

Analyzing Support for U.S. Presidential Candidates in Twitter Polls

STEPHEN SCARANO

VIJAYALAKSHMI VASUDEVAN

University of Massachusetts Amherst, U.S.A.

MATTIA SAMORY

Sapienza University of Rome, Italy

JUNGHWAN YANG

University of Illinois Urbana-Champaign, U.S.A.

PRZEMYSŁAW A. GRABOWICZ

University of Massachusetts Amherst, U.S.A.

Polls posted on social media can provide information about public opinion on a variety of issues from business decisions to support for presidential election candidates. However, it is largely unknown whether the information provided by social polls is useful or not. To enhance our understanding of social polls, we examine nearly two thousand Twitter polls gauging support for U.S. presidential candidates during the 2016 and 2020 election campaigns.

First, we describe the prevalence of social polls. Second, we characterize social polls in terms of the engagement they elicit and the response options they present. Third, leveraging machine learning models, we infer and describe several characteristics, including demographics and political leanings, of the users who author and interact with social polls. Finally, we study the relationship between social poll results, their attributes, and the characteristics of users

Contact author: pgrabowicz@umass.edu

Date submitted: 2024-04-27

interacting with them. Our findings suggest how and to what extent polling on Twitter is biased in terms of content, authorship, and audience. The 2016 and 2020 polls were predominantly crafted by older males and manifested a pronounced bias favoring candidate Donald Trump, whereas traditional surveys favored Democratic candidates. We further identify and explore the potential reasons for such biases and discuss their repercussions.

Keywords: *public opinion, opinion polls, social media*

Social media and the Internet provide an unprecedented opportunity to observe political dialogue on a large scale. Platforms such as Twitter and Facebook have transformed users from passive consumers into active participants who produce and share information. This shift has led to a phenomenon where the buzz generated by users on social media not only serves as an indicator of public interest in various topics but also can potentially influence mainstream media coverage (Wells et al., 2016). Consequently, social media has garnered significant attention from the public, media, and political elites seeking to gauge what people think about key issues in society (McGregor, 2019).

Over the last decade, social media have emerged as an important source of political information (Mason Walker and Katerina Eva Matsa, 2021; Wells et al., 2016; McGregor, 2019). Journalists have employed social media in various ways, from capturing public responses to events such as debates and assessing presidential candidate performances (McGregor, 2019) to amplifying the reach of political messages from Donald J. Trump by featuring his tweets in mainstream news (Wells et al., 2016). For the general public, social media is not only a platform for community building but also serves as a tool for political engagement and voter mobilization (Bond et al., 2012; Rojas and Puig-i-Abril, 2009). It is also an educational tool that can be used to learn about political and social issues (Perrin, 2020; Bail, 2021). Moreover, social signals, such as user responses and comments, can also influence public perceptions about social issues (Muchnik et al., 2013; Kramer et al., 2014; Grabowicz et al., 2020). These pivotal roles social media plays in current politics underscore the need to examine the expressions of public opinion and discourse on social media.

Researchers have been exploring various methods to systematically estimate public opinion using social media data (Beauchamp, 2017; DiGrazia et al., 2013; Zhang et al., 2022a; Li et al., 2023). For example, Zhang et al. (2022a) tracked textual contents expressed by sub-groups on Twitter to understand how politically opposing groups react differently to political events. Other studies used similar techniques to estimate the results of traditional opinion polls from each U.S. state (Beauchamp, 2017) and to approximate individual politicians' electoral performances (DiGrazia et al., 2013). However, these methods have encountered limitations in accuracy and predictive capability, arguably due to the signals of political support used to estimate public opinion, which can be indirect—such as tweet resharing behavior—and complex—such as the natural language content of tweets (O'Connor et al., 2010; Jungherr, 2015; Schober et al., 2016; Klačnja et al., 2018; Dong and Lian, 2021).

Our study focuses on understanding a different indicator of public opinion on social media: *social polls*. These polls, widespread on social platforms, are informal surveys created and shared by users. Elon Musk, the owner of X (formerly known as Twitter) and Tesla, famously used this feature to make key business decisions, e.g., about Twitter's CEO change (Mehta, 2022). Due to their ease of use and quick turnaround time for gauging users' support for political candidates, Twitter polls have gained popularity during election campaigns and it is not uncommon to see social polls amass hundreds of thousands of votes. Such polls can be regarded as a more direct method of sensing public opinion compared to earlier techniques. Yet, little scholarly attention has been paid to understanding social polls.

Social polls encode political information that users express, endorse, and share spontaneously. As such, social polls convey a form of political engagement that traditional surveys often miss. The unique affordances of social media, such as visibility, editability, persistence, and association (Treem and Leonardi, 2013) that are inherent to networked technologies, can facilitate the amplification and spread of both information and social behavior. This new way of social engagement can potentially challenge traditional norms of social interaction and affect the discourses in online public spheres (boyd, 2011). In particular, for individuals disenfranchised from the elite-dominated political culture and mainstream media, the affordances of social media offer opportunities to create and spread their

narratives. Social media platforms emerge as critical arenas where users can express their viewpoints and, importantly, construct and reinforce their social identities (boyd, 2011). Thus, social polls offer a unique outlook on new ways to engage in politics.

However, it is important to understand the quality of social polls as proxies for public opinion. Social media data often disproportionately reflect the expressed views of a reactive, polarized, and highly engaged subset of users (Zhang et al., 2022a) and can be influenced by the activity of bots and astroturfing accounts (Keller et al., 2020a; Ferrara et al., 2016). Furthermore, despite their resemblance to traditional surveys, social polls inherently lack scientific rigor due to non-systematic sampling of the respondents and the absence of demographic information about them. Therefore, these factors may contribute to potential biases in poll outcomes (Auxier and Anderson, 2021; Wojcik and Hughes, 2019). Thus, while social polls provide rich insights into political behaviors on social media, it is paramount to first unpack their characteristics before considering them as indicators of public opinion more broadly.

In this study, we systematically describe social polls, their integrity, and the extent of their biases, to clarify their relevance in the online political landscape and their idiosyncrasies compared to traditional surveys. Our analyses focus on a specific set of Twitter polls posted during the concluding months of the 2016 and 2020 U.S. presidential elections (e.g., “Who has your vote? Biden or Trump?” and “Would you vote for Clinton or Trump in the upcoming election?”). We begin by describing Twitter polls related to the U.S. presidential elections. Then, we address the following research questions.

RQ1: What is the prevalence of social polls and how does it vary over time?

RQ2: What are the characteristics of social polls?

RQ3: What are the characteristics of users who engaged in social polling?

RQ4: What relationships exist between these characteristics and poll outcomes?

To answer these research questions, we assess the extent of bias in Twitter polls by comparing their results with those from traditional election polls and exit polls. Our examination focuses on potential sources of discrepancy between Twitter polls and actual election results, considering factors such as (1) the order effects of vote options in Twitter

polls, (2) the lack of partisan and (3) demographic representativeness among social media users who engage with the polls, and (4) the presence of bots interacting with these polls. Finally, we discuss these results in the context of political participation and public opinion mining.

Related Work

Social polls as political engagement

Public opinion, often defined as an aggregate of individual opinions (Price, 1992), is a crucial part of a well-functioning democracy. Its importance escalates during elections, where campaigners utilize insights from public opinion to shape their campaign strategies and to maximize voter support. One of the most popular ways for assessing public opinion is survey-based opinion polls (Price, 1992). However, critics argue that such polls fail to account for the dynamic interplay of societal groups, overlooking the hierarchical influences and the roles of key individuals and groups in shaping opinions (Blumer, 1948; Herbst, 1998). Similarly, some scholars highlight that relying on survey-based methods might not measure ideal democratic behaviors, such as active participation, open discussion, and thoughtful deliberation (Berelson, 1952; Habermas, 1989).

Social polls, while appearing similar to traditional opinion surveys, serve a distinct function. They should be regarded as indicators of political engagement among specific subsets of the population rather than as accurate reflections of general public opinion.

This distinction is important because the observed opinions on social media are inherently biased. Irrespective of the volume of trending messages, those who actively post political content on these platforms are more likely to have strong political views. Thus, it will be challenging if we try to reproduce the results of traditional opinion polls with social media data.

A historical example illustrating the pitfalls of relying on a biased sample is the Literary Digest's 1936 survey. This survey aimed to forecast the outcome of the U.S. Presidential election using 10 million ballots but failed despite receiving 2.3 million responses. The failure was mainly due to a low response rate and a bias in the demographic of re-

spondents who returned their ballots (Squire, 1988). In a similar vein, the vast number of tweets generated on Twitter every hour should not be misconstrued as a comprehensive representation of the wider public opinion, as they are likely biased towards the views of a more politically vocal subset of Twitter users. Therefore, in this study, we conceptualize social polls as signals of political engagement rather than an accurate reflection of (traditional) public opinion. To understand this under-studied phenomenon, our initial research question asks about the prevalence of this emerging form of political engagement.

RQ1: What is the prevalence of social polls and how does it vary over time?

The new way of political engagement on social media can be attributed to their unique technological affordances. Social media affordances are the perceived and potential properties of social media that emerge from the interplay of technology, social dynamics, and contextual factors (Ronzhyn et al., 2022). As these properties both enable and define specific uses of these platforms, communication scholars categorize social media as a type of networked publics (boyd, 2011). This perspective highlights that the distinctive technological affordances of these platforms enable a variety of actions, such as one-click retweeting of user-generated content and nurturing the formation of connections and networks among users (Treem and Leonardi, 2013) and they can shape social media as a venue where they create and share political information and political actions that potentially lead to online public spheres (Ronzhyn et al., 2022; boyd, 2011). The unique affordances of social media, especially Twitter polls, may shape user responses and their answers to the polls. To further examine the relationship between various affordances of Twitter polls and the votes, we dissect our remaining research questions into nuanced inquiries.

RQ2.1: Are there correlations between the number of votes and the counts of retweets and favorites?

Social media users may perceive platforms like Twitter, particularly Twitter polls, as a means of expressing their political views and influencing political discourse. Social media affordances such as social polls present opportunities to construct and propagate narratives and perspectives independently from the elite-led political discourse and mainstream news media and therefore may be favored for political expression by different groups. This pattern

has been documented in a recent study done by the Pew Research Center. The study reveals systematic differences in the trust of Republicans and Democrats in different types of media. Democrats trust news media more if they perceive them as “mainstream”, whereas Republicans if *not* mainstream (Gottfried, 2021). Similarly, there is a trend that individuals identifying as young Republicans trust social media more than national news media, while young Democrats trust more national news media than social media (Liedke and Gottfried, 2022). Therefore, we expect that these differences in trust can be reflected in the ways that Republican and Democrat users create and interact with political polls on Twitter. We propose RQ 2.2 to examine whether the positioning of response options of Twitter polls reflects the biases of poll authors.

RQ2.2: How are the voting options presented in social polls?

Biases in Twitter poll participants

Social media users in the U.S. tend to skew more liberal in their political affiliations compared to the general population. A Pew Research Center study shows that 36% identify as Democrats among social media users, whereas the corresponding figure in the U.S. general population is 30% (Wojcik and Hughes, 2019). Moreover, the observable behavior of these populations may differ. Research suggests that conservatives on social media are more inclined to engage with liberal discourse on contentious issues such as gun control, compared to their liberal counterparts (Zhang et al., 2022b). Our next research question aims to delve into the political leanings of individuals participating in Twitter polls, seeking to uncover any potential political biases within them.

RQ3.1: What are the political orientations of users who participate in social polling?

On a similar note, social media users do not accurately reflect the demographic makeup and location distribution of the general population. Audiences on Twitter and Facebook are significantly younger (Mellon and Prosser, 2017; Wojcik and Hughes, 2019). Facebook users skew towards a female demographic, while Twitter users are biased towards men (Mislove et al., 2011; Mellon and Prosser, 2017; Wojcik and Hughes, 2019). To examine

these biases in Twitter polls, we propose RQ 3.2.

RQ3.2: What are the demographic characteristics of users who participate in social polling on Twitter?

Finally, social media data could be influenced by inauthentic activities such as astroturfing campaigns and bots. As famously manifested in Russia’s *Internet Research Agency* (IRA) intervention in the 2016 U.S. Presidential election, many organized or sponsored inauthentic social media operations are uncovered and documented (Schoch et al., 2022). Similarly, research has long found that a vast amount of bot accounts generate social media posts daily, and, further, that the prevalence of bot accounts on Twitter has only increased since Elon Musk’s acquisition of the site in 2022 (Ferrara et al., 2016; Hickey et al., 2023). The tools for manipulating social media discourse, now broadly accessible via online services, have expanded their reach beyond governmental institutions to everyday consumers (Al-Rawi and Rahman, 2020). These external influences have the potential to skew metrics and distort the representation of public opinion, particularly as some campaigns with political motivations deliberately exploit them to sway the course of social media dialogues (Schoch et al., 2022). If these kinds of inauthentic information operations and bot activities can influence Twitter activities, they can also create another bias in Twitter polls. Thus, we aim to identify the extent of inauthentic accounts related to Twitter Polls.

RQ3.3: What is the fraction of bot accounts among these users?

Biases in Twitter poll outcomes

The characteristics of social polls and of the users engaging with them could result in biased poll outcomes. To begin with, the positioning of vote options in such polls may be biased and this bias can influence social poll results, similarly to how it influences traditional poll results (Strack, 1992). Further, the opinions of specific demographics can be over-represented because (i) Twitter users are more likely to be male and young (Mislove et al., 2011; Wojcik and Hughes, 2019), and (ii) politically-interested users are non-representative of all users (Hughes and Asheer, 2019; Hughes, 2021). Finally, the prevalence of bot accounts, astroturfing campaigns, and artificially-acquired likes and comments, might distort various aspects of Twitter activity including social polls (Keller et al., 2020a; Ferrara et al.,




Figure 1. Two Twitter polls with the highest number of votes in the dataset.

2016). With the following research questions, our study aims to provide detailed descriptions of the poll outcomes and their biases.

RQ4.1: How do the results of social polls compare to those of mainstream polls?

RQ4.2: How do poll attributes, such as the order or response options, relate to their outcomes?

RQ4.3: How do user attributes relate to the outcomes of social polls?

Data

We collected and labeled a large dataset of Twitter polls gauging support for the 2016 and 2020 presidential elections and users interacting with these polls. To facilitate the reproduction of the results of this study, we share the tweet IDs of all collected polls and user IDs of all users interacting with these polls on Zenodo.¹

Social polls. Provided access to the Twitter API v2, we compiled election polls from the periods of 11/02/2019–11/30/2020 and 11/02/2015–11/30/2016 by making full-archive searches for the queries “vote trump biden” and “vote trump clinton” respectively. This step resulted in a vast number of polls. To identify polls gauging support for presidential candidates (Donald Trump and Hillary Clinton in 2016, and Donald Trump and Joe Biden in 2020), we focused on the polls that explicitly mention them in the post’s text.

¹Link: <https://doi.org/10.5281/zenodo.11049780>

This resulted in 4,551 polls. To identify relevant polls, we manually inspected all of them. To this end, we developed a labeling guideline that defines as relevant the polls gauging support for the 2016 and 2020 presidential election candidates either by (i) directly asking for voting preferences (e.g., "Who has your vote?"), or (ii) asking for election predictions (e.g., "Who do you think will win the presidential election vote?"). This approach identified a total of 1,753 relevant Twitter polls when excluding polls posted after respective election days (see examples in Figure 1). We extracted and normalized poll outcomes for the two major candidates (Trump vs. Clinton in 2016, and Trump vs. Biden in 2020) in polls with greater than 2 options.

User profiles and reactions to social polls. Using Twitter API, we retrieved each poll’s author profile as well as those of their followers. Then, we collected all retweeters and favoriters of each poll. Similarly, we retrieved followees of poll authors, as well as retweeters and favorites of the polls.

Mainstream polls. In addition to the Twitter set, we collected 192 mainstream polls from 2020 and 144 from 2016 aggregated by *FiveThirtyEight* 2 and used national exit poll data distributed by the *Roper Center* (Pool, 2016, 2020).

Methods

We augmented the profiles of poll authors and their followers, as well as poll retweeters and favoriters, by inferring their attributes (age, gender, political affiliation, botness, organization status) using machine learning models and by extracting location information directly from their profiles. To verify the accuracy of the inference, we asked two trained annotators to label a subset of polls and users. Finally, we related potential sources of bias to poll outcomes using regression models.

Gender and Age We classified the gender and age of users, attributes commonly understood to correlate with voting behavior, and used to stratify or poststratify mainstream surveys (Silver, 2021). To do so, we employ the multilingual, multimodal, and multi-label

²<https://projects.fivethirtyeight.com/polls>

machine learning tool *M3-Inference* (Wang et al., 2019). *M3-Inference* is a deep learning text and image model that uses usernames, profiles, and photos to infer age and gender with state-of-the-art accuracy while diminishing algorithmic bias in comparison to other approaches. Since the model additionally infers the likelihood that the given account represents an organization, we exclude from our analysis those users who exceeded an org score of 0.90.

Location Participation in Twitter polls may differ by location, both between U.S. states and globally. To this end, we inferred the location of Twitter users. On Twitter, users can optionally disclose their position, in plain text, using the `location` field of their profile. We resolved such entries to geolocations using Photon, an open-source geocoder built for OpenStreetMap data.³ We could infer the location of 4,737,715 users thusly. To expand coverage, for users whose location could not be geocoded via the previous method, we combined the `location` and `description` plain-text fields of user profiles, and extracted emojis corresponding to national flags. After excluding cases of users displaying flags of multiple countries, we could infer the location of an additional 156,788 users via this second method. In total, 4,894,503 users were mapped to countries, 1,019,610 of which could be resolved to specific U.S. states.

Political Ideology We estimated relative political ideology in Twitter polls using a Markov chain Monte-Carlo approach, following (Barberá et al., 2015; Barbéa, 2015). The approach infers the political ideology of a user based on list of users they follow: users following primarily well-known right-wing accounts (such as `@realDonaldTrump` and `@FoxNews`) will be scored as more right-wing than users following left-wing accounts (such as `@JoeBiden` and `@CNN`). Each user instance is mapped to a continuous political ideology value in the interval $[-3, 3]$, indicating left-to-right. While we provide the raw distribution of these values in Section 3, we later discretize this range into three bins (*Left*, *Moderate*, *Right*), splitting the space evenly for simplicity. Prior work by (Barberá et al., 2015) shows that this approach performs comparably with standard ideological assessment surveys. We inferred political affiliation for all poll authors, retweeters, and favoriters.

³<https://photon.komoot.io/>

Botness We estimated the distribution of human and bot accounts using the machine learning classifier *Botometer*. Trained on a dataset of 5.6 million tweets, *Botometer* is a random forest classifier that evaluates network, user, friend, temporal, content, and sentiment features to label a profile as authentic or artificial (Davis et al., 2016). While *Botometer* provides a bot score within the interval $[0, 1]$ to gauge bot likelihood, it does not specify a cut-off threshold for determining whether an account is a bot or not. We determined the cut-off threshold of 0.83 so that a human annotator and Botometer would classify a similar fraction of users as bots in a random sample. We obtained bot scores for all poll authors and retweeters. We could obtain bot scores for only 35% of favorites and 6% of followers because the Botometer service ceased its operation once Twitter stopped providing free API access to researchers in the middle of 2023 when this study was conducted.

Validation of polls and inferred attributes Although the algorithms we used for estimating demographic and political attributes have undergone prior testing and validation, given the heavy reliance of our study on social poll data from Twitter and estimates of various demographic and political attributes, we took additional steps to validate the relevancy of the data we used and the results obtained from machine learning algorithms. To accomplish this, we rely on human coders to validate the integrity of the poll data and the inferred attributes.

The first validation task determined whether the Twitter polls are relevant to the construct that we aim to measure: public opinion on election outcomes. Three coders independently rated the polls using the following criteria: the poll should at least include both of the top leading candidates of each election (i.e., Hillary Clinton and Donald Trump in 2016 and Joe Biden and Donald Trump in 2020) and the poll instructions should directly ask about users' voting intention or their projection of the election outcomes. We constructed a random sample of 196 polls and asked the coders to code them independently. The three coders' inter-rater reliability (IRR) measured in Fleiss' Kappa was 0.914. Given the high IRR, one of the coders coded the rest of the dataset and identified a final set of 1,753 relevant polls out of 4,000.

The second validation task evaluated the accuracy of the methods for estimating

the demographic and political attributes of Twitter users, which are used in this study to describe potential biases in Twitter polls. Specifically, we validated model accuracy in (1) distinguishing between organizational and personal accounts, (2) discerning bot-like or human-like traits of the account, (3) estimating the political leaning of the account, and for personal accounts, and identifying the (4) age and (5) gender of the account holder. We asked one human coder to review the estimated attributes for a random set of 239 Twitter accounts. To simplify the process, we showed the machine-determined attributes to the coder (who is not a co-author of this work) and asked them to determine whether they agreed with those classifications or not. The methods achieved approximately 93% accuracy in distinguishing between organizational and personal accounts, 91% accuracy in assessing bot-likeness, 93% accuracy in estimating political ideology, 91% accuracy in estimating the age of the account holders, and 88% accuracy in classifying the gender of the account holders.

Regression of poll outcomes To relate potential sources of bias to poll outcomes, we performed ordinary least squares regressions using poll outcomes as the dependent variable, y , operationalized as the fraction of votes cast in favor of Trump in the 2016 and 2020 election seasons. As the potential sources of bias in poll outcomes we considered as independent variables the characteristics of potential poll respondents, such as their gender, age, political ideology, location, and botness. Since poll respondents are anonymous, we considered retweeters and favoriters as proxies for voters. More precisely, as predictors, we used the marginal fractions $p_p^d(g)$ of potential voters in a poll p belonging to a population stratum g . The resulting model is

$$\hat{y}_p = \sum_{d \in \mathcal{D}} \sum_{g \in \mathcal{G}_d \setminus \{g_d^{\text{rf}}\}} \beta_g^d p_p^d(g) + \beta_0, \quad (1)$$

where $\mathcal{G}_d$ stands for all population strata within the dimension $d \in \mathcal{D}$, e.g., $\mathcal{G}_{\text{gender}} = \{\text{male}, \text{female}\}$. The dimensions, $\mathcal{D}$ are the users' gender (male or *female*), age group (*less than 30*, *between 30 and 39*, *greater than 40*), political ideology (democrat, *moderate*, republican), and location (U.S. red state, U.S. blue state, *U.S. swing state*). We also used the bot classification outcome (bot or *not bot*) as another feature to check whether the presence of bots is related to poll outcomes. To account for biases stemming from

response ordering, we also encoded whether Trump or the Democratic candidate is listed first among poll response options (Trump or *not Trump*). The reference strata, g_d^{rf} (marked in *italics* in the text), are excluded to avoid feature collinearity arising from the probability normalization $\sum_{g \in \mathcal{G}_d} p_p^d(g) = 1$. To reduce noise in the dependent variable and the number of missing predictors, we exclude polls with $M = 2$ users or fewer in the union of retweeters and favoriters. We imputed missing values with the mean of the corresponding predictor.

Results

Next, we address each of our four major research questions.

The prevalence of social polls

We first examined the prevalence of social polls over the two election campaigns (RQ1). To get a better understanding of this pattern, we visualize the number of social polls and mainstream polls side by side. Figure 2 suggests that both social and mainstream polls exhibit similar temporal trends, with an increase in the number of polls as election day approaches. The number of both polls peak in the last week of October for both 2016 and 2020. Notably, social polls have a broader temporal range, spanning almost the entire election year. Before September of the respective election years, there were 680 (32%) social polls, while there were only 15 (4%) corresponding mainstream polls in 2016 and none in 2020. Although the sample of mainstream polls from *FiveThirtyEight* might not encompass all pertinent mainstream polls, there is still a strong indication that social polls are significantly more frequent between January and September 2020 than mainstream polls.

As described in Table 1, social polls not only were exposed to a large potential audience of the poll authors' followers, but also elicited significant user engagement through votes, favorites, and retweets. When visualizing the average number of votes (see Figure 3), it becomes apparent that the popularity of social polls assessing support for U.S. presidential candidates follows a heavy-tailed distribution. Although the majority (70%) of polls received fewer than 100 votes, the top 15% polls attracted numbers of votes that are comparable to mainstream political polls (and sometimes two orders of magnitude more votes).

It is important to note that our sample of polls is a small subset of all polls men-




Figure 2. The number of Twitter and mainstream polls published over the course of the 2016 (top) and 2020 (bottom) U.S. presidential election years.

tioning presidential election candidates among response options. Given the popularity of social media polls, they may become a treasure trove of public opinions.

Moreover, there was a substantial increase in social poll participation from 2016 to 2020. The total number of votes for social polls showed an eight-fold increase between the two election cycles as detailed in Table [1](#).

The characteristics of social polls

We turn our attention to describing various characteristics of social polls, focusing on user engagement and poll design (RQ2).

RQ2.1: Are there correlations between the number of votes and the counts of retweets and favorites?

One of the unique aspects of social polls is that people can engage with them by sharing the poll with their followers or by favouriting them to increase visibility. Figures

Table 1: Dataset Summary.

	2016		2020		Total
	Total	Top 15%	Total	Top 15%	
Votes	88,300	75,919	925,421	853,633	1,013,721
Retweeters	2,991	2,669	36,290	34,216	39,281
Favorites	2,623	1,915	26,078	23,519	29,514
Followers	450,612	140,712	17,728,247	12,853,901	18,178,859
Polls	348		1,405		1,753
Authors	298		960		1,258

Note. Quantitative summary of the dataset of Twitter polls gauging support for the 2016 and 2020 U.S. presidential candidates, including the number of polls and votes in these polls, as well as the number of poll authors, their followers, and the retweeters and favoriters of polls. All presented user counts are the numbers of unique users.





Figure 3. Distributions of the number of votes in Twitter and mainstream polls gauging support for the 2016 (top) and 2020 (bottom) U.S. presidential candidates.

[5](#) and [6](#) show that user engagements follow a heavy-tailed distribution: 36,885 of the total 39,281 collected retweets (93%) belong to the top 15% of polls. Likewise, 25,434 of the total of 28,701 favorites (89%) are concentrated within this same fraction (see [Table 1](#)). This suggests that only a relatively small fraction of social polls garner significant engagement from social media users.



(a) Histogram of the number of poll options.

(b) The fraction of head-to-head polls grows over the course of the election year.

Figure 4. Breakdown of the number of Twitter poll options.

How do these user engagement metrics correlate with the number of votes? Retweets spread polls to new users, hence exposing them to more potential voters. Similarly, favorites reflect user interests. The high favorite count may influence other potential voters, possibly alike, to vote due to social influence (Muchnik et al., 2013; Kramer et al., 2014; Grabowicz et al., 2020) or by boosting their exposure via Twitter’s algorithmic news feed (Huszár et al., 2022). Thus, it is plausible that retweets and favorites boost poll visibility, thereby increasing the likelihood of further interactions.

We found that there are strong correlations between the number of votes and the number of retweets. Pearson’s r is 0.84 ($p < 0.001$) for 2016 and 0.60 ($p < 0.001$) for 2020, suggesting that polls attract more votes via retweets. We find similar correlations between the number of votes and the number of favorites: $r = 0.84$ ($p < 0.001$) for 2016 and $r = 0.86$ ($p < 0.001$) for 2020.

These results suggest that there are strong positive correlations between user engagement and the number of votes. However, we do not yet know whether retweets push biased political polls into echo chambers and whether this process exacerbates biases in poll outcomes. We will examine these questions in RQ3 and RQ4, respectively.




Figure 5. Histograms of the number of retweets per poll.




Figure 6. Histograms of the number of favorites per poll.

RQ2.2: How are the voting options presented in social polls?

The selection of response options in a survey has important implications for respondents' behavior and the quality of data. Respondents' choices are related to response bias, response variability, and social desirability bias, among other phenomena (Price, 1992). To answer RQ2.1, we describe the voting options provided in social polls.

By design, the social polls in our dataset must include the two leading candidates of each election campaign (Trump & Clinton in 2016 and Trump & Biden in 2020); however, some polls also include other candidates, providing up to four options to choose from. As shown in Figure 4a, the vast majority (72%) of the dataset are 2-option polls, whereas

3-option and 4-option polls are less frequent (15% and 13%). As the presidential election campaigns develop and the Republican and Democratic party candidates emerge from the presidential primaries over an election year, the fraction of the head-to-head polls increases (see Figure 4b for 2020, data for 2016 was too sparse to observe this pattern).

Another interesting aspect of polls to consider is the “order effect,” which is a well-documented phenomenon that different orders in response options can influence survey outcomes (Strack, 1992). Our data, shown in Figure 7, reveal the pattern where candidates from the two major parties consistently dominate the first and second option placements. Trump and the respective Democratic candidate consistently are ranked first more often than the statistical expectations (i.e., 33% for 3-option polls and 25% for 4-option polls). Notably, regardless of the number of response options, Trump tends to be positioned above the Democratic candidate, suggesting a potential bias in poll design. We will further investigate this bias in response options’ impact on poll outcomes in RQ4.1.

The characteristics of users engaged in social polling

As discussed in Section 1, social poll outcomes may be the result of a non-representative user base engaged in the polling process. In this section, we outline the characteristics of the authors and potential voters in such polls. Considering that voters in social polls can encompass those who retweet or favorite the poll, as well as the followers of the poll authors, we discuss the attributes of these three groups of users. We assess whether the accounts are human versus automated bots and determine whether they resemble a representative sample in terms of their age, gender, and political orientation. These findings are then juxtaposed with data from corresponding exit polls.

RQ3.1: What are the political orientations of users who participate in social polling?

One of the most pertinent voter characteristics to political polls is the ideological makeup of the participants. We estimate the political ideology of the poll authors, retweeters, and favoriters using methods suggested by Barberá et al. (2015). Although an in-depth discussion of this method is out of the scope of the current paper, a brief explanation is provided in the Methods section.




Figure 7. The fraction of Twitter polls listing a given candidate as first among the potential answers.

Note. The majority of polls rank Trump as first. The dashed lines mark the fractions of 50% for 2-option polls, 33.3% for 3-option polls, and 25% for 4-option polls, which correspond to the hypothetical condition where all candidates are ordered randomly (fairly) among poll options.

To compare Twitter users' political ideology scores with the political affiliation data from exit polls, we converted the continuous scale of inferred political ideology of users into a discrete scale. Also, for this discussion, we use the terms political ideology and political affiliation interchangeably. The result of this classification is summarized in Figure 8. Our analysis shows that the distributions of political ideology of poll authors and users interacting with polls are skewed towards the right. This result is consistent with the fact that social poll results are skewed towards Trump (RQ2.2). Interestingly, retweeters and favoriters of social polls are even more likely to be conservative than the poll authors themselves. This asymmetry resembles the one observed in the distribution of political ideology of users who interact with misinformation, which also predominantly leans to the right (Nikolov et al., 2021; González-Bailón et al., 2023).




Figure 8. Political affiliation distribution of poll authors, retweeters, and favoriters

Note. The applied political affiliation inference tool projects users into a continuous partisanship interval $[-3, 3]$, which we discretize into equally-sized partisanship bins (dashed bars of 3 distinct colors) to make comparisons to exit polls, which represent ideology using 3 categories: left, moderate, right.

RQ3.2: What are the demographic characteristics of users who participate in social polling?

We analyze the gender and age of poll authors, retweeters and favoriters, and author followers. First, we find that the fraction of males is about 2 times larger among poll authors than among exit poll respondents (Figure 9) and that the fraction of poll authors below 30 years old is almost 3 times larger (Figure 10). Our results conform to prior research suggesting that Twitter skews heavily male and young (Mellon and Prosser, 2017). Interestingly, these biases are greater among authors and their followers than among retweeters and favoriters, suggesting that while young males are mobilized to author polls, people engaging with the polls are more similar to the general population in terms of age and gender.

RQ3.3: What is the fraction of bot accounts among users who participate in social polling?

We applied the popular bot classifier, *Botometer*, to estimate the fraction of bots among users related to polls. Figure 11 plots the fraction of bots separately for poll authors,




Figure 9. Gender distribution among social poll authors, their followers, as well as the retweeters and favoriters of the polls

Note. For comparison, the rightmost figure shows gender distribution for the exit polls of 2016 and 2020, respectively.




Figure 10. Age distribution of social poll authors, their followers, as well as the retweeters and favoriters of the polls.

Note. For comparison, the rightmost figure shows the age distribution for the exit polls of 2016 and 2020, respectively. The bars are color-coded to mark the correspondence between the age brackets, e.g., the second bin for social polls corresponds to the first two age bins for exit polls.

retweeters, favoriters, and author followers. To contextualize these results, we compared the Botometer scores of the authors of presidential-election polls to that of a random sample of Twitter polls posted over the same period and of similar vote distribution. Poll authors




Figure 11. The fraction of accounts that are likely to be bots among poll authors, retweeters, favoriters, and poll author followers.

are not found to score significantly higher on botness than a random sample of non-election poll authors. In contrast, retweeters of political polls are the highest-scoring user group and indeed are 4 times more likely than authors to be classified as bots (17% vs. 4%). This suggests that there may be some degree of astroturfing, i.e., inauthentic user behavior, involved in political campaigning via social polls (Keller et al., 2020b).

The relationship between social polling characteristics and poll outcomes

With the next set of research questions, we aim to uncover potential sources of social poll biases by comparing the characteristics of Twitter polls with their outcomes. To summarize the results, we regress poll outcomes against the fraction of potential voters of certain political ideology, gender, age, authenticity, and U.S. state (Table 3).

RQ4.1: How do the results of social polls compare to those of mainstream polls?

Using the results of mainstream social polls and actual vote share as the reference, we describe the poll outcomes. Given Trump’s candidacy in both the 2016 and 2020 elections, we present results by referring to the fraction of votes cast in his favor.

Overall, the results of Twitter polls demonstrate that there is a substantial partisan slant, with a consistent leaning toward Trump. When analyzing the average Twitter poll




Figure 12. Average outcome of Twitter polls published before the 2016 and 2020 U.S. presidential election days the average outcomes of mainstream polls, and the official election outcomes.

Note. Gray bars signify the median result.

results across all polls during the 2016 and 2020 presidential elections (see Figure 12), it becomes evident that these polls tend to overestimate support for Trump compared to actual vote shares on election day. In contrast, mainstream polls tend to underestimate Trump’s support. Specifically, the median support for Trump in social polls is 17%p higher than that of mainstream polls (60% to 43% respectively) in 2020. Mainstream polls, on the other hand, exhibit a bias in favor of the Democratic candidate, with a margin of either 6.1%p or 3.8%p when compared to the actual election results. The 2016 election gap between the outcomes of Twitter and mainstream polls is 3%p larger (60% to 40%), largely the consequence of the well-documented inaccuracies in mainstream polls’ election predictions at that time (Silver, 2021; Shirani-Mehr et al., 2018; Gelman, 2021). This contrasting bias pattern suggests that social media polls contain unique information that can be useful for leveling the biases in mainstream polls.

RQ4.2: How do poll attributes, such as the order or response options, relate to their outcomes?

The gap between Twitter poll outcomes and actual election results widens in the

Table 2: Option Placement and Result Correlations.

Pearson ρ	2016	2020
2-option polls	-0.12 ($p < 0.01$)	0.01 ($p = 0.805$)
3-option polls	-0.34 ($p < 0.01$)	-0.46 ($p < 0.001$)
4-option polls	-0.59 ($p < 0.001$)	-0.35 ($p < 0.001$)

Notes. Pearson correlation between the fraction of votes for an option and its position among all potential poll answers, where position 1 is at the top, and position 4 is at the bottom of poll options. The correlation is computed separately for 2-option, 3-option, and 4-option Twitter polls of 2016 and 2020.

most popular polls: Trump’s lead (measured as a median poll result) accelerates from 59% for all polls to 69% for the top 15% of polls in 2016 and from 60% to 86% in 2020, indicating the influence of poll popularity on the bias. This result is consistent with our other findings that the number of votes correlate with the numbers of retweeters and favoriters who are more likely to be Republican than Democrat (Figure 8).

While the majority of Twitter polls position Trump as the first option among the list of potential poll answers, it is not clear whether this influences poll outcomes. Here, we compute the correlation between the positions and the fractions of votes cast for the candidates placed at the respective positions (Table 2), where position 1 is at the top, and position 4 is at the bottom of the poll options. We find a significant negative correlation between the positions and poll outcomes, suggesting that a candidate who is positioned higher on the list of potential poll answers gets more votes. This correlation, however, is smaller for polls with fewer options, and in the case of head-to-head polls from 2020, the correlation is insignificant. In fact, the regression of poll outcomes against user attributes and a binary variable encoding whether Trump is above the respective Democratic candidate among poll options reveals no significant impact of the position of poll options on their outcomes (Table 3). We explain this result with the prevalence of two-option polls and widespread awareness of the two main candidates in presidential election races.

RQ4.3: How do user attributes relate to the outcomes of social polls?

Next, we study whether cross-tabulated results of social polls exhibit similar patterns

Table 3: Regression Correlations.

	2016		2020	
Predictor	coef	P> t	coef	P> t
const	<i>0.41</i>	***	<i>0.36</i>	***
$p(\text{gender}=\text{male})$	<i>-0.02</i>		<i>-0.08</i>	**
$p(\text{age} \geq 40)$	<i>0.06</i>		<i>0.16</i>	***
$p(\text{age} \in [30, 39])$	<i>0.08</i>		<i>0.08</i>	**
$p(\text{ideology}=\text{dem})$	<i>-0.14</i>		<i>-0.14</i>	***
$p(\text{ideology}=\text{rep})$	<i>0.35</i>	***	<i>0.42</i>	***
$p(\text{location}=\text{blue state})$	-0.10		0.005	
$p(\text{location}=\text{red state})$	-0.02		0.007	
$p(\text{bot}=\text{yes})$	0.01		-0.02	
$p(\text{first option}=\text{Trump})$	-0.01		-0.02	
Dependent variable:	% for Trump		% for Trump	
No. observations:	141		569	
Adj. R²:	0.33		0.58	

Notes. Parameters of the linear regression model using as dependent variable percent support for Trump in 2016 and 2020 social polls. Predictors encode characteristics of potential voters and poll option ordering information. We indicate statistical significance at levels $p < 0.001$ (***), $p < 0.01$ (**), and $p < 0.05$ (*). Coefficient values in italic font are significant for at least one of the two datasets and have the same sign for both datasets.

to the ones of exit polls. For instance, are older or male Twitter users more likely to vote for Trump? While we do not have direct information about the voters in Twitter polls, poll retweeters and favoriters likely voted in the poll they interacted with. Furthermore, voters in a social poll will likely have similar characteristics to those of the poll’s author, due to the effect of social network homophily (Byrne, 1971; McPherson et al., 2001). For instance, social media users who are connected and interact tend to live close to each other and be of a similar age (Ugander et al., 2011; Grabowicz et al., 2016) and share political preferences (Conover et al., 2011; Barberá et al., 2015).

Remarkably, the 20220 poll results averaged separately for demographic groups of their authors, that is by age or gender group, reveal the same qualitative demographic cross-tabulation patterns as exit polls (compare blue and orange lines in the left columns of




Figure 13. The average fraction of votes for Trump among Twitter polls authored by a user of certain inferred gender, age, and political affiliation (blue) in comparison to the fraction of votes for Trump in respective exit polls (orange).

Note. The shaded area marks the 95% confidence interval.

Figure 13.⁴ According to the exit polls, older male voters are more likely to vote for Trump. Similarly, the results of polls authored by older or male users are more biased toward Trump than polls authored by younger or female users. The bias of social poll outcomes towards Trump in some cases is relatively more pronounced, e.g., among polls authored in 2020 by users 40 years old or older.

In addition, social poll outcomes show bias toward Trump, i.e., almost all blue points are above the orange one in Figure 13. We explain these results by noting that (i) Republican users have a much higher propensity to author and, even more so, retweet or favorite polls (Figure 8), and (ii) there are biases in poll results with respect to political affiliation of social media users (middle column in Figure 13). These biases can be addressed

⁴For 2016, the respective differences are not significant, due to four times smaller data size.

by poststratification similar to the one performed for Xbox surveys [Wang et al. \(2015\)](#), as we illustrate in the discussion, but its thorough evaluation is beyond the scope of this paper. However, there may be other sources of bias in social media polls, e.g., due to the bias in poll positions, the effect of echo chambers and algorithmic news feeds, or manipulation and campaigning via astroturfing.

We also study how social poll outcomes are related to the locations of users engaged in social polling. To this end, we split U.S. states into blue states, red states, and swing states (Wisconsin, Pennsylvania, New Hampshire, Minnesota, Arizona, Georgia, Virginia, Florida, Michigan, Nevada, Colorado, North Carolina, and Maine). Interestingly, in 2020 we see a larger support for Trump among users in swing states than in other states (bottom row in Figure [13](#)), which may correspond to an intensified campaigning by Trump's camp. However, judging based on the last column of Figure [13](#), it is not clear whether polls authored by bot accounts are biased toward Trump more than polls authored by authentic users.

The regression of social poll outcomes against inferred user attributes and poll option ordering reveals that the two attributes that are the most significantly related to the poll outcomes are the fraction of Republicans and people older than 40 among the potential poll voters (Table [3](#)). Older Republican users are likely to vote for Trump. Democrats and users located in blue states are likely to vote for Biden, although the respective coefficients are significant only for one of the two models, for 2020 and 2016 respectively, likely because of relatively small sample sizes. The regression models achieve high adjusted R^2 values, 0.58 and 0.33 for 2020 and 2016, despite using only a small set of features. Remarkably, all significant coefficients of the regression models maintain the same sign across the two election years, despite sample size differences, lack of temporal validity (e.g., the political ideology classifier was developed in 2020), and any other potential issues that are not considered in this manuscript. We conclude that despite potential manipulation and the impact of algorithmic news feeds and echo chambers, demographic patterns resemble those found in presidential exit polls, suggesting that social polls may contain valuable information about public opinion.

Discussion

Our examination of Twitter polls offers insights at multiple levels. First, this study provides a comprehensive description of how Twitter polls were used during the 2016 and 2020 U.S. presidential elections. Second, our findings suggest various sources of biases present in Twitter polls. Third, this study paves the way for future research in social polling.

This study suggests that Twitter polls could be an effective tool for boosting political engagement. Platforms like Twitter have the potential to cultivate a sense of community and belonging in the online political sphere, thereby stimulating greater political participation and voter mobilization (Bond et al., 2012; Rojas and Puig-i-Abril, 2009). Our study primarily focuses on identifying and understanding the factors that could introduce biases in social media polls, such as the presentation order of poll options, the spread of polls through homogeneous networks, and the influence of algorithmic news feeds. We notice a pattern of increased political engagement among Trump supporters, indicated by their active participation in sharing and interacting with Twitter polls. This pattern of engagement mirrors the pattern observed in 2020 where there was a 20% gap in *strong* support for Trump over Biden (Doherty et al., 2020). The findings are also consistent with what previous research suggests (Wells et al., 2016): Trump voters are more inclined to engage with social media content related to his candidacy.

We highlight the risks associated with undisclosed biases in social poll outcomes, which can act as a source of polarizing and misleading content (Wu et al., 2019). As seen in other studies (Vosoughi et al., 2018; Juul and Ugander, 2021; Zafar et al., 2016; Brady et al., 2019), misleading and polarizing Twitter polls garner more attention and spread more widely than true and balanced information. Such misinformation often resonates more with individuals who have strong partisan identities, exacerbating the spread of polarizing content (Babaei et al., 2019; Pennycook and Rand, 2021; Guess et al., 2021; Allcott and Gentzkow, 2017). Likewise, polarizing misinformation, such as biased polls, can gain more support among people with strong partisan identities. This phenomenon is reflected in the asymmetric distribution of partisanship among users engaging with misinformation (Nikolov et al., 2021; González-Bailón et al., 2023), a pattern also evident in those interacting with

social polls.

Furthermore, social polls offer a common, structured format to compare expressions of political preference across a variety of platforms with diverse user bases, scopes, and values, and thus enable an unprecedentedly nuanced characterization of public opinion and political engagement. While differences exist between social and mainstream polls, their parallels suggest a potential for future studies to apply statistical methods to correct biases in social polling. For instance, Wang et al. (2015) demonstrated that with the appropriate statistical adjustments of demographic and political variables, non-representative polls from Xbox users can yield accurate election forecasts. These findings underline the potential of unconventional data sources in capturing public opinion once biases in data are corrected. For instance, such poststratification could be performed using regression models such as the ones we constructed. The fact that these models provide consistent results that align with the patterns observed offline, adds to the promise. Future user attribute inference and poststratification approaches could establish social polls as a vital data source for estimating public opinion, either independently or in conjunction with mainstream polls. Overall, social media polls may become a valuable source of information about a wide breadth of political preferences. They offer an alternative to measuring public opinion from social media over sentiment analysis, which has had limited success so far (O'Connor et al., 2010; Jungherr, 2015; Schober et al., 2016) and faces multiple challenges due to its reliance on unstructured information in the form of text (Klašnja et al., 2018; Dong and Lian, 2021; Diaz et al., 2016).

Overall, this study opens up multiple avenues for future research. These include exploring the motivations behind publicly sharing political opinions in social polls, comparing audience perceptions of social polls versus institutional surveys, and their impact on offline political mobilization. While our study is focused on U.S. presidential races, future research could expand to other political elections, both domestic and international. Twitter (re-branded as X in 2023) is, at the time of writing, a primary avenue where the public and elite advance the political discourse in the U.S., and thus was a natural setting for studying social polls about presidential candidates. Yet, social polling as a form of political activity is a global phenomenon that is commonplace on several online platforms beyond Twitter. For instance, polls can be created on Facebook Pages and Groups. Correspondingly, the

research opportunities are far greater than what has been covered by the present study. Understanding how social polling varies across platforms, and how it is used across a wide range of topics, is necessary to comprehend social polling as a broader practice in an age of interconnected media.

Limitations

This study performs inference of several unobserved attributes of Twitter users, such as their age, gender, and political leaning. We recognize that, although we used the most accurate inference methods available and evaluated them on our data, such inference is still subject to error. Therefore, the results of the inference are used only in aggregate and should be interpreted as coarse-grained comparisons between user cohorts.

We note also that for a poll to be in our dataset, it must contain both "vote" and the candidates' names in its text, which is not exhaustive. Twitter API changes cut researchers' access to Twitter data in 2023 when this study was conducted. In light of this newly imposed restriction limiting Twitter's transparency, we deemed it important to publish the results of this study despite the potential incompleteness of the data, to keep the public informed about Twitter polls that could have impacted the voting behavior of U.S. citizens and could have reinforced the beliefs of a subset of the U.S. citizens in voting fraud. Subsequent work can identify usable and precise methods to cut through the vast noise in social platforms for a more complete analysis of relevant polls.

Broader Perspective, Ethics, and Competing Interests

This research obtained approval from the pertinent Institutional Review Board. This work concerns itself with the political opinions of Twitter users. As such, it is important to discuss the ethical aspects of the fielding and dissemination of this research.

First, we believe that social media platforms should be transparent about user activity in relation to political issues such as elections. We note that biased social polls can be misleading and can contribute to voter fraud belief. As such, it is important that social media platforms inform users about biases in political polls

Second, the computational model used for inference affords a specific and limited operationalization of gender as a binary construct. This operationalization is admittedly limited, but instrumental in comparing the results of this research to existing research hypotheses and data from mainstream polls that also adopt it. The authors recognize the need for further studies that endorse more complex conceptualizations of user characteristics, and especially that perform at-scale surveying of unobserved user characteristics, although such efforts exceed the scope of this first-comer analysis.

Acknowledgments

We thank Filippo Menczer and Kaicheng Yang for providing us the Botometer scores for millions of Twitter users. We also thank Pablo Barberá for sharing the estimated political leanings of Twitter users. We acknowledge support from the University of Massachusetts Amherst via an Interdisciplinary Research Grant.

References

- Al-Rawi, A. and Rahman, A. (2020). Manufacturing rage: The russian internet research agency’s political astroturfing on social media. *First Monday*, 25(9).
- Allcott, H. and Gentzkow, M. (2017). Social Media and Fake News in the 2016 Election. *Journal of Economic Perspectives*, 31(2):211–236.
- Auxier, B. and Anderson, M. (2021). Social Media Use in 2021. Technical report, Pew Research Center.
- Babaei, M., Chakraborty, A., Kulshrestha, J., Redmiles, E. M., Cha, M., and Gummadi, K. P. (2019). Analyzing Biases in Perception of Truth in News Stories and Their Implications for Fact Checking. In *Proceedings of the Conference on Fairness, Accountability, and Transparency - FAT* ’19*, pages 139–139, New York, New York, USA. ACM Press.
- Bail, C. (2021). *Breaking the Social Media Prism: How to Make Platforms Less Polarising*, volume 23. Princeton University Press.

- Barb era, P. (2015). Birds of the same feather tweet together: Bayesian ideal point estimation using twitter data. *Political Analysis*, 23:76–91.
- Barber a, P., Jost, J. T., Nagler, J., Tucker, J. A., and Bonneau, R. (2015). Tweeting from left to right: Is online political communication more than an echo chamber? *Psychological Science*, 26(10):1531–1542. PMID: 26297377.
- Beauchamp, N. (2017). Predicting and Interpolating State-Level Polls Using Twitter Textual Data. *American Journal of Political Science*, 61(2):490–503.
- Berelson, B. (1952). Democratic theory and public opinion. *Public Opinion Quarterly*, 16:313–330.
- Blumer, H. (1948). Public opinion and public opinion polling. *American Sociological Review*, 13(5):542–549.
- Bond, R. M., Fariss, C. J., Jones, J. J., Kramer, A. D. I., Marlow, C., Settle, J. E., and Fowler, J. H. (2012). A 61-million-person experiment in social influence and political mobilization. *Nature*, 489(7415):295–8.
- boyd, d. (2011). Social Network Sites as Networked Publics: Affordances, Dynamics, and Implications. In Papacharissi, Z., editor, *Networked Self: Identity, Community, and Culture on Social Network Sites*, pages 39–58. Taylor & Francis.
- Brady, W. J., Wills, J. A., Burkart, D., Jost, J. T., and Van Bavel, J. J. (2019). An ideological asymmetry in the diffusion of moralized content on social media among political leaders. *Journal of Experimental Psychology: General*, 148(10):1802–1813.
- Byrne, D. E. (1971). *The Attraction Paradigm*. Academic Press.
- Conover, M., Ratkiewicz, J., Francisco, M., Goncalves, B., Menczer, F., and Flammini, A. (2011). Political Polarization on Twitter. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 5, pages 89–96.
- Davis, C. A., Varol, O., Ferrara, E., Flammini, A., and Menczer, F. (2016). Botornot: A system to evaluate social bots. In *Proceedings of the 25th International Conference Companion on World Wide Web, WWW ’16 Companion*, page 273–274, Republic

- and Canton of Geneva, CHE. International World Wide Web Conferences Steering Committee.
- Diaz, F., Gamon, M., Hofman, J. M., Kıcıman, E., and Rothschild, D. (2016). Online and Social Media Data As an Imperfect Continuous Panel Survey. *PLOS ONE*, 11(1):e0145406.
- DiGrazia, J., McKelvey, K., Bollen, J., and Rojas, F. (2013). More Tweets, More Votes: Social Media as a Quantitative Indicator of Political Behavior. *PLOS ONE*, 8(11):e79449.
- Doherty, C., Kiley, J., Asheer, N., and Jordan, C. (2020). Election 2020: Voters Are Highly Engaged, but Nearly Half Expect To Have Difficulties Voting. Technical report, Pew Research Center.
- Dong, X. and Lian, Y. (2021). A review of social media-based public opinion analyses: Challenges and recommendations. *Technology in Society*, 67:101724.
- Ferrara, E., Varol, O., Davis, C., Menczer, F., and Flammini, A. (2016). The rise of social bots. *Commun. ACM*, 59(7):96–104.
- Gelman, A. (2021). Failure and Success in Political Polling and Election Forecasting. *Statistics and Public Policy*, 8(1):67–72.
- González-Bailón, S., Lazer, D., Barberá, P., Zhang, M., Allcott, H., Brown, T., Crespo-Tenorio, A., Freelon, D., Gentzkow, M., Guess, A. M., Iyengar, S., Kim, Y. M., Malhotra, N., Moehler, D., Nyhan, B., Pan, J., Rivera, C. V., Settle, J., Thorson, E., Tromble, R., Wilkins, A., Wojcieszak, M., de Jonge, C. K., Franco, A., Mason, W., Stroud, N. J., and Tucker, J. A. (2023). Asymmetric ideological segregation in exposure to political news on Facebook. *Science*, 381(6656):392–398.
- Gottfried, J. (2021). Republicans less likely to trust their main news source if they see it as ‘mainstream’; democrats more likely. *Pew Research Center*.
- Grabowicz, P., Kulshrestha, J., Babaei, M., and Weber, I. (2016). The Road to Popularity: The Dilution of Growing Audience on Twitter. In *Proceedings of the 10th International Conference on Web and Social Media, ICWSM 2016*.

- Grabowicz, P. A., Romero-Ferrero, F., Lins, T., Benevenuto, F., Gummadi, K. P., and de Polavieja, G. G. (2020). Bayesian social influence in the online realm.
- Guess, A., Aslett, K., Tucker, J., Bonneau, R., and Nagler, J. (2021). Cracking Open the News Feed: Exploring What U.S. Facebook Users See and Share with Large-Scale Platform Data. *Journal of Quantitative Description: Digital Media*, 1.
- Habermas, J. (1989). *The Structural Transformation of the Public Sphere*. The MIT Press, Cambridge, MA.
- Herbst, S. (1998). *Reading Public Opinion: How Political Actors View the Democratic Process*. The University of Chicago Press, Chicago, IL.
- Hickey, D., Schmitz, M., Fessler, D., Smaldino, P., Muric, G., and Burghardt, K. (2023). Auditing elon musk’s impact on hate speech and bots.
- Hughes, A. (2021). A small group of prolific users account for a majority of political tweets sent by U.S. adults. Technical report, Pew Research Center.
- Hughes, A. and Asheer, Ni. (2019). National Politics on Twitter: Small Share of U.S. Adults Produce Majority of Tweets. Technical report, Pew Research Center.
- Huszár, F., Ktena, S. I., O’Brien, C., Belli, L., Schlaikjer, A., and Hardt, M. (2022). Algorithmic amplification of politics on Twitter. *Proceedings of the National Academy of Sciences of the United States of America*, 119(1):1–6.
- Jungherr, A. (2015). *Analyzing Political Communication with Digital Trace Data*. Springer, Cham.
- Juul, J. L. and Ugander, J. (2021). Comparing information diffusion mechanisms by matching on cascade size. *Proceedings of the National Academy of Sciences*, 118(46):e2100786118.
- Keller, F. B., Schoch, D., Stier, S., and Yang, J. (2020a). Political astroturfing on twitter: How to coordinate a disinformation campaign. *Political Communication*, 37(2):256–280.

- Keller, F. B., Schoch, D., Stier, S., and Yang, J. H. (2020b). Political Astroturfing on Twitter: How to Coordinate a Disinformation Campaign. *Political Communication*, 37(2):256–280.
- Klašnja, M., Barberá, P., Beauchamp, N., Nagler, J., and Tucker, J. A. (2018). Measuring Public Opinion with Social Media Data. In Atkeson, L. R. and Alvarez, R. M., editors, *The Oxford Handbook of Polling and Survey Methods*. Oxford University Press.
- Kramer, A. D. I., Guillory, J. E., and Hancock, J. T. (2014). Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences*, 111(24):8788–8790.
- Li, M., Sun, L., Wang, Y., Sun, Y., Kwon, H., Suk, J., Yang, J., and Shah, D. V. (2023). Computational Approaches to Online Political Expression: A Framework for Research. In Coleman, S. and Sorenson, L., editors, *Handbook of Digital Politics*. Edward Elgar Publishing, Inc., Chltenham, UK, 2nd edition.
- Liedke, J. and Gottfried, J. (2022). Us adults under 30 now trust information from social media almost as much as from national news outlets. *Pew Research Center*.
- Mason Walker and Katerina Eva Matsa (2021). News Consumption Across Social Media in 2021. Technical report, Pew Research Center.
- McGregor, S. C. (2019). Social media as public opinion: How journalists use social media to represent public opinion. *Journalism*, 20(8):1070–1086.
- McPherson, M., Smith-Lovin, L., and Cook, J. M. (2001). Birds of a Feather: Homophily in Social Networks. *Annual Review of Sociology*, 27(1):415–444.
- Mehta, C. (2022). Elon Musk’s Twitter polls - from stake sales to CEO decisions.
- Mellon, J. and Prosser, C. (2017). Twitter and facebook are not representative of the general population: Political attitudes and demographics of british social media users. *Research & Politics*, 4(3):2053168017720008.

- Mislove, A., Lehmann, S., Ahn, Y.-Y., Onnela, J.-P., and Rosenquist, J. N. (2011). Understanding the Demographics of Twitter Users. *Int'l AAAI Conference on Weblogs and Social Media (ICWSM)*, pages 554–557.
- Muchnik, L., Aral, S., and Taylor, S. J. (2013). Social influence bias: A randomized experiment. *Science (New York, N.Y.)*, 341(6146):647–51.
- Nikolov, D., Flammini, A., and Menczer, F. (2021). Right and left, partisanship predicts (asymmetric) vulnerability to misinformation. *Harvard Kennedy School Misinformation Review*, 1(7):1–13.
- O'Connor, B., Balasubramanyan, R., Routledge, B., and Smith, N. (2010). From Tweets to Polls: Linking Text Sentiment to Public Opinion Time Series. In *AAAI Conference on Web and Social Media (ICWSM)*.
- Pennycook, G. and Rand, D. G. (2021). The Psychology of Fake News. *Trends in Cognitive Sciences*, 25(5):388–402.
- Perrin, A. (2020). 23% of users in us say social media led them to change views on an issue; some cite black lives matter. *Pew Research Center*.
- Pool, N. E. (2016). National Election Pool Poll: 2016 National Election Day Exit Poll, 2016 [Dataset]. Roper #31116396, Version 3. Edison Research [producer]. Cornell University, Ithaca, NY: Roper Center for Public Opinion Research [distributor]. Access Date: Oct-05-2023.
- Pool, N. E. (2020). National Election Pool Poll: 2020 National Election Day Exit Poll, 2020 [Dataset]. Roper #31119913, Version 4. Edison Research [producer]. Cornell University, Ithaca, NY: Roper Center for Public Opinion Research [distributor]. Access Date: Oct-05-2023.
- Price, V. E. (1992). *Public opinion*. Sage, Newbury Park.
- Rojas, H. and Puig-i-Abril, E. (2009). Mobilizers Mobilized: Information, Expression, Mobilization and Participation in the Digital Age. *Journal of Computer-Mediated Communication*, 14(4):902–927.

- Ronzhyn, A., Cardenal, A. S., and Batlle Rubio, A. (2022). Defining affordances in social media research: A literature review. *New Media & Society*, 25(11):3165–3188.
- Schober, M. F., Pasek, J., Guggenheim, L., and et al. (2016). Social media analyses for social measurement. *Public Opinion Quarterly*, 80(1):180–211.
- Schoch, D., Keller, F. B., Stier, S., and Yang, J. (2022). Coordination patterns reveal online political astroturfing across the world. *Scientific Reports*, 12(1):4572.
- Shirani-Mehr, H., Rothschild, D., Goel, S., and Gelman, A. (2018). Disentangling Bias and Variance in Election Polls. *Journal of the American Statistical Association*, 113(522):607–614.
- Silver, N. (2021). The Death of Retail is Greatly Exaggerated. *FiveThirtyEight*.
- Squire, P. (1988). Why the 1936 literary digest poll failed. *Public Opinion Quarterly*, 52(1):125–133.
- Strack, F. (1992). “Order Effects” in Survey Research: Activation and Information Functions of Preceding Questions, pages 23–34. Springer New York, New York, NY.
- Treem, J. W. and Leonardi, P. M. (2013). Social Media Use in Organizations: Exploring the Affordances of Visibility, Editability, Persistence, and Association. *Annals of the International Communication Association*, 36(1):143–189.
- Ugander, J., Karrer, B., Backstrom, L., and Marlow, C. (2011). The Anatomy of the Facebook Social Graph. *Arxiv preprint arXiv:1111.4503*.
- Vosoughi, S., Roy, D., and Aral, S. (2018). The spread of true and false news online. *Science*, 359(6380):1146–1151.
- Wang, W., Rothschild, D., Goel, S., and Gelman, A. (2015). Forecasting elections with non-representative polls. *International Journal of Forecasting*, 31(3):980–991.
- Wang, Z., Hale, S., Adelani, D., Grabowicz, P., Hartman, T., Flock, F., and Jurgens, D. (2019). Demographic inference and representative population estimates from multilingual social media data. In *WWW ’19: The World Wide Web Conference*, pages 2056–2067. ACM.

- Wells, C., Shah, D. V., Pevehouse, J. C., Yang, J., Pelled, A., Boehm, F., Lukito, J., Ghosh, S., and Schmidt, J. L. (2016). How Trump Drove Coverage to the Nomination: Hybrid Media Campaigning. *Political Communication*, pages 1–8.
- Wojcik, S. and Hughes, A. (2019). Sizing up twitter users. *PEW research center*, 24:1–23.
- Wu, L., Morstatter, F., Carley, K. M., and Liu, H. (2019). Misinformation in Social Media: Definition, Manipulation, and Detection. *ACM SIGKDD Explorations Newsletter*, 21(2):80–90.
- Zafar, M. B., Gummadi, K. P., and Danescu-Niculescu-Mizil, C. (2016). Message impartiality in social media discussions. *Proceedings of the 10th International Conference on Web and Social Media, ICWSM 2016*, pages 466–475.
- Zhang, Y., Chen, F., and Rohe, K. (2022a). Social Media Public Opinion as Flocks in a Murmuration: Conceptualizing and Measuring Opinion Expression on Social Media. *Journal of Computer-Mediated Communication*, 27(1).
- Zhang, Y., Shah, D., Pevehouse, J., and Valenzuela, S. (2022b). Reactive and Asymmetric Communication Flows: Social Media Discourse and Partisan News Framing in the Wake of Mass Shootings. *The International Journal of Press/Politics*.