How Partisanship and Perceived Political Bias Affect Wikipedia Entries of News Sources

Khonzodakon Umarova
Department of Computer Science
Wellesley College
kumarova@wellesley.edu

Eni Mustafaraj
Department of Computer Science
Wellesley College
emustafaraj@wellesley.edu

ABSTRACT

Increased polarization and partisanship have become a consistent state of politics, media, and society, especially in the United States. As many news publishers are perceived as "biased" and some others have come under attack as being "fake news", efforts to make such labels stick have increased too. In some cases (e.g., InfoWars), the use of such labels is legitimate, because some online publishers deliberately spread conspiracy theories and false stories. Other news publishers are perceived as partisan and biased, in ways that damages their reporting credibility. Whether political bias affects journalism standards appears to be a debated topic with no clear consensus. Meanwhile, labels such as "far-left" or "alt-right" are highly contested and may become cause for prolonged edit wars on the Wikipedia pages of some news sources. In this paper, we try to shine a light into this phenomenon and its extent, in order to start a conversation within the Wikipedia community about transparent processes for assigning political orientation and journalistic reliability labels to news sources, especially to unfamiliar ones, which users would be more likely to verify by looking them up. As more of Wikipedia's content is used outside Wikipedia's "container" (e.g., in search results or by voice personal assistants), the issue of where certain statements appear in the Wikipedia page and their verifiability becomes an urgent one to consider not only by Wikipedia editors, but by third-party information providers too.

CCS CONCEPTS

• **Information systems → Wikis**; *Web search engines*; *Crowdsourcing*; • **Human-centered computing → Empirical studies in collaborative and social computing**; *Interactive systems and tools*;

KEYWORDS

wikipedia, google search, knowledge panels, political bias, editing

ACM Reference Format:
Khonzodakon Umarova and Eni Mustafaraj. 2019. How Partisanship and Perceived Political Bias Affect Wikipedia Entries of News Sources. In *Companion Proceedings of the 2019 World Wide Web Conference (WWW'19 Companion), May 13–17, 2019, San Francisco, CA, USA.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3308560.3316760

WWW '19 Companion, May 13–17, 2019, San Francisco, CA, USA

ACM ISBN 978-1-4503-6675-5/19/05.
https://doi.org/10.1145/3308560.3316760

1 INTRODUCTION

Recently, a new kind of Wikipedia vandalism appeared in the news: the use of an extreme political label, in this case, "nazism", to describe the ideology of the California Republican Party. Interestingly, the public learned about it because Google showed this label in the Knowledge Panel of the California Republican Party's search engine result page (SERP). In fact, Google also received most of the blame for it,[1] despite the fact that the information was extracted from Wikipedia's page of this organization. Knowledge Panels (KP) in SERPs frequently contain information that comes directly from the Wikipedia article of the searched entity (a person, an organization, a company, etc.). Thus, edits in either the infobox[2] or the lead[3] of a Wikipedia entry may quickly find their way to Google's SERPs and therefore to a larger audience. In cases of blatant vandalism, such Wikipedia edits are reverted quickly (although Google search might need several hours to refresh its index that contains the changes). However, other labels describing the perceived political bias of a news publisher (rather than a self-declared political affiliation) or the lack of reliability of news sources are a matter of opinion and don't receive immediate resolution by Wikipedia editors, especially if they belong to pages that don't have a huge Wikipedia traffic. Such lack of clarity about how such labels are assigned, might have consequences in how these sources are viewed by the public.

One of our long-term research goals is to understand and support users' decision making about the credibility of online news sources. Through user studies [7, 10], we have identified that KPs in SERPs play an important role in this process, both through the amount of information they contain and its quality. Explicit references to perceived political bias of news sources, which are sometimes present in the Wikipedia lead of their articles are particularly helpful to users [10]. However, by comparing SERPs of news publishers over time,[4] we have noticed a few interesting patterns: for some news sources, the description swings between overly critic to fact-omitting (see case study in the next section); in others, the description relies on self-representation (citing the primary source) by burying down in the article secondary sources that contest the primary source's self-depiction.

Prior research [8] has established that when users read excerpts from Wikipedia content on Google search results, they often don't click-through to the entire article. Thus, these users might be at risk of creating a distorted or an incomplete impression of a news source based on a few sentences. Given the prominence of KPs

[1] https://www.politico.com/story/2018/05/31/mccarthy-slams-google-in-nazism-flap-616109
[2] A box of information that is displayed in most Wikipedia pages, on the right-side.
[3] The first paragraph in a Wikipedia page.
[4] https://medium.com/@enimust/the-information-panels-on-google-and-facebook-uncovering-their-blind-spots-2e8210b2e697

in SERPs (especially on mobile devices, where they are frequently placed ahead of organic search results, e.g., see Figure 2), the quality of excerpts from the Wikipedia leads, which are displayed in KPs, may potentially "make or break a reputation". Simultaneously, because of this importance, various actors might have an incentive to modify such content with the intention of positively or negatively impacting such reputation. Although issues related to the NPOV (neutral point of view) and notability are nothing new for Wikipedians, the increased political polarization in the society might pose new challenges that need to be addressed consistently.

In this preliminary work, we establish the grounds for studying such challenges systematically. Instead of considering broadly the issue of political bias and its detection in Wikipedia, as performed in [5], we make a deep-dive into a subset of Wikipedia articles of news sources for which we have scores of perceived political bias released by [9]. Initial results indicate that the "political labeling battle" through constant revisions is more fierce for news sources that are not very popular and might not have the same Wikipedia page protections in place as more popular news sources.

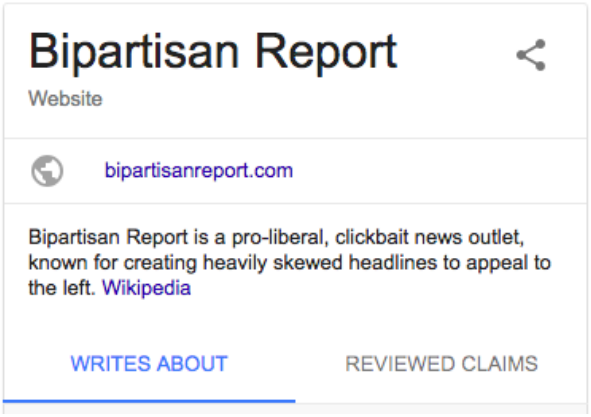

Figure 1: The Knowledge Panel for *Bipartisan Report* in January 2018. It contains a few sentences from the Wikipedia lead of this publication entry. Notice the mentions to perceived political bias and lack of reliability.

2 A CASE STUDY: BIPARTISAN REPORT

Research in news literacy has emphasized the use of "lateral reading" as a strategy for evaluating the reliability of news stories [13]. The first step of this strategy is to Google for the source of a news story. The KP containing information from Wikipedia is often prominent in Google's SERP for a query. One such example is shown in Figure 1, where the screenshot shows the KP for *Bipartisan Report* in January 2018. The brief description addresses this source's perceived political leaning (pro-liberal, left)[5] as well as lack of high journalistic standards (click-bait, heavily-skewed). When the page was created in December 2016, the description used stronger political labels:

> Bipartisan Report is a **far-left**, clickbait news outlet, known for creating heavily skewed headlines to appeal to the **far left**.

However, the author self-corrected, by replacing one instance of "far-left" with "pro-liberal" and removing "far" in the second "far left" expression, settling for the version shown in Figure 1. The complete text of the Wikipedia lead contained more details about this source's perceived lack of reliability, using a biased point of view,[6] quite similar to the language in Figure 1.

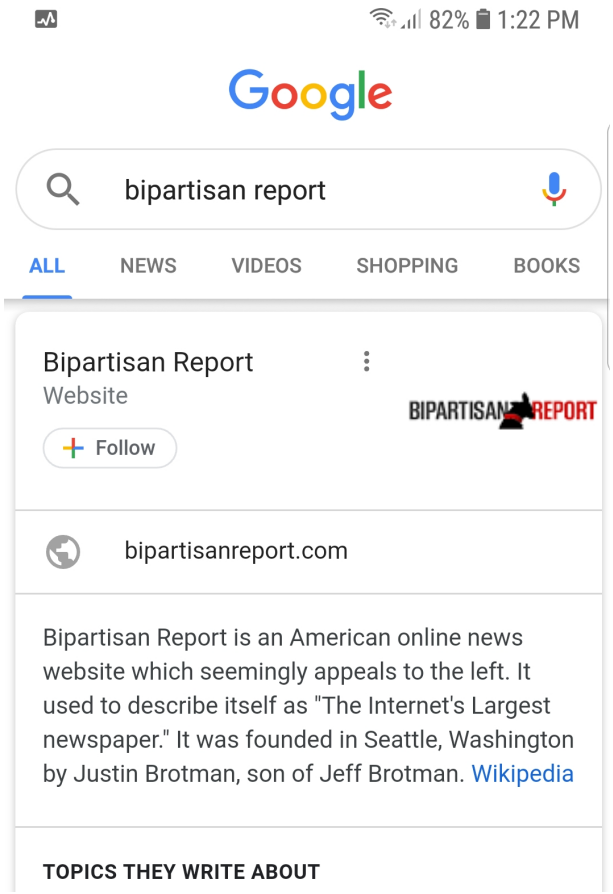

Figure 2: The Knowledge Panel for *Bipartisan Report* as of this writing (March 2019). Compared to the text in Figure 1) the Wikipedia lead was edited and now displays neutral language. Also, this screenshot was taken using a mobile device to indicate that the panel with Wikipedia content is very prominent in the search results shown in such devices.

It took 1.5 years, until May 2018, for the article's lead to be changed into a more neutral-sounding description, which is shown in Figure 2. As a result of a debate in the talk page, many contested parts of the article were removed too. The current version[7] is mostly neutral, but it also omits relevant details, such as several failed fact-checks of this source. Concretely, the fact-checker Snopes has

[5] No citations/references to support these political labels are provided in the Wikipedia article.

[6] https://en.wikipedia.org/w/index.php?title=Bipartisan_Report&oldid=753604666

[7] https://en.wikipedia.org/w/index.php?title=Bipartisan_Report&oldid=875424620

fact-checked four of *Bipartisan Report* stories, finding two of them False, one Mostly False, and one Unproven. The current Wikipedia article doesn't contain any of these facts about possible lack of reliability of *Bipartisan Report*.

If readers were to visit *Bipartisan Report's* Wikipedia page, they would see a big warning box listing multiple issues with the neutrality and quality of the article.[8] If furthermore, users were to visit the Talk page of the article, they would find a 4,600-word discussion, twenty times more text than to the 200-word article (including the warnings). However, few visitors come to this page. In 2018, it had a total of 8,769 page views and 11 edits.

There is an interesting conundrum here. Google continues to display the neutral-sounding description of *Bipartisan Report* in its KP, which doesn't contain any mention of the journalistic shortcomings of this news source, because Wikipedia editors couldn't agree on what to write. Meanwhile both Google and Wikipedia are providing a kind of legitimacy to this news source by "hiding" truly relevant information from users (its lack of reliability as a news source), while showing trivia information (the father of the founder happens to be a successful businessman).

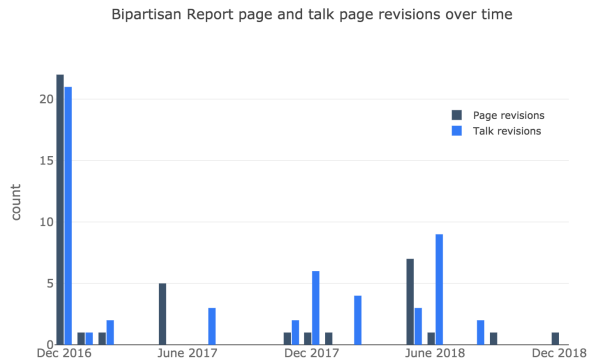

Figure 3: Comparison of edit distributions for the *Bipartisan Report* Wikipedia page and Talk page. Only discussions in May/June 2018 led to changes in the content of the page, by removing most of the contested information.

This case scenario raises many interesting questions worth discussing with respect to the design of information systems and their interfaces, as well as the ethical implications of such design choices:

(1) Why it took more than 1.5 years to take action about the lack of a neutral point of view? A warning label about how "The neutrality of this article is disputed" was added since the page creation in December 2016, and two more warnings about neutrality and bias were added in May 2017, but it took until May 2018 for any changes to happen.
(2) What did the editors discuss in the Talk page and why couldn't they agree on how to make the Wikipedia entry for *Bipartisan Report* more informative, especially by addressing issues of journalistic reliability? As Figure 3 shows, the talk page was active from the start, but that didn't initially affect the content of the page.
(3) Who should be responsible for misleading the public about this news source by omitting relevant information: Wikipedia editors who participated in the deliberation and couldn't agree, or Google that shows a KP based on Wikipedia content without any hints that this content is disputed within Wikipedia.

However, in order to motivate the need to answer such questions, we want to establish first the extent to which cases similar to *Bipartisan Report* occur in Wikipedia. After a discussion of partisanship and bias in the following section, we then explain our approach for discovering other cases potentially similar to *Bipartisan Report*.

3 PARTISANSHIP AND BIAS IN MEDIA

What makes *Bipartisan Report* an interesting study example is that it is perceived by certain readers as a partisan outlet, politically biased to entirely favor one side. The issue of partisanship in media is one with a long history. In the United States, partisan media was the only kind of media in existence for a long time [11]. However, over many decades, newspapers slowly freed themselves from political parties' influence and created a "norm of objectivity" [12] to focus on reporting the facts. Interestingly, this process happened mostly in the United States, and as a result, established newspapers do not identify as having a political orientation (such as left of right). Meanwhile, such labels are very common in Europe, where newspapers continue to be openly aligned with certain political ideologies. In fact, Wikipedia itself is a good place to experience this distinction. While the Wikipedia pages for the list of newspapers in Germany,[9] or the UK[10] contain an explicit column for the political alignment of newspapers, the corresponding page for the United States[11] contains no such column.

However, with the advent of Cable TV, conservative radio, and then later the web, the media landscape has changed in the United States too. Conservative elites (for example, radio and TV hosts such as Rush Limbaugh and Bill O'Reilly) have for years argued that the media in general has a "liberal bias" [3], in order to justify the need for conservative outlets. Researchers have argued that there is even a "Limbaugh effect" that could explain increased polarization in the American public [6]. In such a polarized environment, even though the majority of news sources in the United States remain objective in their reporting, the public perceives them increasingly as biased, as polls have been showing over the years.[12]. Moreover, the establishment of openly partisan web-based outlets has further contributed in the creation of a media ecosystem in which users feel skeptic and want to ascertain the bias of every news publisher.

Researchers have been investigating various methods to characterize the possible bias (or slant) of various news publishers. One method relies on the similarity of language between politicians and newspapers [4], another on a combination of crowdsourcing and machine learning decision-making [2]. In contrast to these content-based approaches, we also find in the literature approaches that

[8] However, this is not true if one visits the Wikipedia page on a mobile device. On it, the warning box is not visible, it's a link to click, further hiding the information about the quality of the page.

[9] https://en.wikipedia.org/wiki/List_of_newspapers_in_Germany
[10] https://en.wikipedia.org/wiki/List_of_newspapers_in_the_United_Kingdom
[11] https://en.wikipedia.org/wiki/List_of_newspapers_in_the_United_States
[12] https://news.gallup.com/poll/225755/americans-news-bias-name-neutral-source.aspx

rely on the ideological alignment of audiences [1]. We rely in one such approach to find news publishers with a score of perceived bias.

4 DATA COLLECTION: SOURCES WITH PERCEIVED BIAS

A common limitation of the content-based approaches mentioned above is the limited size of the news sources they rate. For example, [2] focuses only on 15 news sources. In comparison, the audience-based approach that works by "polling" the users who share news stories in social networks such as Facebook and Twitter have the potential to scale and provide information for a large number of news sources. We explain such an existing approach in the following and then make use of its dataset to discover news sources with Wikipedia pages for our analysis.

4.1 Partisan Audience Bias

In [9], the authors created the Partisan Audience Bias (PAB) dataset[13] from links shared by real users on Twitter. Using voter registration records of US citizens with republican and democratic affiliations, the study identified 519,000 Twitter accounts matching these citizens. Over a certain period of time, a set of 113 million tweets by these accounts were collected and only tweets with URLs were kept. The URLs were processed to extract the second-level domain names. For example, `http://www.bbc.com/news/business-38686568` was converted into `bbc.com`. Then, to reduce noise, since 63% of links were shared only once, the authors kept only the domains that were shared more than 50 times (by different users). This led to a dataset of 19,022 sites. For each site, the authors calculated a bias score between -1 (a site shared only by democratic voters) to +1 (a site shared only by republican voters). Sites that get a bias score between -1 and +1 were shared by a mix of democratic and republican voters. For example, *The Wall Street Journal* had a score of 0.0106, signaling that it is a news source shared almost equally by both sides. Meanwhile, the bias score for *Bipartisan Report* was -0.7469, and the score for *InfoWars* was 0.7817, putting these two sources at opposite ends. However, the study's authors caution the readers that these scores are not absolute but only relative, since there is no fixed zero point in assigning bias.

4.2 SERP pages for the PAB dataset

Using one computer with a Chrome browser in incognito mode, and creating one new instance of the browser for each query, we automatically collected the Google SERPs for all 19,022 site names in the PAB dataset. By parsing the SERPs, we found 3,763 SERPs (or 19.8%) that contained a KP. 89.3% of these KP contained a Wikipedia link, indicating their provenance from Wikipedia. Given that the use of URLs as search queries doesn't guarantee good search results, we decided to rectify this issue with a second collection. We automatically performed approximate string matching of the URLs from the PAB dataset with the list of URLs and their titles from Google's SERP in order to find textual phrases to search. For example, the URL `aflcio.org` from the PAB dataset is matched with three pairs: (`aflcio.org`, *AFL-CIO: America's Unions*), (`wyomingaflcio.org`, *Wyoming State AFL-CIO*), and (`www.massaflcio.org`, *Massachusetts AFL-CIO*). Using this technique, we were able to find matches for 17,497 URLs. We repeated the SERP collection and this time found 3,539 KPs (19.3%), again with the majority with Wikipedia links. At the end of this two-step SERP collection, we had a list of 4,967 matches between (URL, score) pairs from the PAB datasets and Wikipedia links that had appeared in KPs for our two data collections.

4.3 Two lists of online news sites

Inspecting the created dataset of Wikipedia links and KPs description, we noticed that many sites don't are not news sources. As expected, people share on Twitter other links too, such as event descriptions, products, etc. Since our focus is on news sources, we selected a subset of 1,391 URLs that come from the two categories below:

- *Newspaper news sites* - a list that contains all sources, whose KP has a "Newspaper" category. It is made of 695 unique sites.
- *Non-newspaper news sites* - a list that contains all sources, whose KP short description contains the word "news". There are 696 unique sites in this list.

We acknowledge that other links in our dataset might be news sources, and identifying them remains for future work.

4.4 Collecting Wikipedia Leads

We are interested in how news sources are labeled in terms of their perceived bias and reliability and such labeling often happens in the lead paragraph(s) of the Wikipedia entry. Therefore, our current focus is on examining how Wikipedia leads change over periods of time. To that purpose, we collected all revisions of pages corresponding to our news sources dataset. The dataset was collected on 01/12/2019, and hence contains revisions between February 2001 and January 2019. Deleted or inaccessible revisions are excluded. Overall, for 1,391 news-related sources from the two lists, our dataset contains 434,923 revisions, 100,269 of which occur in the lead section.

Using the MediaWiki action API[14], we collected all revisions for each page, including the revision id, time-stamp, user who made the revision, and comments. For every revision, the lead section (or everything that comes before the first section of Wikipedia article) was also extracted. After filtering out Wiki code, such as templates, style tags, and references, we obtained the plain-text of the lead section.

5 RESULTS

In a political atmosphere of polarization, individuals at opposite sides are applying political labels to media outlets, with the purpose of devaluing their reporting (by implicitly considering them as biased and unfair). For example, President Trump tweeted in August 2018 the following: "Republican/Conservative Fair media is shut out. Illegal? 96% of results on "Trump News" is from National Left-Wing Media"[15]. To examine whether this process of political

[13] Available at https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/QAN5VX

[14] https://www.mediawiki.org/wiki/API:Main_page

[15] https://www.politifact.com/truth-o-meter/statements/2018/aug/29/donald-trump/no-96-google-news-stories-trump-arent-left-wing-ou/

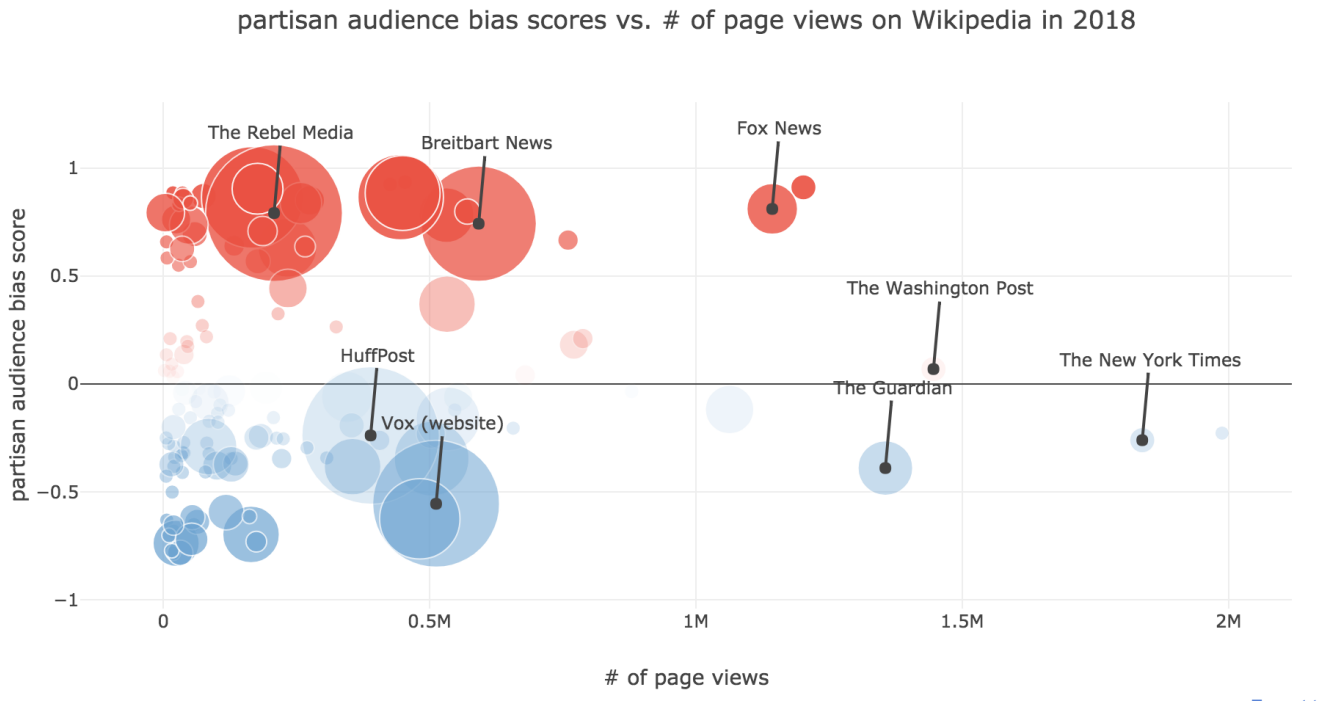

Figure 4: A visual representation of the "political labeling battles" in the Wikipedia pages of news sources that are perceived as politically biased. For visual clarity purposes, we plotted a subset of 127 sources, for which political labels were removed/re-added more than 5 times in 2018. Due to this decision, many sources with perceived bias scores between -0.5 and 0.5 are not displayed, because they didn't pass the threshold.

Table 1: Number of Wikipedia pages for news sources that contain a political label in their latest revision.

Left-leaning	left (46), democrat (6), progressive (18), liberal (35)
Right-leaning	right (49), republican (1), conservative (51)

labeling is happening in Wikipedia too, we compiled a list of target words that are associated with opposite political leanings: 'left' and 'right'; 'democratic' and 'republican', as well as 'liberal', 'progressive' and 'conservative'. By inspecting changes in the lead, we identify revisions in which these words were repeatedly added in or removed from the Wikipedia lead section, an action that indicates lack of consensus about the use of such political labels. For the latest revisions of Wikipedia articles in our dataset, we tabulate the presence of such labels in Table 1. We can notice a somewhat balanced use, 101 are right-leaning labels and 106 are left-leaning. Meanwhile, the total number is small (207) compared to the list of 1,391 sources, indicating that for the majority of news sources, no such political labels are used.

5.1 Changes Related To Political Bias

When Wikipedia editors don't agree about the use of a political label for a news source, they might engage in an edit war to either preserve or remove a label. Are such edit wars correlated with the perceived bias? Although in this paper we do not formally test such an hypothesis, we begin its exploration through the visualization in Figure 4, which illustrates the relationship between (1) the partisan audience bias score of a news source (between -1 and +1, shown in the y-axis), (2) the number of views its English Wikipedia page received in 2018 (shown in the x-axis), and the number of times the political labels from Table 1 were repeatedly removed and added back to the article lead in 2018. The third variable is expressed through the radius of markers.

There are several interesting trends we observe here. On the farther end of x-axis, we have popular websites, such as *Fox News* and *The New York Times*, whose Wikipedia pages received over 1 million views in 2018. These pages have overall a larger number of page revisions (339 and 183 respectively) more than other pages in our list (89.1 in average), however, such revisions are rarely about the political labels in question, and thus these sources appear with a small radius in the graph. Meanwhile, somewhat less popular pages, such as *HuffPost* or *The Rebel Media* that received in the range of 500,000 views, had a large number of changes in political labels, contributing to the depiction in the graph with a larger radius. Such changes might be due to disagreements on the extent to which a certain political label could be applied to these news sources. It is however important to note that both *Fox News* and *The New York Times* pages are under Wikipedia's semi-protection policy, as

opposed to *HuffPost* and *The Rebel Media*, which can be edited by unregistered and newly registered users.

Most importantly, Figure 4 suggests that news sources that are perceived as strongly biased (on both sides of the political spectrum) are targets for these political labelling actions, even though many of these news sources are not popular in terms of page views.

5.2 User participation in revisions

In addition to investigating the nature of edits in the Wikipedia lead section, we were also interested in identifying the involved users. For each revision that was associated with addition and removal of political labels, we identified its editor and when available collected information about their account, such as the registration date, total number of contributions, and permissions using Media Wiki action API. In our analysis we considered two groups of editors: (1) users who added a political label at some point in the lead of a page, and (2) users who removed a political label. Overall there are 3,946 and 3,796 unique users in each group respectively, including unregistered editors (identified by their IP address), and 1,368 users who occur in both groups.

We found that for most editors from both groups, their scope of edits related to political labelling in the lead is focused on a limited number of pages. In fact, about 90-91% of users from each group edit only one Wikipedia page from our lists of newspapers and news-related sources. There are a few users whose edits that span over as many as 30-31 different pages. These are primarily Wikipedia anti-vandalism bots.

Lastly, we analyzed the age of accounts of users who added/removed a political label at the time the revision occurred. Figure 5 shows that most revisions related to political labelling were made by accounts which are 1 to 2 years old. This can be observed in both groups of editors, though "younger" accounts are involved in adding such labels (notice the visible blue spike). We also couldn't find a statistically significant difference between the distributions.

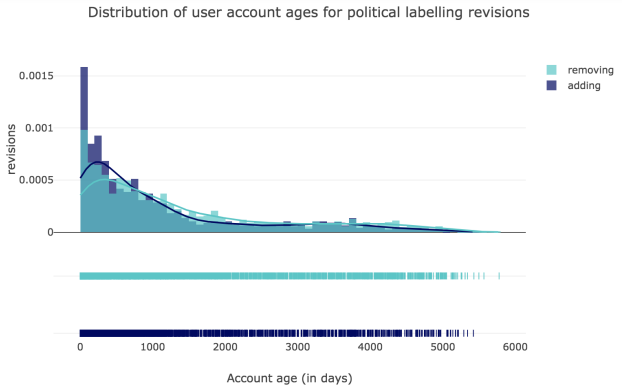

Figure 5: Comparison of distribution of ages for accounts that make edits involving addition/removal of political labels in the Wikipedia article lead.

6 CONCLUSION

Motivated by the presence of Wikipedia content in third-party platforms, such as Google's search results, we seek to identify the extent to which the characterization of news publishers is affected by edit wars in Wikipedia, in the backdrop of a polarized society. Our initial results, which are limited to mostly North American news sources, indicate that many pages that don't have editing protections in place are affected by continuous efforts to add/remove political bias labels to their lead text. That is, this problem is not unique to our case study, *Bipartisan Report*, therefore, it merits further research. Inspecting the editors involved in such edits, we found relatively sizeable groups on both sides of the edit wars, indicating that this is not an isolated issue, and that it should be considered for discussion in the Wikipedia governance deliberation processes. This research has also implications for web literacy efforts. Learning to recognize warning labels, consulting talk pages, and looking at revision histories of Wikipedia pages needs to be one of the skills that is explicitly taught as part of news literacy programs.

ACKNOWLEDGMENTS

The authors are grateful to the members of the Wellesley Cred Lab for their support. This work is partially supported by the National Science Foundation through grant IIS 1751087.

REFERENCES

[1] Eytan Bakshy, Solomon Messing, and Lada A. Adamic. 2015. Exposure to ideologically diverse news and opinion on Facebook. *Science* 348, 6239 (2015), 1130–1132. https://doi.org/10.1126/science.aaa1160 arXiv:http://science.sciencemag.org/content/348/6239/1130.full.pdf

[2] Ceren Budak, Sharad Goel, and Justin M Rao. 2016. Fair and balanced? quantifying media bias through crowdsourced content analysis. *Public Opinion Quarterly* 80, S1 (2016), 250–271.

[3] David Domke, Mark D Watts, Dhavan V Shah, and David P Fan. 1999. The politics of conservative elites and the âĂIJliberal mediaâĂİ argument. *Journal of Communication* 49, 4 (1999), 35–58.

[4] Matthew Gentzkow and Jesse M Shapiro. 2010. What drives media slant? Evidence from US daily newspapers. *Econometrica* 78, 1 (2010), 35–71.

[5] Christoph Hube and Besnik Fetahu. 2018. Detecting Biased Statements in Wikipedia. In *Companion Proceedings of the The Web Conference 2018 (WWW '18)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, 1779–1786. https://doi.org/10.1145/3184558.3191640

[6] David A Jones. 2002. The polarizing effect of new media messages. *International Journal of Public Opinion Research* 14, 2 (2002), 158–174.

[7] Emma Lurie and Eni Mustafaraj. 2018. Investigating the Effects of Google's Search Engine Result Page in Evaluating the Credibility of Online News Sources. In *Web Science*. ACM, 107–116.

[8] Connor McMahon, Isaac L Johnson, and Brent J Hecht. 2017. The Substantial Interdependence of Wikipedia and Google: A Case Study on the Relationship Between Peer Production Communities and Information Technologies.. In *ICWSM*. 142–151.

[9] Ronald E. Robertson, Shan Jiang, Kenneth Joseph, Lisa Friedland, David Lazer, and Christo Wilson. 2018. Auditing Partisan Audience Bias Within Google Search. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 148 (Nov. 2018), 22 pages. https://doi.org/10.1145/3274417

[10] Annabel Rothshild, Emma Lurie, and Eni Mustafaraj. 2019. How the Interplay of Google and Wikipedia Affects Perceptions of Online News Sources. In *Computation + Journalism Symposium*.

[11] Michael Schudson. 1978. *Discovering the news: A social history of American newspapers*. Basic Books.

[12] Michael Schudson. 2001. The objectivity norm in American journalism. *Journalism* 2, 2 (2001), 149–170.

[13] Sam Wineburg and Sarah McGrew. 2017. Lateral Reading: Reading Less and Learning More When Evaluating Digital Information. *Stanford History Education Group Working Paper No. 2017-A1* (2017). https://ssrn.com/abstract=3048994