

A report from the Center on Technology and Society

OCT
2018

Computational Propaganda, Jewish-Americans and the 2018 Midterms:

The Amplification of Anti-Semitic
Harassment Online

Our Mission:

To stop the defamation of the Jewish people and to secure justice and fair treatment to all.

ABOUT

CENTER FOR TECHNOLOGY & SOCIETY AND THE BELFER FELLOWS

In a world riddled with cyberhate, online harassment, and misuses of technology, the Center for Technology & Society (CTS) serves as a resource to tech platforms and develops proactive solutions. Launched in 2017 and headquartered in Silicon Valley, CTS aims for global impacts and applications in an increasingly borderless space. It is a force for innovation, producing cutting-edge research to enable online civility, protect vulnerable populations, support digital citizenship, and engage youth. CTS builds on ADL's century of experience building a world without hate and supplies the tools to make that a possibility both online and off-line.

The Belfer Fellowship program supports CTS's efforts to create innovative solutions to counter online hate and ensure justice and fair treatment for all in the digital age through fellowships in research and outreach projects. The program is made possible by a generous contribution from the Robert A. and Renée E. Belfer Family Foundation. The inaugural 2018-2019 Fellows are:

- **Rev. Dr. Patricia Novick, Ph.D., of the Multicultural Leadership Academy, a program that brings Latino and African-American leaders together.**
- **Dr. Karen Schrier, an associate professor at Marist College and its founding director of the Games and Emerging Media Program.**
- **Samuel Woolley of the Oxford Internet Institute at the University of Oxford, who works with Jigsaw, Google's think tank.**

ADL is a leading anti-hate organization that was founded in 1913 in response to an escalating climate of anti-Semitism and bigotry. Today, ADL is still the first call when acts of anti-Semitism occur and continues to fight all forms of hate. A global leader in exposing extremism, delivering anti-bias education and fighting hate online, ADL's ultimate goal is a world in which no group or individual suffers from bias, discrimination or hate.



Samuel Woolley is the research director of the Digital Intelligence (DigIntel) Lab at IFTF, a research associate at the Oxford Internet Institute, and an associate member of Green Templeton College at the University of Oxford. He specializes in the study of automation/AI, political communication, and information warfare. He is also a co-founder and former research director of the Computational Propaganda (ComProp) research team at the University of Oxford and the University of Washington. He and his collaborators have done foundational research on the topics of online disinformation and political manipulation, coining the terms 'computational propaganda' and 'political bot.'

Computational Propaganda, Jewish-Americans and the 2018 Midterms:

The Amplification of Anti-
Semitic Harassment Online

The Amplification of Anti-Semitic Harassment Online

Jeff, a Jewish-American reporter at a major U.S. news outlet, first experienced anti-Semitic political attacks online during the 2016 U.S. presidential election. Soon after the election, he began to receive harassing messages over social media that used religious slurs or featured photoshopped images of him containing violent or graphic content. Each time one of his stories got traction online—or featured details on topics such as white nationalism, Donald Trump, or libertarianism—he was sent photoshopped images of his face in a gas chamber or was threatened with the public release of his address and contact details. “It has become something that I expect to happen,” he said, “I don’t even think about it anymore.” For Jeff, the 2016 election was a major impetus for allowing the online sphere, and public platforms including Twitter and Reddit, to become openly hostile places for him and people like him.

The number of Jewish people living in the United States is estimated to be between 4.2 million and 12 million, the wide range due to religious versus ethnic distinctions (Steinhardt Social Research Institute, 2016; DellaPergola, 2017). For many of these individuals—especially those in the public eye—social media platforms have become inhospitable for both general communication and as forums for discussing public life. This report explores the ways in which online propaganda, harassment and political manipulation are affecting Jewish People in the runup to 2018 U.S. midterm elections. In the course of our research, members of this group have described a marked rise in the number of online attacks their community is experiencing. This is proving especially true during electoral contests and major political events. Correspondingly, our analyses suggests that tools like social media bots, and tactics including doxing, disinformation, and politically-motivated threats, have been used online during the 2018 midterms to target Jewish Americans. According to interviewees, veiled human users—rather than automated accounts—often deliver the most worrisome and harmful anti-Semitic attacks.

Anonymity and automation are integral features of computational propaganda—the use of algorithms over social media in attempts to manipulate public opinion. Anonymity allows the people who spread digital disinformation and political harassment to do so without fear of reprisal or repercussion. Automation, often in the form of social media bots or automated profiles that look like real users, allows these same individuals to scale their offensives. Both anonymity and automation have been used in online propaganda offensives against the Jewish community during the 2018 midterms. During this contest, political bots—which explicitly focus on political communication online—are playing a significant role in artificially amplifying derogatory content over Twitter about Jewish people. Human users, however, still accounted for the majority of derogatory Twitter traffic. People used the protective power of anonymity over a variety of social media platforms to spread harmful or misleading content about Jewish American people. Many human-led efforts had features of organized propaganda campaigns and made use of twitter bombing—barraging hashtags associated with the Jewish community with highly politicized, and sometimes hateful, content in an effort to demobilize, coopt and interrupt normal communication and organization over social media.

Each of the Jewish American experts interviewed for this study has experienced online attacks in 2018. These respondents stated that many of these onslaughts came from what they considered extremist groups in American politics: the emergent alt-right, white nationalists and neo-Nazi organizations. Those interviewed spoke about the rise of an emboldened anti-Semitic community online and consistently correlated this rise with the election of Donald Trump.

The following report is an analysis of computational propaganda, the Jewish American community, and the 2018 elections. As part of the wider paper series focused on “humanizing the effects of computational propaganda” this empirical work details the ways in

which the Jewish socio-religious population in the U.S. is being disproportionately targeted with disinformation and abuse during this crucial political moment. We use a mixed methods approach in this research, deploying both qualitative and quantitative analysis in order to generate both a culturally deep and statistically broad understanding of how computational propaganda is being leveraged against this community. Interviews with five prominent Jewish Americans reveal themes in the ways in which disinformation and political attacks flow against and within the community. Analysis of 7,512,594 tweets over a period from August 31, 2018 to September 17, 2018 shows the prevalence of political bots in these efforts and highlights groups within the U.S. political spectrum most involved in anti-Semitic attacks. In fact, as many as 30 percent of the accounts messaging using derogatory terms gathered in this data set appear to be highly automated. In the final section, we discuss the implications of our research, as well as policy suggestions for social media platforms, governmental actors, and civil society.

Literature Review

Following revelations about the role of computational propaganda in the 2016 election, there has been an undisputed rise in white supremacist activities and overt anti-Semitism (Woolley & Guilbeault, 2017; Astor, 2018). From 2016 to 2017, the number of established neo-Nazi groups increased from 99 to 121 (SPLC, 2018); twice as many hate-motivated murders were committed by white supremacists (Baynes, 2018); and there was a 258% increase in the number of white supremacist propaganda incidents on college campuses (ADL, 2018). While not all white supremacist groups consider themselves anti-Semitic, anti-Semitism is often a core tenet of white supremacy and, by extension, white nationalism and neo-Nazism (Ferber, 1999). As such, it comes as little surprise that 1,986 anti-Semitic incidents—harassment, vandalism, and assault—occurred in 2017 (“Audit of Anti-Semitic Incidents,” 2018). The 57% increase in such events was the largest escalation in a single year since the Anti-Defamation League (ADL) first began recording incidents in 1979. Schools, from kindergarten through to high school, were the most common locations of anti-Semitic incidents, following a 94% increase. In the words of ADL CEO Jonathan Greenblatt: “Kids repeat what they hear. And so in an environment in which prejudice isn’t called out by public figures, figures of authority, we shouldn’t be surprised when we see young people repeat these same kind of tropes” (Cohen, 2018).

A staggering expansion of online harassment coincided with, and arguably fomented, the increase in offline anti-Semitism. Fringe Internet communities, such as 4chan, 8chan, and Gab allowed for the propagation of such ideas, which quickly spread to Twitter, Reddit, and other mainstream online communities (Glaser, 2017; Marwick & Lewis, 2017). An analysis of over 100 million posts on Gab and 4chan’s Politically Incorrect message board (/pol/) found that, between July 2016 and January 2018, the use of the terms “Jew” and “kike,” a derogatory term for Jewish people, more than doubled on /pol/ and dramatically increased on Gab (Finkelstein et al., 2018). Spikes also occurred in the use of both terms following President Trump’s inauguration and the 2017 Unite the Right Rally in Charlottesville, which is believed to be the largest white supremacist rally in the United States in at least a decade (ADL, 2017).

On Twitter, Jewish journalists have faced an onslaught of online persecution and trolling. Between August 2015 to July 2016, a study featuring 800 journalists found that they received 19,253 “overtly anti-Semitic tweets,” with ten prominent Jewish journalists receiving 83% of the tweets (ADL, 2016). The top five most common words in the harassing accounts’ bios were: Trump, conservative, white, nationalist, and America. Two journalists in particular were heavily targeted: Julia Ioffe and Jonathan Weisman. Ioffe was doxxed¹ by Andrew Anglin, the founder of the world’s biggest neo-Nazi website The Daily Stormer, following the April 2016 publication of her profile on Melania Trump in GQ magazine (Gambino, 2016; O’Brien, 2017). She received threatening, graphic tweets that depicted her face superimposed upon the face of a concentration camp prisoner, and tweets, emails, and phone calls that referenced ovens, lamp shades, coffins, homicide cleanup crews, and Hitler speech recording, among other sordid anti-Semitic threats (Wemple, 2016).

The harassment of Weisman was triggered by a tweet from a Twitter account, @CyberTrump. Which said: “hello, (((Weisman)))” (Weisman, 2016). When asked to explain the three sets of parentheses, @CyberTrump wrote: “It’s a dog whistle, fool. Belling the cat for my fellow goyim.” The torrent of hate that followed—including a photo of Weisman in a gas chamber with a smiling Donald Trump, in Nazi uniform, flipping the switch—brought national attention to a Google Chrome extension called the Coincidence Detector. The extension, drawing from a database of user-generated Jewish names and others deemed “anti-white,” placed an (((echo))) around these

1

Merriam-Webster Dictionary defines “dox” as: “to publicly identify or publish private information about (someone) especially as a form of punishment or revenge” (<https://www.merriam-webster.com/dictionary/dox>)

names when they appeared on the Internet, visible only to those who have the extension (Fleishman & Smith, 2016b). The echo originated from the podcast, *The Daily Shoah*, of *The Right Stuff*, a neo-Nazi blog (Fleishman & Smith, 2016a). The symbol is meant to highlight the reach of “Jewish power” by showing that “all Jewish surnames echo throughout history” (Fleishman & Smith, 2016a).

Harassment of Jewish journalists continued in October of 2016 when attendees at a Trump rally in Cleveland, Ohio, chanted “Lügenpresse,” the German Nazi slur for “lying press” (Fleishman, 2016). Two days later, Trump supporters began #TheList on Twitter—a compilation of journalists who “speak out against Donald Trump, for Hillary Clinton, or other forms of Kikery” (Fleishman, 2016). Journalists were tweeted images with large, red X’s on their faces, alerting them that they had been placed on #TheList due to “their crimes against the American people.” As noted by Mic, which reported on #TheList, the effort “play[ed] into the age-old conspiracy theory of Jewish collusion to control the world’s powerful elite.” Indeed, one user on 8chan, where the #TheList was created, wrote: “Name ‘em and shame ‘em. I look forward to seeing plenty of echoed names” (Fleishman, 2016).

The themes of this online harassment against the Jewish American community, especially against journalists and prominent members of the group, have been carried from the 2016 U.S. presidential election to the 2018 midterm contest. Anti-Semitism is also being carried into offline mainstream politics, with neo-Nazi and holocaust denier Arthur Jones running as the GOP candidate for a congressional seat in the Illinois Third District (Korecki, 2018). The following sections detail the methods used to analyze the prevalence of computational propaganda against the Jewish American population during the 2018 midterms. We explain the use of a mixed method approach to develop understandings of this problem, using both computational and interview-based study.

CyberTrump wrote: “It’s a dog whistle, fool. Belling the cat for my fellow goyim.” The torrent of hate that followed—including a photo of Weisman in a gas chamber with a smiling Donald Trump, in Nazi uniform, flipping the switch—brought national attention to a Google Chrome extension called the Coincidence Detector. The extension, drawing from a database of user-generated Jewish names and others deemed “anti-white,” placed an (((echo))) around these names when they appeared on the Internet,

visible only to those who have the extension (Fleishman & Smith, 2016b). The echo originated from the podcast, *The Daily Shoah*, of *The Right Stuff*, a neo-Nazi blog (Fleishman & Smith, 2016a). The symbol is meant to highlight the reach of “Jewish power” by showing that “all Jewish surnames echo throughout history” (Fleishman & Smith, 2016a).

Methodology

The intention of this report is to better understand the use of online political harassment and disinformation about Jewish Americans during the 2018 U.S. midterm elections. We conducted both interviews and data analysis of tweets gathered in order to understand both the scope of the issue on national scale and the repercussions faced on the individual level. As such, the analyses section of this report is arranged in two parts.

Part A discusses the findings from interviews with five Jewish Americans who are involved in American politics as elected officials, policy makers, journalists, political consultants, and commentators. We analyze the data from our interviews via the creation of thematic memos. During interviews we took intensive timestamped notes. We then used these notes to create empirically-driven interview research memos that reflect upon the experience of the interviewer and stated experience of the interviewee; expose core themes within the discussion of computational propaganda against and within the Jewish community during the 2018 midterms; and highlight paths for continued research in this area.

Part B consists an analysis of 7,512,594 tweets. Selecting for specific hashtags (see Table 4), we collected tweets relating to U.S. politics. Due to constraints upon the collection mechanism, Tweepy—a library of the coding language Python used to access the Twitter API, the tweets were gathered in groups between August 31, 2018 and September 17, 2018. The hashtags were categorized by political leaning: conservative, liberal, extremist, and neutral (e.g. “#vote”). The hashtags were purposively gathered using markers from previous and ongoing research on Twitter conversations (Kollanyi et al., 2016; Woolley and Guilbeault, 2016; Woolley, 2018). We worked to be non-partisan in our selection of hashtags and analyses of data in order to produce the most objective results possible, though we accept the impossibility of true positivism in social scientific research. As this is a study of how Jewish Americans are discussed in the

general Twitter debate concerning U.S. politics, the majority of hashtags studied are associated with liberal and conservative issues, and a minority relate to hate groups, white supremacy and white nationalism.

Tweets were then filtered based upon whether or not the text of the tweet contained a series of terms related to Judaism and/or anti-Semitism. Instances of term use in hashtags, usernames, and shared links were not included. The terms (see Table 5) were categorized as: derogatory, lean derogatory, context dependent (meaning they could be derogatory or not depending on the context), lean context dependent, neutral (e.g. Jew, Orthodox, Israeli), and other, which consisted of derogatory terms historically used by Jews to describe other ethnicities, non-Jewish individuals, and to criticize other Jews (e.g. kushi, shiksa, kapo). The accounts that tweeted five or more derogatory or lean derogatory terms during the time period were then run through Botometer to ascertain whether or not they were automated. Botometer is a machine learning dashboard that works to classify Twitter bots, created by the Observatory on Social Media (OSoME), a joint project through the Network Science Institute (IUNI) and Center for Complex Networks and Systems Research (CNetS) at Indiana University. It builds upon a previous iteration of the tool, known as BotorNot, to scan accounts for automation using a variety of measures and signals and is arguably the academic industry standard in Twitter bot detection (Davis et al, 2016).

Analysis

A. Qualitative: Interviews with Expert Representatives of the Community

Despite the diversity of the five interview subjects in profession, age, denomination of Judaism practiced, and perspective on Israel, there were several consistent themes that came up. First, there is a consistent pattern of harassment against the Jewish American community by extremists: white supremacists, neo-Nazis, and aspects of the fragmented group generally known as the alt-right. Some of this harassment is bleeding into, and being taken up by, more mainstream U.S. conservatism. There was little evidence from interviews pointing to left-wing anti-Semitism. Second, much of the trolling and disinformation about Jewish Americans relies upon archaic stereotypes and conspiracies. These outmoded,

and often hateful, clichés are commonly used to stoke fear and mistrust between other U.S. minority groups and American Jewish people. Third, there are platform specific differences in how harassment happens and what it looks like. This is due not only to trolls' apparent obsession with anonymity, but also due to the demographics and affordances of particular platforms.

The interview subjects stated that, while they were familiar with the use of bots in spreading online propaganda, they were more concerned—and had more frequently experienced—human-based attacks on social media. Three of the five interview subjects, those most in the public eye, have been doxxed by prominent white supremacist and neo-Nazi leaders. All three attacks occurred following the publication of a news article. The two interviewees who work as politicians, one previously a candidate and the other about to take office, had their personal information and addresses leaked by known white nationalists following positive profiles on them and their work in the Jewish Telegraphic Agency (JTA). They told us that anti-Semitic individuals and groups seemed to be following reports from JTA and using them to choose Jewish American leaders to attack online. Angie, the soon to be elected 2018 political candidate, was doxxed by white supremacist leader David Duke. “I was getting random hate messages and tweets and didn't realize it was because [Duke] had re-tweeted [the JTA profile],” she said. She illuminated the tactics of barraging her with hate messages: “I still don't have that many followers, so I think they figured it would be overwhelming for me and it was.”

Ellen, a former political candidate, was similarly outed by Don Black. The Stormfront creator called her a “disgusting hissing weasel” and told his supporters to attack her. The next morning, her Gmail, Facebook, Instagram, and voicemail were flooded with thousands of messages. They used horrible slurs and contained threats of physical and sexual violence. There were dozens of such messages in her email. Because her phone number had been released she received a cascade of similarly harassing voicemails. The FBI got involved and police had to patrol the area surrounding her house during the election cycle. Nearly two years later, she continues to receive many threatening messages from trolls online. Although the number of messages has decreased, the graphic nature of the content has remained intense.

The other interviewees highlighted the fact that they get attacked online after they say or write anything public—this is especially true if the writing contains discussion of Judaism or anti-Semitism. Mark, a prominent economist and writer, said that his attackers “are not automated, they are [a] non-bot twitter mob.” He made it clear that although many of the accounts used messages and tactics similar to those of computational propagandists, further inspection revealed them to be real people. They weren’t tweeting at computationally enhanced levels, as bots often do, but they seemed to be organized as groups rather than as individuals. Zane, a journalist, was doxxed on 4chan after he interviewed an attendee of the Unite the Right 2 Rally in 2018 and wrote the first name and home city of his interview subject. He was subsequently attacked via his digital wedding guestbook, where users wrote threats and used references to Hitler and the Holocaust (e.g. offering to buy him an oven for his cremation).

While each interview subject spoke of not wanting to let threats of the trolls impact their online activity, political campaigns, academic research or news reporting, they all admitted the threats of violence and deluges of anti-Semitism had become part of their internal equations. For some, it drove them to speak out louder and more vigorously, defying the trolls; for others, often citing concern over the harassment of family members, friends and romantic partners, sought to make adjustments. Although we only spoke to a sample of five individuals, it was clear that, although anti-Semitic harassment has become almost normalized and expected following the election of Donald Trump, it has a chilling effect on Jewish Americans’ involvement in the public sphere.

B. Quantitative

Analysis of hashtags: First, it is evident that conservative hashtags were the main conduits of political conversation in our sample of 7,512,594 tweets (8,183,545 hashtags). Although we collected data from fewer conservative hashtags—94, as opposed to 106 liberal hashtags (Table 4)—conservative hashtags were tied to more tweets than all of the other categories combined, including six times the number of tweets linked to liberal hashtags (Table 1). Thus, it comes as no surprise that of the ten most popular hashtags in our sample, 80.26 percent of the related tweets pertained to conservative causes:

#MAGA (2,300,281; conservative); #QAnon (989,277; conservative); #WWG1WGA (745,311; conservative); #Trump (739,803; conservative); #WalkAway (724,495; conservative); #Resist (540,973; liberal); #KAG (524,752; conservative); #FBRParty (341,096; liberal); #VoteThemOut (311,341; liberal); and, #TheResistance (288,067; liberal). Although, it is worth noting that while #Trump and #WalkAway (as in walk away from the liberal party) are categorized as conservative hashtags, a minority of liberals and other non-conservatives use them.

Significantly, the top ten hashtags reflect divide in discourse between supporters and opponents of President Trump. While #MAGA (Make America Great Again), Trump’s 2016 presidential campaign slogan, and #KAG (Keep America Great), Trump’s 2020 motto, were extremely prolific, producing 37.6% of the total tweets in the sample, their predominance is conventional. The abundance of #QAnon and #WWG1WGA is far more surprising. QAnon, is a conspiracy theory that emerged on the 4chan message board /pol/ on October 28, 2017. An anonymous figure, “Q,” named in reference to Q-level security clearance (Department of Energy authorization to access “Top Secret” information) posted a cryptic message reference Hilary Clinton and National Guard, and then a series of questions about Trump (Coaston, 2018). QAnon is based on the idea that Trump is in control of everything, and he is bringing “the storm” to disrupt the “deep state” history of U.S. presidential involvement in a global criminal empire hellbent on pedophilia (Coaston, 2018). There are strong anti-Semitic undertones, as followers decry George Soros and the Rothschild family as puppeteers. While there are several oft-repeated refrains, including “trust the plan,” “the great awakening,” “follow the white rabbit,” and “walk away,” the central saying is “where we go one, we go all” (WWG1WGA).

There is coordination surrounding the most popular liberal hashtags as well. #FBRParty stands for “follow back resistance,” and the hashtag is a somewhat codified way for liberals, progressives, and other members of “the resistance,” to both find other members and promote messages on a large scale. Users who share the hashtag are added to mass lists, and all members follow each other and tweet to gain more followers, creating a chain of message promotion and coordination. Members of this movement often have an emoji of a crashing blue wave in their username or the body of their tweet to

signify their affiliation with the democratic party. Due to the speed at which they add each followers with these “follow back tweets,” individuals’ abilities to follow are occasionally slowed by Twitter or their accounts are taken down. Although not covered in this analysis, this chain of communication appears likely to have been targeted by malicious bot attacks.

Breaking down the sample by category of hashtag, tweets containing conservative hashtags consisted of 58.09 percent of the entire conversation, liberal hashtags tweets made up 31.01 percent, tweets with neutral hashtags were 9.51 percent, extremist hashtags comprised 0.88 percent. These percentages are notable for the sake of understanding our sample, but due to an arguably biased sampling of hashtags, they are not a perfect representation of American political conversation on Twitter. For instance, tweets were gathered from only 20 extremist hashtags, several of which were overly niche (#ProudOfYourBoy had 4 tweets and #Hammerskins had 1 tweet).

Significantly, the top ten hashtags reflect divide in discourse between supporters and opponents of President Trump. While #MAGA (Make America Great Again), Trump’s 2016 presidential campaign slogan, and #KAG (Keep America Great), Trump’s 2020 motto, were extremely prolific, producing 37.6% of the total tweets in the sample, their predominance is conventional. The abundance of #QAnon and #WWG1WGA is far more surprising. QAnon, is a conspiracy theory that emerged on the 4chan message board /pol/ on October 28, 2017. An anonymous figure, “Q,” named in reference to Q-level security clearance (Department of Energy authorization to access “Top Secret” information) posted a cryptic message reference Hilary Clinton and National Guard, and then a series of questions about Trump (Coaston, 2018). QAnon is based on the idea that Trump is in control of everything, and he is bringing “the storm” to disrupt the “deep state” history of U.S. presidential involvement in a global criminal empire hellbent on pedophilia (Coaston, 2018). There are strong anti-Semitic undertones, as followers decry George Soros and the Rothschild family as puppeteers. While there are several oft-repeated refrains, including “trust the plan,” “the great awakening,” “follow the white rabbit,” and “walk away,” the central saying is “where we go one, we go all” (WWG1WGA).

There is coordination surrounding the most popular liberal hashtags as well. #FBRParty stands for “follow back resistance,” and the hashtag is a somewhat codified way for liberals, progressives, and other members of “the resistance,” to both find other members and promote messages on a large scale. Users who share the hashtag are added to mass lists, and all members follow each other and tweet to gain more followers, creating a chain of message promotion and coordination. Members of this movement often have an emoji of a crashing blue wave in their username or the body of their tweet to signify their affiliation with the democratic party. Due to the speed at which they add each followers with these “follow back tweets,” individuals’ abilities to follow are occasionally slowed by Twitter or their accounts are taken down. Although not covered in this analysis, this chain of communication appears likely to have been targeted by malicious bot attacks.

Breaking down the sample by category of hashtag, tweets containing conservative hashtags consisted of 58.09 percent of the entire conversation, liberal hashtags tweets made up 31.01 percent, tweets with neutral hashtags were 9.51 percent, extremist hashtags comprised 0.88 percent. These percentages are notable for the sake of understanding our sample, but due to an arguably biased sampling of hashtags, they are not a perfect representation of American political conversation on Twitter. For instance, tweets were gathered from only 20 extremist hashtags, several of which were overly niche (#ProudOfYourBoy had 4 tweets and #Hammerskins had 1 tweet).

Analysis of terms within hashtags:

Observing term occurrence by hashtag category, extremist tweets were more likely than any other category to contain derogatory and lean derogatory terms, as well as more likely to contain terms that can be derogatory depending on the context (Table 1). Given that words such as “shoah,” “nazi,” and “oy vey” are categorized as context-dependent, it is unlikely that these words are being used in non-derogatory ways on extremist channels of communication. In fact, these traditionally Jewish words, have been co-opted by extremists to mock Jewish people. For example, anti-Semites often write tweets saying things such as, “Oy vey the goyim know. Shut it down!” They also

TABLE 1: TERM PREVALENCE BY HASHTAG CATEGORY

Hashtag Category	Total Tweets	Derogatory + Lean Derogatory (%)	Context Dependent + Lean Dependent (%)	Neutral (%)	Other (%)	Contain Any Term (%)
Conservative	4802207	0.82	1.65	0.32	0.01	2.79
Liberal	786225	0.53	1.26	0.16	0.006	1.95
Neutral	2563308	2.02	0.89	0.13	0.01	3.05
Extremist	72406	2.51	11.86	0.66	0.03	15.06

often write in poor imitation of Yiddish accent, ridiculing the Holocaust by claiming small inconveniences to be “annuda shoah” (or another shoah). Notably, around 15 percent of extremist tweets contained any term relating to anti-Semitism or Judaism. Given that the many extremists, especially white nationalists, target a number of minority groups and dedicate energy to condemning Antifa, liberals, and safe spaces, this is an interesting metric.

Certain hashtags appear to be correlated with the prevalence of specific terms. On a more general level, the ten hashtags that contained the highest percentages of tweets with derogatory or lean derogatory words were: #ReligiousRight (22.81%; conservative); #SCOTUS (5.63%; neutral); #NWO (5.60%; extremist); #WhiteGenocide (4.21%; extremist); #Libertarian (3.63%; conservative); #LiarInChief (3.29%; liberal); #FollowTheWhiteRabbit (2.55%; conservative); #Dems (2.3%; liberal); and, #TrumpTrain (2.21%; conservative). Both #ReligiousRight and #Uhrh had very few tweets (114 and 275 tweets respectively), so the large percentage of derogatory or lean derogatory words is more likely due to chance. Oddly the only term that was used in #ReligiousRight tweets was “NWO” (New World Order), which relates to a conspiracy that elites, often Jewish elites, are going to submit the entire world to servitude under totalitarian governance. It is tied to The Protocols of the Elders of Zion, an anti-Semitic treatise that was published in Russia in early 1900s. “NWO” was also by far the most common term found in #SCOTUS tweets, while “illuminati” and “globalist” (which is often an anti-Semitic slur) were the most popular terms in #NWO. For #WhiteGenocide, there was a spread of terms, but the most common was unsurprisingly “non-white.” Interestingly, far and away the most popular term used in #LiarInChief tweets was “soap.” Because

extremists use soap in reference to the Holocaust, it is classified as lean derogatory, but upon further inspection of the term usage in #LiarInChief tweets, it is often used in reference to soap operas, likening Donald Trump’s presidency to a scripted televised drama. As for #Dems, the most common words were “Aryan” and “Nazi.” Not surprisingly, most accounts used these terms to accuse Trump and other republicans of being Nazis or obsessed with Aryan ideals, not espousing those values themselves. Lastly, the most popular term, by one to two orders of magnitude, used in #TrumpTrain tweets was “Soros,” in reference to George Soros, and the second most popular term was “NWO.” Many of the tweets are in reference to QAnon or George Soros as a “puppeteer” paying for Antifa, Hilary Clinton, or any other person challenging far-right conservatism or extremists

Looking at the relationship between specific terms and hashtags (Table 6), the five terms that were most prevalent within specific hashtags were: “Nazi” (71.80% of #ProudBoys tweets, extremist hashtag); “Jew” (56.92% of #JCOT tweets, conservative hashtag); and “Hitler” (44.14% of #GoodbyeDemocrats tweets, conservative hashtag).

Automated “bot” accounts:

To assess whether or not automated accounts impacted the U.S. political conversation surrounding Jewish Americans, particularly regarding the use of demeaning language, we analyzed accounts that tweeted five or more tweets during the collection time period that contained derogatory or lean derogatory terms. In all, 3733 accounts satisfied the requirements, but Botometer only returned the results for 3,060 accounts. It is unknown what happened to the 727 missing accounts. Although only speculative, it is possible they were removed by Twitter.

Botometer provides a number of scores, including the “universal score,” which is a language-independent CAP – essentially, the probability that the account is automated. We classified accounts with a universal score of greater than .5 as highly automated or botlike, and any accounts with lower scores as likely human. Previous work measuring computational propaganda using BotoMeter suggests that the program is more likely to rank bot accounts as false negatives than false positives. Accounts with a score of 50% or more bot-like demonstrate high degrees of automated behavior, though it is important to note that the distinction between bot (or automated account) and human account exists on a scale rather than as a binary (Woolley and Guilbeault, 2017). Of accounts that tweeted five or more derogatory or lean derogatory terms during the collection time period, 28.14 percent of the accounts are likely to be automated and 67.45 percent of the accounts are likely to be human (Table 2). Automated accounts produced 43.14 percent of tweets in this category and human accounts produced

56.86 percent. Preliminary examination suggests that between 30 and 40 percent of the accounts using derogatory terms were bots.

Interestingly, humans are more likely to tweet all categories of terms relating to Jewish Americans, including derogatory and lean derogatory terms (Table 3). This finding is in alignment with the qualitative interviews in Part A. Each of the interview subjects said that their most significant interactions with harassment were either due to doxxing or were from accounts that did not act or look like automated accounts. This is deeply concerning. If this problem was driven largely by automated accounts, there would be avenues for recourse, such as improved detection and dismantlement of bots, that would decrease the level of harassment. Given that this is largely a human problem, the path to a safe internet and public space for all is far more complicated, calling into question laws regarding hate speech and constitutional rights.

TABLE 2: AUTOMATION OF ACCOUNTS

Category	Number of Accounts	Percentage of Total (%)	Number of Tweets
Likely Bot	861	28.14	426023
Likely Human	2199	67.45	561433

TABLE 3: COMPARING TERM USAGE BY AUTOMATED AND HUMAN ACCOUNTS

Category	Total Tweets	Derogatory + Lean Derogatory (%)	Context Dependent + Lean Dependent (%)	Neutral (%)	Other (%)	Contain Any Term (%)
Bot	426023	2.52	2.52	2.52	0.00	5.69
Human	561433	4.05	4.05	4.05	0.00	6.81

Conclusion

The online public sphere—now a primary arena for communication about American politics—has become progressively inhospitable for Jewish Americans. Prior to the election of President Donald Trump, anti-Semitic harassment and attacks were rare and unexpected, even for Jewish Americans who were prominently situated in the public eye. Following his election, anti-Semitism has become normalized and harassment is a daily occurrence. The harassment, deeply rooted in age-old conspiracies such as the New World Order, which alleges that an evil cabal of Jewish people have taken autocratic control of the globe, and Holocaust imagery—faces placed inside Nazi concentration camp ovens or stretched on lampshades—shows no signs of abating. Unfortunately, the more minority or vulnerable groups one identifies with (e.g. Jewish Latina), the more targeted one becomes. The harassment is largely perpetrated by individuals, as opposed to mass automated “bot” networks, and hypocritically, these trolls are often obsessed with anonymity, lashing out with doxxing attacks when they feel their secret identities are threatened. In the words of one interview subject: “They themselves would not be able to stand the level of cruelty and scrutiny that they treat others with. They are afraid of being identified as someone who acts in a hateful way.” The platforms are key facilitators of this anti-Semitic harassment.

Twitter and Facebook: In our qualitative research, every interview subject said that harassment was worse on Twitter than on Facebook. On Facebook, they were sent harassing messages, but because private pages are very difficult to infiltrate, they were able to take more control of such spaces. When they were sent harassing messages, they were often sent from seemingly real accounts and did not contain the anti-Semitic vitriol that is often found in tweets. One interviewee said that when she was targeted by Fox News viewers, these users more often reached out to her on Facebook as opposed to Twitter and were “less focused on [her] as a filthy Jewess, [her] as a hissing weasel, or images of [her] lying dead with Nazis and SS.” Instead they sent messages containing “a lot of unbelievable sexism.” They called her “stupid, dumb, waste of air,” and said her “mother should have gotten an abortion.” Their trolling relied more upon patriarchal and sexist ideas, conservative values, and her inadequacy due to being a “libtard.”

On Twitter, on the other hand, many interviewees spoke of massive, coordinated, attacks by trolls—sometimes at the behest of white nationalist or hate-group leaders such as Andrew Anglin and David Duke—that were impossible to filter or staunch. They all felt that they had to be on Twitter, as politicians, journalists, public facing academics, and consultants, but that Twitter did not make it easy for them to exist in that ecosystem peacefully. The ease of attack was highlighted several times—with minimal effort, an anonymous harasser could mention one in a tweet or comment on a post, and without forewarning or consent, the target will receive an automatic notification and be subjected to disturbing imagery or threats. As one interviewee remarked: “Twitter does an awful job. An awful, awful, awful job policing discourse on the site. It is ideally designed for organized or non-organized harassment. They haven’t reckoned with it.”

Recommendations for the platforms: Our quantitative research shows that human-run accounts make up the majority of anti-Semitic accounts on Twitter. This means that simply removing bot accounts is insufficient. Twitter needs to make changes to their interface and harassment response efforts if they desire to improve user experience. One recommendation is to allow users more nuanced control over who can see their tweets. For many public-facing individuals, changing their profiles to private is not feasible, but neither is blocking each and every troll. One of our interview subjects mentioned that ever since he changed location in his bio to Germany (although he is based in the U.S.), he has not been harassed. He suspected that due to stricter data rules in Germany, Twitter either wasn’t showing him malicious tweets, or wasn’t showing his tweets to anti-Semitic profiles.

The later hypothesis is especially interesting in that it evades some constitutional rights issues, as the user is willfully limiting his own freedom of speech. A second recommendation is to allow for greater ease in filtering notifications and direct messages from unfamiliar accounts. Although one can filter notifications, the process is not intuitive and can be time intensive. A third recommendation that arose from the interviews, is the proposal to block accounts that actively facilitate trolling. One such example is the conservative “news” website Twitchy, which serves to curate Tweets and call

for readers to attack or support the tweets, which are imbedded for ease of interaction. Although Twitchy is more mainstream conservative and is thus not typically anti-Semitic, one interviewee remarks that “It is an organized hate platform, directing hate and mobs towards people, and that is an inevitable part of how it functions.” He recommended that the Twitter accounts of people who work at Twitchy should not be allowed to operate.

Lastly, a fourth recommendation came from an interviewee who was egregiously doxxed on Instagram. Facing thousands of violent and disturbing messages, she needed an immediate solution to stem the flow of messages, but she did not know whom to contact. She ended up reaching out a friend who worked in the fashion department to help her deal with the torrent of harassment. She remarked: “There are some of us who are made the poster children by these trolls and there needs to be a chain of command that we can go to counteract trolling during a deluge.” Although the experience occurred on Instagram, it is reasonable to apply these findings to Twitter and all social media platforms. Platforms need clearer mechanisms, that involve human facilitators and not just automated or online complaint systems, for identifying serious harassment and trolling.

Recommendations for civil society: Social media platforms face mounting scrutiny from government, civil society, and users. There are numerous opportunities for negotiation and reshaping of social media platforms—both internally, through self-regulation, and externally, through government regulation and civil society pressure. Employees of technology companies and regular social media users also have power to create change, but they need help to organize and communicate concerns and desired changes. Groups like Coworker.org, Tech Solidarity, the Tech Workers Coalition, and the Center for Human Technology are working to organize and give voice to tech workers, but they face challenges in connecting with employees due to strict non-disclosure agreements and company cultures that penalize and isolate those that speak out.

Governments, social media platforms, and website-hosting companies can and should hold hateful and defamatory websites accountable. Following the Daily Stormer’s mockery of Heather Heyer’s death at the Unite the Right Rally in Charlottesville, Virginia in 2017, GoDaddy and Google Domains both evicted the Daily Stormer for violating the companies’ terms of service (Mettler & Selk,

2017). Several social media firms have moved to ban Infowars, a known conspiracy and hate site, from their platforms. One of our interviewees argued that 4chan and 8chan, websites on which many mass harassment attacks are coordinated and anti-Semitic memes generated, should also be more heavily scrutinized. He said that companies and platforms hosting hateful or harassing content “should be a liability [for web hosts], just like it is to have a neo-Nazi publication.”

Our interviewees suggested that the business models of many social media platforms incentivize the companies to allow disinformation and harassment. They highlighted problematic political advertising practices, which can be gamed by bots and taken advantage of by anonymous groups. Interviewees also underscored the idea that companies seem reluctant to remove both bot-driven and harassing content not only because they are cautious about regulating speech, but also because they are beholden to share-holders and thus do not want to effect user growth or the overall number of users on sites—metrics used to determine company success and worth. One interview subject said, “the obligation is on the tech platforms not to specifically engineer themselves to promote harassment, not so much that they have the duty to prevent harassment.”

Governments, including the German parliament, have instituted monetary consequences if harassment and hate speech are left on platforms for particular periods of time. While massive fines for hosting hate-speech or propaganda may not get to the heart of the problem—i.e. a need for systematic, and ethically-minded, redesign of trending algorithms and features that prioritize easily co-opted anonymity and automation—they do place pressure on the companies to take responsibility as curators of content. Rather than fines for harmful online content, governments should set sensible benchmarks for social media companies that work towards instituting socially and technologically informed changes to gameable trending algorithms, advertising mechanisms, defamation policies, and bot-policies. When benchmarks are not met they should result in fines or other penalties.

At present regulation is slow in coming. It is also complicated by technical nuance and a need for systematic redesign of certain aspects of social media. It is important, with this in mind, for civil society to aid both technology company works and general users in organizing and negotiating for positive democratic

change. Journalists, in particular, are well-positioned to negotiate changes to Twitter. Twitter is a crucial tool for journalists, but journalists are also integral to the fabric of Twitter and produce much of the high-quality content on that platform. Journalists are under constant harassment on Twitter, but they are also extremely valuable to the Twitter landscape. This unique position provides opportunity for organization and negotiation.

State sponsored trolling of journalists and democratic figures using social media is a new form of human rights abuse, and national and international law should treat it as such (Monaco & Nyss, 2018). State linked accounts used for attacking or defaming the public must be expeditiously detected, identified and deleted. As Monaco and Nyss note:

For many people around the world, the Internet has become the key medium through which their free speech rights can be exercised. The weaponization of information in the form of state sponsored trolling attacks thus constitutes an interference with individuals' right to freedom of expression and opinion, which encapsulates a right not only to impart, but also to seek and receive, information and ideas of all kinds, regardless of frontiers. (pp. 49-50)

There is a challenge in using national law to hold governments accountable for trolling that they perpetuate and sponsor. In these circumstances international law and corporate policy must work to protect democracy. According to Monaco and Nyss, international human rights law requires States to take action to prohibit "hate speech". They write, "article 20 of the International Covenant on Civil and Political Rights states that 'any advocacy of national, racial or religious hatred that constitutes incitement to discrimination, hostility or

violence shall be prohibited by law'." International law must work to hold non-compliant states accountable both for state-perpetuated trolling but also for hate speech and defamation that occurs online, spread by any number of groups, during everyday moments but also throughout elections, security crises and other political events.

Recommendations for combatting bots: To combat the use of bots in amplifying both online hate speech and disinformation companies must develop a labeling system for accounts that demonstrate high-levels of automation. It is true that it would be impossible to delete all social bots on Twitter, let alone on the internet—where they make up as much as 50 percent of all online traffic (Incapsula, 2016). It is also true that not all bots are used for nefarious purposes. In fact, most bot accounts online and on social media are used for routine tasks and web searches. Automation of Twitter accounts is then, not a metric of defamatory use in and of itself, but it would be helpful for users and researchers to know what accounts (or pages) on a site like Twitter or Facebook demonstrate high levels of automation and why because bots can and are used to amplify harmful political attacks and hate speech.

Beyond this, companies and researchers must do more work to identify and score bot-driven, or highly automated, political accounts on Twitter—but also on sites from Reddit to YouTube to Facebook and beyond—with a higher degree of confidence. When automation is used to inflate social media metrics of political candidates, social issues/perspectives or, simply put, hate, it can give widespread illusion of popularity to ideas or people that really have no backing. We call this automatically amplified support manufactured consensus and initial research suggests that this phenomena could have myriad implications for bandwagon politics and reporting of disinformation or misinformation in the news media. Also, the usage of social media bots to suppress, attack or spam activist or opposition hashtags is parallel problem that must be addressed.

Society—but more specifically governments and technology firms who regulate and self-regulate automation and harmful speech online—must ask several questions about the political use of bots: Is it within the bounds of democratic communication to allow a person or small group of people to use hundreds or thousands

of accounts to amplify their perspective? What about one person using one or two accounts to automatically send thousands of political or derogatory messages a day? If this activity is mostly consigned to well-resourced actors—who have time, money and computational skill—then what are the implications for fairness during elections? Overall, what are the implications, when regulating bot communications, for free speech? In considering this latter question, Lamo and Calo (2018) point out that:

Ultimately, bots represent a diverse and emerging medium of speech. Their use for mischief should not overshadow their novel capacity to inform, entertain, and critique. We conclude by urging society to proceed with caution in regulating bots, lest we inadvertently curtail a new, unfolding form of expression.

This said, the manipulative usage of bots on social media remains a problem—especially in the political sphere and increasingly in conversations about vaccination/healthcare, climate science and the business section—and must be addressed.

It is time for technology companies to design their products for democracy. This does not just mean a facile attempt to protect or prioritize free speech. It means protecting minority ethnic and religious groups, which are integral to successful democracy but also disproportionately the targets of computational propaganda. Social media companies cannot escape responsibility by claiming not to be “arbiters of truth”. They are curators of information, their algorithms prioritize which content and ads people see. This includes which news and what political communication citizens engage with during elections. It is time for companies to inject ethics, and—more strongly—human rights, into the heart of product design.

TABLE 4: CATEGORIZATION OF HASHTAGS

Category	Hashtags
<p>Extremist (20)</p>	<p>"#GothRight", "#NewRight", "#AltRight", "#ItsOkToBeWhite", "#proudboys", "#POYB", "#POYG", "#ProudBoysMedia", "#proudboy", "#ProudBoysGirls", "#proudfyourboy", "#8Chan", "#Uhuru", "#WhiteGenocide", "#nazis", "#Hammerskins", "#WhitePower", "#WhitePride", "#NewWorldOrder", "#NWO"</p>
<p>Conservative (94)</p>	<p>"#TheDonald", "#DonaldTrump", "#TrumpTrainPortal", "#BuildTheWall", "#KeepAmericaGreat", "#JewsForTrump", "#MAGA", "#trump", "#Republicans", "#GOP", "#Libertarian", "#Anarchists", "#tcot", "#silentmajority", "#NRA", "#RedWave", "#RedWaveRising", "#RedWave2018", "#RedWaveRising2018", "#KAG2018", "#KAG2020", "#VoteRedToSaveAmerica", "#TrumpTrain", "#RedNation", "#TrumpNation", "#TrumpsArmy", "#FakeNewsMedia", "#FakeNewsCNN", "#Redhat", "#GoodbyeDemocrats", "#PresidentTrump", "#potus", "#LatinosForTrump", "#BlacksForTrump", "#KAG", "#MakeAmericaGreatAgain", "#Trump2020", "#KAG2020", "#VoteRed2020", "#TrumpSupporters", "#WalkAway", "#RedNationRising", "#VoteRedOrAmericalsDead", "#snowflake", "#AmericaFirst", "#draintheswamp", "#VoteRed", "#VoteDemsOut", "#fuckHillary", "#freealexjones", "#CrookedHillary", "#Neoliberals", "#TeamTrump", "#ReligiousRight", "#WWG1WGA", "#VoteDemOut2018", "#liberalmedia", "#MediaBias", "#EnemyOfThePeople", "#PromisesKept", "#ccot", "#tlot", "#TPOT", "#GayConservative", "#NeverVoteDemocratAgain", "#ItsOkToBeMale", "#MuellerWitchHunt", "#PJNET", "#BetterDeadThanRed", "#QAnon", "#Q", "#QArmy", "#TheStormIsHere", "#Qanon8chan", "#FollowTheWhiteRabbit", "#EnjoyTheShow", "#WeThePeople", "#DeepStateInPanic", "#QanonArmy", "#TrustThePlan", "#TrustSessions", "#TheGreatAwakening", "#TheStorm", "#ResistanceUnited", "#PedoGate", "#InsigniaGate", "#PatriotsUnited", "#jcot", "#FuckAntifa", "#Molonlabe", "#Soros", "#GothsforTrump", "#TeaParty", "#Teabaggers"</p>
<p>Liberal (106)</p>	<p>"#Dem", "#Dems", "#Liberals", "#ProgressDems", "#DNC", "#BlueWave2018", "#AntiFas", "#FightingFascism", "#takemoneyoutofpolitics", "#AWAG", "#TrumpKnew", "#TrumpDerangementSyndrome", "#TrumpCult", "#ImpeachTrump", "#Impeach45", "#VOTEBLUE", "#VoteBlueToSaveAmerica", "#VoteBlue2018", "#votebluenomatterwho", "#Resist", "#theResistance", "#Cult45", "#TRE45ON", "#UnfitToBePresident", "#FoxFakeNews", "#TrumpLies", "#complicitgop", "#corruptGOP", "#PutinsPuppet", "#RussiaGate", "#PutinsPoodle", "#RussianPuppet", "#KremlinAnnex", "#TrumpRussiaConspiracy", "#TrumpRussiaCollusion", "#crazytown", "#TrumpRussia", "#ProtectMueller", "#MarchForTruth", "#WhatsAtStake", "#ProtectOurCare", "#TakeltBack", "#CultureOfCorruption", "#RedToBlue", "#riggedGOP", "#NeverAgain", "#BlueWave", "#NeverTrump", "#NotMyPresident", "#Progressive", "#FatNixon", "#TrumpColluded", "#25thAmendment", "#LunaticInChief", "#LockTrumpUp", "#GOPCorruption", "#GPOTraitors", "#VoteOutGOP", "#StrongerTogether", "#dividedwefall", "#FuckTrump", "#CultureOfCorruption", "#DemForce", "#GangOfPuppets", "#TrumpTreason", "#TrumpLiesMatter", "#LiarInChief", "#DumpTrump", "#TraitorInChief", "#FakePresident", "#TrumpResign", "#TrumpCrimeFamily", "#illegitimatePOTUS", "#OurRevolution", "#Treason", "#ctl", "#p2", "#UniteBlue", "#AlternativeFacts", "#CommonSenseGunLaws", "#MAGAts", "#Deplorables", "#FuckTheNRA", "#AggressiveProgressives", "#tytlive", "#SecondCivilWar", "#ANTIFA", "#ElizabethWarren", "#StillBerning", "#feelthebern", "#HindsightIs2020", "#FollowBackResistance", "#FBR", "#FBRparty", "#BlueTsunami", "#BlueWave2018", "#BlueTsunami2018", "#BlueWaveComing2018", "#VoteThemOut", "#VoteThemOut2018", "#MuellerTime", "#1Voice", "#OnEveryCorner", "#jlot", "#WakeUpAmerica", "#UnionStrong"</p>
<p>Neutral (13)</p>	<p>"#vote", "#PrimaryDay", "#2018Midterms", "#Nov6", "#1A", "#2A", "#GodBlessAmerica", "#jewishamerican", "#MuellerInvestigation", "#Mueller", "#WhiteHouse", "#SCOTUS", "#WomensRights"</p>

TABLE 5: CATEGORIZATION AND OCCURRENCE OF TERMS

Term	Label	Count
soros	context dependent	36916
nazi	context dependent	31735
nwo	lean derogatory	27318
globalist	context dependent	20084
jew	neutral	15312
hitler	context dependent	13216
israeli	neutral	8445
globalism	context dependent	7984
kek	derogatory	7258
heil	derogatory	6928
soap	lean derogatory	4656
jewish	neutral	4153
apartheid	lean dependent	4120
holocaust	context dependent	3940
kiki	derogatory	3856
aryan	derogatory	3796
zionist	neutral	2517
illuminati	derogatory	2420
semitic	neutral	2389
rothschild	lean derogatory	2024
shekel	context dependent	1936
anti-semitic	neutral	1787
alt-right	lean dependent	1783
new world order	lean derogatory	1578
neo-nazi	lean dependent	1491
concentration camps	lean dependent	1385
jq	lean derogatory	989
goy	context dependent	948
1488	derogatory	699
revisionist	lean derogatory	686
glue	lean derogatory	606
yid	context dependent	588
moch	derogatory	558
kushi	other	423
anti-white	derogatory	372
kapo	other	350
heeb	derogatory	347

TABLE 5: CATEGORIZATION AND OCCURRENCE OF TERMS (CONT.)

Term	Label	Count
antisemitic	neutral	306
cultural marxism	lean derogatory	297
echoes	context dependent	266
orthodox	neutral	236
judas	lean derogatory	235
kike	derogatory	168
kekistan	derogatory	160
shlomo	derogatory	157
gentile	neutral	151
white genocide	derogatory	125
shiksa	other	77
zios	derogatory	74
goyim	context dependent	67
kosher	neutral	65
jews will not replace us	derogatory	64
jewess	derogatory	58
diaspora	neutral	39
shoah	context dependent	30
cultural enrichment	lean derogatory	27
dogwhistle	dependent	26
non-jew	neutral	24
cultural marxists	lean derogatory	18
hymie	derogatory	17
zionazi	derogatory	17
kapos	other	16
oy vey	context dependent	16
gas chambers	lean dependent	14
jewish question	lean derogatory	8
idolaters	other	7
labor camps	lean dependent	6
jewboy	derogatory	5
jewboy	derogatory	5
shylock	derogatory	3
haredi	neutral	3
conservative jew	neutral	2
mocky	derogatory	2
reform jew	neutral	2
marrano	derogatory	1
alt-light	lean derogatory	1

TABLE 6: TERM PREVALENCE BY HASHTAG

Hashtag	Term	Category of Term	Term Prevalence in Total Tweets (%)
#Dems	Aryan	Derogatory	2.05
#IsOkaytoBeWhite	KEK	Derogatory	0.79
#AmericaFirst	KEK	Derogatory	0.80
#ReligiousRight	Kiki	Derogatory	2.63
#ResistanceUnited	Kiki	Derogatory	1.98
#NewWorldOrder	Illuminati	Derogatory	3.67
#Libertarian	Illuminati	Derogatory	3.47
#Snowflake	Heil	Derogatory	0.93
#VoteRed2020	Globalism	Dependent	6.49
#KAG2018	Globalism	Dependent	1.42
#AmericanFirst	Globalism	Dependent	1.36
#TrumpTrain	Globalism	Dependent	1.07
#Soros	Globalist	Dependent	7.37
#8Chan	Globalist	Dependent	6.68
#PedoGate	Globalist	Dependent	2.01
#NWO	Globalist	Dependent	1.58
#ProudBoys	Nazi	Dependent	71.80
#ProudBoy	Nazi	Dependent	16.83
#Antifas	Nazi	Dependent	13.49
#LiberalMedia	Nazi	Dependent	10.95
#Antifa	Nazi	Dependent	9.03
#Nazis	Nazi	Dependent	8.99
#WhitePower	Nazi	Dependent	5.61
#GoodbyeDemocrats	Hitler	Dependent	44.14
#Nazis	Hitler	Dependent	7.13
#RedWaveRising	Hitler	Dependent	2.52
#NewWorldOrder	Soros	Dependent	21.24
#NeverTrump	Soros	Dependent	18.64
#Soros	Soros	Dependent	8.71
#Antifa	Soros	Dependent	3.40
#JewishResistance	Jew	Neutral	77.16
#JCOT	Jew	Neutral	56.92
#IfNotNow	Jew	Neutral	41.67
#LiarInChief	Soap	Lean Derog.	2.77
#ReligiousRight	NWO	Lean Derog.	20.18
#Uhuru	NWO	Lean Derog.	6.18
#SCOTUS	NWO	Lean Derog.	5.39
#NewWorldOrder	Rothschild	Lean Derog.	2.58
#NWO	Rothschild	Lean Derog.	1.36
#AIPAC	Rothschild	Lean Derog.	1.17

REFERENCES

- ADL. (2016). Anti-Semitic Targeting of Journalists During the 2016 Presidential Campaign. New York, NY: Anti-Defamation League. Retrieved from https://www.adl.org/sites/default/files/.../CR_4862_Journalism-Task-Force_v2.pdf
- ADL. (2017, August 7). "Unite the Right" Rally Could Be Largest White Supremacist Gathering in a Decade. Retrieved October 5, 2018, from <https://www.adl.org/blog/unite-the-right-rally-could-be-largest-white-supremacist-gathering-in-a-decade>
- ADL. (2018, February 1). ADL Finds Alarming Increase in White Supremacist Propaganda on College Campuses Across U.S. Retrieved October 5, 2018, from <https://www.adl.org/news/press-releases/adl-finds-alarming-increase-in-white-supremacist-propaganda-on-college-campuses>
- Astor, M. (2018, July 18). Anti-Semitic Incidents Surged 57 Percent in 2017, Report Finds. The New York Times. Retrieved from <https://www.nytimes.com/2018/02/27/us/anti-semitism-adl-report.html>
- Baynes, C. (2018, January 8). Murders by white supremacists in America more than doubled last year. Retrieved August 6, 2018, from (O'Brien, 2017)
- Coaston, J. (2018, August 1). #QAnon, the scarily popular pro-Trump conspiracy theory, explained. Retrieved October 19, 2018, from <https://www.vox.com/policy-and-politics/2018/8/1/17253444/qanon-trump-conspiracy-theory-reddit>
- Davis, C., Varol, E., Ferrara, E., Flammini, A., & Menczer, F. (2016). BotOrNot: A System to Evaluate Social Bots. In Proceedings of the 25th International Conference Companion on World Wide Web (pp. 273–274). New York, NY: Association for Computing Machinery (International World Wide Web Conferences Steering Committee). doi:10.1145/2872518.2889302.
- DellaPergola, S. (2017). World Jewish Population, 2016. In American Jewish Year Book 2016: The Annual Record of North American Jewish Communities (Vol. 116, pp. 253–332).
- Ferber, A. L. (1999). *White Man Falling: Race, Gender, and White Supremacy*. Lanham, MD: Rowman & Littlefield Publishers.
- Finkelstein, J., Zannettou, S., Bradlyn, B., & Blackburn, J. (2018). A Quantitative Approach to Understanding Online Antisemitism. arXiv preprint arXiv:1809.01644.
- Fleishman, C. (2016, October 24). #TheList: alt-right Donald Trump trolls have found a new way to attack journalists. Retrieved August 6, 2018, from <https://mic.com/articles/157543/the-list-alt-right-donald-trump-trolls-harass-jewish-journalists-8chan-raid>
- Fleishman, C., & Smith, A. (2016a, June 1). (((Echoes))), Exposed: The Secret Symbol Neo-Nazis Use to Target Jews Online. Retrieved August 6, 2018, from <https://mic.com/articles/144228/echoes-exposed-the-secret-symbol-neo-nazis-use-to-target-jews-online#.vMn9fLWMg>
- Fleishman, C., & Smith, A. (2016b, June 2). "Coincidence Detector": The Google Chrome Extension White Supremacists Use to Track Jews. Retrieved August 6, 2018, from <https://mic.com/articles/145105/coincidence-detector-the-google-extension-white-supremacists-use-to-track-jews#.pp4GSkxCv>
- Flynn, K. (2017, December 28). Why Twitter banned the bot that unmasked neo-Nazis. (Hint: Neo-Nazis). Retrieved August 6, 2018, from <https://mashable.com/2017/12/28/twitter-shuts-down-imposter-buster-bot-nazis/>
- Gambino, L. (2016, April 28). Journalist who profiled Melania Trump hit with barrage of antisemitic abuse. Retrieved August 6, 2018, from <https://www.theguardian.com/us-news/2016/apr/28/julia-ioffe-journalist-melania-trump-antisemitic-abuse>
- Glaser, A. (2017, August 30). The alt-right Wants to Build Its Own Internet. Slate.
- Korecki, N. (2018, June 29). "I snookered them": Illinois Nazi candidate creates GOP dumpster fire. Retrieved October 5, 2018, from <https://politi.co/2KfRhJY>
- Lamo, M. & Calo, R. (2018, August 14). Regulating bot speech. UCLA Law Review (2019). Retrieved October 30, 2018, from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3214572
- Marwick, A., & Lewis, B. (2017). *Media Manipulation and Disinformation Online*. New York, NY: Data & Society Research Institute.
- Mettler, K., & Selk, A. (2017, August 14). GoDaddy — then Google — ban neo-Nazi site Daily Stormer for disparaging Charlottesville victim. Retrieved October 19, 2018, from <https://www.washingtonpost.com/news/morning-mix/wp/2017/08/14/godaddy-bans-neo-nazi-site-daily-stormer-for-disparaging-woman-killed-at-charlottesville-rally/>

Monaco, N., & Nyss, C. (2018). State Sponsored Trolling: How Governments are Deploying Fake News as Part of Broader Harassment Campaigns. Institute for the Future Working Research Papers.

O'Brien, L. (2017, December). The Making of an American Nazi. The Atlantic. Retrieved from <https://www.theatlantic.com/magazine/archive/2017/12/the-making-of-an-american-nazi/544119/>

Rosenberg, Y. (2017, December 27). Confessions of a Digital Nazi Hunter. Retrieved August 6, 2018, from <https://www.nytimes.com/2017/12/27/opinion/digital-nazi-hunter-trump.html>

SPLC. (2018, February 21). The Year in Hate: Trump buoyed white supremacists in 2017, sparking backlash among black nationalist groups. Retrieved October 5, 2018, from <https://www.splcenter.org/news/2018/02/21/year-hate-trump-buoyed-white-supremacists-2017-sparking-backlash-among-black-nationalist>

Steinhardt Social Research Institute. (2016). American Jewish Population Project. Waltham, MA: Cohen Center for Modern Jewish Studies, Brandeis University. Retrieved from <http://ajpp.brandeis.edu/aboutestimates.php>

Weisman, J. (2016, May 26). Opinion | The Nazi Tweets of "Trump God Emperor." The New York Times. Retrieved from <https://www.nytimes.com/2016/05/29/opinion/sunday/the-nazi-tweets-of-trump-god-emperor.html>

Wemple, E. (2016, April 29). Why we can no longer laugh about the Trumps' media obsession. Retrieved August 6, 2018, from https://www.washingtonpost.com/blogs/erik-wemple/wp/2016/04/29/why-we-can-no-longer-laugh-about-the-trumps-media-obsession/?utm_term=.20c6ea6ba33b

Woolley, S., & Guilbeault, D. (2017). Computational Propaganda in the United States of America: Manufacturing Consensus Online (Computational Propaganda Working Paper Series No. 2017.5) (p. 37). Oxford, United Kingdom: Oxford Internet Institute, University of Oxford.

Support

This work is made possible in part by the generous support of:

**Robert A. and Renée E. Belfer Family
Charitable Fund**

**Joyce and Irving Goldman
Family Foundation**

The Grove Foundation

Walter & Elise Haas Fund

Craig Newmark Philanthropies

Omidyar Network Fund, Inc.

Lisa and John Pritzker Family Fund

Righteous Persons Foundation

Alan B. Slifka Foundation

ADL Leadership

Marvin D. Nathan
National Chair

Glen S. Lewy
President,
Anti-Defamation League
Foundation

Jonathan A. Greenblatt
CEO and National Director

Tech Advisory Board Members

Danielle Citron
Law Professor,
University of Maryland

Brad Hamm
Dean, Medill School of Journalism,
Northwestern University

Shawn Henry
Former FBI Executive Assistant
Director; President, CrowdStrike

Steve Huffman
Founder, CEO, Reddit

James Joaquin
Founder, Managing Director,
Obvious Ventures

Aileen Lee
Founder, Partner, Cowboy Ventures

Craig Newmark
Founder, Craigslist

Matt Rogers
Founder, Chief Product Officer, Nest

Guy Rosen
Vice President, Product, Facebook

Jeffrey Rosen
President of the National
Constitution Center

Jeffrey Saper
Vice Chair, Wilson Sonsini

Katie Jacobs Stanton
Chief Marketing Officer, Color
Genomics

Anne Washington
Public Policy Professor,
George Mason University

Whitney Wolfe
CEO, Bumble

The ADL Center for Technology & Society

Brittan Heller
Director

Daniel Kelley
Associate Director

For additional and updated
resources please see:
www.adl.org

Copies of this publication are available in
the Rita and Leo Greenland Library and
Research Center.

ADL by the numbers | OUR IMPACT

EDUCATE

5.2M
STUDENTS 

impacted through our Holocaust education program, Echoes & Reflections, since inception in 2005

28k
EDUCATORS 

acquired skills to teach anti-bias concepts and support students in using ally behaviors to challenge bias and identity-based bullying

30,000
EDUCATORS 

receive ADL Education Updates and Curriculum resources

ADVOCATE

45
STATES + D.C. 

have enacted Hate Crimes legislation based on, or similar to, the ADL model produced in 1981

INVESTIGATE

11k 
CASES

in which ADL has provided extremist related information to law enforcement, including critical, up-to-the-minute background on extremist threats

MONITOR

2.6M 
TWEETS

Tweets containing anti-Semitic language were identified in our 2016 report about online harassment of journalists

250 
WHITE SUPREMACISTS

who attended the 2017 "Unite the Right" Rally in Charlottesville were identified by ADL in cooperation with local law enforcement

PARTNERS

300+ 
MAYORS

pledged to join ADL through the Mayors' Compact to Combat Hate, Extremism and Discrimination, a new partnership with the U.S. Conference of Mayors

TRAIN

100% 

Of all new FBI agents have been trained by ADL since 2001

+150k 
LAW ENFORCEMENT PROFESSIONALS

were trained by ADL in the last 10 years, helping them to fight extremism and build trust with the communities they serve

How ADL is Responding


ADL has a comprehensive approach to address anti-Semitic incidents and to counter all forms of hate.

ADL's Center on Extremism is a foremost authority on extremism, terrorism and all forms of hate.


The COE's team of investigators and analysts strategically monitors and exposes extremist movements and individuals, using cutting-edge technology to track real

time developments and provide actionable intelligence and data-based analysis to law enforcement, public officials, community leaders and technology companies.

For a database of reports and resources on hate, extremism and more, visit [ADL.org](https://www.adl.org).

 Anti-Defamation League

 @ADL

 @adl_national

ADL
FIGHTING HATE FOR GOOD