How Accurate Are Survey Responses on Social Media and Politics?*

Andrew Guess[†]
Princeton University

Kevin Munger[‡]
New York University

Jonathan Nagler[§]
New York University

Joshua Tucker[¶]
New York University

August 7, 2018

Abstract

How accurate are survey-based measures of social media use, in particular about political topics? We answer this question by linking original survey data collected during the U.S. 2016 election campaign with respondents' observed social media activity. We use supervised machine learning to classify whether this Twitter and Facebook account data is content related to politics. We then benchmark our survey measures on frequency of posting about politics and the number of political figures followed. We find that, on average, our self-reported survey measures tend to correlate with observed social media activity. At the same time, we also find a worrying amount of individual-level discrepancy and problems related to extreme outliers. Our recommendations are twofold. First, for survey questions about social media use to provide respondents with options covering a wider range of activity, especially in the long tail. Second, for survey questions to include specific content and anchors defining what it means for a post to be "about politics."

*This study has been approved by the New York University Institutional Review Board (IRB-12-9058, IRB-FY2017-150). We are grateful to Samantha Luks at YouGov for facilitating survey data collection, Yvan Scher for research assistance, and Jona Ronen for development and technical assistance. Thanks to Rachel Gibson, workshop participants at the University of Manchester, the editors of this special issue, and to three anonymous reviewers for their helpful comments and feedback.

[†]Assistant Professor, Department of Politics and Woodrow Wilson School, Princeton University. Email: `aguess@princeton.edu`

[‡]Wilf Family Department of Politics and SMaPP Lab, New York University.

[§]Professor, Wilf Family Department of Politics and SMaPP Lab, New York University.

[¶]Professor, Wilf Family Department of Politics and SMaPP Lab, New York University.

1 Introduction

In 2016, the Pew Research Center surveyed Americans about their social media use during the election campaign. One question asked: "In the past week, did you yourself use a social networking site to share news or information about the presidential campaign or candidates, such as by posting or replying to or commenting on a post?"[1] This is a vital question for understanding the changing ways in which people consume and interact with the news about political developments (e.g., Mutz and Young 2011; Duggan 2015; Duggan and Smith 2016; Gottfried and Shearer 2016). Yet, so far, we know very little about how accurately people's responses to such a question reflect their observed behavior on social media.

This paper provides the first systematic attempt to benchmark such questions, focusing on a set of items developed to study Americans' social media use during the 2016 U.S. presidential election campaign. It is important to know how well such questions perform when compared against a known "ground truth," as well as whether there are systematic biases that could potentially be corrected. In the future, this will aid both studies of social media use as well as research on the effects of social media on politics (e.g., Bartels 1993; Bohdanova 2014). At the same time, we do not claim that the current paper is the final word on the subject; there are thorny measurement issues involved, and studies of social media are inevitably aiming at an ever-shifting target (Karpf 2012).

We combine survey data from a panel fielded over several waves in 2016 with data on respondents' Twitter and Facebook posts, classified as political or not using supervised learning techniques to generate measures of real-world posting behavior: both in general and on political topics. We find that, on average, there is a correspondence between the average frequency of tweets and Facebook posts and the self-reporting by our respondents. However, individual-level discrepancies are common, and there remain persistent issues related to outliers and top-coding of response categories. In particular, our model of individual-level discrepancies of Twitter behavior performs poorly. This is due to the mismatch between

[1] http://www.journalism.org/2016/02/04/the-2016-presidential-campaign-a-news-event-thats-hard-to-miss/

the non-normal distribution of true rates of tweeting and the response categories in the traditional survey questions about media use. A crucial observation for future research is the need to develop survey measures designed to capture the activity of both normal and power users.

The next section briefly covers some related literature. We then provide an overview of our data sources, survey questions, and coding procedures. In Section 4, we present our results on tweeting and posting frequency. Section 4.4 discusses results on our survey measures of political follow networks. We conclude with a discussion of issues with the kinds of survey questions studied here.

2 Measuring Media Exposure and Social Networks

Social scientists have long been interested in the effects of exposure to information from the mass media and interpersonal interactions (Lazarsfeld, Berelson and Gaudet 1944; Iyengar and Kinder 2010; Mutz and Martin 2001). However, scholars have consistently found that their favorite tool in this context — self-reported survey data — is flawed, as individuals often misreport their true exposure to these information sources. This insight has taken creative research design, as objective measures of exposure to traditional mass media such as television and radio are rarely linked to individual-level survey data (Prior 2009, 2012; Scharkow 2016; Jerit et al. 2016).

The rapidly evolving and increasingly fragmented information environment adds urgency to the important task of benchmarking measures before widespread deployment (Niederdeppe 2014; Guess 2015; de Vreese and Neijens 2016; Taneja 2016). Simply put, can we trust responses to questions about what people do on social media sites? This paper provides the first comprehensive benchmarking test of self-reported social media exposure measures. By comparing self-reports of activity on Twitter and Facebook to observed activity, we are able to measure the accuracy of the self-reports that researchers have previously had to take at

face value.

There is a long history of asking for self-reports about political activity, everything from "how often do you talk to a neighbor about politics" to whether or not the respondent voted or gave money to a candidate (Prior 2009; Anderson and Silver 1986; Schwarz 1999). Responses to a large set of behavior questions cannot be verified, and thus scholars have had to rely on self-reports for both simply tabulating the behavior and for trying to understand the correlates of the reported behavior (Romantan et al. 2008). However, social media data, by nature of its "digital footprints," gives us the opportunity to observe political behavior.[2] That means we can, in principle, see if self-reports of activity are valid. But, since observation of the actual behavior can be difficult, we want to know if the self-report corresponds to observed behavior, and if we can rely on surveys when direct observation of behavior is not possible.

The literature on the accuracy of behavioral self-reports suggests a core set of expectations (e.g., Stone et al. 1999). Questions about sensitive behaviors are known to produce systematic underreporting, but it is unlikely that most respondents consider their social media activity — by definition public or semi-public in nature — to fall into this category. As a result, we are more concerned about patterns of overreporting. There are a number of possible mechanisms for systematic inflation of self-reported behavioral tendencies, including imperfect recall (Revilla, Ochoa and Loewe 2017) or inattentive responding (McKibben and Silvia 2017), faulty calibration to population frequencies (Prior 2009), and social desirability (Krumpal 2013). Some of these errors may cancel out in the aggregate, while others (such as social desirability) lead in a predicted direction. This implies that survey measures could be substantially correlated with ground truth even in the presence of substantial biases at the individual level. Perhaps for this reason, scholars have sought to identify demographic or political correlates of over- or underreporting in surveys. Some hypothesized factors, such as

[2]At the same time, the issue of what defines "political" behavior in the context of Facebook is just beginning to receive scholarly attention. Below, we discuss the issue of subjectivity of what is and is not political behavior.

the role of political interest and partisan strength in driving earnest or expressive respondent enthusiasm, are more plausible than others (Prior 2013); predictors such as age or gender may be associated with overreporting for some questions in certain contexts, but it is not clear whether effects of this kind should be expected to generalize.

Taken together, this discussion implies that the misreporting we uncover should tend to inflate rather than deflate social media use frequency, and that this tendency to overreport may be associated with measures of political interest and partisan allegiance (Prior 2009). At the same time, we have some expectation that our respondents' self-selection into social media use in the first place may be related some of the same characteristics. Thus we are left with a puzzle: Individuals who we would expect to be overrepresented in social media discussions about politics may be the very types of people who are predisposed to overreport that activity, due to partisan misperceptions, flawed recall, or other cognitive biases. It is unclear whether this might cause more extreme errors in self-reporting or, instead, whether social media use could be a case in which some subjects' observed behaviors happen to match their generally inflated response patterns.

We also face a classic measurement issue in benchmarking posts about politics. We can only truly determine if respondents' reports of behavior of a specific type are accurate if the behavior we are asking about can be objectively measured. If we ask people whether or not they voted, there is a true answer. But when we ask if they posted "about politics" there is an inherent subjectivity to the question. Even if all respondents had perfect recall and desired to be as accurate as possible, their varying conceptions the definition of a "political" post would present a challenge to any benchmarking effort. To use the example developed in detail in Settle (2018), consider a Facebook post in which someone says that they are having lunch at Chik-Fil-A. This appears to be a non-political post, but during a period around 2012, there was considerable controversy about Chik-Fil-A's funding of anti-LGBT organizations. Some liberals announced a boycott of Chik-Fil-A, and some conservatives began a counter-boycott with associated posts on their Facebook feeds. Many other people

continued to eat (or not) at Chik-Fil-A and post the occasional picture of a fried chicken sandwich.

If our coders do not interpret posts about where someone ate lunch as political, then the supervised machine learning model we use will not be likely to classify posts with a picture of a chicken sandwich at Chik-Fil-A as political. But a conservative activist who answers our survey question may think that their posting the sandwich is political, and we would thus conclude that they have overreported their true rate of posting about politics. On the other hand, if our coders are attuned to the politics surrounding Chik-Fil-A, they may think the posting of a chicken sandwich is political. This would lead us to be likely to classify such postings as political, and we would conclude that respondents who do not view postings of their lunch decisions as political were underreporting their rate of posting about politics.[3]

We note that the high degree of inter-coder reliability we achieved suggests that there is in fact generally a shared sense of what it means for a post to be "about politics." As a result, for the survey questions we study, we refrain from claiming that gaps between our baseline measures and self-reports of posting political content represent intentional distortions. But they do represent the discrepancy between self-reports and what our coders believe to be posts about politics. We note that this particular issue could be resolved by asking respondents a more specific question that gives them guidance about what it means for a post to be "about politics": we could, for instance, spell out that we mean a post mentioning a campaign, a candidate, or a public policy issue.

[3]The contours of this subjectivity have only begun to be explored by political scientists. Settle (2018) asks a sample of undergraduates and a sample of MTurk coders to label Facebook posts as "political" or not. She reports that "Users who say they more frequently encounter political content on their News Feeds do in fact have a broader conception of content that could be considered to be political" (p 132). Settle's argument is that the more frequently people use Facebook, the more likely they are to identify posts as being political. We do not try to test this claim directly, but our data in Appendix C are consistent with it.

3 Data Overview

3.1 Panel Survey

The self-reports against which we benchmark objective measures comes from a panel survey conducted during the 2016 U.S. presidential election.[4] These surveys were designed explicitly to understand the relationship between social media use and changes in political beliefs and preferences; this was accomplished by pairing a panel survey design with information about respondents' social media accounts.

The U.S. presidential election survey took place over the course of three waves: the first was April 9–May 1, 2016 (3,500 respondents); the second was September 9–October 9 (2,635 respondents); and the third was October 25–November 7 (2,628 respondents). The first wave of the survey was estimated by YouGov to take 15 minutes to complete, and contains a wealth of demographic, social media use and political preference questions. This paper focuses on questions in which respondents self-report several aspects of their use of social media.

In wave 1, respondents were asked to select all of the social networking sites on which they had accounts (options provided: Twitter, Facebook, Instagram, LinkedIn, Snapchat, Other); if they indicated that they had a Facebook and/or Twitter account, we asked them about their frequency of posting on that platform both in general and specifically about politics. Response options ranged from "Never" to "Several times a day." We also asked respondents to describe the number and types of people they were connected to on the accounts they have — both in general (for Twitter only) and "elected officials, candidates for office, or other political figures."[5] Tables 1 and 2 show the raw responses in each category for these questions (for those who reported using the relevant social media platform). The modal response for the number of politicians followed on both Facebook and Twitter is zero; likewise, the modal respondent who uses either social media platform reports never posting

[4]The survey was designed by the authors and conducted by the polling firm YouGov.

[5]See Appendix D for full question wordings and response options.

about politics.

Table 1: Self-Reported Social Media Posting Behavior (Counts and Proportions)

	Twitter	Twitter (politics)	Facebook	Facebook (politics)
Never	219 (0.09)	992 (0.44)	102 (0.04)	874 (0.33)
Less often	673 (0.29)	473 (0.21)	329 (0.12)	562 (0.21)
Every few weeks	361 (0.15)	257 (0.11)	362 (0.13)	373 (0.14)
1 to 2 days a week	325 (0.14)	178 (0.08)	391 (0.14)	245 (0.09)
3 to 6 days a week	195 (0.08)	95 (0.04)	356 (0.13)	152 (0.06)
About once a day	257 (0.11)	124 (0.05)	453 (0.17)	172 (0.07)
Several times a day	304 (0.13)	160 (0.07)	711 (0.26)	256 (0.10)

Table 2: Self-Reported Social Media Following Behavior (Counts and Proportions)

	# politicians followed: Twitter	# politicians followed: Facebook
0	1047 (0.47)	1245 (0.49)
1-10	892 (0.40)	1067 (0.42)
More than 10	269 (0.12)	254 (0.10)

The sampling frame included members of the YouGov panel who had previously agreed to supply YouGov with their Twitter IDs. There were 2,163 respondents in the first wave (62% of the overall sample) who elected to enter some text when prompted for their Twitter ID. We were able to verify 1,816 of these with the Twitter API.[6] We scraped the tweets by these respondents on November 20, 2016.[7] We also retrieved the IDs of the accounts followed by these respondents. Ultimately, we were able to collect a non-zero number of tweets from 1,421 respondents. This raises understandable concerns about sample selection. For a full discussion of the attrition process in our data, including descriptive statistics suggesting that it did not seriously bias the final sample, see Appendix B. These results are similar to the finding by Munger et al. (2016) that a sample of U.K. Twitter users who shared their account information with YouGov is representative of U.K. Twitter users generally.

[6]Some people simply provided something other than a valid Twitter ID to YouGov, and it is possible that some of these accounts were deleted by their owners or banned/suspended by Twitter.

[7]Because we only accessed each respondents' tweets at one point, the number of tweets per subject in our dataset is capped at 3,200, the maximum allowed by the Twitter API.

We also collected Facebook profile data from our respondents. Using a web application that we developed, we asked respondents if they would be willing to supply information about their own past Facebook activity and told them that we could get this information directly from their Facebook accounts if they agreed. This was done via a separately administered survey question post-election that linked respondents to the web app connecting to the Facebook API. 1,221 of our respondents agreed to let us retrieve their Facebook information in return for compensation in the form of $5 in YouGov "points." Specifically, we requested their public profile information, Timeline posts (including text and links if available), page likes, and what Facebook saves as religious and political views. If a respondent chose to log into Facebook after the survey prompt, they were asked what specific pieces of information they were willing to share. They could approve sharing all of the types of information, selectively approve only some of these types of information, or approve nothing.[8]

Despite the seeming high level of attrition for Facebook users, we highlight that 45% of the 2,711 respondents who said they have a Facebook account chose to share data with us — all the more reassuring given that unlike Twitter, Facebook data is not public. Appendix B also provides information about this form of attrition. These populations are indeed somewhat different: the respondents who shared their information were more likely to be Democrats, were more educated, and were heavier Facebook users. Of course, individuals who are politically interested and more habitual social media users may have been more likely to select into our sample. This suggests caution is warranted in interpreting our results, if, for example, people who tend to "lurk" without posting (e.g., Steinmetz 2012) are also less likely to share account data with researchers. The quantity we aim to measure—agreement between subjective and objective measures of social media behavior—is latent, and it is certainly possible that this quantity covaries with selection into sharing account

[8]No data on News Feed content or exposure was shared with researchers, since such access is no longer allowed via the API (for a rare example of data collected with the earlier permissions, see Wells and Thorson 2017). Data access was temporary and lasted only 2 months after permission was granted. All respondents who agreed to share information consented to a privacy policy that specified, in part, "This application will not access the profile information of any friends, groups, or other information associated with your profile page."

information. This problem is a serious one, and not easily solved given the ethical necessity of acquiring consent before accessing private social media data. We can, however, get an idea of the scope of the problem by examining distributions of observable variables in the different samples, in Appendix B.

3.2 Coding Social Media Networks

To create an objective measure about the networks of our respondents, we used friend/follower data. The first question to benchmark — how many people do you follow on Twitter? — was straightforward to measure: Twitter API data always contains this information.[9] To create a measure of how many *politicians* respondents follow on Twitter/Facebook, we combined existing public lists of official accounts. Appendix G describes this procedure in more detail.

3.3 Coding Social Media Posts

The simplest measures of social media activity are how often respondents post to the respective platforms. However, we are not only interested in how often they post, but also in how often they post about politics. To examine the calibration of self-reported posts "about politics," we trained a machine learning model using a combination of supervised and semi-supervised methods.

We go into significant detail in Appendix A about our iterative method of generating labeled training data and validating separate classifiers for Facebook and Twitter posts out of sample. To summarize briefly, we produced a set of Facebook posts and tweets hand-coded as being "about politics" or not, and predicted this attribute for the rest of the data using a Naive Bayes classifier trained on the labeled benchmark data. Ultimately, we achieve a 50-fold cross-validated accuracy of 89% for a human-coded subset of the Facebook posts and an 83% accuracy for tweets. We estimate that 27% of subjects' tweets were political,

[9] Unfortunately, the Facebook data our app gave us access to does not contain information on the number of friends users have, so we could not create an analogous measure.

compared to only 5% of Facebook posts.

We emphasize that despite an inherently subjective component to assessments of what counts as "political," our coders had fairly high levels of agreement in their categorizations of tweets and Facebook posts. While this suggests some uncertainty about our chosen benchmark, it is nevertheless uncertainty that we can characterize and quantify. Our survey respondents, by contrast, may or may not have approached the survey with an accuracy motivation in mind, and we cannot obtain intercoder reliability statistics on the survey responses themselves. Thus we proceed on the assumption that our labeled training data, and the classifiers they enable, are useful as a benchmark for validating our self-reported survey responses.

4 Results

4.1 Twitter

First, we compare respondents' self-reported tweet frequency to the overall number of tweets they posted. Figure 1 displays the raw individual-level data by each response category, with overlaid box plots. We see that there is a correspondence: respondents are more likely to report a greater frequency of tweeting if they actually tweet more ($r = 0.47$).[10] We reiterate that the individual-level tweet posting data is limited by the Twitter API to approximately 3,200 per person, which can be seen in the figure. This and subsequent figures also display the marginal distribution of responses within each category. Here, the marginals indicate that a plurality of respondents (28.8%) said that they tweet "Less often" (where "less often" means less often than every few weeks).

Figure 2 illustrates a similar relationship between self-reported *political* tweet frequency

[10]Pearson's correlations — here and in subsequent results — are computed by converting self-reported frequency to a numeric variable and treating it as if on an interval or ratio scale. For per-day measures, we translate response categories as follows: "Never" to 0, "Less often" to 1/50, "Every few weeks" to 1/21, "1 to 2 days a week" to 2/7, "3 to 6 days a week" to 6/7, "About once a day" to 1, and "Several times a day" to 3.

and the number of such tweets posted by respondents with linked accounts ($r = 0.45$). In this case, we identify tweets as "political" according to the procedure outlined in Coding Social Media Posts And while the relationship is about equally strong, we note that the plurality of users (43.5%) say they never tweet about politics.

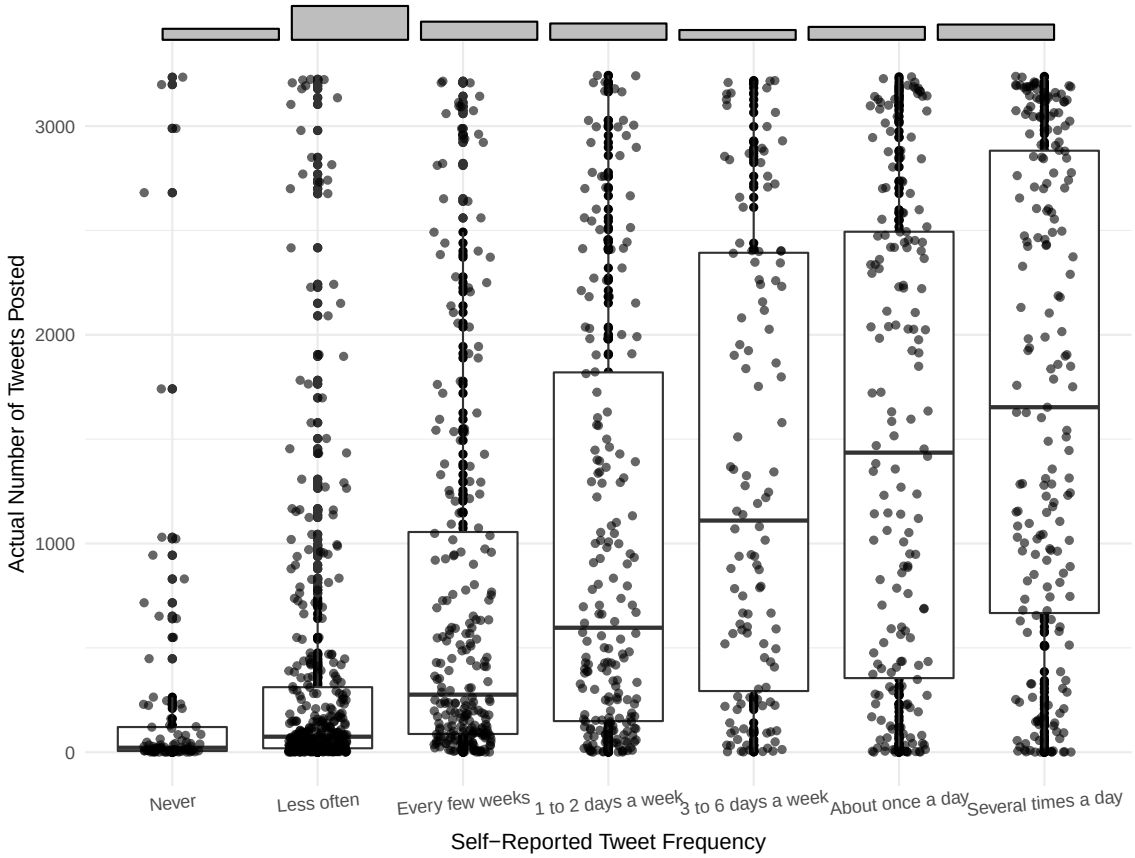

Figure 1: Total number of tweets posted (using linked data from respondents' Twitter accounts) plotted against self-reported tweet frequency.

For tweeting, then, there at least appears to be a rough correspondence between overall post volume and perceived frequency of posting. But are these perceptions well-calibrated to the wording of the survey response options, which specify frequencies such as "Every few weeks" and "About once a day"? To answer this question, we construct measures of tweet frequency by day. This is simply the number of tweets (all or political) posted in the month prior to when we queried the Twitter API for users' posts divided by the number of

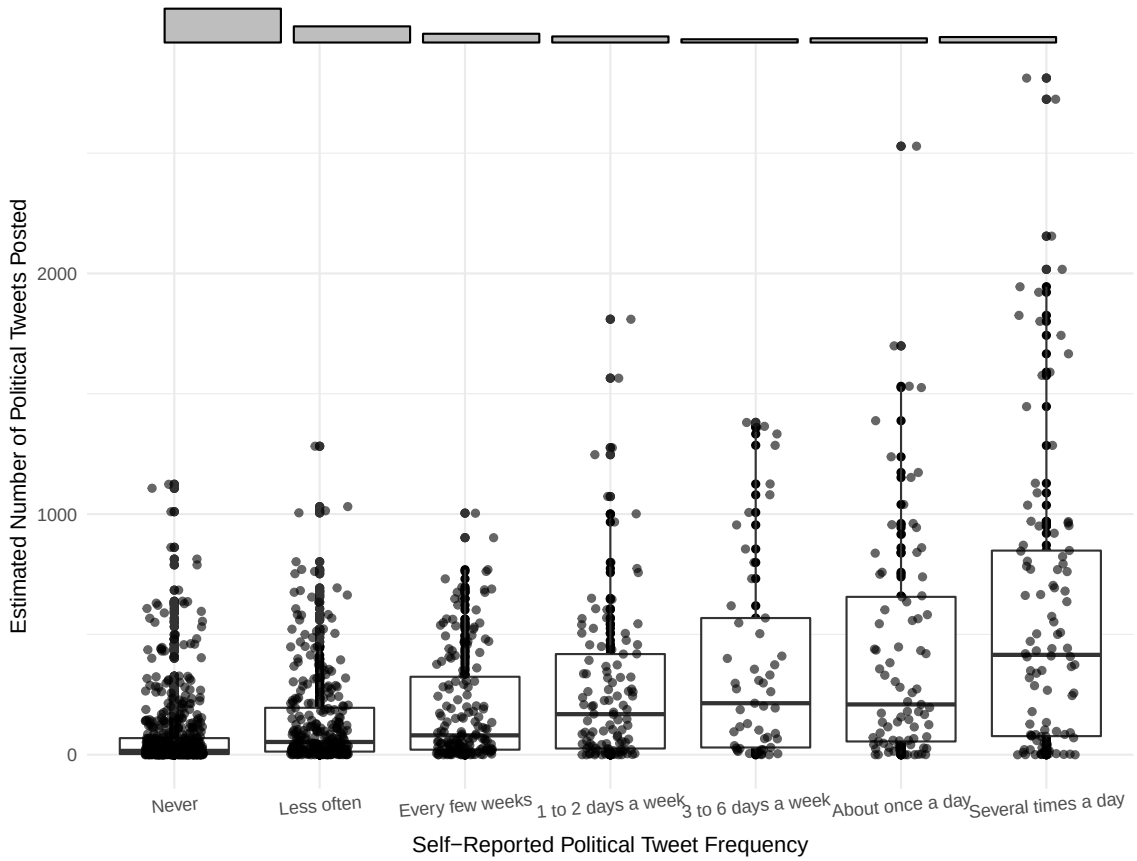

Figure 2: Total number of political tweets posted (using linked data from respondents' Twitter accounts) plotted against self-reported political tweet frequency. Tweets are categorized as "political" via the supervised learning technique discussed in Coding Social Media Posts

days (31).[11] When we plot this daily tweet measure against the same self-reported survey responses (Figure 3), the correspondence is difficult to detect visually but still present, since we are effectively normalizing tweet volume (all: $r = 0.36$; political: $r = 0.43$). One notable feature of both the total and political versions of the graph is that there are a nontrivial number of tail-end outliers, especially in the "Several times a day" category. This likely reflects right censoring in our data: people who tweet 50 to 100 times a day, although a minority, are placed in the same category as people who tweet 2 to 3 times per day.

Since the average difference across survey response categories can be obscured by the

[11] We also created versions of this measure for the period of Sept. 9–Oct. 9, 2016, which coincided with the second wave of survey data collection and was farther away from Election Day. The correlation between the two measures is 0.76 for tweets per day and 0.71 for political tweets per day, giving us confidence in their reliability.

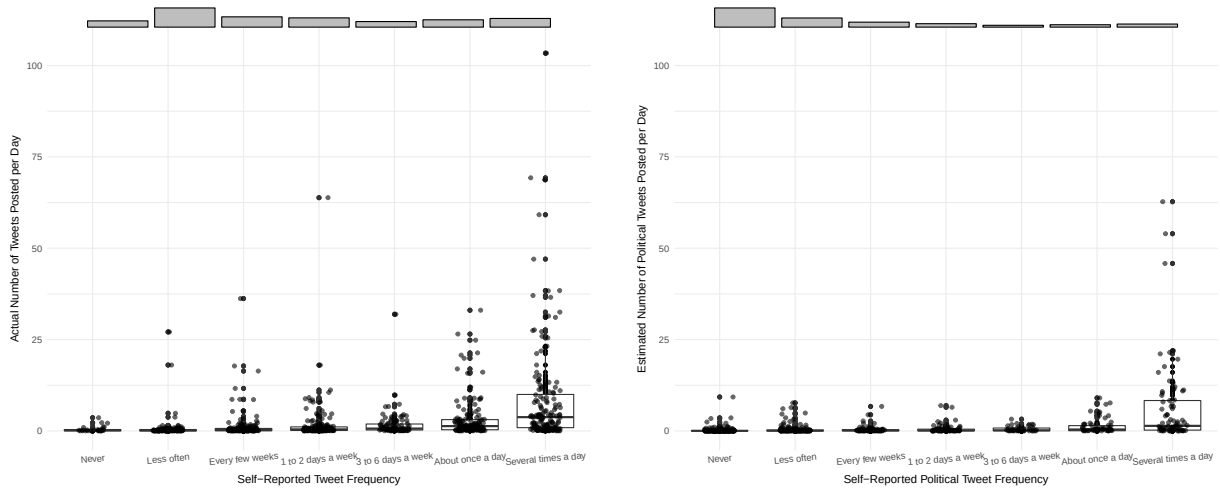

Figure 3: Number of total (left) or political (right) tweets per day posted (using linked data from respondents' Twitter accounts) plotted against self-reported total or political tweet frequency. Tweets are categorized as "political" via the supervised learning technique discussed in Coding Social Media Posts

outliers in the figures, we report the mean number of tweets per day in Table 3 below. Here, for both total tweets and political tweets, we generally see a positive relationship between a given tweet frequency in each survey category and the mean number of tweets posted by respondents. Strikingly, the calibration appears to be almost perfect in the case of political tweets: an average of 1.23 political tweets were posted per day by respondents who said they tweeted "About once a day." The survey categories representing frequencies less than one tweet a day are also associated with an average number of tweets posted per day of less than one. And there were an average of 6.24 tweets per day posted by those who said they tweet "Several times a day." It is important to keep in mind, however, that these calculations reflect specific decisions about how to define the total number of days counted in the denominator. Using only weekdays, for instance, would somewhat inflate these numbers. Of course these means are obscuring substantial variation in individual behavior. Note that when we examine the column for the mean number of all tweets per day we see that the mean number of tweets per day even for people who claim they tweet only "1 to 2 days a week" is 1.76; thus the mean is higher than it should be if all respondents were accurately reporting. However, much of this is driven by outliers who report tweeting only 1 to 2 days

Survey response	Mean Number of (all) Tweets Per Day	Mean Number of Political Tweets Per Day
Never	0.53	0.16
Less often	0.63	0.40
Every few weeks	1.17	0.40
1 to 2 days a week	1.76	0.51
3 to 6 days a week	1.71	0.55
About once a day	3.24	1.23
Several times a day	8.60	6.24

Table 3: Mean number of tweets per day for each survey response category. Tweets are categorized as "political" via the supervised learning technique discussed in Coding Social Media Posts

per week, but are in fact tweeting as often as 10 times per day.[12]

4.2 Facebook

Turning to Facebook, we see a similar pattern: a correspondence between self-reported post frequency and the mean number of total posts as determined by linked profile data. For all types of posts, the correlation is somewhat higher ($r = 0.38$) than for political posts ($r = 0.32$). In terms of general posts, we see in Figure 4 that a plurality of respondents who shared Facebook data (26.3%) said, perhaps unsurprisingly, that they do so "Several times a day." This drops to less than 10% when we look at posts related to politics in Figure 5 (where the most commonly chosen category was "Never," at 33.2%).

Per-day measures of Facebook posting frequency are also correlated with the corresponding survey measures (all: $r = 0.26$; political: $r = 0.20$), but they are heavily mismatched; the mean number of political posts per survey category is off by nearly two orders of magnitude, as can be seen in Figure 6. This may be due to the fact that, lacking API limitations, we used the entire history of users' Facebook posts to construct the measures: We divided the total number of posts by the total number of days active on Facebook.

[12]Respondents may have interpreted the question to mean 'how many times a day' they post things, rather than 'how many items they post' - and it may be that people post 'about once a day', but each time they post 5 items.

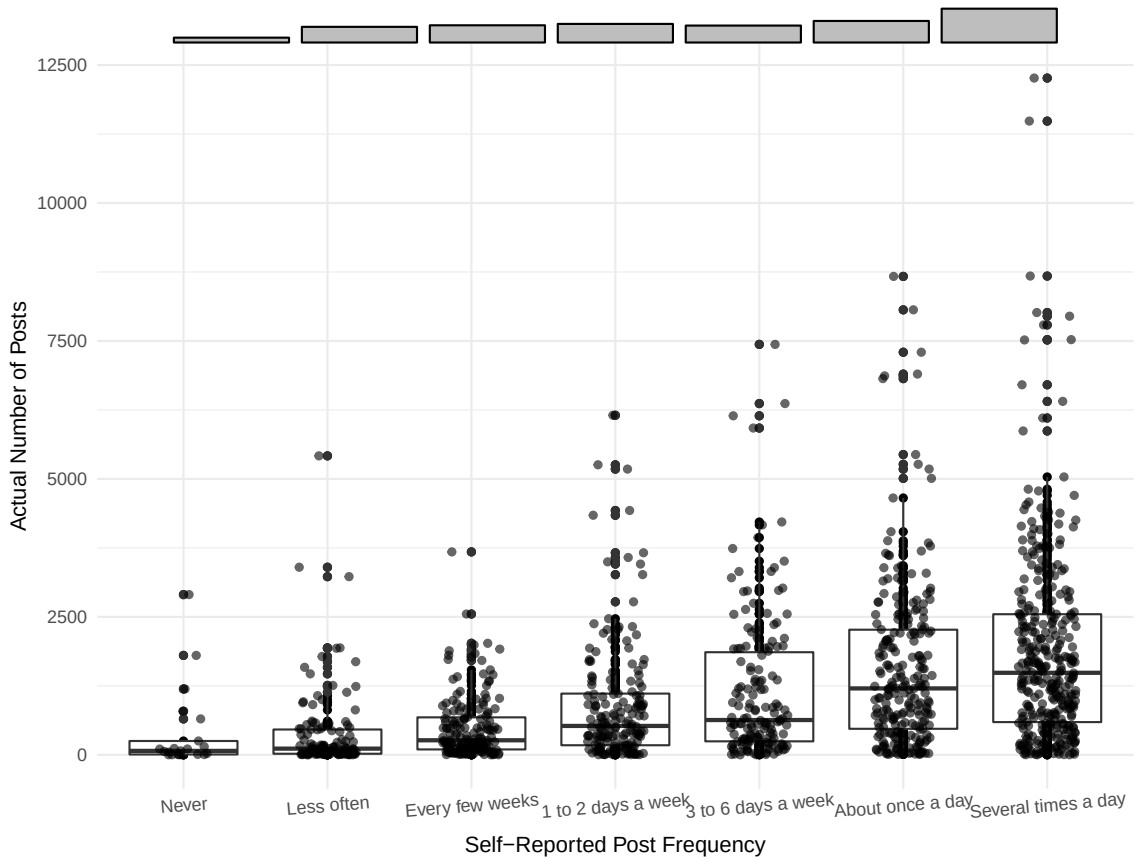

Figure 4: Total number of posts (using linked data from respondents' Facebook profiles) plotted against self-reported post frequency.

4.3 Characterizing Individual-level Discrepancies

While on average there is a correspondence between survey responses and observed social media posting behavior, it is critical to further explore discrepancies on the individual level. Do people tend to over- or underestimate their rates of crafting tweets and Facebook posts? Do certain characteristics systematically predict these discrepancies?

To investigate these questions, we calculate individual-level reporting discrepancies in the following way. First, we convert the survey categories to numerical magnitudes in two different ways.[13] Then we subtract the daily tweet measures from each of those numerically translated survey measures and keep the quantity with the smallest absolute value (given

[13]In the first version, we equate "Never" to 0, "Less often" to 1/50, "Every few weeks" to 1/21, "1 to 2 days a week" to 2/7, "3 to 6 days a week" to 6/7, "About once a day" to 1, and "Several times a day" to 3. In the second version, "Less often" is 1/40, "1 to 2 days a week" is 1/7, and "3 to 6 days a week" is 3/7.

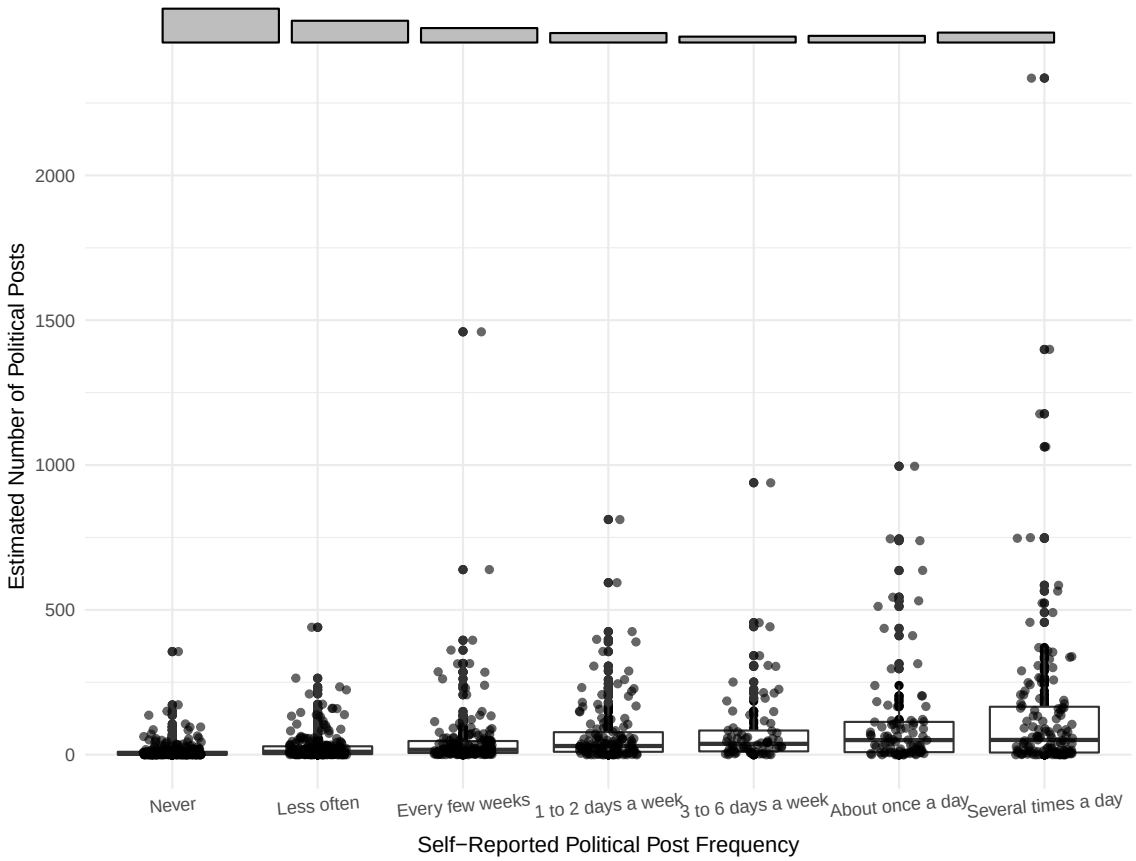

Figure 5: Total number of political posts (using linked data from respondents' Facebook profiles) plotted against self-reported political post frequency. Posts are categorized as "political" via the supervised learning technique discussed in Coding Social Media Posts

the inherent imprecision in survey categories). Note that we are explicitly measuring *over-reporting*. We then run several linear regression models including a number of hypothesized predictors from the literature on self-reported media consumption, and a set of demographic characteristics. Most importantly, we include a measure of interest — whether or not a respondent voted in the 2016 primaries.[14] We also include both party identification and a measure of party strength: whether or not someone identifies as a "Strong Democrat" or "Strong Republican." The results are shown in Table 5.

There is a conceptual difference between the discrepancies in the rates of *overall* posting and posting *about politics*. The former is objective, while the latter is subjective; for some people, the definition of what makes a post "about politics" is more expansive than for others.

[14]The survey did not have a more direct measure of political interest.

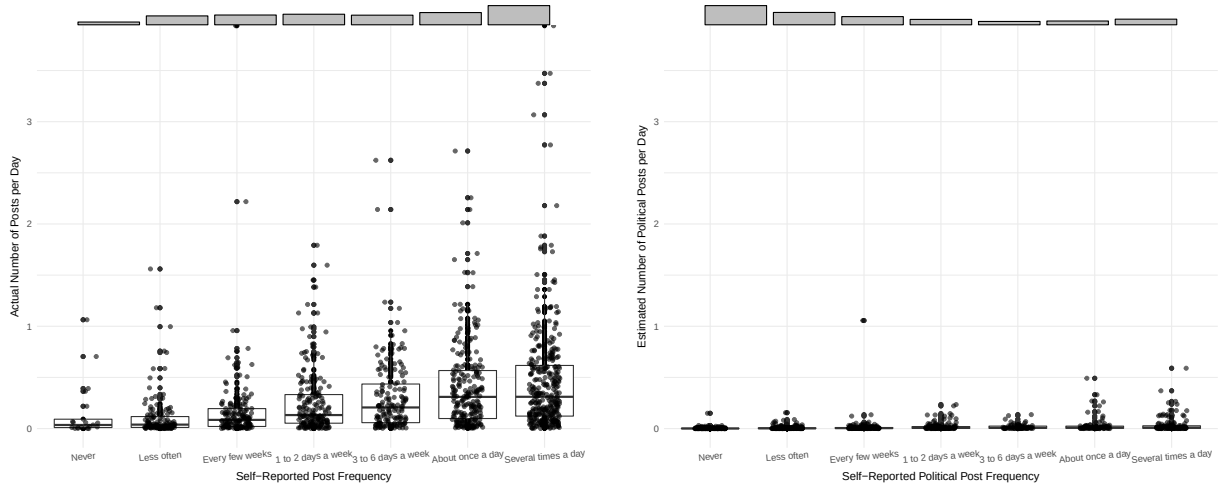

Figure 6: Number of total (left) or political (right) posts per day (using linked data from respondents' Facebook profiles) plotted against self-reported total or political post frequency. Posts are categorized as "political" via the supervised learning technique discussed in Coding Social Media Posts

One could interpret the results in columns 1 and 3 as "errors" while those in columns 2 and 4 could be some mix of errors and of measures of what describes the types of people with these more expansive conceptions of the political.

Patterns in reporting social media activity are quite different on Facebook and Twitter. In the first two columns, it is evident that none of our predictors are related to discrepancies in survey self-reports of overall tweet frequency or of tweeting about politics. The constant in both models is also insignificant, suggesting no general bias toward overreporting or underreporting among our respondents. Indeed, the adjusted R^2 of these models is below 0, indicating that they are quite poor at explaining the variation. A major explanation for this problem is that there is a massive left tail for the dependent variable: our transformed measure has a maximum of 3, a median near 0, and a minimum of approximately -100. The models in the first two columns have had their dependent variables log-transformed (symmetric around zero) to account for this fact.

This is due to the true variability in the way that people use Twitter. With some people posting hundreds of times a day and others posting twice a month, 6-point questions like

Table 4: Predictors of Individual-level Social Media Over-Reporting

	Tweets/day (All)	Tweets/day (Pol)	FB/day (All)	FB/day (Pol)
	(1)	(2)	(3)	(4)
Voted in primary	−0.069	0.036	0.223**	0.196***
	(0.106)	(0.103)	(0.090)	(0.075)
Party: Democrat	−0.066	0.043	0.166	0.229**
	(0.132)	(0.128)	(0.108)	(0.089)
Party: Republican	−0.108	0.009	0.123	0.116
	(0.260)	(0.252)	(0.229)	(0.192)
Party: Independent	−0.153	0.120	0.172	0.249
	(0.212)	(0.214)	(0.210)	(0.174)
Party: Not sure	−0.095	−0.085	−0.046	0.032
	(0.115)	(0.113)	(0.098)	(0.081)
Strong partisan	0.066	0.098	0.221**	0.294***
	(0.114)	(0.111)	(0.093)	(0.077)
Age	−0.003	−0.003	0.009***	0.009***
	(0.003)	(0.003)	(0.002)	(0.002)
High school graduate	0.181	0.147	0.025	0.084
	(0.148)	(0.146)	(0.126)	(0.104)
Some college	−0.035	−0.119	0.138	−0.028
	(0.153)	(0.151)	(0.126)	(0.105)
2-year college	−0.216	−0.087	0.382	0.194
	(0.368)	(0.375)	(0.247)	(0.202)
College graduate	0.054	0.067	−0.207	−0.115
	(0.169)	(0.165)	(0.137)	(0.114)
Postgrad	0.066	0.021	0.092	0.036
	(0.146)	(0.144)	(0.123)	(0.102)
Female	−0.019	−0.040	0.126*	−0.153***
	(0.084)	(0.083)	(0.072)	(0.059)
Nonwhite	0.117	0.084	−0.049	0.080
	(0.102)	(0.100)	(0.084)	(0.070)
Constant	−0.319	−0.606***	−0.050	−0.566***
	(0.225)	(0.221)	(0.193)	(0.160)
N	645	627	1,084	1,063
Adjusted R^2	-0.007	-0.006	0.028	0.048

*p < .1; **p < .05; ***p < .01
OLS models. Reference category for party is "Other."
Columns 1 and 2 take as their DV the log-transformed error, to deal with the wide range of the data.

the ones we employ are simply insufficient. We did not attempt to develop a more Twitter-appropriate survey measure because this insufficiency was not apparent before we conducted this research. We encourage researchers to experiment with novel survey measures to address this problem.

On Facebook, by contrast, a more familiar pattern emerges: Those who voted in the primaries, as well as people who are more partisan, are significantly more likely to overreport their posting frequency. People also seem to systematically underestimate the number of political posts they make to Facebook, as illustrated by the negative constant in the fourth column. Finally, age appears to be positively associated with errors in reporting Facebook posting frequency. The only coefficient that significantly switches signs between columns 3 and 4 is gender: women over-report their rate of posting on Facebook but significantly (at $p < .01$) *under-report* their rate of posting about politics. We do not have a strong theory for why this might be.

Appendix C expands this analysis, using survey responses about the use of Twitter to predict discrepancies in the reported/objective use of Facebook (and vice versa). This analysis is only possible on the subset of subjects who use *both* networks, but the results are informative: these questions are extremely powerful predictors. Also unexpected is the asymmetric effect of self-reported *viewing* the two services. People who report viewing Twitter more often have lower levels of overreporting *all four* behaviors. The effect of viewing Facebook is the opposite: those who report viewing Facebook more often have *higher* levels of overreporting all four behaviors. We discuss possible explanations in Appendix C.

4.4 Friends and Political Sources

We extend our focus beyond tweets and posts to people's social networks in Appendix F. Similar to the previous results, we find a correspondence between the number of political figures people say they follow and the number they actually do, and note that most respondents appear to follow either no or relatively few such figures on social media. The calibration for

Twitter political follow networks seems to be especially accurate on average.

5 Discussion

We take advantage of a unique linked dataset which allows us to validate self-reports of political activity. The good news is that self-reports are correlated with observed behavior. People who follow political actors on Twitter are more likely to say they do than are respondents who do not follow political actors on Twitter. However, at the individual level there are substantial discrepancies in reporting, discrepancies that covary with demographic variables of interest (on Facebook).

Our data also allows us to make comparisons about how people behave across different social media platforms. The most striking difference is that a much higher percentage of tweets in our dataset (20%) were coded as political than were Facebook posts (5%). This result is somewhat surprising; although Twitter users are less representative of the population than Facebook users (in particular they tend to be more highly educated), users of each platform report roughly equivalent rates of seeing (and posting) political content (Duggan and Smith 2016). Pew's most recent survey reports that 8% of Twitter users and 6% of Facebook users claim the highest possible category of political posting — that "a lot" of what they post is related to politics. This is at odds with our finding of very different proportions of content being about politics.[15]

We have demonstrated five important points in this paper. First, self-reports of social media use are meaningful; they are (perhaps surprisingly) accurate and correlated with our objective measure at the aggregate level. Second, there are substantial discrepancies between objective and self-reported posting behavior at the individual level. The nature of these discrepancies is difficult to summarize, and it is not clear whether the predictors we observe, such as age and political interest for Facebook posting, will always be correlated

[15] As mentioned above, adjusting these proportions as suggested by Bachl and Scharkow (2017) gives estimates of 6% political tweets and 2% political Facebook posts. Although the gap in percentage points terms is lower, the percentage gap is similar to the unadjusted gap. See Appendix E for details.

with overreporting.

Third, questions about the number of political actors a given respondent follows or likes appear to be particularly well-calibrated; this is a particularly explicit question, suggesting that people find it easier to answer. Fourth, contrary to what self-reports suggest and given the caveats above, a much higher proportion of posts on Twitter are about politics than are posts on Facebook.

Finally, we have shown that researchers interested in measuring online behavior should consider modifications to standard survey question wordings. In the first place, they should be more liberal in their top-coding of survey responses and allow for the behavior of heavy users. We recommend that an adjustment be made to all survey questions about social media use to reflect the nearly exponential distribution of behavior we observe with our objective measures. There is simply more information on the tails of the distribution than is being recorded by current survey questions.

Additionally, given persistent questions about differing subjective perceptions of what constitutes "political" content, it would be useful to experiment with anchors that explicitly provide respondents with guidance on the kinds of political topics researchers wish to focus on. This could be as simple as "posts about policy proposals and government affairs" or "debates about collective responses to societal problems," and could additionally encompass examples illustrating the scope of such discussions.

Our findings point the way forward for future substantive work in political communication. By helping to validate and refine survey-based measures of social media use for politics, we enable more confident inferences from self-reported data on online behavior — the empirical bedrock of research in this growing area — and point toward individual-level correlates to take into account in order to reduce measurement error. The specific features that we find are associated with misreporting, such as age, partisan strength, and possibly the choice of social platform itself, are moreover of substantive interest themselves. Finally, we hope that our findings on the need for response categories to capture behavior at the long tail of social

posting will enable research into the most influential users on social media, about whom we know relatively little despite their likely disproportionate impact on the contours of political discourse.

The opportunity to compare self-reported behavior with observed behavior is rare (Guess 2015; Scharkow 2016; Revilla, Ochoa and Loewe 2017). Such cases provide us with a reference standard that is extremely valuable in validating the use of self-reports, as well as in providing appropriate cautions.

References

Anderson, Barbara A. and Brian D. Silver. 1986. "Measurement and Mismeasurement of the Validity of the Self-reported Vote." *American Journal of Political Science* 30:771–785.

Bachl, Marko and Michael Scharkow. 2017. "Correcting Measurement Error in Content Analysis." *Communication Methods and Measures* 11(2):87–104.

Bartels, Larry M. 1993. "Messages Received: The Political Impact of Media Exposure." *American Political Science Review* 87(2):267–285.

Bohdanova, Tetyana. 2014. "Unexpected revolution: the role of social media in Ukraine's Euromaidan uprising." *European View* pp. 1–10.

de Vreese, Claes H. and Peter Neijens. 2016. "Measuring Media Exposure in a Changing Communications Environment." *Communication Methods and Measures* 10(2-3):69–80.

Duggan, Maeve. 2015. *The Demographics of Social Media Users*. Pew.

Duggan, Maeve and Aaron Smith. 2016. "The political environment on social media." *Pew Research Center* 25.

Gottfried, Jeffrey and Elisa Shearer. 2016. "News use across social media platforms 2016." *Pew Research Center* .

Guess, Andrew M. 2015. "Measure for Measure: An Experimental Test of Online Political Media Exposure." *Political Analysis* 23(1):59–75.

Iyengar, Shanto and Donald R Kinder. 2010. *News that matters: Television and American opinion*. University of Chicago Press.

Jerit, Jennifer, Jason Barabas, William Pollock, Susan Banducci, Daniel Stevens and Martijn Schoonvelde. 2016. "Manipulated vs. Measured: Using an Experimental Benchmark to Investigate the Performance of Self-Reported Media Exposure." *Communication Methods and Measures* 10(2-3):99–114.
URL: *https://doi.org/10.1080/19312458.2016.1150444*

Karpf, David. 2012. "Social science research methods in Internet time." *Information, Communication & Society* 15(5):639–661.

Krumpal, Ivar. 2013. "Determinants of social desirability bias in sensitive surveys: a literature review." *Quality & Quantity* 47(4):2025–2047.

Lazarsfeld, Paul F., Bernard Berelson and Hazel Gaudet. 1944. *The People's Choice: How the Voter Makes Up His Mind in a Presidential Campaign*. New York: Columbia U. Press.

McKibben, William Bradley and Paul J Silvia. 2017. "Evaluating the distorting effects of inattentive responding and social desirability on self-report scales in creativity and the arts." *The Journal of Creative Behavior* 51(1):57–69.

Munger, Kevin, Patrick Egan, Jonathan Nagler, Jonathan Ronen and Joshua A Tucker. 2016. "Learning (and Unlearning) from the Media and Political Parties: Evidence from the 2015 UK Election." *Working Paper* .

Mutz, D.C. and L. Young. 2011. "Communication and Public Opinion." *Public Opinion Quarterly* 75(5):1018–1044.

Mutz, Diana C and Paul S Martin. 2001. "Facilitating Communication Across Lines of Political Difference: The Role of Mass Media." *American Political Science Review* 95(1):97.

Niederdeppe, Jeff. 2014. "Conceptual, empirical, and practical issues in developing valid measures of public communication campaign exposure." *Communication Methods and Measures* 8(2):138–161.

Prior, Markus. 2009. "Improving media effects research through better measurement of news exposure." *The Journal of Politics* 71(03):893–908.

Prior, Markus. 2012. "Who watches presidential debates? Measurement problems in campaign effects research." *Public Opinion Quarterly* 76(2):350–363.

Prior, Markus. 2013. "Media and political polarization." *Annual Review of Political Science* 16:101–127.

Revilla, Melanie, Carlos Ochoa and Germn Loewe. 2017. "Using Passive Data From a Meter to Complement Survey Data in Order to Study Online Behavior." *Social Science Computer Review* 35(4):521–536.
URL: *https://doi.org/10.1177/0894439316638457*

Romantan, Anca, Robert Hornik, Vincent Price, Joseph Cappella and K Viswanath. 2008. "A comparative analysis of the performance of alternative measures of exposure." *Communication Methods and Measures* 2(1-2):80–99.

Scharkow, Michael. 2016. "The Accuracy of Self-Reported Internet Use—A Validation Study Using Client Log Data." *Communication Methods and Measures* 10(1):13–27.

Schwarz, Norbert. 1999. "Self-reports: How the questions shape the answers." *American Psychologist* 54(2):pp. 93–105.

Settle, Jaime. 2018. *Frenemies: How Social Media Polarizes America*. Cambridge, Cambridge University Press.

Steinmetz, Kevin F. 2012. "Message received: Virtual ethnography in online message boards." *International Journal of Qualitative Methods* 11(1):26–39.

Stone, A.A., C.A. Bachrach, J.B. Jobe, H.S. Kurtzman and V.S. Cain. 1999. *The Science of Self-report: Implications for Research and Practice*. Taylor & Francis.
URL: *https://books.google.com/books?id=k9p4AgAAQBAJ*

Taneja, Harsh. 2016. "Using commercial audience measurement data in academic research." *Communication Methods and Measures* 10(2-3):176–178.

Wells, Chris and Kjerstin Thorson. 2017. "Combining big data and survey techniques to model effects of political content flows in Facebook." *Social Science Computer Review* 35(1):33–52.