as “Russian disinformation” (Sanger, 2021a). This stance persisted even after mainstream media validated the authenticity of the emails (Times, 2022).

The bias is further evident in Wikipedia’s list of reliable sources. Openly left-leaning media outlets are seen as reliable, while openly right-leaning sources are not. Jon Weiss, a prominent Wikipedian, has also observed this trend, noting that while Wikipedia excels in areas like science and sports, it shows significant bias in its coverage of current political events (Weiss, 2022).

Administrators, who have significant power over content, often openly identify as socialists or Marxists. They use their authority to protect left-leaning content and suppress right-leaning edits. For example, descriptions of the Antifa movement are minimized in terms of violence, and attempts to label it as a far-left movement are swiftly removed (Sanger, 2021a). This administrative bias extends to the portrayal of historical topics. Articles on socialism and communism have historically downplayed the atrocities committed under these regimes, emphasizing any perceived benefits while minimizing or ignoring the significant human costs (Weiss, 2022).

Attempts to correct these biases are frequently met with resistance. Edits that introduce balanced perspectives or highlight leftist extremism are often quickly reverted. This creates a hostile environment for editors who do not align with the dominant ideological stance, discouraging broader participation and reinforcing the existing bias (Sanger, 2021a). Critics say that the root of Wikipedia’s bias lies in the ideological homogeneity of its most active editors and administrators. The overrepresented presence of left-leaning contributors leads to a natural skew in the content (Weiss, 2022).

The difference between neutrality and objective truth is fundamental to understanding this issue. Objective truth is generally a point of contention, as people often disagree on what constitutes objective truth. Sanger says that it would be ideal to have a reference containing only objective truths. Although he acknowledges that this would be impossible because no two people will ever agree on everything (Sanger, 2021b). Neutrality, on the other hand, attempts to explain all different points of view on a subject with sufficient detail and evidence for readers to form their own opinions. Neutrality involves presenting a wide range of views rather than deciding the facts for the reader. Writing that promotes a single point of view resembles propaganda, especially to those who do not share that perspective. Propaganda aims to alter beliefs without considering alternative points of view, making people less informed and less objective (Sanger, 2021b).

To conclude, while Wikipedia was founded on principles of neutrality and collective knowledge, its current trajectory shows a tilt toward left-leaning

bias. If the reader is convinced by the arguments above, they should approach politically charged topics on Wikipedia with a healthy dose of skepticism.

C. Counterpoints: Defending Wikipedia's Neutrality

Accusations that Wikipedia has a left-leaning political bias often arise, but these claims fail to account for the mechanisms and policies that Wikipedia employs to maintain neutrality. The Neutral Point of View (NPOV) policy is central to Wikipedia's editorial guidelines, mandating that all articles must be written without bias, fairly representing all significant viewpoints based on reliable and verifiable sources (Wikipedia, 2024h). Examples can always be found at specific points in time of potential bias, although by and large, those examples do not last long and are promptly fixed by users.

Wikipedia's open-editing model, which allows anyone to contribute to the overwhelming majority of articles, is a fundamental aspect that supports its neutrality. This approach makes sure that a diverse range of perspectives are considered. When biases are introduced, they can be quickly identified and corrected by other editors. This continuous peer-review process creates a balanced representation of information. Wikipedia's community is known to be vigilant and proactive in addressing biases, engaging in discussions and reaching consensus to resolve disagreements and ensure articles reflect a neutral point of view (Wikipedia, 2024b).

The requirement for verifiability means that all information in Wikipedia articles must be supported by sources. Wikipedia's guidelines say that editors should avoid stating opinions as facts and should attribute conflicting viewpoints to their sources (Wikipedia, 2024k). This ensures that articles do not present opinions as truths, maintaining an objective tone.

Wikipedia's handling of contentious topics further demonstrates its commitment to neutrality. For example, during the COVID-19 pandemic, the discussion around the efficacy of masks was highly debated. Wikipedia navigated this by presenting the prevailing scientific consensus alongside significant dissenting opinions, accurately reflecting the state of the debate (Wikipedia, 2024c). This approach prevents the platform from giving weight to fringe theories or marginal perspectives, ensuring that mainstream viewpoints are proportionately represented (Wikipedia, 2024f).

Critics often cite single entries at specific moments to argue that Wikipedia is biased. However, these usually involve topics where mainstream and fringe views clash. Wikipedia's policy is to represent mainstream views proportionately while acknowledging significant minority opinions (Wikipedia, 2024e). This gives way to the idea of due weight, which requires that the prominence of each viewpoint in the article reflects its prevalence in reliable sources. Thus,

fringe theories are not given the same level of coverage as widely accepted views, preventing a distorted representation of the subject matter (Wikipedia, 2024e).

Misunderstandings about how Wikipedia operates often create accusations of bias. The platform’s editorial structure and policies are designed to prioritize factual accuracy and neutrality over individual beliefs. While editors may have personal biases, the collective editing process acts as a counterbalance (Wikipedia, 2024j). This ensures that no single viewpoint can dominate, preserving the integrity of the content.

Moreover, Wikipedia has mechanisms to address and resolve biases. The Dispute Resolution process allows editors to collaboratively resolve conflicts through discussion and consensus (Wikipedia, 2024d). This process, combined with the supervision of experienced editors and administrators, helps maintain the neutrality of articles. Additionally, the NPOV Noticeboard is a platform where concerns about potential biases can be raised and addressed by the community, further reinforcing the commitment to neutrality (Wikipedia, 2024i).

The broader context of how truth is perceived and evolves over time also plays a role in understanding Wikipedia’s approach. Wikipedia is not the ultimate arbiter of truth but rather a repository of the best knowledge available at any given time, reflecting the evolutionary and unstable nature of “truth”. Historically, ideas that were once considered fringe, like Copernicus’s heliocentric theory, took centuries to gain acceptance. Showing that truth is often contested and evolves with new evidence and perspectives (Wikipedia, 2024a).

Wikipedia’s model acknowledges this complexity. It operates on the principle that what we know now is based on the best available evidence and is subject to change as new information emerges. This is critical in a world where public understanding and consensus are continually evolving. By presenting information that is verifiable and sourced from reliable references, Wikipedia provides a balanced view that accommodates the diversity of perspectives (Wikipedia, 2024g).

In conclusion, Wikipedia’s mechanisms and policies work together to prevent political bias. The platform’s commitment to neutrality and verifiability, along with processes for addressing and resolving biases, make sure that it remains a balanced and reliable source of information.

III. Data

This study uses data sourced from Wikipedia and the Wayback Machine, accessed through the Wayback Machine API. The Wayback Machine is a digital archive of the World Wide Web, allowing users to access historical

snapshots of web pages. It provides a vast collection of archived web content, including old versions of Wikipedia pages, which can be used to track changes over time. Wikipedia was chosen as the main data source for several reasons:

- It has been a primary source of information for the past decade (Santana, 2010).
- It offers comprehensive coverage of numerous politicians.
- The availability of historical snapshots allows for a temporal analysis.

Typical Wikipedia entries for politicians contain information on their early life, political career, and potentially post-political career activities.³ These entries often include achievements, community impacts, and other publicly available relevant information.

A. Data Collection and Processing

To gather this data, two lists were created: one containing all politicians who switched political parties between 2004 and 2024, and another containing a random selection of politicians from the United States, Canada, and the United Kingdom who were active during a large portion of that time frame and did not switch parties. The datasets included the date of the party switch, the old party and the new party, as well as names, countries, and states of politicians.⁴

Python functions were created to retrieve the Wikipedia URLs for each politician using their name, country, state, and party, as well as for other data-cleaning purposes. The main analysis script was used to scrape Wikipedia entries through the WayBack machine for each politician from 2004 onward, with chunking used to perform sentiment analysis on each snapshot. Finally, a dataset was obtained with the percentage of the article’s positive, neutral, or negative sentiment of each available snapshot for each politician. However, many challenges were faced during this process.⁵

B. Variables and Measures

The dataset was matched by name to include the date of each observation, the politician’s name, the URL used, the overall sentiment score (calculated as positive sentiment percent minus negative sentiment percent)⁶ the word count of each Wikipedia page,⁷ indicators of a political move, as well as all the above

³For a graphical representation of sentiment over for all politicians, see Figure 8

⁴For further information, refer to Appendix Sections A.1

⁵For further information, refer to Appendix Sections A.2

⁶Neutral was ignored as we are looking to focus on the bias. See more on Sections A.4

⁷To see the average word count of politicians Wikipedia pages split by country refer to figure 13

information. Additionally, a move from any party to being independent was categorized as a shift to the right.⁸

C. Data Statistics

This study uses data on 1,399 different politicians from the United States, Canada, and the United Kingdom, covering the period from 2004 to 2024. The sample consists of:

- A control group of 1,078 politicians who never switched parties
- A treatment group of 321 politicians who switched parties, further divided into:
 - 43 Politicians who moved more toward the left
 - 256 Politicians who moved more toward the right
 - 23 politicians who made equal moves (neither more left nor more right)⁹

Group	N	Average Score	Word Count	Party Alignment
All Politicians	1399	-0.02 (0.00)	3724.60 (102.06)	2.49 (0.04)
Pre Treatment	225	-0.01 (0.01)	2664.26 (251.41)	2.61 (0.09)
Treatment	321	-0.05 (0.01)	2525.12 (181.14)	2.56 (0.08)
Control	1078	-0.02 (0.00)	4112.29 (118.76)	2.47 (0.04)
Right Wing	677	-0.04 (0.00)	3689.00 (132.89)	3.87 (0.01)
Left Wing	704	-0.01 (0.00)	3716.71 (150.92)	1.16 (0.01)
USA	800	-0.02 (0.00)	4562.59 (138.37)	2.51 (0.05)
Canada	228	0.03 (0.01)	1774.65 (120.21)	2.28 (0.06)
UK	371	-0.07 (0.01)	2722.39 (174.24)	2.58 (0.08)

Table 1: Summary Statistics and Balance Table

D. Dataset Structure

Finally, the dataset was split into two: one combining the control group with only the politicians who moved to the right and the other combining the control group with only the politicians who moved to the left.

⁸All politicians who became independents due to candles in their personal lives were removed to no bias in the results. For further information, refer to Appendix Section A.3. To see the average sentiment over time of the treated and control group, see Figure 9

⁹Party Alignment $\in [1, 4]$ with 1 being a far-left party and 4 being a far-right party. In these cases, if a politician moved from a party with an alignment score equal to the new party score, it was considered an equal move. To see sentiment over time split by politician party alignment see 12 and Figure 14

E. Data Limitations and Considerations

While Wikipedia provides a rich source of data on politicians, it’s important to note potential limitations.

- Possible biases in Wikipedia coverage of politicians, which may vary based on notability or other factors.
- Potential biases in the sentiment analysis, which depend on the training data and methodology of the sentiment analyzer used.
- The reliability and completeness of Wikipedia entries may vary.

These caveats should be considered when interpreting the results of analyses based on this dataset.

F. Ethical Considerations

This study uses publicly available data from Wikipedia. While this mitigates many privacy concerns, care has been taken to use this data responsibly and to not misrepresent the politicians in the study.

IV. Research Design and Econometric Methods

An Event study/Staggered Difference-in-Difference can provide a framework for evaluating the causal effect of specific events on an outcome variable with panel data, which is useful in settings where the timing of the event varies across individuals (Callaway and Sant’Anna, 2021).

A. Applying Event Study to Political Bias in Wikipedia

I attempt to gauge whether a causal relationship exists regarding political bias on Wikipedia. Specifically, looking at estimating the effects post-political party switch, considering that sentiment should not change significantly if unbiased (Callaway and Sant’Anna, 2021). The control group consists of politicians who have never switched parties. The treatment group includes politicians who have switched parties, further divided into those moving more to the left and those moving more to the right. This distinction aims to identify any asymmetry in sentiment changes related to the direction of the switch.

B. Specification of the Model

$$Y_{itl} = \alpha + \sum_{k=-K}^K \beta_k D_{itl}^k + \gamma X_{itl} + \delta_i + \tau_t + \phi_l + \epsilon_{itl} \quad (1)$$

- Y_{itl} represents the sentiment score for politician i at time t , in location l

¹⁰The sentiment is obtained through a multilingual sentiment analysis API designed to process and classify sentiments across various languages (Yuan, 2023).

- α represents the baseline sentiment score for the reference group, assuming no party switch and controlling for other variables set to zero.
- D_{itl}^k are indicator variables for k periods before and after the party switch.
- β_k is our variable of interest and represents the change in sentiment score associated with the k periods before and after a politician switches parties.
- X_{itl} includes control variables such as the word count of a politician’s Wikipedia page and their original party alignment.
- δ_i represents fixed effects for individual politicians, controlling for their intrinsic characteristics that do not vary over time.
- τ_t represents time-fixed effects, accounting for global or national factors affecting all politicians at a particular time.
- ϕ_l is location fixed-effects, controlling for country and state influences.
- ϵ_{itl} is the error term, representing random noise affecting the sentiment.

Here the model incorporates fixed effects to control for unobserved heterogeneity that could bias our estimates (Abadie, 2005). The treatment effect in our context is defined by the party-switching event. We analyze the sentiment change by comparing politicians who switched parties to those who did not, using a staggered treatment design. This approach is particularly useful in settings like this, where treatment—here, the party switch—is homogeneous to all units but occurs at various times across the treated units (Callaway and Sant’Anna, 2021).

C. Motivation for Using Event Study Designs

The nature of the intervention drives the decision to employ an event study design. Political party switches are discrete events that potentially alter public and media perceptions instantaneously. The event study approach is normally used for analyzing the effects of such time-stamped events, providing a clear before-and-after comparison across many periods (Freyaldenhoven et al., 2019).

D. Methodological Considerations and Challenges

The model accounts for fixed effects to control for unobserved heterogeneity across politicians and time. Robust standard errors are clustered by politicians to address potential autocorrelation issues. The main challenge lies in the irregular timing of observations, which gave way to the need to use observational intervals rather than fixed time intervals.¹¹

¹¹Observational intervals are based on when data points are actually collected, varying irregularly, unlike fixed time intervals that are consistent, like weekly or monthly. Unfortunately, research on this in Difference-in-Differences frameworks is not available; most studies focus on time series contexts.

V. Results

This section presents the findings from our event study. It looks to examine the sentiment changes on Wikipedia for politicians who switch political parties to determine whether Wikipedia exhibits political bias.

A. Impact of Shifting to a More Right-Wing Party

The results, summarized in Table 2 and illustrated in Figure 1, reveal a consistent, negative sentiment shift for politicians moving to more conservative political parties. Immediately upon the party switch, there is a statistically significant decline of about 0.02 in sentiment scores. This maps to an overall drop of 2%.¹² This decline not only persists but also amplifies; looking at coefficients B1 to B5, it increases and then stabilizes. These changes suggest a bias against these politicians, with the largest decrease occurring at the fifth interval (B5)¹³.

Table 2: Staggered DiD Sentiment Analysis for Shifting Right

	Model 1 (No Controls)	Model 2 (Word Count)	Model 3 (Party Alignment)	Model 4 (Both Controls)
B0	-0.01888*** (0.00467)	-0.02017*** (0.00533)	-0.01871*** (0.00479)	-0.02059*** (0.00547)
B1	-0.05175*** (0.00565)	-0.04931*** (0.00612)	-0.05222*** (0.00582)	-0.05003*** (0.00634)
B2	-0.05766*** (0.00625)	-0.05759*** (0.00676)	-0.05902*** (0.00643)	-0.05898*** (0.00697)
B3	-0.05719*** (0.00628)	-0.05651*** (0.00697)	-0.05829*** (0.00647)	-0.05767*** (0.00720)
B4	-0.06493*** (0.00690)	-0.06561*** (0.00759)	-0.06605*** (0.00709)	-0.06699*** (0.00784)
B5	-0.22886*** (0.02700)	-0.22150*** (0.02223)	-0.22945*** (0.02718)	-0.22059*** (0.02156)
WC		-0.00000312* (0.00000102)		-0.00000356* (0.00000114)
PA			-0.00838*** (0.00002)	-61.48216*** (0.22977)

Examination of pre-trends in Figure 1 suggest that the model assumptions hold, with pre-shift sentiment levels remaining stable and centred around

¹²This is because the sentiment analysis originally outputs a percentage of positive, negative and neutral sentiment in a Wikipedia page; we can move forward looking at all the results in percentage change as well as the marginal effect.

¹³Upon analyzing the data points with five post periods, it was evident that the reason for this massive drop in sentiment was the small sample size exacerbated by one very negative observation. The reader may ignore this result at their discretion.

zero. This stability supports the potential reliability of the findings regarding post-right-wing shifts in sentiment changes.




Figure 1: Staggered DiD Sentiment Analysis for Shifting Right

Given the observed sentiment declines associated with politicians switching to more right-wing parties, proponents of the view that Wikipedia has a left-wing bias might interpret these results as proof of such bias.

However, it is important to consider alternative explanations that can also have plausible reasoning. One hypothesis could be that switching parties, regardless of direction, may inherently provoke a negative sentiment. This could be attributed to public surprise or displeasure over a politician’s realignment. This could reflect a general preference for political stability and loyalty. The following section further explores this idea by analyzing the sentiment changes for politicians who shift to the left, providing a comparative analysis to assess if the negative sentiment is a general response to party switching rather than an indication of ideological bias.

B. Impact of Shifting to a More Left-Wing Party

Here, in contrast, shifts to the left, analyzed in Table 3, and illustrated in Figure 2 do not exhibit the same pattern. The coefficients for politicians who move to more left-wing parties’ are negative but much smaller than those

who move to the right; they also lack statistical significance, suggesting that moving parties that are more left wing does not affect the sentiment of a politician’s Wikipedia page negatively. This difference in sentiment changes between right and left shifts aligns with the belief of a left-wing bias within Wikipedia. However, the results suggest that this perceived bias is driven more by negative sentiment towards right-leaning political opinions rather than positive sentiment for left-wing politics.

Table 3: Staggered DiD Sentiment Analysis for Shifting Left

	Model 1 (No Controls)	Model 2 (Word Count)	Model 3 (Party Alignment)	Model 4 (Both Controls)
B0	-0.01147 (0.00985)	-0.01269 (0.01081)	-0.01216 (0.01058)	-0.01341 (0.01165)
B1	-0.00687 (0.01252)	-0.00298 (0.01252)	-0.00665 (0.01356)	-0.00274 (0.01362)
B2	-0.01236 (0.01289)	-0.00900 (0.01322)	-0.01233 (0.01389)	-0.00827 (0.01429)
B3	-0.01579 (0.01295)	-0.01337 (0.01287)	-0.01589 (0.01399)	-0.01351 (0.01394)
B4	-0.00925 (0.01455)	-0.00829 (0.01469)	-0.00865 (0.01583)	-0.00755 (0.01600)
WC		-0.00000312** (0.00000102)		-0.00000356** (0.00000114)
PA			-0.02428*** (0.000005)	-43.52602*** (0.34420)

Examining Figure 2, in contrast to the results shown in Figure 1, it becomes apparent that the sentiment shifts for politicians moving to a left-wing party are substantially smaller, with standard errors including zero in their range. This suggests that any changes in sentiment are negligible when compared to those observed for shifts towards right-wing parties. Moreover, the pre-trend analysis in Figure 2 reinforces our parallel pre-trends assumption, as it depicts trends that closely hover around zero, indicating a stable pre-event sentiment across our sample.

The comparative analysis of shifts to the right and left reveals a significant asymmetry. While rightward shifts lead to pronounced negative sentiment, leftward shifts do not seem to incur those same sentiment penalties. This observation challenges the general perception of Wikipedia as unbiased, as the evidence above suggests a bias against right-wing politics.¹⁴

¹⁴Please refer to section C in the Appendix for individual tables and graphs of these results




Figure 2: Staggered DiD Sentiment Analysis for Shifting Left

C. Robustness Checks

C.1. Excluding Politicians Who Became Independent

To address potential selection bias, a robustness check was conducted by excluding politicians who shifted to an independent. Here the claim that politicians who move to more right-wing parties explicitly have lower sentiment score coefficients is tested. This adjustment left us with 89 individuals, compared to the over 200 initially included in the main analysis. The results showed a consistent pattern of negative sentiment scores for these individuals, with statistically significant coefficients for post-treatment periods. For instance, the coefficient for B4 was -0.0775 ($p < 0.001$), indicating a substantial negative shift in sentiment. These results very closely mirror the results from the main analysis. This consistency suggests that the negative sentiment associated with rightward shifts remains robust even when excluding independents. The robustness of these results also supports the claim that becoming independent is often associated with a rightward ideological shift.¹⁵

¹⁵For these results, Refer to Figure 7 and Table 10 in the Appendix

C.2. Only Including Politicians Who Became Independent

Further analysis was conducted on a subset of 167 politicians who transitioned to independent status, which constituted the majority of the treatment group. The findings revealed a similar trend of negative sentiment, though with less statistical significance and lower magnitude compared to those moving to right-wing parties. For example, the coefficient for B3 was -0.037 ($p = 0.058$), which is only statistically significant at the 10% level. This decrease in significance and magnitude may indicate a less pronounced sentiment change for those becoming independent. Nonetheless, the overall negative trend reinforces the argument that a rightward shift, including becoming independent, generally results in a decrease in favourable sentiment, matching and strengthening the main findings.¹⁶

C.3. Centrality Shift Analysis

The reader may notice that every estimation coefficient so far, whether it be statistically significant or not, was near zero or negative. This could lead them to believe that perhaps the true value of these coefficients is negative for any party switch. Here, this is disputed by looking at politicians who move to a more central-leaning party. (e.g., from far-right to right-leaning or far-left to left-leaning). This analysis included 23 politicians. The results showed a positive but statistically insignificant effect on sentiment scores, with coefficients ranging from 0.004 to 0.053 (this maps to 0.4% to 5.3%). For instance, the coefficient for B4 was 0.053 ($p = 0.764$), indicating a slight positive shift in sentiment, though without strong statistical backing.

This outcome suggests that it is possible that centrism may not evoke the same negative sentiment associated with more polarized shifts. The lack of statistical significance, possibly due to the small sample size, prevents definitive conclusions but hints at a potentially different narrative for centrists compared to more extreme ideological movements. While also disputing any potential claims that all estimations are negative.¹⁷

C.4. Investigation of Pretrends

The next robustness check focused on investigating the pre-trends. The pre-trends appear to be very close to zero, with near-zero confidence intervals. This is rarely observed in real-world data, though there is a reasonable explanation and analysis to address these concerns. Firstly, for many politicians, each subsequent observation, despite occurring at a different time, has a sentiment score similar or identical to the previous observation. This phenomenon arises because politicians' Wikipedia pages typically change minimally; updates usually occur only when a new event transpires. By examining the average

¹⁶For these results Refer to Figure 6 and Table 9 in the Appendix

¹⁷For these results Refer to Figure 5 and Table 8 in the Appendix

change in sentiment scores observation over observation, grouped by individual politicians, an average change of -0.002 with standard errors of 0.0008 was observed. This analysis encompasses all politicians over the entire time frame, the estimation could be even smaller if only politicians who have never been treated or have not yet been treated are considered. In the absence of major events or shifts, there is little change in sentiment scores from one observation to the next. This is particularly evident when one limits the number of observations per period to four pre- and four post-treatment periods. This analysis helps explain why the pre-periods exhibit such negligible changes: without significant events, the alterations in Wikipedia pages for politicians are minimal.

C.5. Correlational Analysis on Right-Wing Alignment and Sentiment

The final robustness check involved a correlational analysis to ensure that the overall picture of the data was accurately captured. Examining the relationship between being right-wing and sentiment scores. According to Figure 15 correlation matrix, there is a negative correlation between right-wing party alignment and sentiment scores. Additionally, analyzing political alignment more broadly Using a standard OLS regression with sentiment score as the dependent variable and party alignment as the independent variable and found a statistically significant coefficient of -0.0058 (Figure 16 and Table 11). This indicates that as a politician’s alignment becomes more right-wing, there is an average marginal effect of a half-percentage point decrease in sentiment. These negative correlations are consistent with the figures presented in the appendix and align with existing literature, reinforcing the notion that right-wing political alignment is associated with a decrease in favourable sentiment on Wikipedia.

D. Final Results Statements

These results, together with the results from (Greenstein and Zhu, 2012), who found that there was a tendency towards a more neutral point of view over time in Wikipedia articles, suggest that while immediate biases are evident, the platform may self-correct over time. These ever-changing adjustments might reflect the evolving nature of Wikipedia’s content and editorial practices.

In their ongoing efforts to achieve a non-biased perspective, Wikipedia continues to have challenges in mitigating the short-term biases that some of its most prominent editors may exhibit. Over time, however, these biases appear to diminish. This could be attributed to Wikipedia’s slow but ever-present self-correcting mechanisms. However, this might also be explained by the fading affect bias, where the emotional intensity of memories decreases over time, particularly those associated with negative emotions. This psychological phenomenon suggests that as time progresses, past events are recounted with

less emotional bias, allowing Wikipedia articles to evolve towards a more neutral point of view (Walker and Skowronski, 2009).

VI. Conclusion

I attempt to address the question of whether Wikipedia exhibits a political bias, particularly in favour of left-wing politics and against right-wing politics. Utilizing an event study approach, I look to estimate the effect of switching political parties on the sentiment of politicians' Wikipedia pages. Deriving the causal effect of being right-wing on the sentiment of a Wikipedia page. The analysis showed that politicians switching to right-wing parties experienced a statistically significant decrease in sentiment score of up to 6.7% four periods after the party switch. While shifts to left-wing parties show small and not statistically significant sentiment changes.

A. Lines for Further Research

Future research could expand this study by incorporating a more diverse dataset with added controls and more politicians and observations. It could utilize more post-treatment observations to explore whether the biases diminish over time as observed by Greenstein and Zhu, 2012.

One can also examine the economic implications of political biases on Wikipedia, such as impacts on political careers and post-political outcomes, which could provide insights into some of the real-world effects of this potential bias.

Future studies could also include a more diverse set of countries for improved generalizability. Additionally, using the politician's page in different languages could complement studies that have suggested Wikipedia has a Western bias Callahan and Herring; Temperton, 2011; 2015. Analyzing different language versions of the same pages could serve as control groups or provide a comparative perspective on how political biases manifest differently across linguistic contexts. This approach could also help to identify whether certain biases are specific to the English version or prevalent across multiple languages.

Using the methods in this research, future studies could explore the impact of various events on the outlook and perception of politicians. For example, researchers could examine events that are generally perceived to positively enhance a politician's public image, such as being appointed to a specific cabinet position. In the Canadian context, this could involve analyzing the sentiment change associated with becoming a minister of a particular portfolio, which typically signifies a more significant role than merely representing a jurisdiction. Such analyses could provide valuable insights into how different milestones and achievements influence the public's perception, as reflected in

Wikipedia sentiment. This is just one example of the potential avenues for future research that could further elucidate the factors affecting the sentiment of Wikipedia pages about politicians.

B. Limitations and Caveats

One limitation of this study is the relatively small sample size of politicians who switch parties, particularly from major English-speaking countries, with only 43 individuals moving to the left. This small number makes it difficult to observe any statistically significant changes. It may affect the generalizability of the results.¹⁸ Another shortcoming is the inability to filter for only relevant information from the sentiment analysis. Furthermore while controls can limit certain biases, there is a chance that the Large Language Model (LLM) and the formula used to calculate sentiment may themselves be biased. This could mean that observed sentiment changes are not actual shifts in sentiment but rather reflect changes in the topics being discussed, as politicians may alter their political focus after switching parties. This introduces a potential bias in the analysis, suggesting that caution should be exercised when interpreting these results as definitive proof of bias, as the inherent sentiment may not be accurately represented.

C. Broader Implications

The findings suggest potential short-term biases on Wikipedia against right-wing politics, which could influence public perception and contribute to societal biases that arise from a small group of people who control this information. This bias is particularly concerning as it can feed into large language models trained on this content, potentially perpetuating editor biases on a larger scale. Efforts to ensure neutrality and reduce bias in user-generated content are needed to mitigate these effects.

In conclusion, while the immediate findings indicate a bias against right-wing politics on Wikipedia, the long-term implications and the evolution of content toward neutrality remain areas that need further investigation.

References

- Abadie, A. (2005). Semiparametric difference-in-differences estimators. *Review of Economic Studies*, 72(1):1–19.
- Ackerly, B. A. and Michelitch, K. (2022). Wikipedia and political science: Addressing systematic biases with student initiatives. *PS: Political Science & Politics*, 55:429–433.

¹⁸This, of course, may also not be the case, Particularly if there actually is no effect. As seen in the B5 example from the first results, we found statistical significance with a much smaller number of politicians.

- Barnard, M. (2020). No, wikipedia isn't biased, except toward reality. *The Future is Electric*.
- Bhattacharya, A. (2016). Wikipedia's not as biased as you might think. *Quartz*.
- Callahan, E. and Herring, S. C. (2011). Cultural bias in wikipedia content on famous persons. *Journal of the American Society for Information Science and Technology*, 62(10):1899–1915.
- Callaway, B. and Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230.
- Cohen, N. (2011). Define gender gap? look up wikipedia's contributor list. *The New York Times*.
- Dan Bernhardt, Stefan Krasa, M. P. (2008). Political polarization and the electoral effects of media bias. *Journal of Public Economics*, 92(5-6):1092–1104.
- de Quetteville, H. (2023). How wikipedia became too powerful. *The Telegraph*.
- DiStaso, M. W. and Messner, M. (2010). Forced transparency: Corporate image on wikipedia and what it means for public relations. *Public Relations Journal*, 4(2).
- Ford, H. and Wajcman, J. (2017). 'anyone can edit', not everyone does: Wikipedia's infrastructure and the gender gap. *Social Studies of Science*, 47(4):511–527.
- Freyaldenhoven, S., Hansen, C., and Shapiro, J. M. (2019). Pre-event trends in the panel event-study design. *American Economic Review*, 109(9):3307–38.
- Greenstein, S. and Zhu, F. (2012). Is wikipedia biased? *American Economic Review: Papers & Proceedings*, 102(3):343–348.
- Greenstein, S. and Zhu, F. (2016). Can wikipedia be trusted? *Kellogg Insight*.
- Greenstein, S. and Zhu, F. (2018). Do experts or crowd-based models produce more bias? evidence from encyclopedia britannica and wikipedia. *MIS Quarterly*, 42(3):945–960.
- Hafner, K. (2007). Corporate editing of wikipedia revealed. *The New York Times*.
- Hamborg, F. (2022). *Media Bias Analysis*, chapter 2. Springer.
- Hube, C. and Fetahu, B. (2018). Detecting biased statements in wikipedia. In *Companion of the The Web Conference 2018*, pages 1779–1784.
- Johnson, B. (2007). Rightwing website challenges 'liberal bias' of wikipedia. *The Guardian*.

- Jr., J. M. R. (2005a). Can you trust the wikipedia?
- Jr., J. M. R. (2005b). A case of mutual aid: Wikipedia, politeness, and perspective taking. *Proceedings of Wikimania 2005*.
- Kalla, J. L. and Aronow, P. M. (2015). Editorial bias in crowd-sourced political information. *PLOS ONE*, 10(9):e0136327.
- Lih, A. (2009). *The Wikipedia Revolution: How a Bunch of Nobodies Created the World's Greatest Encyclopedia*. Hyperion.
- Mastrine, J. (2024). Is wikipedia biased? *AllSides*.
- Natapov, A. (2023). Former editor explains wikipedia's anti-russia bias. *RT*.
- Noor, P. (2018). The five wikipedia biases: Pro-western, male-dominated. *The Guardian*.
- Rehman, I. u. (2019). Facebook-cambridge analytica data harvesting: What you need to know. *Library Philosophy and Practice (e-journal)*, 2497.
- Reynolds, M. (2016). Wikipedia 'facts' depend on which language you read them in. *New Scientist*.
- Richards, J. (2009). Wikipedia — don't trust, and verify. *AEIdeas*.
- Rozado, D. (2024a). Is wikipedia politically biased? Available at: <https://www.manhattan-institute.org>.
- Rozado, D. (2024b). Trump bad, obama good: Wikipedia's bias revealed. *Manhattan Institute*.
- Sanger, L. (2005). The early history of nupedia and wikipedia: A memoir.
- Sanger, L. (2021a). The bias on wikipedia. <https://www.youtube.com/watch?v=10P4Cf0UCwU>. Accessed: 2024-07-16.
- Sanger, L. (2021b). Interview on wikipedia's neutrality and bias. <https://www.youtube.com/watch?v=McoEd6VqijY>. Accessed: 2024-07-16.
- Sanger, L. (2021c). Wikipedia co-founder says site is now 'propaganda' for left-leaning 'establishment'. *New York Post*.
- Santana, A. (2010). The new york times: Quality newspaper in the age of compromise. *Journal of Journalism Studies*.
- Stossel, J. (2022). Wikipedia's left-wing bias. *The Daily Signal*.
- Sumi, R., Yasseri, T., Rung, A., Kornai, A., and Kertész, J. (2012). Edit wars in wikipedia. In *Proceedings of the 2012 International Conference on Weblogs and Social Media*. AAAI.

- Temperton, J. (2015). Wikipedia's world view is skewed by rich, western voices. *WIRED*.
- Times, N. Y. (2022). Emails on hunter biden's laptop are authentic. Accessed: 2024-07-16.
- Walker, W. R. and Skowronski, J. J. (2009). The fading affect bias: But what the hell is it for? *Applied Cognitive Psychology*, 23(9):1122–1136.
- Weiss, J. (2022). Observations on wikipedia's political bias. <https://www.youtube.com/watch?v=kiRgJYMw6YA>. Accessed: 2024-07-16.
- Wikipedia (2024a). Truth and evolution. https://en.wikipedia.org/wiki/History_of_scientific_thought.
- Wikipedia (2024b). Wikipedia:about. <https://en.wikipedia.org/wiki/Wikipedia:About>.
- Wikipedia (2024c). Wikipedia:covid-19 pandemic. https://en.wikipedia.org/wiki/COVID-19_pandemic.
- Wikipedia (2024d). Wikipedia:dispute resolution. https://en.wikipedia.org/wiki/Wikipedia:Dispute_resolution.
- Wikipedia (2024e). Wikipedia:due and undue weight. https://en.wikipedia.org/wiki/Wikipedia:Due_and_undue_weight.
- Wikipedia (2024f). Wikipedia:fringe theories. https://en.wikipedia.org/wiki/Wikipedia:Fringe_theories.
- Wikipedia (2024g). Wikipedia:model of knowledge. <https://en.wikipedia.org/wiki/Wikipedia:About>.
- Wikipedia (2024h). Wikipedia:neutral point of view. https://en.wikipedia.org/wiki/Wikipedia:Neutral_point_of_view.
- Wikipedia (2024i). Wikipedia:neutral point of view/noticeboard. https://en.wikipedia.org/wiki/Wikipedia:Neutral_point_of_view/Noticeboard.
- Wikipedia (2024j). Wikipedia:systemic bias. https://en.wikipedia.org/wiki/Wikipedia:Systemic_bias.
- Wikipedia (2024k). Wikipedia:verifiability. <https://en.wikipedia.org/wiki/Wikipedia:Verifiability>.
- Wikipedia contributors (n.d.a). Wikipedia: Countering systemic bias.
- Wikipedia contributors (n.d.b). Wikipedia: Neutral point of view.

- Worku, Z., Bipat, T., McDonald, D. W., and Zachry, M. (2020). Exploring systematic bias through article deletions on wikipedia from a behavioral perspective. In *OpenSym 2020*, Madrid, Spain. Association for Computing Machinery.
- Yuan, L. (2023). Distilbert base multilingual cased sentiments. <https://huggingface.co/lxyuan/distilbert-base-multilingual-cased-sentiments-student>. Accessed: 2024-07-27.
- Zeller, S. (2007). Conservapedia: See under "right". *CQ Politics*.
- Zhu, F. (2015). Is wikipedia more biased than encyclopædia britannica? *Harvard Business School Working Paper*.

VII. Appendix

A. More on data

A.1. Explaining Party Shifts and State Clarification

The left-leaning dataset includes politicians who shifted from a slight left-leaning party to a hard left-leaning party or from a right-wing party to a left-wing party. Likewise, politicians who moved more to the right are included in the right-leaning dataset. For example, states such as Vermont and California were included for the USA, Ontario and Quebec for Canada, and regions such as England, Northern Ireland, Scotland, and Wales for the UK.

A.2. Data Standardization and Challenges Encountered

A subsequent function standardized all dates to a single format, removed duplicate entries, and checked missing observations for data integrity.

The challenges that were encountered and their solution:

1. **Rate Limits with the API:** A proxy service (Bright Data) was used to redirect traffic through different IP addresses.
2. **Language Model Token Capacity:** A sliding window was implemented to process text in chunks, and results were aggregated to produce the final output.
3. **Large Scraped Data File:** Data streaming with Dask was implemented to handle large files without loading them entirely into memory.
4. **Computational Expense of Analysis Scaling:** A GPU was rented from Hugging Face to speed up processing time due to the computational expense of language processing.

Given the size of the data (hundreds of gigabytes), to manage costs and computational resources, only a random sample of about half the total politicians was included in the final dataset. The total cost of this project was about \$300 CAD.

A.3. Dataset Details

The complete dataset, now containing both control and treatment groups, was matched by name to include various details: the date of each observation, the politician’s name, the URL used, positive, negative, and neutral sentiment scores, the overall sentiment score (calculated as positive sentiment minus negative sentiment), the country, the state, the date they switched parties, their first party, and their new party.

A column was created to align parties as either right or left-wing based on their political platforms, with values ranging from 1 (very far left) to 4 (very far right). In a Canadian context, moving from the Conservative Party to the People’s Party was marked as moving more to the right, while moving from the Liberal Party to the NDP was marked as moving more to the left.

This was based on a case study that observed politicians who moved to independent status. The move to independence was almost always due to disagreements with their party and a tendency towards more right-wing views. In this case study, all politicians who became independents due to scandals in their personal lives (which will very negatively affect their sentiments) were removed to not bias in the results.

A column called Treated Dummy was created, where a value of 0 was used for observations before a party switch and 1 for observations after. Another column, Ever Treated, indicated whether a politician had ever switched parties. The political alignments were matched to each individual politician, and further data cleaning was performed.

A.4. Sentiment Analysis Methodology

The sentiment analysis for this paper was conducted using the *distilbert-base-multilingual-cased-sentiments-student* model, which is a distilled version of the *mDeBERTa-v3-base-mnli-xnli* teacher model. This model was fine-tuned on a multilingual sentiment dataset, allowing it to classify text into positive, neutral, or negative sentiment categories across multiple languages. The model operates using a pipeline from the Hugging Face library, which processes the text data and returns sentiment scores for each classification label. The model works by first encoding the input text into a series of numerical representations, capturing the semantic meaning of the text. These representations are then passed through the model’s layers, which have been trained to recognize patterns

associated with positive, neutral, and negative sentiments. Finally, the model outputs a probability distribution across these sentiment categories.

B. More on the model

- Sentiment outputs are categorized into either positive, neutral, or negative, with the total adding up to one. The overall score is calculated by subtracting the negative score from the positive
- This intercept captures the general sentiment level before any specific adjustments or events.

C. Additional Tables and Figures

Table 4: Impact of Party Switching on Sentiment (Right, Including Word Count and Party Alignment)

Coefficient	Estimate	Std. Error	t value	p-value
B0	-0.02059***	0.00547	-3.76	0.00018
B1	-0.05003***	0.00634	-7.90	<2e-15
B2	-0.05898***	0.00697	-8.46	<2e-16
B3	-0.05767***	0.00720	-8.01	<2e-16
B4	-0.06699***	0.00784	-8.54	<2e-16
B5	-0.22059***	0.02156	-10.23	<2e-16
Word Count	-0.00000356**	0.00000114	-3.11	0.00194
Party Alignment	-61.48216***	0.22977	-267.58	<2e-16
RMSE:	0.04548			
Adj. R2:	0.8507			
Within R2:	0.03065			

Table 5: Impact of Party Switching on Sentiment (Left, Including Word Count and Party Alignment)

Coefficient	Estimate	Std. Error	t value	p-value
B0	-0.01341	0.01165	-1.15	0.2502
B1	-0.00274	0.01362	-0.20	0.8408
B2	-0.00827	0.01429	-0.58	0.5628
B3	-0.01351	0.01394	-0.97	0.3327
B4	-0.00755	0.01600	-0.47	0.6372
Word Count	-0.00000356	0.00000114	-3.11	0.00194
Party Alignment	-43.52602	0.34420	-126.46	<2e-16
RMSE:	0.04540			
Adj. R2:	0.8507			
Within R2:	0.02823			

Table 6: Impact of Party Switching on Sentiment (Right, No Controls)

Coefficient	Estimate	Std. Error	t value	p-value
B0	-0.01888***	0.00467	-4.04	0.00006
B1	-0.05175***	0.00565	-9.17	<2e-16
B2	-0.05766***	0.00625	-9.22	<2e-16
B3	-0.05719***	0.00628	-9.11	<2e-16
B4	-0.06493***	0.00690	-9.42	<2e-16
B5	-0.22886***	0.02699	-8.48	<2e-16
RMSE:	0.04509			
Adj. R2:	0.8484			
Within R2:	0.00298			




Figure 3: DiD for Politicians Moving Right with No Controls

Table 7: Basic Model (Left, Without Controls)

Coefficient	Estimate	Std. Error	t value	p-value
B0	-0.01147	0.00985	-1.16	0.2445
B1	-0.00687	0.01252	-0.55	0.5835
B2	-0.01236	0.01289	-0.96	0.3376
B3	-0.01579	0.01295	-1.22	0.2229
B4	-0.00925	0.01455	-0.64	0.5250
RMSE:	0.04502			
Adj. R2:	0.8485			
Within R2:	0.00002152			




Figure 4: DiD Impact of Party Switching on Sentiment (Left, No Controls)

Table 8: Politicians That Moved Towards the Center

	Estimate	Std. Error	t value	Pr(> t)
B0	0.004412	0.009788	0.450726	0.65660
B1	-0.001925	0.061444	-0.031330	0.97529
B2	0.031454	0.062007	0.507268	0.61701
B3	0.036906	0.134084	0.275243	0.78570
B4	0.053081	0.174572	0.304065	0.76394
RMSE:	0.010549			
Adj. R2:	0.805146			
Within R2:	0.078707			




Figure 5: Politicians that moved towards the center

Table 9: Politicians that moved to independent

	Estimate	Std. Error	t value	Pr(> t)
B0	-0.02585349	0.01021033	-2.532092	0.0124342 *
B1	-0.03847316	0.01172659	-3.280848	0.0013043 **
B2	-0.04232851	0.01394040	-3.036391	0.0028528 **
B3	-0.03759556	0.01963940	-1.914292	0.0576086 .
B4	-0.04766407	0.01855274	-2.569112	0.0112350 *
RMSE:	0.015434			
Adj. R2:	0.940589			
Within R2:	0.089909			




Figure 6: Politicians that moved to independent

Table 10: Politicians that moved to the Right not including Independent

	Estimate	Std. Error	t value	Pr(> t)
B0	-0.02024050	0.01047959	-1.93142	5.3738e-02 .
B1	-0.06001132	0.01304965	-4.59869	4.8460e-06 ***
B2	-0.06941109	0.01484098	-4.67699	3.3472e-06 ***
B3	-0.06390545	0.01508976	-4.23502	2.5150e-05 ***
B4	-0.07749185	0.01706013	-4.54228	6.3051e-06 ***
B5	-0.22549316	0.02319787	-9.72043	< 2.2e-16 ***
RMSE:	0.04546			
Adj. R2:	0.850525			
Within R2:	0.029495			

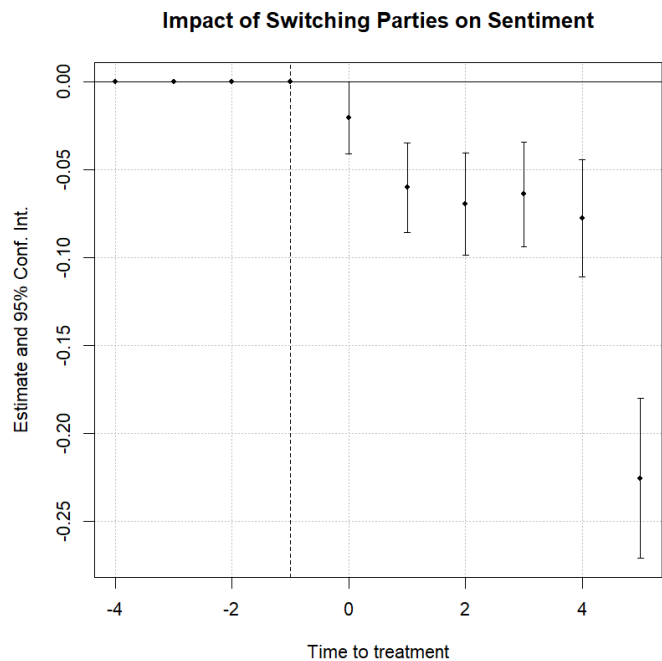


Figure 7: Politicians that moved to the Right not including Independent

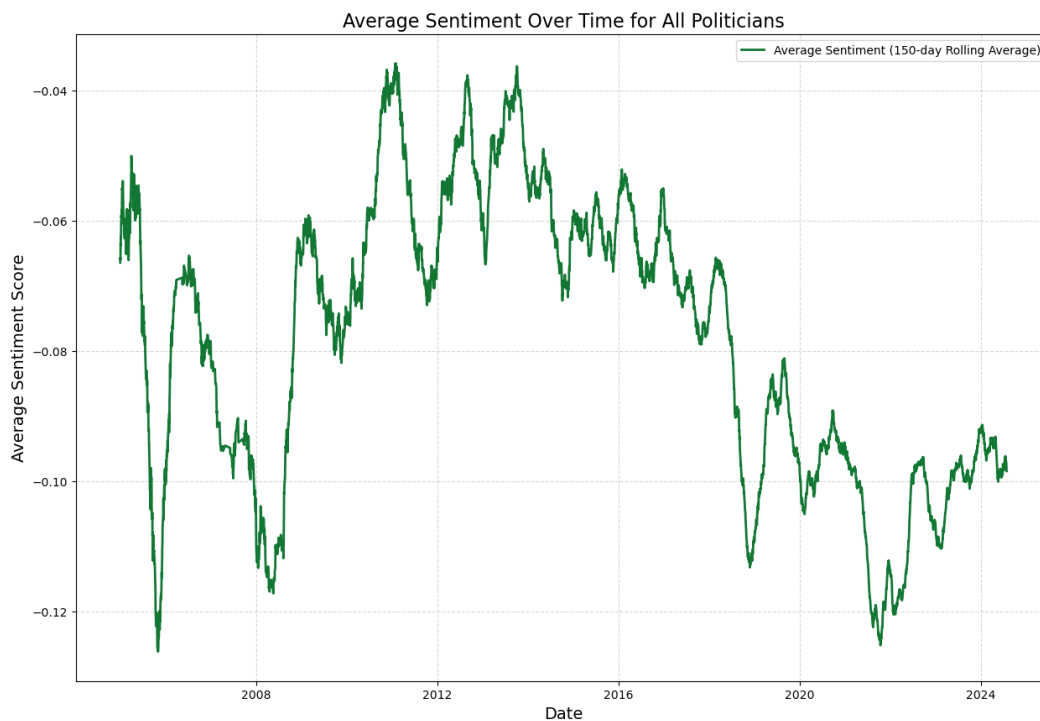


Figure 8: Average Sentiment of All Politicians Over Time



Figure 9: Average Sentiment of All Politicians Over Time Split by Treatment Status

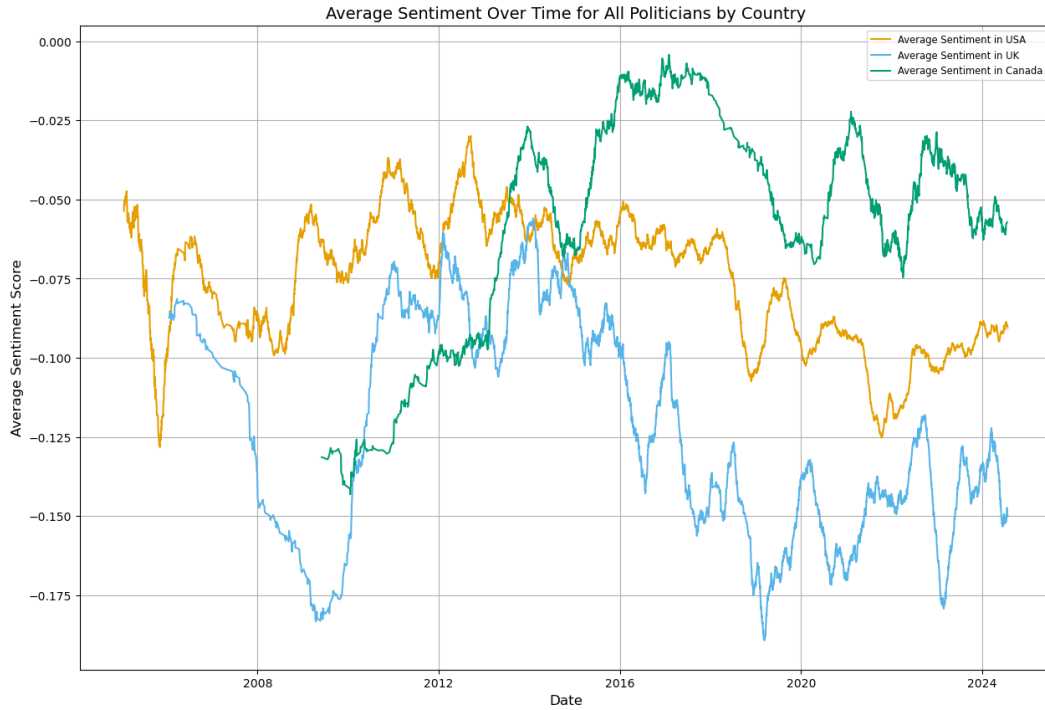


Figure 10: Average Sentiment of All Politicians Over Time Split by Country




Figure 11: Average Sentiment of All Politicians Over Time Split by Right-Wing

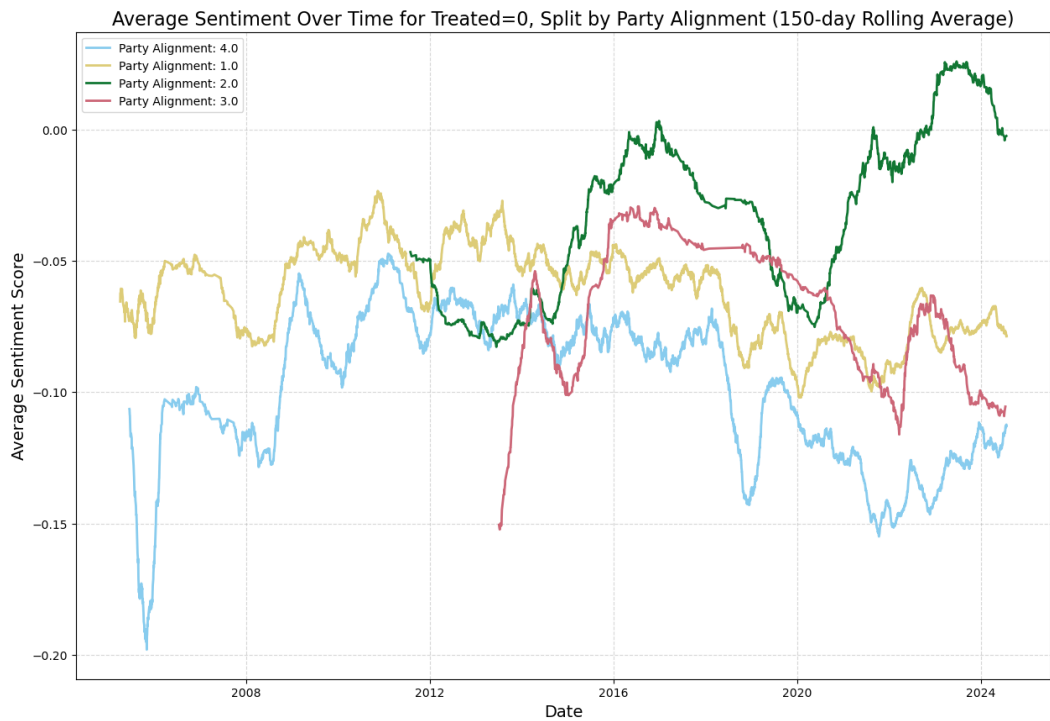


Figure 12: Average Sentiment of All Politicians Over Time Split by Alignment

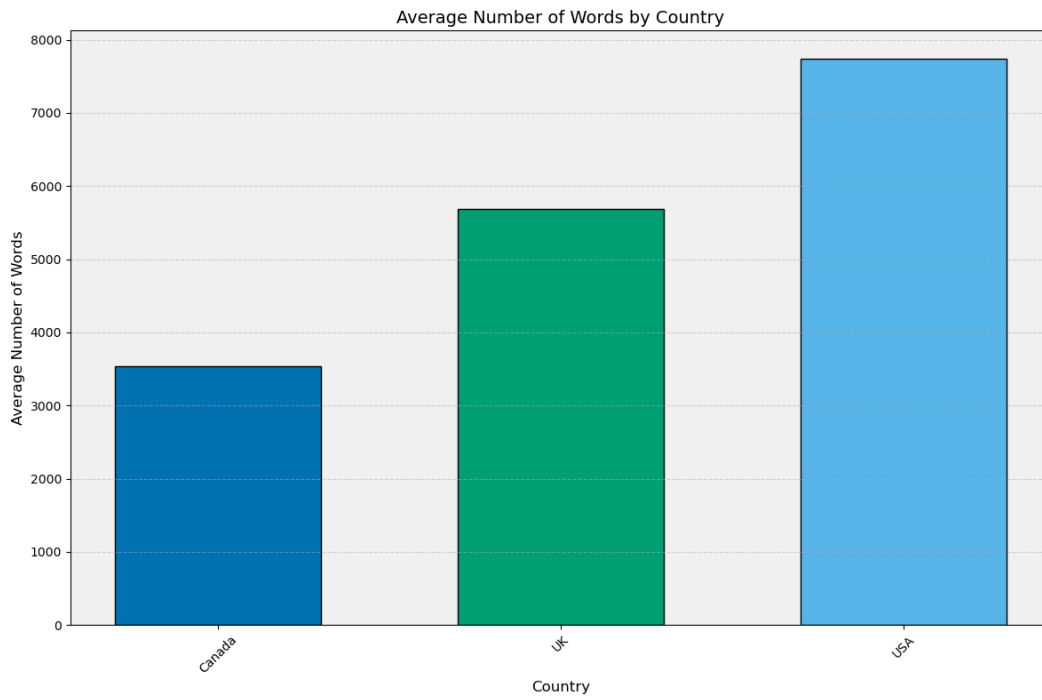


Figure 13: Average Number of Words for Politicians Wiki Per Country

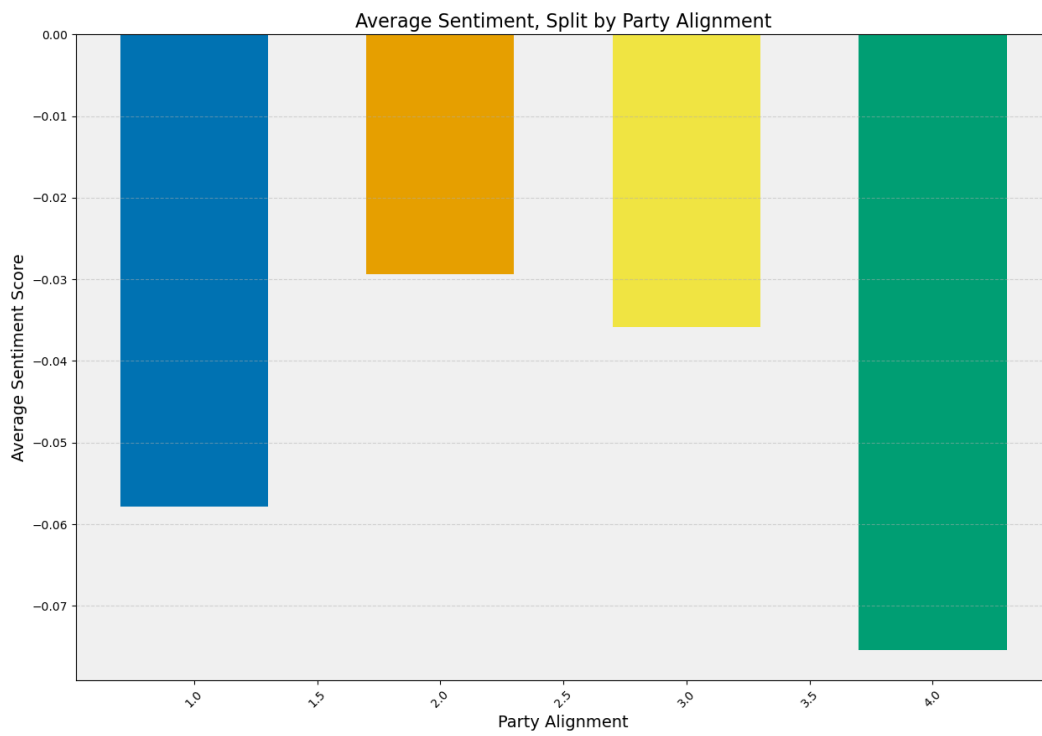


Figure 14: Average Sentiment by Party Alignment

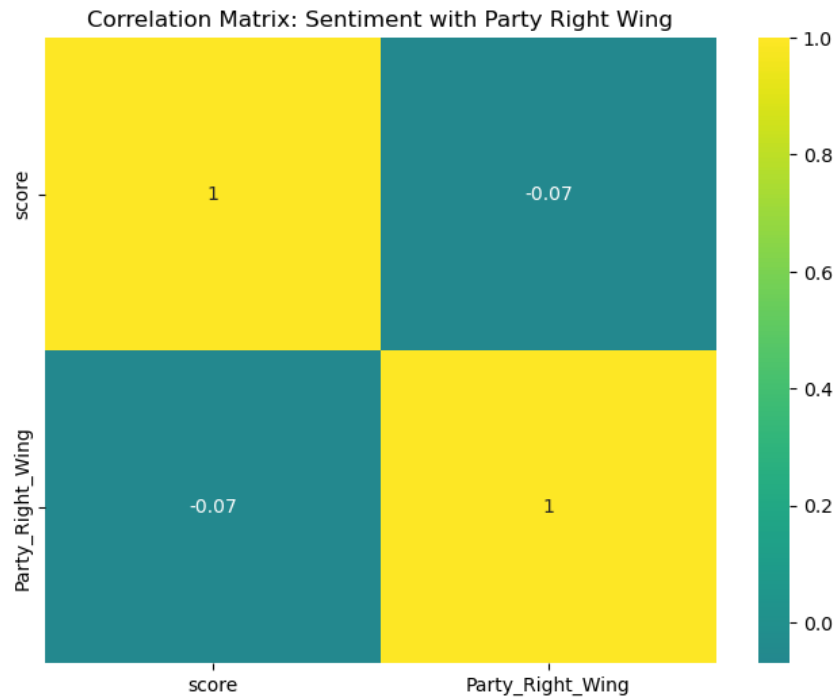


Figure 15: Correlation matrix for Party Right-Wing and Sentiment Score

Table 11: OLS Regression Results

	Coef.	Std. Err.	tstat	$P > t $	0.025	0.975
const	-0.0503***	0.001	-98.883	0.000	-0.051	-0.049
Party_Alignment	-0.0058***	0.000	-32.346	0.000	-0.006	-0.005
R-squared:	0.005					
Adj. R-squared:	0.005					
F-statistic:	1046.					

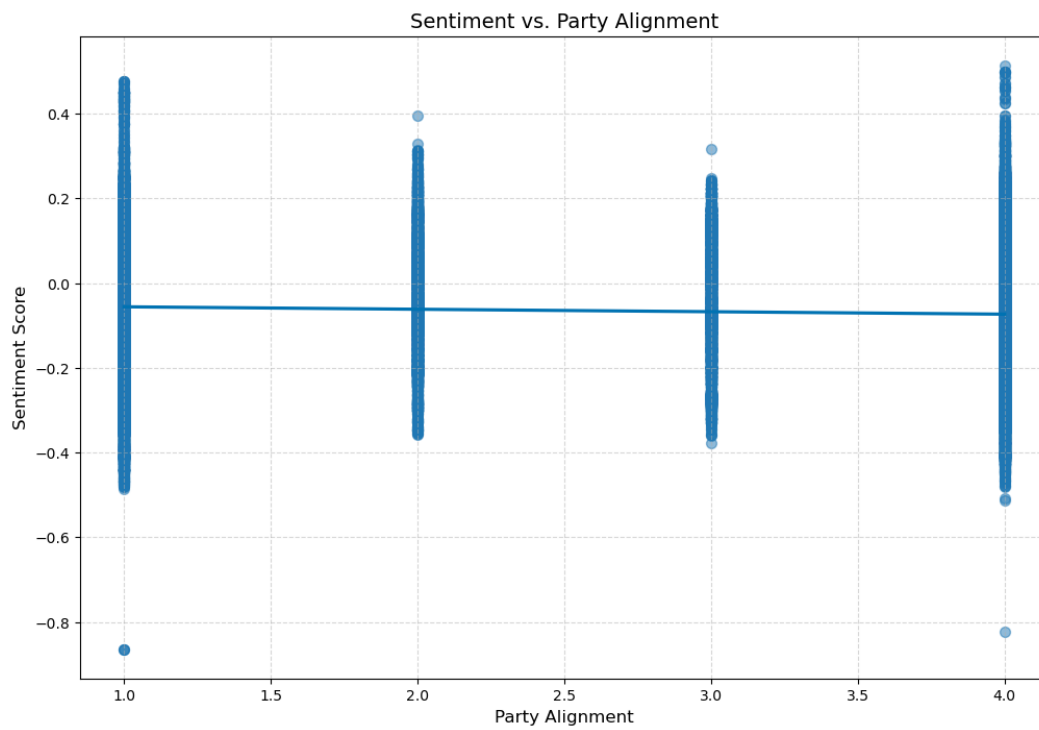


Figure 16: Sentiment vs Party Alignment OLS Regression