Political Resources and Online Political Hostility: How and Why Hostility Is More Prevalent Among the Resourceful

STIG HEBBELSTRUP *Aarhus University, Department of Political Science*

MATHIAS OSMUNDSEN *Aarhus University, Department of Political Science*

MICHAEL BANG PETERSEN *Aarhus University, Department of Political Science*

***T**oxicity and hostility permeate political debates on social media, but who is responsible? Canonical theories of political engagement equate political resources with being a “model democratic citizen.” In contrast, we argue that in the current polarized political climate, those same resources make people behave hostile in political debates. Combining two years of survey and behavioral Twitter data, we provide empirical support for a link between political resources and online political hostility. This link is pronounced among citizens high in affective polarization – a trait held by many resourceful citizens in current US society. Worryingly, resourceful but hostile social media users do not simply cater to fringe audiences. Rather, they dominate online debates by tweeting more frequently, having more friends and followers, and by occupying powerful positions in their online networks. The present findings thus shed important light on the causes and consequences of online political hostility.*

.Postdoc at Aarhus University (stighebbelstrup@ps.au.dk)
Postdoc at Aarhus University
Professor at Aarhus University

INTRODUCTION

No more than two decades ago, political debate took place in closed social circles, town hall meetings, and the opinion pages of newspapers. Today, the debate has moved to the public network societies of online social media (Castells 2011). This radical transformation of citizens' political participation was believed to hold enormous potential for the quality and scope of dialogue as it democratized access to the public sphere (Dahlberg 2001). Even without formal connections or the ability to write in prose that satisfied the editor of a newspaper, everyone now had a voice and could engage with others online, including politicians. Yet lowering the entry barriers to the public sphere seemingly came with a cost. As any regular user of social media will know, many online political conversations are anything but friendly. Citizens view online political conversations as far more negative and hostile than offline political conversations (Bor and Petersen 2021). Even more concerning, hate speech directed at minority groups, incitements of violence against political enemies, and other uses of hateful and toxic language are commonplace on online platforms (Bliuc et al. 2018; Keipi et al. 2016; Siegel 2020; Kim et al. 2021).¹

In explaining the prevalence of online political hostility, one key argument is that the Internet's democratization of access is in itself to blame. Thus, a range of studies converge on the notion that online hostility reflects a lack of resources in the form of low socioeconomic status, low cognitive ability, or low emotion-regulation abilities. In psychological research, for example, online hostility, aggression, and "trolling" are often attributed to a cocktail of anti-social personality traits (e.g., sadism) and situational triggers of negative mood and frustration (Cheng et al. 2017). Among political scientists, a dominant theory is that socioeconomic deprivation and perceptions of being economically and socially "left behind" have fueled deep-seated frustrations with the political system among certain segments of the population (e.g., rural voters with low incomes and no college degrees) (Pettigrew 2017).

Highlighting a lack of resources as a key trigger of hostility seemingly resonates well with one of the most canonical models of political behavior: the resource model of political participation (Brady,

1. Although the proportion of hateful tweets out of the total number is small according to some studies (Siegel et al. 2021; Gagliardone et al. 2016), this does not imply they are not influential in social media networks (Hawdon, Oksanen, and Räsänen 2017; Kaakinen et al. 2018), or perceived as problematic (Celik 2018).

Verba, and Schlozman 1995). Building on research accumulated since the 1960s, this model has identified important features of the “model democratic citizen” (Nie, Junn, and Stehlik-Barry 1996) and argues that *political engagement*, in the form of political participation and political interest, is “at the heart of democracy” (Verba, Schlozman, and Brady 1995, p1). Some of the most important precursors of political engagement are education and internal efficacy, which grants citizens the capacity and inclination to engage in politics.

We examine the relationship between the traditional resource-related construct of political engagement and its antecedents and online political hostility. In contrast to established intuitions, we argue and demonstrate that the factors typically associated with democratic participation in political activities are associated with more – not less – online hostility. In short, those who are more hostile in online political debates are those who have resources, rather than those lacking resources. Importantly, this finding extends rather than challenges the original formulation of the resource model of political participation. As envisioned by Nie, Junn, and Stehlik-Barry (1996), the democratic effects of political engagement depend on *political enlightenment*, that is, the “understanding of and commitment to the rules of the democratic game,” which “tempers the unbridled pursuit of self-interested political engagement” (p. 6). Whether the development of political resourcefulness (e.g., via education) leads to the internalization of political enlightenment depends on features of the wider political system (Thomsen and Olsen 2017). In a highly polarized country with low levels of political trust and high levels of affective polarization, such as the present-day United States (Levi and Stoker 2000; Citrin and Stoker 2018), hostility may be the default way of engaging with politics. We therefore also investigate whether *affective polarization* (Iyengar et al. 2019) acts as a moderator of the relationship between political engagement and online political hostility.

We investigate the overall argument using a unique dataset that combines behavioral data on actual online political hostility from Twitter with panel survey data on traditional measures of political engagement. We first replicate the classical finding of a positive association between education and internal efficacy on the one hand and (offline) political engagement in the form of political participation and political interest on the other hand. Second, in a cross-sectional analysis, we demonstrate that political engagement and its antecedents, such as education and internal efficacy, also predict

higher levels of online political hostility, even after adjusting for personality traits and additional socio-demographic variables. Third, we show that this effect is moderated by measures of political enlightenment such that resources for politics only predict political hostility among those who are affectively polarized and, hence, *not* politically enlightened in Nie, Junn, and Stehlik-Barry (1996)'s sense of the term. Fourth, we use the panel element of our infrastructure to investigate the link between individual-level changes in engagement and actual hostile online behaviors over time while keeping many dispositional and time-invariant factors constant. To our knowledge this is the first study to explicitly study the sources of online political hostility using this unique data infrastructure.

The association between political resources and political hostility is both important and sobering. It suggests that the hostility of online platforms is not caused by the fact that online platforms expand access to the public sphere to previously underrepresented and resource-deprived groups. Yet, at the same time, the resourcefulness of hostile individuals may help explain why hostility dominates discussions on online platforms. As our fifth and last set of analyses demonstrate, politically hostile users tend to tweet more about politics, have larger numbers of friends and followers, and occupy more central positions in their social media networks. Thus, not only is political engagement associated with online political hostility; those who are hostile online are also some of the most powerful and influential actors in the network.

ONLINE POLITICAL HOSTILITY: DEFINITION AND EXISTING RESEARCH

We define *online political hostility* as the expression of degrading, malicious, and intimidating online behavior towards individuals or collectives based on political differences. Online political hostility shares commonalities with other forms of online extremism, like hate speech (Siegel 2020) and cyberbullying (Feinberg and Robey 2009), that have increasingly come under scrutiny by researchers across different scientific disciplines. It involves hostile behaviors expressed through various means – comments on visual or auditory content, memes, tweets and retweets, and postings on social media platforms – that transgress the boundaries of what ordinary citizens consider acceptable input into political conversations. Of critical concern, the key distinguishing feature of online political hostility is the explicit focus on hostile content that enters the realm of *politics*. Accordingly, toxic online

exchanges in the aftermath of, say, a sports game fall outside the purview of what we examine here.

As conversations about politics and other social issues rapidly move into the online world, concern is growing that online political hostility is running rampant (Duggan 2017). Social media users increasingly shy away from online discussions perceived to be intensely hostile and uncivil (Bor and Petersen 2021). Facebook, Twitter, and other tech companies are devoting increasing amounts of resources to detecting and removing hateful content from their platforms. And research teams across the globe are focusing on developing counter-strategies that can shield against the detrimental effects of hostility. As a consequence of these developments, our knowledge of the dynamics of online hostility is constantly expanding. We now know that social media users typically become *victims* of online hostility based on a host of personal attributes, like gender, social class, ethnicity, or sexual orientation (Bliuc et al. 2019; Costello et al. 2016; Saha, Chandrasekharan, and De Choudhury 2019; Silva et al. 2016). We also know that people who are visible on social media (Costello, Rukus, and Hawdon 2019; ElSherief et al. 2018) and who confront the hostile behavior of others (Costello, Rukus, and Hawdon 2019; Hawdon, Oksanen, and Räsänen 2014; Mathew et al. 2019; Ribeiro et al. 2017) often fall victim to online hostility themselves.

At the same time, we still lack a clear understanding of the characteristics of the *perpetrators* of online hostility; that is, the set of features that may propel social media users to express hostility towards others in online political discussions (Bor and Petersen 2021). As Siegel (2020, 60) succinctly summarizes: “Less is known about the actors in informal communities dedicated to producing harmful content, or the accounts that produce hate speech on mainstream platforms.”

From different perspectives, existing studies converge on a particular argument: Online hostility reflects a lack of resources either in the form of socioeconomic resources or in the form of cognitive or personality resources. For example, a central finding in the psychological literature on cyberbullying and so-called trolling is that the most severe cases of hate can be attributed to a (relatively) small segment of social media users with extreme personality traits, like those scoring high on the “Dark Triad” (Moor and Anderson 2019) and sadism (Buckels, Trapnell, and Paulhus 2014), which are intimately coupled with a lack of emotion regulation. Other research similarly highlights how the challenges of emotion regulation, under certain circumstances, imply that “anyone can become a

troll” (Cheng et al. 2017). All of us, often due to undesirable or stressful events, can occasionally feel overwhelmed by anger and negativity. In turn, these small bursts of negative emotions – perhaps amplified by the possibility of hiding behind an online veil of anonymity (Ribeiro et al. 2017) or belonging to a network of friends and followers who are accepting of hostile comments (Cheng et al. 2017; ElSherief et al. 2018) – will occasionally lead us to lash out at online strangers who happen to cross our path. The notion that grievances can trigger spurts of online hostility is also represented in research within political science. One focal topic has been to explain the public’s endorsement of Donald Trump, and, in particular, the rising support for hateful rhetoric against marginalized groups (Siegel 2020). To this end, researchers have argued that expressions of hate speech flow from the anger and frustrations felt by certain demographic groups – especially rural, white males with low incomes and little formal education – who perceive themselves to be economically and socially threatened by immigration, globalization, and new social norms (Mutz 2018; Petersen, Osmundsen, and Arceneaux 2018; Pettigrew 2017; Silver 2016).

The importance of these studies notwithstanding, it is relevant to note that these studies typically focus on destructive online behaviors that are not strictly about politics (i.e., cyberbullying, trolling), and the literature that does focus on disruptive political behavior tends to be narrow in its scope (e.g., right-wing hate speech targeting marginalized groups). The present aim is to establish a more general theory of online political hostility where the focus is on understanding the precursors of hostility as expressed through various forms of conversations, debates, and online exchanges that take place in the political realm. In building such a theory, we contend that online political hostility is best understood as a contemporary expression of political *engagement* that particular citizens may find attractive.

POLITICAL ENGAGEMENT – HOSTILE AND NON-HOSTILE – IS FOR THE RESOURCEFUL

While research on online hostility is in its infancy, there is a solid literature about the sources and conceptualization of political engagement (Nie, Junn, and Stehlik-Barry 1996; Verba, Schlozman, and Brady 1995). Political engagement, typically conceptualized as the motivation to become and stay

informed about political affairs and pursue one's self-interests through political participation (Nie, Junn, and Stehlik-Barry 1996; Verba, Schlozman, and Brady 1995), is a cornerstone of the democratic process. (Verba, Schlozman, and Brady 1995). Accordingly, answering the question of when and how citizens become politically engaged and pursue enlightened deliberation over political choices is of critical importance for understanding the well-being of any democracy.

The classical answer from the resource model of political participation (Verba, Schlozman, and Brady 1995; Brady, Verba, and Schlozman 1995) is that civic engagement rests on two key building blocks: the *motivation* and the *capacity* to engage in politics. Two of the most studied, and important, constructs affecting the motivation and capacity to participate in politics are (1) formal education and (2) political efficacy. Traditionally, education has been linked to a host of "participation-enhancing skills such as cognitive ability, civic skills, and civic orientations" (Kam and Palmer 2011, p. 613). It has been termed the "universal solvent" (Converse 1972) for its ability to influence most constructs deemed important for citizens' political activities, such as the motivation to become politically interested, and cognitive factors such as knowledge of politics and cognitive sophistication (cf. also (Delli Carpini and Keeter 1996; Nie, Junn, and Stehlik-Barry 1996; Neuman and Neuman 1986)). And although it seems that the effect size of education on political engagement was likely overstated in early studies (Kam and Palmer 2011; Persson 2015; Campbell 2009; Rasmussen and Nørgaard 2018), its continued research focus testifies to its ongoing importance. In a similar vein, political efficacy is one of the the strongest predictors of political engagement and one of the most studied constructs in the literature on political engagement (Sullivan and Riedel 2001). A particularly important form of efficacy is *internal efficacy*, which is the subjective feeling of competence in politics (Balch 1974; Craig 1979; Craig and Maggiotto 1982; Craig, Niemi, and Silver 1990; Niemi, Craig, and Mattei 1991).

Within this literature, political engagement has been seen as a boon to democracy. Formal education and a strong sense of internal efficacy potentially equip citizens with the motivation and capacity to partake in respectful conversations about topics of political importance online. Further, a formal education may even confer a strong safety net against the perceived threats of immigration and globalization (Stubager 2006) and thus shield against the frustrations potentially fuelling online hostility.

At the same time, it is key to note that the literature underlying the resource model explicitly highlights that political engagement does not in itself ensure pro-democratic participation. In addition, it requires *political enlightenment*. According to Nie, Junn, and Stehlik-Barry (1996, p. 6), “...democratic enlightenment is basically consensual in nature. Enlightenment signifies an understanding of and commitment to the rules of the democratic game and tempers the unbridled pursuit of self-interested political engagement.” Thus, although political engagement is a virtue to aspire to and foster, it also contains the seeds of political *inequality* and conflict. The motivation and capacity to engage politically will lead to respectful exchanges of political views only to the extent that people trust that others are playing by the same democratic rules and are interested in arriving at a consensus through deliberation. Yet, as some researchers have argued, many Western democracies are increasingly affected by processes of political dissatisfaction (Arceneaux et al. 2021) and political polarization (Iyengar et al. 2019). In such contexts, having politically engaged citizens is not necessarily a boon to democracy. In a society characterized by dissatisfaction and polarization, politically efficacious and highly educated citizens may use their political skills not to reduce, but rather to *increase* the hostility of political exchanges on social media. This danger of political engagement is the key concern we address here.

As is detailed below, the present study takes place in the United States, a highly polarized country (Iyengar et al. 2019), which for many years has experienced declining levels of political trust (Levi and Stoker 2000; Citrin and Stoker 2018). In such a setting, with hatred towards political out-groups, low trust in the political system and its political representatives, and a general skepticism towards others, democratic enlightenment may be lower than the ideal, and the potential for engaged conflict may be high. In essence, in contrast to the hundreds of studies showing the benefits of education, internal efficacy, and political engagement, there are reasons to believe that the current political conditions of the United States imply that political engagement and its antecedents facilitate *more* rather than less online political hostility.

PREDICTIONS

Our key claim about the effects of political engagement and its antecedents (i.e., political efficacy and education) on online political hostility may appear counterintuitive. For this reason, an important

first step is to show that we can reproduce findings regarding more traditional measures of political engagement. Therefore, we first investigate whether we can reproduce classical findings.

We next explicitly state our main hypothesis. Given the low levels of democratic enlightenment (i.e. high levels of affective polarization and low levels of trust) in the United States, we expect political engagement and its antecedents (i.e. internal efficacy and education) to predict online political hostility. Accordingly, we predict that

H1 Political resources are associated with higher levels of hostility in online political discussions.

Our argument is that any positive association between political resources and hostility reflects an absence of democratic enlightenment. Accordingly, we do not expect the relationship between political resources and online political hostility to operate unconditionally.

Affective polarization reflects the tendency to view one's in-group favorably and out-groups unfavorably (Iyengar et al. 2019) (see also (Uscinski et al. 2021)). Affective polarization has previously been found to make both the political system and people's everyday lives less "consensual." For example, previous research has demonstrated that economic pricing decisions can be affected by whether the customer is an in-group or out-group partisan (McConnell et al. 2018); Hetherington and Rudolph (2015) have demonstrated that affective polarization can negatively affect political trust over time; and Druckman et al. (2021) have demonstrated that out-party animus can even lead to a reluctance to follow life-saving health guidelines during a global pandemic.

There is already some evidence that dissatisfaction can lead to increased political engagement because those who are dissatisfied want to make their voices and preferences heard (Citrin and Stoker 2018). Similarly, increasing levels of affective polarization have been found to increase political participation (Iyengar and Krupenkin 2018). In this perspective, education and internal efficacy simply grant citizens the capacity and inclination to be active in the political realm: They are neither good nor bad, but can lead to democratically problematic behavior given the right circumstances. On this basis, the last prediction is:

H2 The effect of political resources on online political hostility *increases* as the level of affective polarization increases.

DATA AND METHODS

Data

We rely on a novel data infrastructure combining panel survey data with online behavioral Twitter data. While a handful of other papers have a similar data infrastructure (e.g., Mosleh et al. 2021), we are, to our knowledge, the first team to use it to study the antecedents of online political hostility.

To obtain our data, we commissioned the survey agency YouGov to invite panelists from their standing web panel to participate in a series of panel surveys. We recruited the original set of panelists in January 2019. Here, participants were asked to complete a survey with measures of our set of potential causes of online political hostility and, importantly, to provide their Twitter ID and give consent to let us scrape their publicly available Twitter data and link it to their survey answers. These participants, along with an additional set of panelists who were included to increase the study's power in September 2020, were then asked to participate in five additional survey waves, administered at various time points until the final survey in January 2021. The total number of panelists where we were able to link survey responses to their Twitter handle in this period was 2567. For the analyses where we only use Twitter data the maximum number of respondents is 3643. Full number of respondents can be seen in tables in the main text or the appendix. The six surveys did not include the same questions, and, consequently, not all measures described below are available in each wave. Further, as participants were free to opt in and out of the different surveys, the final sample size for each wave differs. Tables 10-13 in the appendix outline in greater detail which measures and how many participants were available for each wave.

The Twitter scraping commenced March 4, 2019 and stopped in January 2021. During this period, we managed to scrape 9,250,180 tweets from our panelists.

Ethical considerations

The studies were conducted in accordance with the Code of Conduct for Research Integrity of the primary investigator's university and associated policies of the National Committee for Health Research Ethics.

All studies were conducted online and involved voluntary and informed consent with the possibility of ending participation at any time. Potential respondents were initially selected based on their willingness to share their Twitter id and link their Twitter profile to their survey responses. Before survey respondents were asked to participate in the actual survey they were informed about the content of the project, who the project responsible was and how to contact them, how to withdraw their consent at a later stage, that participation was voluntary, and they were given the option to decline to participate in the survey. Before given a chance to provide their Twitter id, further in the survey, a detailed description of the research project was again presented to the respondents, such that they could make an informed choice about whether they wanted to participate again. Finally the respondents were informed that none of the Twitter data would be presented in any form outside the research project that would make them personally identifiable.

Participants were paid 500 YouGov points (which roughly translates into \$6) for completing the 15-minutes YouGov surveys, which is consistent with standard pay rates of the YouGov survey agency.

Methods

As we discuss in more detail below, the different analyses rely on slightly different measures and approaches. Accordingly, for ease of exposition and coherence, we present the analyses' methods individually before presenting the results.

Generally speaking, we use Mplus (Muthén and Muthén 2017) for all statistical models and the statistical computing environment R (R Core Team 2020) to connect to Mplus using the MplusAutomation package (Hallquist and Wiley 2018). Most graphs use the plotting library ggplot2 (Wickham 2016), except the network graphs, where we use R's built-in plotting functionality using the igraph package (Csardi and Nepusz 2006). When analyzing the panel data, we use a specific form of cross-lagged panel models, which account for time-invariant confounding at the individual level – the so-called Random-Intercept Cross-Lagged Panel Model (RI-CLPM). By using a RI-CLPM, we account for stable (time-invariant) dispositional features such as stable socioeconomic factors and personality traits. If (within-person) changes in political engagement lead to (within-person) changes in hostility, this is also strong evidence that our hypothesized effect of political engagement on hostility is not driven by, for

example, differences in (negative) character traits.

All tables report 95 % confidence intervals, and intervals that are significantly different from zero in this interval are indicated with a *. Unless otherwise noted, all scales are rescaled to range from zero to one.

Measures of online hostility

Political Hostility Our key measure of online political hostility is based on a word embeddings approach. The gist of the method is that through their use of language in tweets, respondents *themselves* decide what constitutes political hostility (Rasmussen et al. 2021). This method has been validated using expert raters. The method as used here consists of four steps. First, we create *word vectors* based on tweets from our panelists using the word2vec method (Mikolov et al. 2013; Rheault and Cochrane 2020). Word vectors are numerical representations of a word in a multidimensional space. Word vectors can therefore be used to measure the distance between words; are the words "political" and "hate," for example, closely related? Second, we combine the word vectors "political" and "hate" into a combined "political hate" vector. Third, we use this combined measure to calculate the distance between individual tweets and the political hate word vector: If a tweet contains many words that are close to the "political hate" word vector,² it is more politically hostile; if the tweet is further away from the word vector, it is less politically hostile. Fourth and finally, we create a measure of a respondent's general level of political hate by averaging the political hate scores of all her tweets; and for the panel models we do this averaging per time unit investigated. This measure of overall political hostility constitutes our primary dependent variable.

In the appendix and in the main text, however, we also discuss and show results using related constructs that also measure online political hostility to validate and elaborate on the findings. First, we use a classical way of measuring sentiment using a dictionary approach, i.e. the so-called AFINN dictionary ranging from -5 to 5 (Nielsen 2011), which is a standard way of measuring sentiment (Stine 2019). Second, we also report results from Google's Perspective API to detect online toxicity (JIGSAW). Google's Perspective API was developed using Wikipedia's editorial discussions. Although

2. Our distance measure is the cosine similarity measure

it has been criticized (Salminen et al. 2018; Hosseini et al. 2017; Gröndahl et al. 2018), it is one of the most prominent examples of developing an authoritative approach towards measuring online toxicity.

It is of course important to stress that although using a general notion of online hostility or hate speech works in some cases, in many situations a much greater conceptual and empirical delineation of types and targets of hate speech is called for (Siegel and Badaan 2020; Saha, Chandrasekharan, and De Choudhury 2019). We thus do not mean to imply that the ways of measuring online hostility presented here are by any means the only ways to measure it.

Measures of political resources

In order to make the investigation of the influence of political resources on online political hostility as comprehensive as possible, we investigate a broad range of measures of political resources. In essence, we can distinguish between more distal measures of political resources, such as education and internal efficacy, and more proximate measures of political resources such as online and offline political participation.

Internal efficacy Internal efficacy is measured using the standard Center for Political Studies (CPS) questions (Niemi, Craig, and Mattei 1991): "Sometimes politics seem so complicated that a person like me can't really understand what's going on" and "I feel that I have a pretty good understanding of the important political issues facing our country." We also use a third item from the European Social Survey that asks whether "I often think it is difficult to make up my mind about political issues." Cronbach's alpha is 0.70 for our measure of internal efficacy.

Education Education is measured using six different educational categories: No HS, High school graduate, Some college, 2-year, 4-year, Post-grad.

Political participation To fully capture the many ways citizens engage in politics, we include measures of both online and offline political participation. To measure traditional *offline* participation, we created an 11-item battery based on previous research (Oser, Hooghe, and Marien 2012; Zúñiga, Molyneux, and Zheng 2014), asking participants on scales ranging from "1 = Never" to "7 = Several times a day"

how often they participate in eleven offline political activities, like “Contact a government official in person, by phone, or by letter,” “Work or volunteer for a political party, candidate or a social cause,” and “Wear clothes or use items in support of a political or social cause.” To measure *online* participation, we asked participants on the same response scale how often they engaged in ten online political activities (Feezell 2016; Valenzuela et al. 2019), including “Comment on others’ online posts about a political or social cause (including online news articles),” “Like or follow a page of a political group, party or candidate on social media,” and “Start or join a political group or group supporting a cause on a social networking site.”

Measures used for statistical control

Personality traits In assessing the empirical support for these predictions, it is key to account for alternative explanatory factors. As was described earlier, a widespread account stresses that online political hostility reflects the absence of particular cognitive or personality resources, i.e., that online hostility reflects the online presence of emotional, aggressive, and volatile individuals. Accordingly, we control for the Big Five personality traits, which is arguably the standard model to understand stable individual differences in motivational, emotional, and cognitive dispositions to think, act, and feel (Denissen and Penke 2008; DeYoung 2020; John, Naumann, and Soto 2008). Neuroticism, with its proneness to experience negative emotionality, anxiety, and sensitivity towards punishment, as well as extroversion, with its tendency to experience positive emotionality and reward sensitivity, have a strong affective component (Krueger and Johnson 2008; Watson and Clark 1992). Both neuroticism and extroversion have consistently been found to relate to mood (DeYoung 2010b; Watson and Clark 1992; Marcus et al. 1995; DeYoung 2010a).

We rely on the Big Five Inventory (BFI) consisting of 44 items to measure the five major personality traits John, Naumann, and Soto (2008). The alpha reliability was 0.82 for openness, 0.84 for conscientiousness, 0.87 for extroversion, 0.82 for agreeableness, and 0.87 for neuroticism.

Party identification In addition to the measures described above, we also include a measure of party identification in order to rule out the possibility that our findings on political engagement reflect

underlying partisan dispositions. To measure partisanship, we used a standard question from the American National Election Study, asking “Generally speaking, do you usually think of yourself as a Republican, Democrat, Independent, Other?” Participants who answered either “Republican” or “Democrat” were then asked whether they saw themselves as “Strong” or “Not very strong” supporters of their party, while those answering “Independent” or “Other” were asked whether they saw themselves as closer to one of the parties. Based on participants’ responses, we then created our partisanship measure on a seven-point scale, ranging from “Strong Democrat” to “Strong Republican.”

Demographics Standard demographics of age, gender and income were obtained from YouGov’s background information about its panelists. Age is measured in age brackets of five-year intervals starting with 18-22 and ending with 83+. Our gender variable “Female” is a binary variable containing the categories “Male” and “Female” and indicates whether someone self-reported as female or not. Household income is measured in intervals of \$10,000 starting from “less than \$10,000” and ending with “\$500,000 and more.” Controlling for income in particular is important, as alternative theoretical perspectives highlight the lack of more tangible resources (such as money) as a main driver of political engagement.

Measure of Democratic Enlightenment

To measure democratic enlightenment, we use a measure of affective polarization. Affective polarization typically relates to animosity towards opposing political groups, especially party elites (Druckman and Levendusky 2019), and is here measured using the standard approach of measuring the difference between self-reported feeling thermometer scores of the in-party vs. the out-party (Iyengar et al. 2019). A higher score means the person is *more* polarized and vice versa. We asked questions about three Republican politicians (Ted Cruz, Donald Trump, and Mike Pence) and three Democratic politicians (Joe Biden, Elizabeth Warren, and Bernie Sanders). We then use the respondents’ answers on party identification to distinguish whether they self-identify as Republicans or Democrats. The question asked was: “We’d like to get your feelings toward some of our political leaders and other people who are in the news these days on a ‘feeling thermometer.’ A rating of zero degrees means you feel as cold

and negative as possible. A rating of 100 degrees means you feel as warm and positive as possible. You would rate the politician at 50 degrees if you don't feel particularly positive or negative toward the politician."

Descriptive statistics for all measures can be found in Appendix 1, Tables 10-13.

DO POLITICAL RESOURCES PREDICT ONLINE POLITICAL HOSTILITY – AND FOR WHOM?

Methods and measures

To investigate the impact of political resources on online political hostility, we first present results from a series of ordinary least squares (OLS) regressions where we investigate the effect of political resources on political hostility. We conduct the analyses across our measure of political hostility and our two robustness measures, i.e., Afinn scores and toxicity scores. We use both education and internal efficacy as measures of political resources, as well as direct measures of political engagement such as political participation and political interest. Thinking in terms of the "funnel of causality" (Campbell et al. 1980), education and internal efficacy are more distant predictors of political engagement, whereas political engagement in the form of political participation is the *outcome* of being politically engaged. We might thus expect the effect of participation on our measures of online hostility to be greater for this measure compared to education and internal efficacy.

We report cluster robust standard errors to take into account that a single Twitter user can participate across multiple waves. This dataset is thus organized as a so-called Time-Series Cross-Section (TSCS) dataset using pooled OLS. All models control for the Big Five personality traits, age, gender, and income. As can be seen in Table 9 in Appendix 1, not all measures are available in each survey and we therefore use multiple imputation and full information maximum likelihood to estimate missing values in these analyses and the analyses that follow. We elaborate on this in Appendix 5.

Results

Figure 1 makes clear that political engagement and its antecedents are indeed associated with online political hostility, and we can thus confirm Hypothesis **H1**. Inspecting the results more closely, we notice that the largest effects of our measures of political engagement are found for our measure of political hostility, compared to either AFINN scores or toxicity scores. The effect of online participation and political interest are, for example, 0.136 (95 % CI [0.103;0.169]) and 0.118 (95 % CI [0.096;0.140]) on political hate, whereas they are 0.055 (95 % CI [0.019;0.091]) and 0.025 (95 % CI [0.006;0.044]) for AFINN scores; the effects of political engagement on toxicity are also generally smaller compared to political hate. This is also indirect confirmation that it is indeed political resources that are driving the relationship, since it takes more political effort to engage in hostile *political* debates than simply uttering negatively valenced words. It is important to stress that these results regarding political resources persist after controlling for some of the most important individual-level variations in personality.

Interestingly, we find no support for the claim that a negative or bad mood should lead to increased hostility (see Table 4 in Appendix 1). In general, neither extroversion nor neuroticism are positively or negatively related to our measures of online hostility; a partial exception is the effect of neuroticism, which is a significant predictor of toxicity but not of political hate or AFINN scores. These results of course do not imply that some politically hateful tweets cannot be written in a heated moment, but simply that in general the proclivity to write politically hateful tweets is not driven by the emotionally charged dispositions of extroversion or neuroticism.

To investigate Hypothesis 2 about the conditional effect of democratic enlightenment, we conducted a series of interaction models where the effect of education and internal efficacy, the two classical correlates of political engagement, are moderated by affective polarization. These results are presented in Figure 2.

Figure 2 displays the coefficient from the *interaction* between the resources from Figure 1 and the measure of affective polarization; if the interaction term is significant, the resource is moderated by levels of affective polarization. A positive value indicates that the effect of the resource on online hostility increases when affective polarization increases. From Figure 2 it is clear that affective polarization moderates the relationship between political engagement and online hostility across our

FIGURE 1. The effect of political engagement and its antecedents on political hostility. All models control for the Big Five personality traits, income, age, gender, and party identification. The full results can be found in Appendix 1 Table 4

three measures; eight out of fifteen moderation effects are significant and in the expected direction. In general, the resources exerting the largest influence on online hostility, i.e. online participation and offline participation, are also those that are consistently significant across model specifications. This is not surprising since these also have the highest statistical power. This corroborates **H2**.

A sense of the size of the interaction effect can also be assessed from the graph (full marginal effect

FIGURE 2. The effect of political engagement and its antecedents moderated by affective polarization. All models control for age and gender but not the full set of controls to reduce problems of multicollinearity. The figure illustrates the *interaction* effect from each model. Full results can be found in Table 5 in appendix 1

plots with more information is available in Figure 13 in Appendix 1). For instance, the effect of offline participation on toxicity (scaled from zero to one) changes by 0.593 (95 % CI [0.194;0.991] on a zero to one scale. This is a large effect. The effect goes from being insignificant when affective polarization is zero to exerting an effect on toxicity of 0.462 (95 % CI [0.221;0.702]). Not all effects are of this size, but the moderated effects of online and offline participation across our measures of online hostility are

in the same range and thus quite large.

INCREASING INTERNAL AND MEASUREMENT VALIDITY: RANDOM INTERCEPT CROSS-LAGGED PANEL MODELS FOR BEHAVIORAL MEASURES OF ONLINE ENGAGEMENT

The above results support the presented hypotheses. At the same time, they are not without limitations, as they rely on observational data and are influenced by problems of confounding and reverse causality (Angrist and Pischke 2008). Furthermore, the analyses rely on self-reported explanatory variables, which can be biased due to social desirability or measurement error (Holbrook and Krosnick 2010).

To increase the internal validity of the conclusions, we turn to so-called Random Intercept Cross-Lagged Panel Models (RI-CLPM) (Hamaker, Kuiper, and Grasman 2015) by using the panel structure of the dataset. The RI-CLPM allow us to investigate the reciprocal relationships between political engagement and online hostility over time more thoroughly than our previous analysis. That is, we can ask: Do increases in political engagement lead to online hostility, or does online hostility drive citizens' engagement?

Why use the RI-CLPM? Classical cross-lagged models have been criticized for not distinguishing clearly among variation *between* individuals and *within* individuals since the parameter estimates from such a model are a *mix* of the two (Hamaker, Kuiper, and Grasman 2015). The RI-CLPM, in contrast, split the variation into separate estimates for the between variance and within variance. This approach is akin to a fixed effect model – which subtracts the (within-group) mean to isolate the within-individuals variation – and is accomplished by estimating a random intercept for each construct so that the cross-lagged parameter estimates are based only on within-individual variation. The result is a model that can investigate reciprocal relationships, which is not possible in a fixed effect model, as well as hold (time-invariant and possibly unobserved) dispositional traits constant. Time-varying confounding cannot be ruled out (Imai and Kim 2019).

Furthermore, to increase the measurement validity of the conclusions, we turn towards direct measures of political engagement rather than merely analyzing the antecedents of political engagement.

Making fuller use of the available data, we are able to repeat these analyses using two *behavioral* measures of online political engagement. First, we exploit the fact that the measure of “political hate” can be split into two (Rasmussen et al. 2021): A word vector of words associated with “political” and a vector of words associated with “hate.” We use these two vectors to examine whether people who often tweet about politics – i.e., those who are politically engaged – are also more likely to disseminate hateful content. Again we also use the measures of toxicity and Afinn scores as a replacement for the hate vector to gauge the robustness of the findings.

Methods and measures

The previous analyses have used the survey waves as the unit of analysis. Here, we are not limited to that since we have daily data on hundreds of Twitter users. To make full use of this, we therefore organize the dataset in terms of months to ensure reliable measures of “political” and “hate” that are not overly influenced by any single tweet but are still highly fine-grained.

The RI-CLPM are estimated in Mplus version 8. Since the RI-CLPM rest on the assumption of normally distributed variables (Usami, Murayama, and Hamaker 2019), and in order to reduce the influence of outliers, we remove the most extreme observations (see Figures 7-10 in Appendix 1).³ The model includes measures of both stability and change. Positive stability scores indicate that those who were (hateful or engaged) at $time_1$ were also likely to be more hateful or engaged relative to their own expected mean at $time_{+1}$. Conversely the reciprocal effects, i.e. the effects of online hostility on political engagement (online hostility $\rightarrow$ political) or the effect of engagement on online hostility (political $\rightarrow$ online hostility) indicate that a (positive) deviation from an individual’s expected mean in e.g. political engagement or online hostility also leads to a (positive) deviation in the expected mean of political engagement or online hostility.

The measures of toxicity scores and Afinn scores are based on data from (Rasmussen et al. 2021) and thus based on the time period from March 2019 until March 2020. Although we have full data for both our hostility and political measures, we only focus on the period March 2020 until January 2021 since this is the period during which we collected measures of affective polarization. Since we are also

3. We investigate the robustness of these choices (see Figures 4-6) in Appendix 1

interested in investigating whether these behavioral measures of political engagement are moderated by subjective feelings of discontent, we include the average score of affective polarization in a second set of analyses to investigate whether the relationships are moderated by political enlightenment as above. We follow Mulder and Hamaker (2021) in using a multiple group approach towards investigating moderation hypotheses for the RI-CLPM. Put simply, we create a RI-CLPM separately for those high and low on affective polarization to investigate whether the reciprocal relationships between engagement and hostility change as a function of affective polarization.

Results

Figure 3 illustrates the reciprocal relationships between our three measures of online hostility (hate, AFINN scores, and toxicity scores) and our measure of online engagement, i.e. our political vector. We first notice that there is a fair amount of individual-level stability over time, indicated by the positive and significant stability estimates (political → political and online hostility → online hostility). This means that Twitter users who have been hostile or politically engaged at a previous point in time tend to remain hostile or politically engaged at later time points. More importantly, we find that political engagement exerts a positive and significant influence on our three measures of online hostility. An increase in political engagement from zero to one, for example, leads to an increase of 0.088 (95 % CI [0.062;0.113]) in hate on a zero to one scale. This finding supports (but does not definitely prove) the claim that increased political engagement leads to increased hostility.

However, we also find that online hostility, measured using hate and toxicity scores, predicts political engagement (although hate measured by AFINN scores does not). Still, the estimated effects are generally smaller compared to the reverse effect. This finding is perhaps not surprising. The findings above demonstrated that the effect of resources on political hostility is moderated by affective polarization. We would thus expect the effect of online hostility on political engagement to matter most for those who are most affectively polarized. This is also generally confirmed as shown in Figures 11-12 in the Appendix. The effect of online hostility (measured using toxicity scores and hate) on political engagement is only significant for those high on affective polarization but not for those low on

FIGURE 3. RI-CLPM for online hostility. Full results can be found in Table 3

affective polarization.⁴

In sum, these findings from the cross-lagged models bolster our previous findings. Over-time increases in political engagement are associated with increased political hostility on social media. And while we observe occasional hints that hostility can lead to increased political engagement, this association is much weaker.

IMPLICATIONS: THE POLITICAL LEVERAGE OF THOSE WHO ARE HOSTILE

The traditional finding in the literature on political engagement is that more resourceful individuals have more political leverage (Brady, Verba, and Schlozman 1995; Verba, Schlozman, and Brady 1995). Given the finding that more resourceful and engaged individuals are more hostile in their online conversations, this may entail the potentially problematic implication that more hostile individuals

4. The effects are illustrated for the average effect using cut points in the range [0,5;0.65] where affective polarization is measured on a zero to one scale, since we have no a priori guess for which cutpoint to choose.

have greater influence within their online networks.

We investigate the topic of the influence of politically hostile users in two different ways: We first show results from count models that examine whether those who are politically hostile (a) tweet more, (b) have more friends, and (c) have more followers.⁵ A "friend" in Twitter's usage is someone I follow, and a "follower" is someone who follows me. The second analysis views the data as a *network* and investigates (i) whether politically hostile/non-hostile Twitter users tend to bundle together in "echo chambers" or not, and (ii) who the powerful actors in the network are.

Results

In the Appendix 4 we demonstrate that those who are politically hostile are also likely to be influential in the network insofar as they both tweet more and likely also reach more followers (and slightly more friends) in their Twitter feeds. These effects are not trivial. A change from -2SD to +2SD on political hostility means an increase in the number of tweets from around 150 to around 1150 and an increase in the number of followers from around 120 to around 420.

Secondly Figure 4 suggest that those who are politically hostile are dominating the network, and thus exert a large influence in the network, no matter whether you are high or low on political hate; i.e. they do not exhibit separate echo chambers. Further analyses in Appendix 3, where we include formal measures of network centrality, also corroborate this finding: Those who are most hostile also exert the largest influence in the network (see Figure 1 in Appendix 3).

5. In the network analysis literature, these two last features are named indegree and outdegree of the persons (vertices)

FIGURE 4. Network visualization of those with high levels of political hate and low levels of political hate. The size of vertices indicates the number of friends/followers and the color the level of political hate of the respondent using the "magma" color palette. The layout used is a "ring" layout

CONCLUSION AND DISCUSSION

The advent of the Internet has transformed contemporary political life. Social media platforms like Twitter, Facebook, and YouTube have dramatically expanded the potential for citizens' democratic engagement. Yet many believe this increased democratization has carried with it a steep cost: a hostile environment in which toxic exchanges between political opponents are commonplace. Arriving at a better understanding of the factors that drive social media users to engage hostile online has been our goal with this study.

Building on canonical theories about political engagement (Brady, Verba, and Schlozman 1995; Nie, Junn, and Stehlik-Barry 1996), we argue and demonstrate empirically that political resources in the form of political engagement, and its antecedents education and internal efficacy, are associated with increased political hostility on social media. This finding, we contend, turns commonly-held beliefs about the perpetrators of online hostility on their heads. Hostile users of social media do not suffer from political apathy or an inability to navigate the political landscape on social media. Instead, it is those who *do* possess the ability and political motivation to engage politically who turn hostile in online political debates.

At first glance, our key finding may appear paradoxical. Surely a well-educated citizenry ready to engage in political conversations is a boon to democracy? Yet the paradox evaporates once we account for *political enlightenment* – that is, the extent to which citizens trust each other to play by the same democratic rule book. Accordingly, our analyses show that the link between political resources and online political hostility critically depends on people's animosity and distrust towards opposing political groups. Only when people become affectively polarized and stop trusting politicians and voters from opposing political groups will political resources lead to hostility. This effect of (the lack of) democratic enlightenment may also explain the high levels of political hostility witnessed in contemporary politics. In America, but elsewhere in Western politics as well, fear and loathing across party-political lines have worsened for years (Iyengar and Westwood 2014). Such a climate of intense political distrust and polarization may lead broad groups of resourceful citizens to view hostile acts and denigration of political antagonists as the only viable means of promoting their political goals.

Not only are highly politically engaged and polarized individuals more politically hostile; they also

tend to dominate online debates by tweeting more frequently, having more friends and followers, and by taking up more central positions in their social networks. This finding has important ramifications. Hostile social media users are not solitary actors, catering only to small audiences of like-minded friends and followers. Instead, their high degree of “connectivity” may amplify the scope and speed of the spread of hostility, making it visible and accessible to users far removed in the social media sphere Bor and Petersen 2021.

These findings have implications for how to best combat online political hatred. If online hostility resulted from certain features of social media – like online anonymity, lack of emotional cues, or algorithmic promotions of hostile content – the solution would be (relatively) straightforward: platform-based interventions aimed at creating a better social media environment by, for instance, changing ranking systems to discourage sharing of hostile content. But if online hostility instead feeds on a cocktail of strong political resources and out-party animosity, finding viable solutions appears more challenging. Political polarization had been worsening for decades prior to the inception of Facebook, Twitter and other social media platforms. Accordingly, an exclusive focus on platform-based solutions to the problem of political hostility may distract us from the root problem: the deeper societal divides that breed low political enlightenment and the key responsibility of political elites in this regard. The real solution to polarization – and, by extension, online hostility – seems to require difficult policy reforms that can turn the tide of decades of increasing political animosity. Unfortunately, this is much easier said than done.

REFERENCES

- Angrist, Joshua D, and Jörn-Steffen Pischke. 2008. *Mostly harmless econometrics: An empiricist's companion*. Princeton university press.
- Arceneaux, Kevin, Timothy B Gravelle, Mathias Osmundsen, Michael Bang Petersen, Jason Reifler, and Thomas J Scotto. 2021. “Some people just want to watch the world burn: the prevalence, psychology and politics of the ‘Need for Chaos’.” *Philosophical Transactions of the Royal Society B* 376 (1822): 20200147.

-
- Balch, George I. 1974. "Multiple indicators in survey research: The concept" sense of political efficacy"." *Political Methodology*, 1–43.
- Bliuc, Ana-Maria, John Betts, Matteo Vergani, Muhammad Iqbal, and Kevin Dunn. 2019. "Collective identity changes in far-right online communities: The role of offline intergroup conflict." *New Media Society* 21, no. 8 (August): 1770–1786. issn: 1461-4448. <https://doi.org/10.1177/1461444819831779>.
- Bliuc, Ana-Maria, Nicholas Faulkner, Andrew Jakubowicz, and Craig McGarty. 2018. "Online networks of racial hate: A systematic review of 10 years of research on cyber-racism." *Computers in Human Behavior* 87:75–86.
- Bor, Alexander, and Michael Bang Petersen. 2021. "The Psychology of Online Political Hostility: A Comprehensive, Cross-National Test of the Mismatch Hypothesis." *American Political Science Review*, 1–18. <https://doi.org/10.1017/S0003055421000885>.
- Brady, Henry E, Sidney Verba, and Kay Lehman Schlozman. 1995. "Beyond SES: A resource model of political participation." *American political science review* 89 (2): 271–294.
- Buckels, Erin E, Paul D Trapnell, and Delroy L Paulhus. 2014. "Trolls just want to have fun." *Personality and individual Differences* 67:97–102.
- Campbell, Angus, Philip E Converse, Warren E Miller, and Donald E Stokes. 1980. *The american voter*. University of Chicago Press.
- Campbell, David E. 2009. "Civic engagement and education: An empirical test of the sorting model." *American Journal of Political Science* 53 (4): 771–786.
- Castells, Manuel. 2011. *The rise of the network society*. Vol. 12. John wiley & sons.
- Celik, Serkan. 2018. "Tertiary-level internet users' opinions and perceptions of cyberhate." *Information Technology People* 31, no. 3 (January): 845–868. issn: 0959-3845. <https://doi.org/10.1108/ITP-05-2017-0147>.
- Cheng, Justin, Michael Bernstein, Cristian Danescu-Niculescu-Mizil, and Jure Leskovec. 2017. "Anyone can become a troll: Causes of trolling behavior in online discussions." In *Proceedings of the 2017 ACM conference on computer supported cooperative work and social computing*, 1217–1230.
- Citrin, Jack, and Laura Stoker. 2018. "Political trust in a cynical age." *Annual Review of Political Science* 21:49–70.

- Converse, Philip E. 1972. "Change in the American electorate." In *The human meaning of social change*, edited by Angus Campbell and Philip E. Converse, 263–339. New York: Russel Sage Foundation.
- Costello, Matthew, James Hawdon, Thomas Ratliff, and Tyler Grantham. 2016. "Who views online extremism? Individual attributes leading to exposure." *Computers in Human Behavior* 63 (October): 311–320. ISSN: 0747-5632. <https://doi.org/10.1016/j.chb.2016.05.033>.
- Costello, Matthew, Joseph Rukus, and James Hawdon. 2019. "We don't like your type around here: Regional and residential differences in exposure to online hate material targeting sexuality." *Deviant Behavior* 40, no. 3 (March): 385–401. ISSN: 0163-9625. <https://doi.org/10.1080/01639625.2018.1426266>.
- Craig, Stephen C. 1979. "Efficacy, trust, and political behavior: An attempt to resolve a lingering conceptual dilemma." *American Politics Quarterly* 7 (2): 225–239.
- Craig, Stephen C, and Michael A Maggiotto. 1982. "Measuring political efficacy." *Political Methodology*, 85–109.
- Craig, Stephen C, Richard G Niemi, and Glenn E Silver. 1990. "Political efficacy and trust: A report on the NES pilot study items." *Political behavior* 12 (3): 289–314.
- Csardi, Gabor, and Tamas Nepusz. 2006. "The igraph software package for complex network research." *InterJournal Complex Systems*:1695. <https://igraph.org>.
- Dahlberg, Lincoln. 2001. "The Internet and Democratic Discourse: Exploring The Prospects of Online Deliberative Forums Extending the Public Sphere." *Information, Communication & Society* 4, no. 4 (January): 615–633. <https://doi.org/10.1080/13691180110097030>. <https://doi.org/10.1080/13691180110097030>.
- Delli Carpini, Michael X, and Scott Keeter. 1996. *What Americans Know about Politics and Why It Matters*. Yale University Press. ISBN: 9780300062564. <http://www.jstor.org/stable/j.ctt1cc2kv1>.
- Denissen, Jaap JA, and Lars Penke. 2008. "Motivational individual reaction norms underlying the Five-Factor model of personality: First steps towards a theory-based conceptual framework." *Journal of research in personality* 42 (5): 1285–1302.
- DeYoung, Colin G. 2010a. "Mapping personality traits onto brain systems: BIS, BAS, FFFS and beyond."
- . 2010b. "Toward a theory of the Big Five." *Psychological Inquiry* 21 (1): 26–33.

-
- DeYoung, Colin G. 2020. "Intelligence and personality."
- Druckman, James N, Samara Klar, Yanna Krupnikov, Matthew Levendusky, and John Barry Ryan. 2021. "Affective polarization, local contexts and public opinion in America." *Nature human behaviour* 5 (1): 28–38.
- Druckman, James N, and Matthew S Levendusky. 2019. "What Do We Measure When We Measure Affective Polarization?" *Public Opinion Quarterly* 83 (1): 114–122. <https://doi.org/10.1093/poq/nfz003>. <https://doi.org/10.1093/poq/nfz003>.
- Duggan, Maeve. 2017. "Online harassment 2017."
- ElSherief, Mai, Shirin Nilizadeh, Dana Nguyen, Giovanni Vigna, and Elizabeth Belding. 2018. "Peer to Peer Hate: Hate Speech Instigators and Their Targets." ArXiv: 1804.04649, *arXiv:1804.04649 [cs]* (April). <http://arxiv.org/abs/1804.04649>.
- Feezell, Jessica T. 2016. "Predicting Online Political Participation: The Importance of Selection Bias and Selective Exposure in the Online Setting." *Political Research Quarterly* 69 (3): 495–509. <https://doi.org/10.1177/1065912916652503>.
- Feinberg, Ted, and Nicole Robey. 2009. "Cyberbullying." *The education digest* 74 (7): 26.
- Gagliardone, Iginio, Matti Pohjonen, Zenebe Beyene, Abdissa Zerai, Gerawork Aynekulu, Mesfin Bekalu, Jonathan Bright, Mulatu Alemayehu Moges, Michael Seifu, and Nicole Stremlau. 2016. *Mechachal: Online Debates and Elections in Ethiopia - From Hate Speech to Engagement in Social Media*. May. <https://doi.org/10.2139/ssrn.2831369>. <https://papers.ssrn.com/abstract=2831369>.
- Gröndahl, Tommi, Luca Pajola, Mika Juuti, Mauro Conti, and N Asokan. 2018. "All You Need is" Love" Evading Hate Speech Detection." In *Proceedings of the 11th ACM Workshop on Artificial Intelligence and Security*, 2–12.
- Hallquist, Michael N., and Joshua F. Wiley. 2018. "MplusAutomation: An R Package for Facilitating Large-Scale Latent Variable Analyses in Mplus." *Structural Equation Modeling*, 1–18. <https://doi.org/10.1080/10705511.2017.1402334>. <https://www.tandfonline.com/doi/full/10.1080/10705511.2017.1402334>.
- Hamaker, Ellen L, Rebecca M Kuiper, and Raoul PPP Grasman. 2015. "A critique of the cross-lagged panel model." *Psychological methods* 20 (1): 102.

- Hawdon, James, Atte Oksanen, and Pekka Räsänen. 2014. "Victims of online hate groups: American youth's exposure to online hate speech." *The causes and consequences of group violence: From bullies to terrorists*, 165–182.
- . 2017. "Exposure to Online Hate in Four Nations: A Cross-National Consideration." *Deviant Behavior* 38, no. 3 (March): 254–266. ISSN: 0163-9625. <https://doi.org/10.1080/01639625.2016.1196985>.
- Hetherington, Marc J, and Thomas J Rudolph. 2015. *Why Washington won't work*. University of Chicago Press.
- Holbrook, Allyson L, and Jon A Krosnick. 2010. "Social desirability bias in voter turnout reports: Tests using the item count technique." *Public Opinion Quarterly* 74 (1): 37–67.
- Hosseini, Hossein, Sreeram Kannan, Baosen Zhang, and Radha Poovendran. 2017. "Deceiving google's perspective api built for detecting toxic comments." *arXiv preprint arXiv:1702.08138*.
- Imai, Kosuke, and In Song Kim. 2019. "When should we use unit fixed effects regression models for causal inference with longitudinal data?" *American Journal of Political Science* 63 (2): 467–490.
- Iyengar, Shanto, and Masha Krupenkin. 2018. "The strengthening of partisan affect." *Political Psychology* 39:201–218.
- Iyengar, Shanto, Yphtach Lelkes, Matthew Levendusky, Neil Malhotra, and Sean J. Westwood. 2019. "The Origins and Consequences of Affective Polarization in the United States." *Annual Review of Political Science* 22, no. 1 (May): 129–146. ISSN: 1094-2939. <https://doi.org/10.1146/annurev-polisci-051117-073034>.
- Iyengar, Shanto, and Sean J. Westwood. 2014. "Fear and Loathing across Party Lines: New Evidence on Group Polarization." 59, no. 3 (December): 690–707. <https://doi.org/10.1111/ajps.12152>. <https://doi.org/10.1111/ajps.12152>.
- JIGSAW. <https://perspectiveapi.com/#/home>.
- John, Oliver P, Laura P Naumann, and Christopher J Soto. 2008. "Paradigm shift to the integrative Big Five trait taxonomy: History, measurement, and conceptual issues."
- Kaakinen, Markus, Pekka Räsänen, Matti Näsi, Jaana Minkkinen, Teo Keipi, and Atte Oksanen. 2018. "Social capital and online hate production: A four country survey." *Crime, Law and Social Change* 69, no. 1 (February): 25–39. ISSN: 1573-0751. <https://doi.org/10.1007/s10611-017-9764-5>.

-
- Kam, Cindy, and Carl Palmer. 2011. "Rejoinder: Reinvestigating the Causal Relationship between Higher Education and Political Participation." *Journal of Politics* 73 (7): 659–663.
- Keipi, Teo, Matti Näsi, Atte Oksanen, and Pekka Räsänen. 2016. *Online hate and harmful content: Cross-national perspectives*. Taylor & Francis.
- Kim, Jin Woo, Andrew Guess, Brendan Nyhan, and Jason Reifler. 2021. "The distorting prism of social media: How self-selection and exposure to incivility fuel online comment toxicity." *Journal of Communication* 71 (6): 922–946.
- Krueger, Robert F, and Wendy Johnson. 2008. "Behavioral genetics and personality: A new look at the integration of nature and nurture."
- Levi, Margaret, and Laura Stoker. 2000. "Political trust and trustworthiness." *Annual review of political science* 3 (1): 475–507.
- Marcus, George E, Elizabeth Theiss-Morse, John L Sullivan, and Sandra L Wood. 1995. *With malice toward some: How people make civil liberties judgments*. Cambridge University Press.
- Mathew, Binny, Ritam Dutt, Pawan Goyal, and Animesh Mukherjee. 2019. "Spread of hate speech in online social media." In *Proceedings of the 10th ACM conference on web science*, 173–182.
- McConnell, Christopher, Yotam Margalit, Neil Malhotra, and Matthew Levendusky. 2018. "The economic consequences of partisanship in a polarized era." *American Journal of Political Science* 62 (1): 5–18.
- Mikolov, Tomas, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. "Distributed representations of words and phrases and their compositionality." In *Advances in neural information processing systems*, 3111–3119.
- Moor, Lily, and Joel R Anderson. 2019. "A systematic literature review of the relationship between dark personality traits and antisocial online behaviours." *Personality and individual differences* 144:40–55.
- Mosleh, Mohsen, Gordon Pennycook, Antonio A. Arechar, and David G. Rand. 2021. "Cognitive reflection correlates with behavior on Twitter." *Nature communications* 12 (1): 1–10.
- Mulder, Jeroen D, and Ellen L Hamaker. 2021. "Three extensions of the random intercept cross-lagged panel model." *Structural Equation Modeling: A Multidisciplinary Journal*, 1–11.

- Muthén, LK, and BO Muthén. 2017. *Mplus User's Guide. Eighth Edition*.
- Mutz, Diana C. 2018. "Status threat, not economic hardship, explains the 2016 presidential vote." *Proceedings of the National Academy of Sciences* 115 (19): E4330–E4339.
- Neuman, W Russell, and WR Neuman. 1986. *The paradox of mass politics: Knowledge and opinion in the American electorate*. Harvard University Press.
- Nie, Norman H., Jane Junn, and Kenneth Stehlik-Barry. 1996. *Education and democratic citizenship in America*. Chicago: The University of Chicago Press.
- Nielsen, Finn Årup. 2011. "A new ANEW: Evaluation of a word list for sentiment analysis in microblogs." *arXiv preprint arXiv:1103.2903*.
- Niemi, Richard G, Stephen C Craig, and Franco Mattei. 1991. "Measuring internal political efficacy in the 1988 National Election Study." *The American Political Science Review*, 1407–1413.
- Oser, Jennifer, Marc Hooghe, and Sofie Marien. 2012. "Is Online Participation Distinct from Offline Participation? A Latent Class Analysis of Participation Types and Their Stratification." *Political Research Quarterly* 66, no. 1 (March): 91–101. <https://doi.org/10.1177/1065912912436695>. <https://doi.org/10.1177/1065912912436695>.
- Persson, Mikael. 2015. "Education and Political Participation." *British Journal of Political Science* 45 (03): 689–703.
- Petersen, Michael Bang, Mathias Osmundsen, and Kevin Arceneaux. 2018. "The "Need for Chaos" and Motivations to Share Hostile Political Rumors" (September). <https://doi.org/10.31234/osf.io/6m4ts>. <https://doi.org/10.31234/osf.io/6m4ts>.
- Pettigrew, Thomas F. 2017. "Social psychological perspectives on Trump supporters." *Journal of Social and Political Psychology* 5 (1): 107–116.
- R Core Team. 2020. *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing. <https://www.R-project.org/>.
- Rasmussen, Stig H R, Alexander Bor, Mathias Osmundsen, and Michael Bang Petersen. 2021. *Super-unsupervised text classification for labeling online political hate*, June. <https://doi.org/10.31234/osf.io/8m5dc>. psyarxiv.com/8m5dc.

-
- Rasmussen, Stig Hebbelstrup Rye, and Asbjørn Sonne Nørgaard. 2018. "When and why does education matter? Motivation and resource effects in political efficacy." *European Journal of Political Research* 57 (1): 24–46.
- Rheault, Ludovic, and Christopher Cochrane. 2020. "Word embeddings for the analysis of ideological placement in parliamentary corpora." *Political Analysis* 28 (1): 112–133.
- Ribeiro, Manoel Horta, Pedro H Calais, Yuri A Santos, Virgílio AF Almeida, and Wagner Meira Jr. 2017. "'Like Sheep Among Wolves': Characterizing Hateful Users on Twitter." *arXiv preprint arXiv:1801.00317*.
- Saha, Koustuv, Eshwar Chandrasekharan, and Munmun De Choudhury. 2019. "Prevalence and psychological effects of hateful speech in online college communities." In *Proceedings of the 10th ACM Conference on Web Science*, 255–264.
- Salminen, Joni, Hind Almerexhi, Milica Milenkovic, Soon-Gyu JUNG, Haewoon KWAK, and Bernard J JANSEN. 2018. "Anatomy of online hate: developing a taxonomy and machine learning models for identifying and classifying hate in online news media."
- Siegel, Alexandra A. 2020. "Online hate speech." In *Social Media and Democracy: The State of the Field, Prospects for Reform*, edited by Nathaniel Persily and Joshua A Tucker. Cambridge University Press.
- Siegel, Alexandra A., and Vivienne Badaan. 2020. "No2Sectarianism: Experimental Approaches to Reducing Sectarian Hate Speech Online." *American Political Science Review* 114, no. 3 (August): 837–855. ISSN: 0003-0554, 1537-5943. <https://doi.org/10.1017/S0003055420000283>.
- Siegel, Alexandra A., Evgenii Nikitin, Pablo Barberá, Joanna Sterling, Bethany Pullen, Richard Bonneau, Jonathan Nagler, and Joshua A. Tucker. 2021. "Trumping Hate on Twitter? Online Hate Speech in the 2016 U.S. Election Campaign and its Aftermath." *Quarterly Journal of Political Science* 16, no. 1 (January): 71–104. ISSN: 1554-0626, 1554-0634. <https://doi.org/10.1561/100.00019045>.
- Silva, Leandro, Mainack Mondal, Denzil Correa, Fabrício Benevenuto, and Ingmar Weber. 2016. "Analyzing the Targets of Hate in Online Social Media." *Proceedings of the International AAAI Conference on Web and Social Media* 10, no. 11 (March). ISSN: 2334-0770. <https://ojs.aaai.org/index.php/ICWSM/article/view/14811>.
- Silver, Nate. 2016. *Education, Not Income, Predicted Who Would Vote For Trump*, November. <https://fivethirtyeight.com/features/education-not-income-predicted-who-would-vote-for-trump/>.

- Stine, Robert A. 2019. "Sentiment analysis." *Annual review of statistics and its application* 6:287–308.
- Stubager, Rune. 2006. *The Education Cleavage: New Politics in Denmark*. Politica.
- Sullivan, J.L., and Riedel. 2001. "Efficacy: Political." In *International encyclopedia of the social and behavioral sciences*, edited by J. Smelser. Pergamon.
- Thomsen, Jens Peter Frølund, and Mark Olsen. 2017. "Re-examining socialization theory: How does democracy influence the impact of education on anti-foreigner sentiment?" *British Journal of Political Science* 47 (4): 915–938.
- Usami, Satoshi, Kou Murayama, and Ellen L Hamaker. 2019. "A unified framework of longitudinal models to examine reciprocal relations." *Psychological methods* 24 (5): 637.
- Uscinski, Joseph E, Adam M Enders, Michelle I Seelig, Casey A Klofstad, John R Funchion, Caleb Everett, Stefan Wuchty, Kamal Premaratne, and Manohar N Murthi. 2021. "American Politics in Two Dimensions: Partisan and Ideological Identities versus Anti-Establishment Orientations." *American Journal of Political Science*.
- Valenzuela, Sebastián, Daniel Halpern, James E. Katz, and Juan Pablo Miranda. 2019. "The Paradox of Participation Versus Misinformation: Social Media, Political Engagement, and the Spread of Misinformation." *Digital Journalism* 7, no. 6 (June): 802–823. <https://doi.org/10.1080/21670811.2019.1623701>.
- Verba, Sidney, Kay Lehman Schlozman, and Henry E Brady. 1995. *Voice and equality: Civic voluntarism in American politics*. Harvard University Press.
- Watson, David, and Lee Anna Clark. 1992. "On traits and temperament: General and specific factors of emotional experience and their relation to the five-factor model." *Journal of personality* 60 (2): 441–476.
- Wickham, Hadley. 2016. *ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag New York. ISBN: 978-3-319-24277-4. <https://ggplot2.tidyverse.org>.
- Zúñiga, Homero Gil de, Logan Molyneux, and Pei Zheng. 2014. "Social Media, Political Expression, and Political Participation: Panel Analysis of Lagged and Concurrent Relationships." *Journal of Communication* 64, no. 4 (June): 612–634. <https://doi.org/10.1111/jcom.12103>. <https://doi.org/10.1111/jcom.12103>.

Part

Appendix for Political Resources and Online Political Hostility: How and Why Hostility Is More Prevalent Among the Resourceful

Table of Contents

A	Appendix 1: Reproducing traditional relationships regarding political resources	3
A.1	Methods	3
A.2	Results	3
B	Appendix 2: Additional results	5
C	Appendix 3: Additional results for network analyses	23
C.1	Measures	23
C.2	Results	23
D	Appendix 4: Additional results for count models	27
D.1	Methods and measures	27
D.2	Results	28
E	Appendix 5: Missing data	29

APPENDIX 1: REPRODUCING TRADITIONAL RELATIONSHIPS REGARDING POLITICAL RESOURCES

Methods

In this section we examine the traditionally expected patterns of correlations regarding political resources: Do education and internal efficacy propel citizens to become more likely to engage in traditional forms of (offline) political participation?

Our first set of results consists of a simple correlation table to reproduce the results from the numerous studies demonstrating the strong relationships between measures of political engagement and measures of political participation (Nie, Junn, and Stehlik-Barry 1996). We report Spearman correlations to avoid assuming that all measures are necessarily normally distributed. The correlations are based on a dataset grouped at the level of the individual Twitter user and thus do not take the panel aspect into account. We use pairwise deletion to include as many observations as possible for the full table.

Results that are not statistically significant at a .05 level are reported in gray and results that are statistically significant are reported in white.

Results

As expected, there is a strong relationship between internal efficacy and a classical measure of political engagement, namely "offline" participation ($r=0.34$). Higher levels of education are also associated with higher levels of "offline" participation ($r=0.24$). Both education ($r=0.22$) and internal efficacy ($r=0.53$) are also positively correlated with political interest. Further, the figure also shows that both variables correlate positively with *online* forms of political participation, providing the first clue that these factors may also shape political behaviors in social media interactions.

It is also worth noting that income, openness, and extroversion are all significantly and positively correlated with offline participation, which is also in line with previous studies (Mondak 2010; Verba, Schlozman, and Brady 1995).

Generally, the findings demonstrate that this sample of participants is similar to the samples from numerous other studies suggesting that internal efficacy and education are strong predictors of political participation.

FIGURE 1. Correlations with political engagement

APPENDIX 2: ADDITIONAL RESULTS

TABLE 1. Correlations between measures of online hostility and self-reported measures in survey

TABLE 2. Correlations between measures of online hostility and self-reported measures in survey for those self-identifying as Democrats

TABLE 5. Regression table for models for figure 2

[illegible]

*The table illustrates the regression output used to create figure 1 in the main text, for the three dependent variables Political Hostility, Afinn scores and Toxicity scores. Standard errors are in parentheses. # indicates significant at .1, * indicates significant at .05 and ** indicates significant at .01

TABLE 6. Fit statistics for count models for number of followers

Model	LL	BIC
Poisson regression	-562743.898	1125504.106
Negative Binomial Regression	-21443.396	42911.255
Zero Inflated Negative Binomial	-21443.396	42927.564

TABLE 7. Fit statistics for count models for number of friends

Model	LL	BIC
Poisson regression	-782805.547	1565627.38
Negative Binomial Regression	-24141.228	48306.884
Zero Inflated Negative Binomial	-24123.117	48286.948

TABLE 8. Fit statistics for count models for number of tweets

Model	LL	BIC
Poisson regression	-552085.897	1104186.823
Negative Binomial Regression	-12024.56	24071.665
Zero Inflated Negative Binomial	-12024.561	24086.696

TABLE 9. Constructs available in different waves

Construct	Wave1	Wave2	Wave3	Wave4	Wave5	Wave6
Female	yes	yes	yes	yes	yes	yes
Age	yes	yes	yes	yes	yes	yes
Education	yes	yes	yes	yes	yes	yes
Household Income	yes	yes	yes	yes	yes	yes
Internal efficacy	no	no	no	yes	yes	yes
Party identification	yes	yes	yes	yes	yes	yes
Political interest	yes	no	no	no	no	no
Online participation	yes	no	no	no	no	no
Affective polarization	no	no	yes	yes	yes	yes
Openness	no	no	no	no	yes	yes
Conscientiousness	no	no	no	no	yes	yes
Extraversion	no	no	no	no	yes	yes
Agreeableness	no	no	yes	yes	yes	yes
Neuroticism	no	no	yes	yes	yes	yes

TABLE 10. Descriptive statistics for constructs

Construct	1	2	3	4	5	6
Affective polarization _{max}	1	1	1	1	1	1
Affective polarization _{mean}	0.668	0.669	0.674	0.677	0.678	0.679
Affective polarization _{min}	0	0	0	0	0	0
Affective polarization _{sd}	0.225	0.224	0.222	0.227	0.227	0.227
Affective polarization _{count}	1403	1131	1195	1449	1413	1267
Age _{max}	1	1	1	1	1	1
Age _{mean}	0.347	0.411	0.426	0.423	0.437	0.448
Age _{min}	0	0	0	0	0	0
Age _{sd}	0.224	0.219	0.218	0.221	0.219	0.216
Age _{count}	2212	1446	1406	1695	1643	1469
Agreeableness _{max}			1	1	1	1
Agreeableness _{mean}			0.671	0.681	0.682	0.684
Agreeableness _{min}			0	0.111	0	0.074
Agreeableness _{sd}			0.143	0.144	0.15	0.145
Agreeableness _{count}	0	0	1406	1695	1643	1469
Conscientiousness _{max}					1	1
Conscientiousness _{mean}					0.667	0.668
Conscientiousness _{min}					0.029	0
Conscientiousness _{sd}					0.187	0.192
Conscientiousness _{count}	0	0	0	0	1643	1469
Education _{max}	1	1	1	1	1	1
Education _{mean}	0.658	0.676	0.686	0.695	0.692	0.704
Education _{min}	0	0	0	0	0	0
Education _{sd}	0.274	0.27	0.27	0.265	0.266	0.257
Education _{count}	2212	1446	1406	1695	1643	1469

TABLE 11. Descriptive statistics for constructs continued

Construct	1	2	3	4	5	6
Extraversion _{max}					1	1
Extraversion _{mean}					0.492	0.491
Extraversion _{min}					0	0
Extraversion _{sd}					0.225	0.223
Extraversion _{count}	0	0	0	0	1643	1469
Female _{max}	1	1	1	1	1	1
Female _{mean}	0.499	0.486	0.494	0.483	0.481	0.474
Female _{min}	0	0	0	0	0	0
Female _{sd}	0.5	0.5	0.5	0.5	0.5	0.5
Female _{count}	2212	1446	1406	1695	1643	1469
Hate _{max}	0.851	0.801	0.734	1	0.852	0.849
Hate _{mean}	0.482	0.475	0.479	0.493	0.491	0.506
Hate _{min}	0.122	0	0.152	0.135	0.19	0.191
Hate _{sd}	0.09	0.09	0.083	0.093	0.081	0.083
Hate _{count}	1443	934	796	972	916	700
Household income _{max}	1	1	1	1	1	1
Household income _{mean}	0.384	0.399	0.409	0.418	0.416	0.418
Household income _{min}	0	0	0	0	0	0
Household income _{sd}	0.234	0.233	0.233	0.234	0.234	0.237
Household income _{count}	2071	1366	1327	1609	1553	1389

TABLE 12. Descriptive statistics for constructs continued

Construct	1	2	3	4	5	6
Internal efficacy _{max}				1	1	1
Internal efficacy _{mean}				0.768	0.769	0.768
Internal efficacy _{min}				0.056	0	0
Internal efficacy _{sd}				0.207	0.204	0.21
Internal efficacy _{count}	0	0	0	1695	1643	1469
Neuroticism _{max}			1	1	1	1
Neuroticism _{mean}			0.436	0.451	0.445	0.433
Neuroticism _{min}			0	0	0	0
Neuroticism _{sd}			0.218	0.224	0.227	0.221
Neuroticism _{count}	0	0	1406	1695	1643	1469
Online participation _{max}	1					
Online participation _{mean}	0.233					
Online participation _{min}	0					
Online participation _{sd}	0.183					
Online participation _{count}	2212	0	0	0	0	0

TABLE 13. Descriptive statistics for constructs continued

Construct	1	2	3	4	5	6
Openness _{max}					1	1
Openness _{mean}					0.685	0.688
Openness _{min}					0.051	0
Openness _{sd}					0.168	0.168
Openness _{count}	0	0	0	0	1643	1469
Party identification _{max}	1	1	1	1	1	1
Party identification _{mean}	0.631	0.639	0.639	0.643	0.639	0.649
Party identification _{min}	0	0	0	0	0	0
Party identification _{sd}	0.353	0.359	0.36	0.359	0.362	0.351
Party identification _{count}	2212	1446	1406	1695	1643	1469
Politica hate _{max}	1	0.887	0.836	0.975	0.841	0.891
Politica hate _{mean}	0.513	0.516	0.517	0.535	0.533	0.551
Politica hate _{min}	0.082	0.115	0.103	0	0.123	0.115
Politica hate _{sd}	0.109	0.113	0.105	0.112	0.103	0.107
Politica hate _{count}	1443	934	796	972	916	700
Political interest _{max}	1					
Political interest _{mean}	0.746					
Political interest _{min}	0					
Political interest _{sd}	0.283					
Political interest _{count}	2212	0	0	0	0	0
Political _{max}	0.926	0.892	0.858	0.894	1	0.925
Political _{mean}	0.487	0.499	0.497	0.511	0.511	0.526
Political _{min}	0.094	0.079	0.045	0	0.094	0.135
Political _{sd}	0.107	0.116	0.108	0.109	0.111	0.106
Political _{count}	1443	934	796	972	916	700

TABLE 14. Descriptive statistics for number of tweets in the panel

Median tweets	Mean tweets	Min tweets	Max tweets	Sd tweets
845.0	2539.166	1.0	136146.0	6380.274

FIGURE 2. Density plot for number of tweets in the panel **FIGURE 3. Density plot for number of tweets in the panel excluding those with more than 10,000 tweets**

FIGURE 4. RI-CLPM for toxicity exploring the robustness of the cutpoints

FIGURE 5. RI-CLPM for AFINN exploring the robustness of the cutpoints

FIGURE 6. RI-CLPM for Hate exploring the robustness of the cutpoints **FIGURE 7. Distributions for Afinn before and after removing outliers**

FIGURE 8. Distributions for political before and after removing outliers

FIGURE 9. Distributions for toxicity before and after removing outliers

FIGURE 10. Distributions for hate before and after removing outliers

FIGURE 11. RI-CLPM estimates for the reciprocal effects of hate and political across those high and low on affective polarization

The results represent averages of models estimated using cutpoints in the range [0.5;0.65]. Full results can be seen in table

FIGURE 12. RI-CLPM estimates for the reciprocal effects of hate and toxicity across those high and low on affective polarization

The results represent averages of models estimated using cutpoints in the range [0.5;0.65]

TABLE 15. Output for figures 3 and 11-12

Parameter	Estimate	Standard error	N	Group	Dependent variable	Parameter	5% CI	95 % CI	Estimate	N
Online hostility → Online hostility	0.114	0.014	2377.0	HIGH	Afinn	Online hostility → Online hostility	0.09	0.16	0.12	1409.0
Political → Online hostility	0.052	0.016	2377.0	HIGH	Afinn	Online hostility → Political	-0.03	0.02	-0.0	1409.0
Online hostility → Political	0.029	0.01	2377.0	HIGH	Afinn	Political → Online hostility	0.01	0.12	0.07	1409.0
Political → Political	0.112	0.013	2377.0	HIGH	Afinn	Political → Political	0.11	0.18	0.15	1409.0
Online hostility → Online hostility	0.114	0.012	2812.0	HIGH	Political hate	Online hostility → Online hostility	0.09	0.18	0.13	1166.0
Political → Online hostility	0.06	0.024	2812.0	HIGH	Political hate	Online hostility → Political	0.01	0.1	0.05	1166.0
Online hostility → Political	0.016	0.007	2812.0	HIGH	Political hate	Political → Online hostility	0.0	0.08	0.04	1166.0
Political → Political	0.124	0.011	2812.0	HIGH	Political hate	Political → Political	0.11	0.2	0.15	1166.0
Online hostility → Online hostility	0.088	0.013	2785.0	HIGH	Toxicity	Online hostility → Online hostility	0.11	0.18	0.14	1424.0
Political → Online hostility	0.081	0.028	2785.0	HIGH	Toxicity	Online hostility → Political	0.0	0.05	0.03	1424.0
Online hostility → Political	0.0	0.007	2785.0	HIGH	Toxicity	Political → Online hostility	-0.04	0.06	0.01	1424.0
Political → Political	0.141	0.011	2785.0	HIGH	Toxicity	Political → Political	0.09	0.15	0.12	1424.0
				LOW	Afinn	Online hostility → Online hostility	-0.04	0.08	0.02	1409.0
				LOW	Afinn	Online hostility → Political	-0.03	0.05	0.01	1409.0
				LOW	Afinn	Political → Online hostility	-0.01	0.19	0.09	1409.0
				LOW	Afinn	Political → Political	0.2	0.33	0.26	1409.0
				LOW	Political hate	Online hostility → Online hostility	0.09	0.24	0.16	1166.0
				LOW	Political hate	Online hostility → Political	-0.12	0.04	-0.04	1166.0
				LOW	Political hate	Political → Online hostility	-0.01	0.12	0.06	1166.0
				LOW	Political hate	Political → Political	0.12	0.26	0.19	1166.0
				LOW	Toxicity	Online hostility → Online hostility	0.06	0.17	0.12	1424.0
				LOW	Toxicity	Online hostility → Political	-0.02	0.06	0.02	1424.0
				LOW	Toxicity	Political → Online hostility	0.06	0.22	0.14	1424.0
				LOW	Toxicity	Political → Political	0.17	0.29	0.23	1424.0

FIGURE 13. The moderated effect of political resources on various measures of online hostility

TABLE 16. Output for figures 4-6. The output tables are, from left to right: Political hate, Afinn scores and Toxicity scores

Parameter	Cut point	Estimate	Standard error	N	Parameter	Cut point	Estimate	Standard error	N	Parameter	Cut point	Estimate	Standard error	N
Hate → Hate	0,0,7	0.117	0.014	2378.0	Afinn → Afinn	-0.55,0.75	0.087	0.013	2792.0	Toxicity → Toxicity	0,0.55	0.113	0.013	2810.0
Political → Hate	0,0,7	0.035	0.011	2378.0	Political → Afinn	-0.55,0.75	0.059	0.021	2792.0	Political → Toxicity	0,0.55	0.043	0.019	2810.0
Hate → Political	0,0,7	0.042	0.015	2378.0	Afinn → Political	-0.55,0.75	0.003	0.009	2792.0	Toxicity → Political	0,0.55	0.018	0.009	2810.0
Political → Political	0,0,7	0.116	0.013	2378.0	Political → Political	-0.55,0.75	0.154	0.011	2792.0	Political → Political	0,0.55	0.135	0.011	2810.0
Hate → Hate	0.01,0.69	0.117	0.014	2378.0	Afinn → Afinn	-0.54,0.74	0.084	0.013	2789.0	Toxicity → Toxicity	0,0.56	0.109	0.012	2810.0
Political → Hate	0.01,0.69	0.035	0.011	2378.0	Political → Afinn	-0.54,0.74	0.058	0.021	2789.0	Political → Toxicity	0,0.56	0.049	0.019	2810.0
Hate → Political	0.01,0.69	0.042	0.015	2378.0	Afinn → Political	-0.54,0.74	0.004	0.009	2789.0	Toxicity → Political	0,0.56	0.019	0.009	2810.0
Political → Political	0.01,0.69	0.116	0.013	2378.0	Political → Political	-0.54,0.74	0.155	0.011	2789.0	Political → Political	0,0.56	0.136	0.011	2810.0
Hate → Hate	0.02,0.68	0.117	0.014	2378.0	Afinn → Afinn	-0.53,0.73	0.086	0.013	2786.0	Toxicity → Toxicity	0,0.57	0.11	0.012	2811.0
Political → Hate	0.02,0.68	0.036	0.011	2378.0	Political → Afinn	-0.53,0.73	0.058	0.021	2786.0	Political → Toxicity	0,0.57	0.05	0.019	2811.0
Hate → Political	0.02,0.68	0.042	0.015	2378.0	Afinn → Political	-0.53,0.73	0.001	0.009	2786.0	Toxicity → Political	0,0.57	0.018	0.009	2811.0
Political → Political	0.02,0.68	0.116	0.013	2378.0	Political → Political	-0.53,0.73	0.157	0.011	2786.0	Political → Political	0,0.57	0.136	0.011	2811.0
Hate → Hate	0.03,0.67	0.115	0.014	2378.0	Afinn → Afinn	-0.52,0.72	0.086	0.013	2785.0	Toxicity → Toxicity	0,0.58	0.111	0.012	2811.0
Political → Hate	0.03,0.67	0.037	0.011	2378.0	Political → Afinn	-0.52,0.72	0.062	0.021	2785.0	Political → Toxicity	0,0.58	0.05	0.019	2811.0
Hate → Political	0.03,0.67	0.04	0.014	2378.0	Afinn → Political	-0.52,0.72	0.001	0.009	2785.0	Toxicity → Political	0,0.58	0.02	0.009	2811.0
Political → Political	0.03,0.67	0.117	0.013	2378.0	Political → Political	-0.52,0.72	0.155	0.011	2785.0	Political → Political	0,0.58	0.136	0.011	2811.0
Hate → Hate	0.04,0.66	0.116	0.014	2378.0	Afinn → Afinn	-0.51,0.71	0.087	0.013	2783.0	Toxicity → Toxicity	0,0.59	0.112	0.012	2811.0
Political → Hate	0.04,0.66	0.039	0.012	2378.0	Political → Afinn	-0.51,0.71	0.059	0.022	2783.0	Political → Toxicity	0,0.59	0.052	0.018	2811.0
Hate → Political	0.04,0.66	0.037	0.014	2378.0	Afinn → Political	-0.51,0.71	0.0	0.009	2783.0	Toxicity → Political	0,0.59	0.021	0.009	2811.0
Political → Political	0.04,0.66	0.117	0.013	2378.0	Political → Political	-0.51,0.71	0.156	0.011	2783.0	Political → Political	0,0.59	0.136	0.011	2811.0
Hate → Hate	0.05,0.65	0.114	0.014	2377.0	Afinn → Afinn	-0.5,0.7	0.089	0.013	2783.0	Toxicity → Toxicity	0,0.6	0.115	0.013	2811.0
Political → Hate	0.05,0.65	0.042	0.012	2377.0	Political → Afinn	-0.5,0.7	0.056	0.022	2783.0	Political → Toxicity	0,0.6	0.054	0.018	2811.0
Hate → Political	0.05,0.65	0.036	0.014	2377.0	Afinn → Political	-0.5,0.7	-0.001	0.009	2783.0	Toxicity → Political	0,0.6	0.022	0.009	2811.0
Political → Political	0.05,0.65	0.117	0.013	2377.0	Political → Political	-0.5,0.7	0.156	0.011	2783.0	Political → Political	0,0.6	0.138	0.011	2811.0
Hate → Hate	0.06,0.64	0.115	0.014	2377.0	Afinn → Afinn	-0.49,0.69	0.096	0.013	2781.0	Toxicity → Toxicity	0,0.61	0.11	0.013	2812.0
Political → Hate	0.06,0.64	0.043	0.012	2377.0	Political → Afinn	-0.49,0.69	0.053	0.022	2781.0	Political → Toxicity	0,0.61	0.052	0.018	2812.0
Hate → Political	0.06,0.64	0.035	0.013	2377.0	Afinn → Political	-0.49,0.69	0.0	0.009	2781.0	Toxicity → Political	0,0.61	0.025	0.009	2812.0
Political → Political	0.06,0.64	0.117	0.013	2377.0	Political → Political	-0.49,0.69	0.157	0.011	2781.0	Political → Political	0,0.61	0.136	0.011	2812.0
Hate → Hate	0.07,0.63	0.11	0.014	2377.0	Afinn → Afinn	-0.48,0.68	0.1	0.013	2781.0	Toxicity → Toxicity	0,0.62	0.109	0.013	2812.0
Political → Hate	0.07,0.63	0.048	0.013	2377.0	Political → Afinn	-0.48,0.68	0.06	0.022	2781.0	Political → Toxicity	0,0.62	0.041	0.018	2812.0
Hate → Political	0.07,0.63	0.033	0.013	2377.0	Afinn → Political	-0.48,0.68	0.001	0.009	2781.0	Toxicity → Political	0,0.62	0.024	0.009	2812.0
Political → Political	0.07,0.63	0.118	0.013	2377.0	Political → Political	-0.48,0.68	0.157	0.011	2781.0	Political → Political	0,0.62	0.134	0.011	2812.0
Hate → Hate	0.08,0.62	0.116	0.014	2377.0	Afinn → Afinn	-0.47,0.67	0.102	0.013	2780.0	Toxicity → Toxicity	0,0.63	0.107	0.013	2812.0
Political → Hate	0.08,0.62	0.055	0.013	2377.0	Political → Afinn	-0.47,0.67	0.06	0.022	2780.0	Political → Toxicity	0,0.63	0.045	0.018	2812.0
Hate → Political	0.08,0.62	0.034	0.012	2377.0	Afinn → Political	-0.47,0.67	0.0	0.009	2780.0	Toxicity → Political	0,0.63	0.023	0.009	2812.0
Political → Political	0.08,0.62	0.122	0.013	2377.0	Political → Political	-0.47,0.67	0.158	0.011	2780.0	Political → Political	0,0.63	0.135	0.011	2812.0
Hate → Hate	0.09,0.61	0.116	0.014	2377.0	Afinn → Afinn	-0.46,0.66	0.104	0.013	2776.0	Toxicity → Toxicity	0,0.64	0.108	0.013	2812.0
Political → Hate	0.09,0.61	0.054	0.014	2377.0	Political → Afinn	-0.46,0.66	0.074	0.022	2776.0	Political → Toxicity	0,0.64	0.04	0.018	2812.0
Hate → Political	0.09,0.61	0.033	0.012	2377.0	Afinn → Political	-0.46,0.66	0.005	0.009	2776.0	Toxicity → Political	0,0.64	0.023	0.009	2812.0
Political → Political	0.09,0.61	0.121	0.013	2377.0	Political → Political	-0.46,0.66	0.163	0.011	2776.0	Political → Political	0,0.64	0.134	0.011	2812.0
Hate → Hate	0.1,0.6	0.117	0.014	2377.0	Afinn → Afinn	-0.45,0.65	0.105	0.013	2776.0	Toxicity → Toxicity	0,0.65	0.11	0.013	2812.0
Political → Hate	0.1,0.6	0.056	0.014	2377.0	Political → Afinn	-0.45,0.65	0.075	0.023	2776.0	Political → Toxicity	0,0.65	0.039	0.017	2812.0
Hate → Political	0.1,0.6	0.034	0.011	2377.0	Afinn → Political	-0.45,0.65	0.007	0.009	2776.0	Toxicity → Political	0,0.65	0.024	0.01	2812.0
Political → Political	0.1,0.6	0.12	0.013	2377.0	Political → Political	-0.45,0.65	0.163	0.011	2776.0	Political → Political	0,0.65	0.134	0.011	2812.0

APPENDIX 3: ADDITIONAL RESULTS FOR NETWORK ANALYSES

This appendix presents additional results for the network part of the analysis. First, we use a set of descriptive measures to elaborate on Figure 4. Second, we elaborate on the “network centrality” discussed in the main text.

Measures

We use the *density*, *diameter*, and the *largest component* to measure network *cohesion*. Density refers to (Kolaczyk and Csárdi 2014, 55) “the frequency of realized edges relative to potential edges,” where edges simply refers to whether there is an “arrow” pointing to or from a person (“vertice”) in our Twitter panel. This measure is scaled to range from 0 to 1, where a higher score indicates that a person has more edges and therefore a denser network. The *diameter* of the graph measures the longest possible path between two vertices (Kolaczyk and Csárdi 2014); if the path is longer, more people in the network are connected. And, finally, we use the *largest component* to extract the largest number of connected Twitter friends/followers among those low and high on political hate: The larger this metric, the more connected the network is. We scale it to represent the *proportion* out of the total number of vertices. If social media users cluster together in echo chambers, we should observe multiple “centers of gravity” in the networks, such that those who are low on political hostility should have one or more centers of gravity and those who are high on political hostility should have one or more centers of gravity as well. To investigate this, we calculate so-called “hub scores” and “authority scores,” which indicate how influential a given vertice is. Loosely speaking, an authority is a vertice with many arrows pointing towards it and a hub is a vertice that points to many authorities (Kolaczyk and Csárdi 2014). We present these graphs and analyses in the appendix.

Results

The density of the network is much higher for those who are politically hateful ($6.59e-4$) compared to those low on political hate ($5.59e-5$). The diameter – that is, the longest path – is also higher for those

high on political hate, and the largest component among those who are high on political hate amounts to roughly 10 % of the total number of vertices, whereas it only accounts for roughly 1 % among those low on political hate. All in all, these results suggest that politically hateful Twitter users often reside in networks that are densely populated and diverse (in terms of hate), and where large audiences can potentially be exposed to their hateful tweets.

TABLE 1. Measures of connectedness for networks of political hate followers

Density	Largest component share of total	Diameter	Level of political hate
6.59e-04	0.09	9.0	High
8.20e-05	0.01	3.0	Low

We reach the same conclusion when turning to the graph in Figure 1. It illustrates those who are the hubs and authorities in the network. Here, the coloring of the vertices – where darker colors indicate more hateful Twitter users – makes abundantly clear that those who are either hubs or authorities are all relatively high on political hate. The ten highest hub scores are shown, along with their score on political hate (Tables 2 and 3). It thus seems that those who exert influence in the network are those with higher scores on political hate.

FIGURE 1. Network visualization of hubs and authorities in the network containing friends and followers who are also survey respondents. The color indicates the level of political hate and the size of the nodes represent their authority/hub score. The layout is based on the Kamada-Kawai algorithm (Kamada, Kawai, et al. 1989). The grey color represents the connections ("edges") between the network.

TABLE 2. Top ten hub scores

Political hate score	Hub score
0.66	1.0
0.67	0.66
0.54	0.4
0.59	0.34
0.55	0.24
0.65	0.24
0.67	0.23
0.66	0.19
0.7	0.17
0.62	0.17

TABLE 3. Top ten authority scores

Political hate score	Authority score
0.66	1.0
0.66	0.8
0.69	0.62
0.65	0.59
0.54	0.46
0.72	0.4
0.66	0.32
0.55	0.32
0.71	0.31
0.71	0.3

APPENDIX 4: ADDITIONAL RESULTS FOR COUNT MODELS

Some relevant evidence already exists for the assertion that hostile users are also influential users. Hateful users have been found to form cohesive networks that connect them to each other (Mathew et al. 2019; Ribeiro et al. 2017). Furthermore, such users occupy central positions within their networks and their social media posts are shared further in the network (Ribeiro et al. 2017; Mathew et al. 2019). Finally, Ribeiro et al. (2017) have found that hateful users tend to tweet more than typical users and have more friends, but fewer followers. Importantly, none of these studies investigates *political* hostility and, hence, it is unclear whether such findings travel to the domain of politics.

Methods and measures

The models do not take the time aspect into account but are classical count models. To investigate whether those who are more politically hostile also tweet more and have more friends/followers, we estimated, for each dependent variable, a Poisson regression, a negative binomial regression, and a zero-inflated negative binomial regression. The model with the lowest BIC was chosen (Muthén, Muthén, and Asparouhov 2017); see Appendix 1 Tables 6-7. All analyses were implemented in Mplus version 8. Before running the analyses we removed the five percent most extreme counts for each dependent variable separately in order to not make the analyses hinge on a few extreme observations.⁶

For the network analyses, we used *igraph* (Csardi and Nepusz 2006) to analyze and graph the results. If those who are most and least hostile tend to cluster together with their like-minded peers, we should observe relatively distinct networks – that is, echo chambers. This has two implications. First, it would imply that our measures of network cohesion are quite high if we break down our network into those high and low on political hate and investigate the measures separately for these two groups. Accordingly, we split our Twitter panelists into two equally sized groups based on their level of political hate. For each group, we calculate our descriptive measures of network cohesion. We also present more formal measures of “network centrality” in Appendix 3.

6. For the investigation of whether those who are politically hostile also tweet more, we only used the respondents from the first wave and only investigate the year 2019 since we were scraping these respondents continually in this period and there is thus not a problem with “exposure.”

Results

Counts of tweets, friends, and followers The first set of results concerns the count models. Table 4 gives results from models that regress participants' number of followers, friends, and tweets on our measure of political hostility. These models show that political hate predicts both the number of followers and number of tweets and, to a lesser extent, number of friends.

TABLE 4. Negative binomial regression models for political hate

Model	parameters	Number of followers	Number of tweets	Number of friends
Count model	Political Hate	2.575	3.663	1.320
		[2.147;3.003]*	[3.105;4.221]*	[0.961;1.679]*
	Intercept	4.109	4.159	5.458
		[3.886;4.332]*	[3.860;4.457]*	[5.269;5.647]*
Inflation model	Political Hate			-11.022
				[-14.317;-7.728]*
	Intercept			0.180
				[-0.976;1.336]
	Number of individuals	3479	3439	1835

In order to illustrate the impact of political hate on the number of tweets, followers, and friends, Figure 4 calculates predicted values when varying the level of political hate above and below two standard deviations from the mean. The effect of political hate is very large. As the left panel of Figure 4 shows, changing from -2SD to +2SD on political hate means an increase in the number of tweets from around 150 to around 1150 – an increase of roughly seven and half times the number of tweets. The middle panel shows that increased political hate is associated with an increase in the number of followers from around 120 to around 420; that is, roughly 3.5 times the number of followers. As the right panel shows, the effect of political hate on the number of friends is more modest, with an increase from roughly 300 to 600 – a doubling of the number of friends.

These results thus stress that those who are politically hostile are also likely to be influential in the network insofar as they both tweet more and likely also reach more followers (and slightly more friends) in their Twitter feeds.

FIGURE 2. Predicted counts as a function of political hate

APPENDIX 5: MISSING DATA

Missing values are a source of potential bias in the estimation of any statistical parameter. The classical solution of simply “ignoring” this source of bias by dropping the missing values (listwise deletion) is probably the worst solution to this problem (Enders 2010). Our main goal is thus to minimize the bias introduced by missing values, while at the same time acknowledging that no perfect solutions to this problem exist. Using listwise deletion rests on the assumption of missing completely at random (MCAR), whereas newer approaches toward handling missing values rest on the assumption of missing at random (MAR), which is less restrictive. Simplifying somewhat, MCAR means that missing values are completely random, and thus e.g. not influenced by levels of education or other measured or unmeasured factors, whereas MAR rests on the assumption that missing values of a variable are not related to the values of that variable after taking into account (measured) variables that are predictors of missingness (Enders 2010).

Two main techniques exist to deal with missing values: full information maximum likelihood (FIML) and multiple imputation (MI) (Enders 2010), each with its own advantages and challenges. For most models we use FIML, since this is the most efficient in terms of having the lowest standard errors, but for the results in Figure 1 the number of missing values is quite large for some constructs, and we

thus run into convergence problems. We therefore use MI for these analyses.

Before the final multiple imputation, we used autocorrelation plots to investigate at which point the data become independent ⁷. We use the default convergence settings in Mplus for MI and created 50 datasets. Each regression is then run for each dataset and the results averaged automatically in Mplus. These 50 datasets contained all variables used across all models for Figure 1, and we are thus using the same data across all models and only do one MI for this figure.

We use MI/FIML to deal with unit nonresponse, missing data across waves, and item-nonresponse (a person missing data within a wave for a given item).

7. For a fuller exposition of the steps used in MI see (Enders 2010, Chapter6-10)