CSMaP | CENTER FOR SOCIAL MEDIA AND POLITICS

Debate Twitter:

Mapping User Reactions to the 2020 Democratic Presidential Primary Debates

March 3, 2020

Authors: Megan A. Brown, Zhanna Terechshenko, Niklas Loynes, Tom Paskhalis, Jonathan Nagler

Acknowledgments

This research was carried out by the Center for Social Media and Politics at New York University, directed by Richard Bonneau, Jonathan Nagler, and Joshua A. Tucker.

The Center is supported by funding from the John S. and James L. Knight Foundation, the Charles Koch Foundation, Craig Newmark Philanthropies, the William and Flora Hewlett Foundation, the Siegel Family Endowment, the Bill and Melinda Gates Foundation, and the National Science Foundation.

This report is part of the Center's ongoing data report series, which provides key insights into the relationship between social media and politics.

The report's lead author, Megan A. Brown, is a research scientist at the Center. She collected tweets and debate transcripts, merged datasets, calculated candidate mentions, and contributed writing to this report. Zhanna Terechshenko is a postdoctoral fellow at the Center. She conducted policy topic classification for tweets and debate transcripts, and contributed data visualizations and writing to this report. Niklas Loynes is a visiting research fellow at the Center. He estimated the location of Twitter users, calculated user political affinity by candidate mention and topic mention, and contributed data visualizations and writing to this report. Tom Paskhalis is a postdoctoral fellow at the Center. Jonathan Nagler is a co-director of the Center. He supervised research and edited a draft of this report. He conducted user gender estimation and contributed data visualizations and writing to this report. Gregory Eady is an affiliate of the Center. He provided user gender estimation for the first two debates. Andreu Casas is an affiliate of the Center. He created tweet topic models for the first two debates.

The authors would like to thank Venuri Siriwardane, the Center's Researcher/Editor, for writing the introduction and editing a draft of this report; Jen Rosiere Reynolds, the Center's lab manager, for project management support; and Judy Zhang for design and production support.

Introduction

The 2020 Democratic Presidential field was the largest in modern political history. It peaked at nearly 30 candidates in 2019 and narrowed to just five who will compete in the Super Tuesday primaries on March 3. [1]

Twenty-three of these candidates participated in at least one Democratic presidential primary debate, including all five active major candidates. They headed to the stage to argue their case for the nomination, sparring over topics such as climate change, gun control, and "Medicare for All." These policy discussions took place over the course of 10 debates, with two more on the schedule: a March 15 debate in Phoenix and a final, unannounced debate in April.

The debates are organized by the Democratic National Committee, sponsored and moderated by news organizations, and hosted on network or cable television. A record 20 million viewers tuned into the ninth debate hosted on February 19 by NBC News, MSNBC, and Telemundo, while millions more watched online via live-streaming platforms. [2]

There is a large body of literature exploring the impact of general presidential debates on voters. Far less scholarly attention, however, is paid to primary debates — despite their tremendous impact on partisan voters who choose a candidate to represent them in the general election (Benoit et al., 2002).[3] Even fewer studies examine social media users' behavior during primary debates (Jennings et al., 2017),[4] despite platforms such as Twitter becoming one of the most important communication arenas for modern electoral politics.

Twitter chief executive Jack Dorsey has described the platform as a "public square," where many voices gather to "see what's happening and have a conversation about what they see."[5] This idea has taken root in elite spheres, with journalists[6] and elected officials turning to the platform for real-time updates and reactions to the day's events.

This research report explores Twitter's "debate discussion ecosystem" over the course of the first nine debates, which spanned across 11 nights from June 26 to

February 19. We collected and analyzed a total of *11,286,346* tweets by *1,724,306* unique Twitter users to:

- Compare candidates' discussions on the debate stage to Twitter users' discussions
- Find out if candidates discussed topics that resonated with users
- Compare the behavior of users in key battleground states to that of other users
- Observe how users' behavior changes across the debates, especially when candidates they were most interested in dropped out of the race

Twitter's user demographics are not representative of the American voting age population: Just 22% of U.S. adults use the platform, according to a [2019 Pew Research Center survey](#).[7] Those users tend to be younger and more educated. They earn higher incomes and are more likely to identify as Democrats.[8] And if Twitter is a modern public square, it should be noted the square is owned and governed by a private company.

Bearing these limitations in mind, what follows is a comprehensive account of our findings.

11,286,346
1,724,306

Methodology

Our goal was to find patterns in the way Twitter users discussed Democratic presidential candidates and their policy platforms during the debates. To achieve this, we collected tweets using Twitter's Streaming API,[9] filtering for hashtags related to the debates and the primary season.

We tracked the hashtags

#DemDebate
#DemocraticDebate
#Democrat
#2020Election
#Campaign2020
#2020Candidates
#POTUS2020
#DNC
#FlipTheWhiteHouse
#DNCDebates
#PresidentialDebate

for 24 hours after the start of each debate. We also included a hashtag specific to the debate number (e.g., "#DemDebate2") for each debate after the first.[10]

We used a keyword-based method[11] to calculate the number of candidate mentions in the corpus of tweets. We counted a candidate mention if a tweet contained a term — usually a candidate's name or username — from a dictionary of terms associated with each candidate. For example, we counted a mention for Elizabeth Warren if the tweet contained "@ewarren," "Elizabeth Warren," or "Warren."

We classified the tweets[12] into 22 policy issue categories identified by the Comparative Agendas Project (CAP), which assembles and codes information about the policy processes of governments around the world.[13] We narrowed our analysis to the eight most popular policy issue categories across the debates: civil rights, the economy, education, the environment, healthcare, immigration, international affairs, and law and crime.

To identify how much time the candidates spent on these issues during the debates, we used the transcripts of the debates and extracted the candidates' comments and responses to the questions. We then classified each response/comment using the same procedure described above.

To predict the gender of each user, we used the method proposed by Blevins and Mullen (2015).[14] This approach compares the first names of Twitter users to U.S. Social Security Administration (SSA) baby name data, allowing us to calculate the probability a given user's first name is associated with someone in the database who is male or female.[15] We predicted the gender of 63% of users in our corpus by matching their first names to baby names in the SSA database.

We noted the SSA database relies on a binary gender classification, which has been critiqued by many in scholarly research and social activism. We believe our estimates largely correspond to user self-identification based on available evidence, but were not able to measure self-identification using Twitter display name metadata. We were therefore unable to measure users whose self-reported Twitter display names do not correspond to a strictly "male" or "female" gender, according to SSA data. Analyzing the representativeness of this method is outside the scope of this report.

Twitter user profiles contain scant information about users' individual attributes, with the username being the only required field. To fill this void, researchers developed methods to estimate user-level traits using Twitter metadata and content published by users. We adopted the Bayesian ideal-point estimation method (Barberá, 2015)[16] to classify a given user's ideological ideal point into one of three political affinity groups: liberal, conservative, and moderate. [17]

We used Barberá's method to generate a one-dimensional, left-to-right political affinity score based on relevant Twitter accounts a user follows, including media outlets, elected officials, pundits, political organizations, and non-profit organizations. Together, these politically salient accounts paint a picture of a user's political affinity. For example, a user will be classified as conservative if they follow 12 such politically salient accounts, the majority of which are known to be conservative. We define conservative as being to the right of Fox News and liberal as being to the left of the New York Times.

Approximately 40% of users provided an accurate, parsable location in their profiles, such as "Queens, New York," "Istanbul," or "Philippines." We classified these users' locations by matching them to U.S. Census list data. We achieved this by extracting the text from their location field, performing some minimal text cleaning (e.g., removing non-alphanumeric characters), and using the GeoNames API to parse their locations into a three-level (five levels if latitude-longitude is available), machine-readable data point. When GeoNames returned more than one result for a parsing query, we assigned the top result to a given user. We parsed the locations of users in all U.S. states, but narrowed our focus in the analysis to six swing states: Arizona, Florida, Michigan, North Carolina, Pennsylvania, and Wisconsin.

See the appendix for more details about Twitter tracking terms, a full dictionary for candidate mentions, categorization of political affinity, and location string parsing.

Twitter's Debate Discussion Ecosystem

Table 1 shows the number of tweets collected for each debate, the number of unique users who tweeted about each debate, and the number of new unique users who tweeted about the debate.

The number of tweets and unique users declined between the first debate and the debate before the New Hampshire primary (February 7). However, tweets about the debates spiked by 38% during the debate on February 19, 2020, which followed the Iowa Caucus (February 3) and New Hampshire Primary (February 11). It was also the first debate Bloomberg participated in, which may have contributed to a surge in viewership. And while debates 5 through 8 brought in between 70,000 and 99,000 new unique users per debate, the 9th debate brought in 126,944 new unique users.

Table 1: Number of tweets over the course of the debates

Debate	Date	Tweet Count	Unique Users	New Unique Users Added[18]
1	June 26-27, 2019	2,155,370	571,816	n/a
2	July 30-31, 2019	2,387,688	667,097	425,693
3	September 12, 2019	1,115,821	365,226	154,441
4	October 15, 2019	1,057,735	332,777	121,599
5	November 20, 2019	790,141	241,012	75,409
6	December 19, 2019	823,788	253,895	80,161
7	January 14, 2020	954,644	299,755	98,529
8	February 7, 2020	839,396	242,444	69,714
9	February 19, 2020	1,161,763	390,221	126,944

Candidate Mentions

We used a keyword-based method to determine whether each tweet about the debates mentioned a particular candidate. Our keyword dictionaries included the candidates' last names, Twitter usernames, and common nicknames such as "Mayor Pete."[19] Table 2 gives the percentage of mentions of each candidate as a percentage of the total candidate mentions for each debate. The total number of candidate mentions is calculated using the sum of times any candidate is mentioned, so a tweet that mentions both Elizabeth Warren and Bernie Sanders would be counted twice.

Table 2: Percentage of candidate mentions by debate

debate	1	2	3	4	5	6	7	8	9
Joe Biden	12.8	12.2	18.9	14.6	10.1	14.4	8.3	14.2	6.6
Michael Bloomberg	0	0	0	0	0	0.1	0.8	0.8	28.6
Cory Booker	8.5	4.2	4.2	3.5	4.9	0.9	0.3	0.1	0
Pete Buttigieg	9.1	8.2	13.3	9.2	11.2	16.9	8.8	22	10.8
Tulsi Gabbard	6.9	11.3	1.1	11.1	15.5	0.6	1.2	0.5	0.2
Kamala Harris	19.2	16.8	11.5	10	17	1.3	0.3	0.5	0.3
Amy Klobuchar	3.1	1.1	2.7	4.1	3.9	7	6.4	8.9	5.6
Bernie Sanders	11.7	17.7	17	15.2	13.3	20.6	34.3	22.9	22
Tom Steyer	0	0	0	2.2	2.3	3.6	7.6	5.3	0.1
Elizabeth Warren	25.5	19.3	16	16.4	8.1	13.6	24.8	11.3	24.7
Andrew Yang	3.3	9.2	15.1	13.6	13.6	21.1	7.1	13.6	1
Total Candidate Mentions	916,798	1,271,619	654,993	640,268	656,433	623,622	711,904	707,446	1,054,904

Policy Issues

To analyze the policy topics people tweeted about during the debates, we classified tweets into eight topics using the CAP coding scheme: civil rights, the economy, education, the environment, healthcare, immigration, international affairs, and law and crime.

Table 3 shows the most common terms associated with each policy issue: For example, tweets classified into the civil rights category discussed women's rights, reproductive rights, voting rights, LGBTQIA+ rights, or race and ethnicity. Tweets classified into the law and crime category on the other hand discussed crime, violence, gun policy, and the justice system. (See appendix for details about the classifier and classification process.)

Just 10% of the approximate 10 million tweets from debates 1 to 8 in our corpus discussed one or more of these eight topics, while 63% did not discuss any policy topics. The remaining tweets discussed policy topics outside our focus, such as transportation or government operations, and were excluded from our analysis.

Table 3: Twitter terms associated with policy topics

Civil Rights	women, rights, black, abortion, reproductive, trans, justice, equal, voting, racial
Law and Crime	gun, violence, guns, assault, police, control, weapons, mass, ban, background
Environment	climate, change, crisis, threat, water, fossil, environmental, clean, existential, issue
Education	student, college, public, education, debt, loan, schools, school, free, kids
Healthcare	health, healthcare, care, insurance, private, plan, pay, companies, drug, universal
International Affairs	foreign, war, policy, military, nuclear, troops, endless, voted, wars, weapons
Immigration	immigration, immigrants, border, illegal, undocumented, children, open, immigrant, borders, legal
Economy	tax, taxes, wealth, pay, economy, middle, income, class, debt, economic

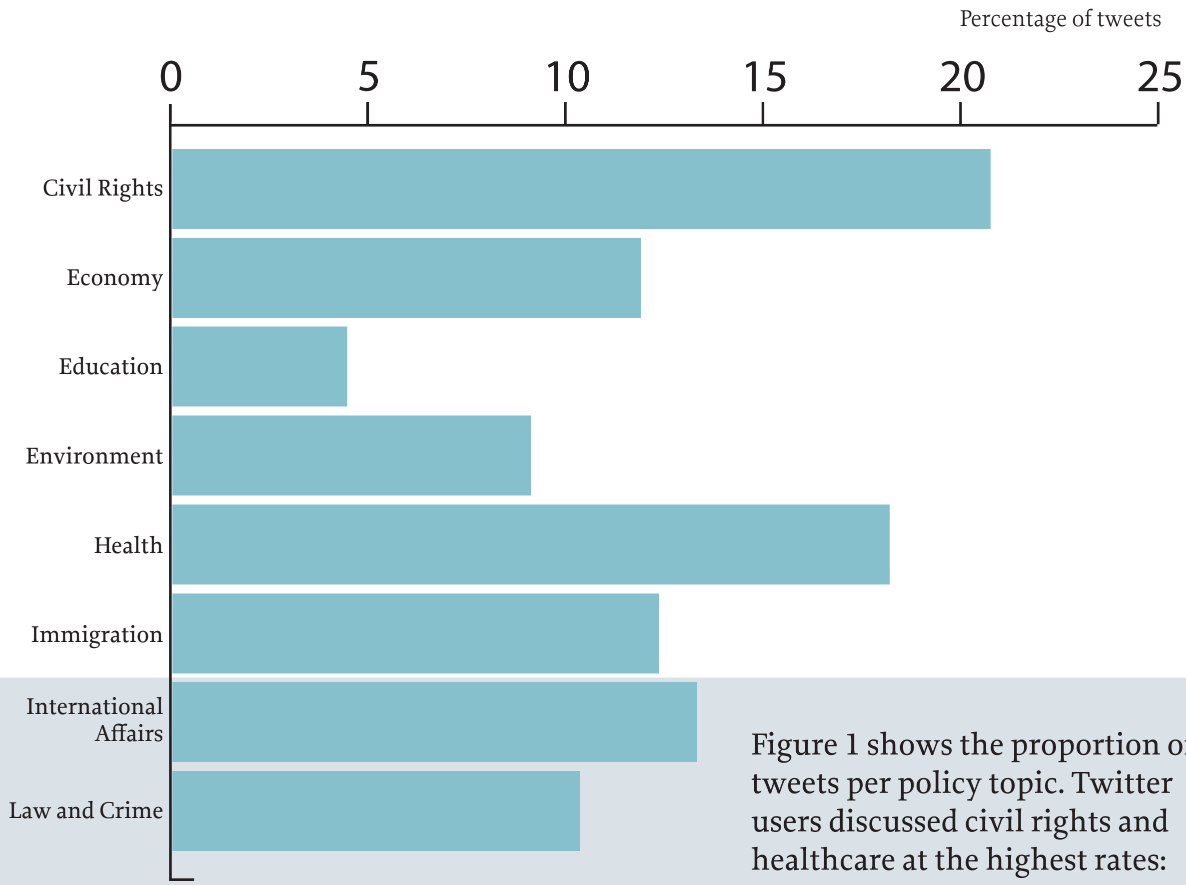

Figure 1: Percentage of tweets by topic aggregated across debates

Figure 1 shows the proportion of tweets per policy topic. Twitter users discussed civil rights and healthcare at the highest rates: about 20% and 18% respectively. Users also showed interest in international affairs (13%), immigration (12%), and the economy (12%). The environment comprised just 9% of the discussion on Twitter, ranking just below law and crime (10%). Education (4%) trails other policy topics in the amount of attention it drew from users across the debates.

Gender

We identified 63% of about 1.5 million unique Twitter users in our corpus as either women or men. We identified the first word of each user's display name as a possible first name and matched them to U.S. Census name distribution lists. Figure 2 shows a remarkable agreement between women and men in their prioritization of policy topics, with slight variations: Women focused more on civil rights and law and crime, while men paid more attention to healthcare, the economy, the environment, and international affairs. We observed near parity between women and men's focus on education and immigration. And if we were to rank the issues most talked about by each gender, the lists would be almost identical.

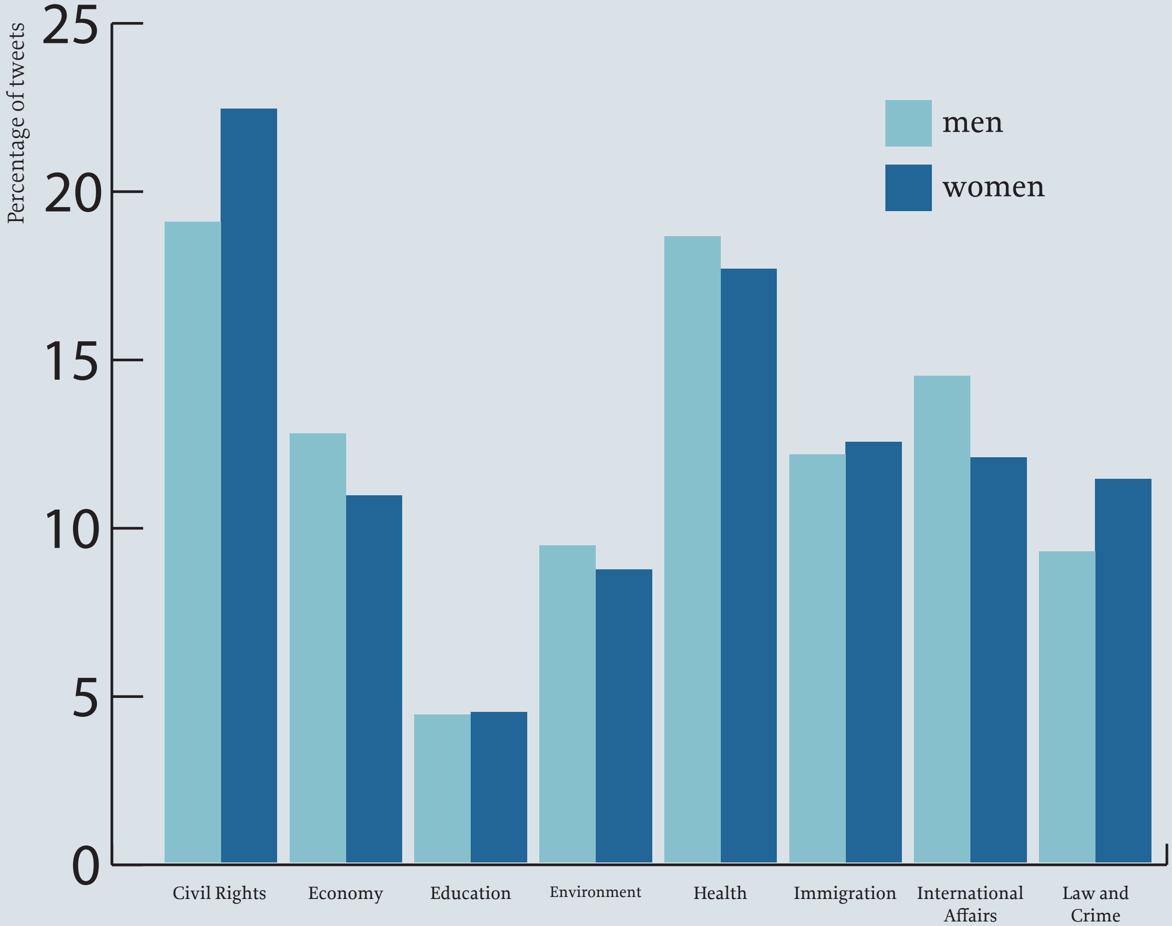

Figure 2: Percentage of tweets per topic, by gender.

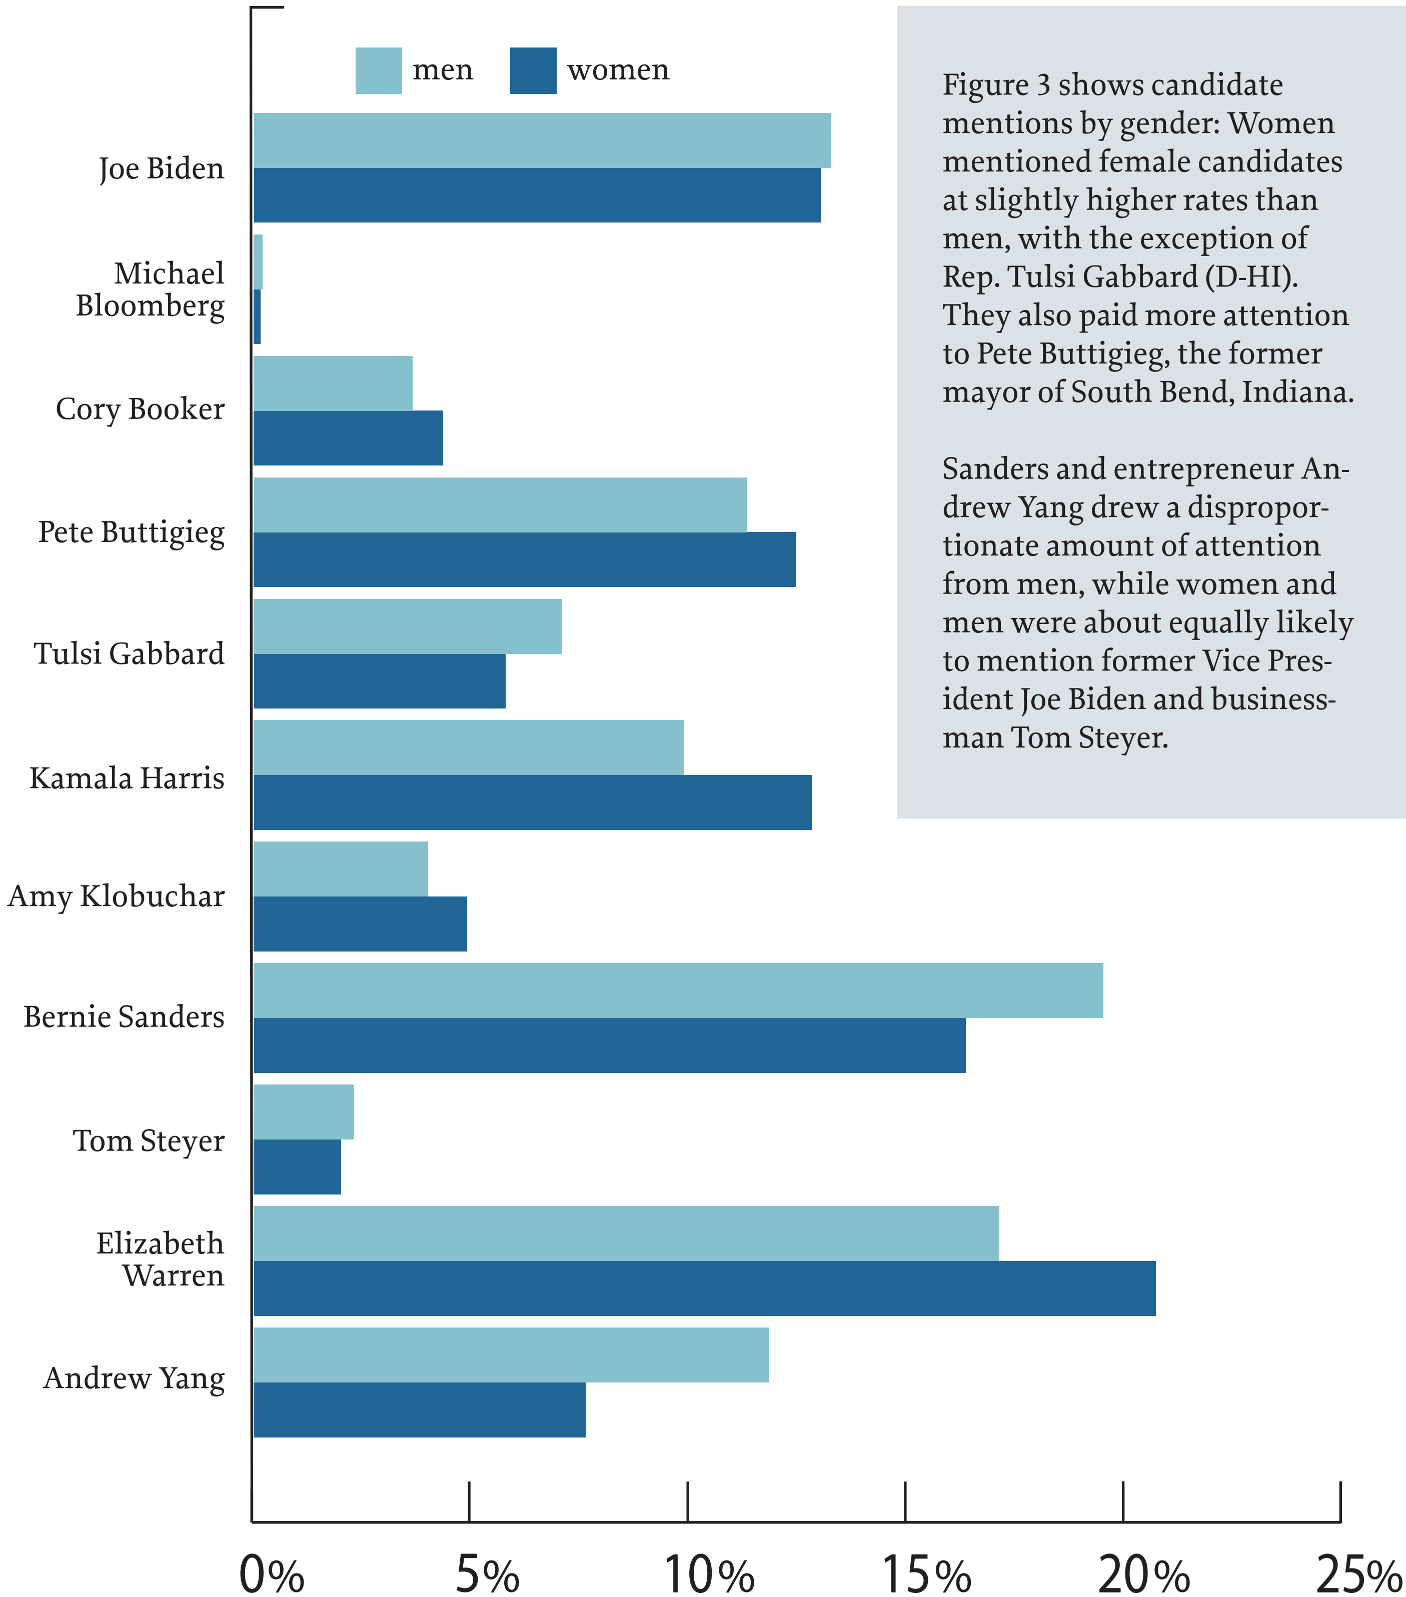

Figure 3: Percentage of tweets about candidates by gender

Figure 3 shows candidate mentions by gender: Women mentioned female candidates at slightly higher rates than men, with the exception of Rep. Tulsi Gabbard (D-HI). They also paid more attention to Pete Buttigieg, the former mayor of South Bend, Indiana.

Sanders and entrepreneur Andrew Yang drew a disproportionate amount of attention from men, while women and men were about equally likely to mention former Vice President Joe Biden and businessman Tom Steyer.

Political Affinity

To find out how different political affinity groups reacted to the candidates, we computed mention ratios for each candidate, for each debate, by each of three affinity groups: liberal, moderate, and conservative. Ratios were calculated by taking the number of candidate mentions by a group, and dividing it by the number of expected candidate mentions by a group. We computed expected mentions as the number of mentions a candidate would get if the group tweeted equally about each candidate who participated in the debate.

For example, if there were 100 tweets by liberal users and 10 candidates, the number of expected mentions would be 10 per candidate. If a candidate received 20 mentions by liberals rather than 10, they would have a candidate mention ratio of 2 (or 20/10). Higher values indicate more mentions proportionally by the group. Because the computed candidate-level ratio by political affinity group was contingent on the number of candidates in each debate, this metric allowed us to compare relative candidate mention volume by political affinity groups across debates.

Figure 4 shows candidate mention ratios by political affinity group, for each individual debate (see the appendix for a description of how users' political affinity scores were calculated and how users were assigned a group). Candidates considered frontrunners at various times during the debates — such as Biden, Sanders, and Sen. Elizabeth Warren (D-MA) — consistently outperformed their baseline expectation among all groups. Biden, considered further right than Sanders and Warren, consistently received a larger share of conservative users' mentions than he did of liberals or moderates until the seventh debate -- when his share among all groups dropped. While Buttigieg, despite his centrist reputation, received fewer mentions from conservatives than he did from the other two groups.

Candidates considered less likely to perform well in the primaries show fairly volatile mention patterns across groups — sometimes outperforming their baseline, sometimes underperforming it. This volatility is likely linked to specific events during the debates, such as a "viral moment," that made them the focus of discussion on Twitter. Overall, Figure 4 suggests the candidates who are considered most popular offline also attract the most attention on Twitter, across groups.

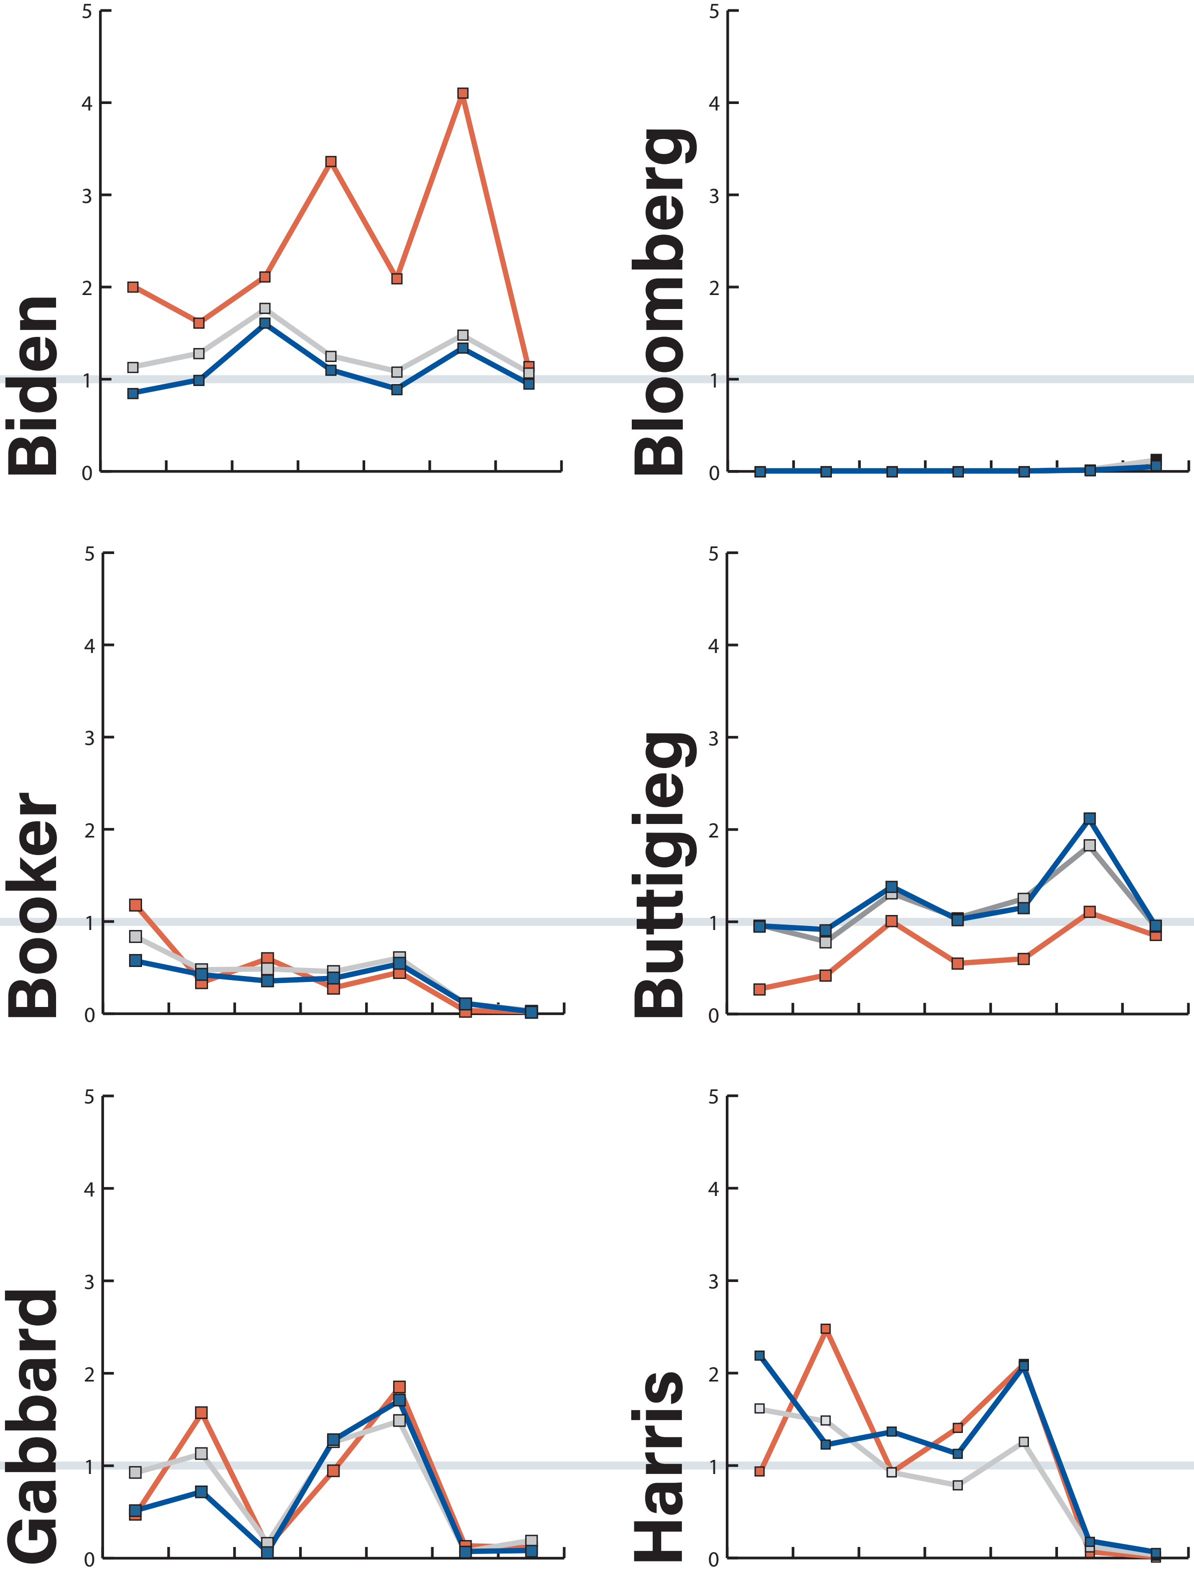
Biden
Bloomberg
Booker
Buttigieg
Gabbard
Harris
5
4
3
2
1
0

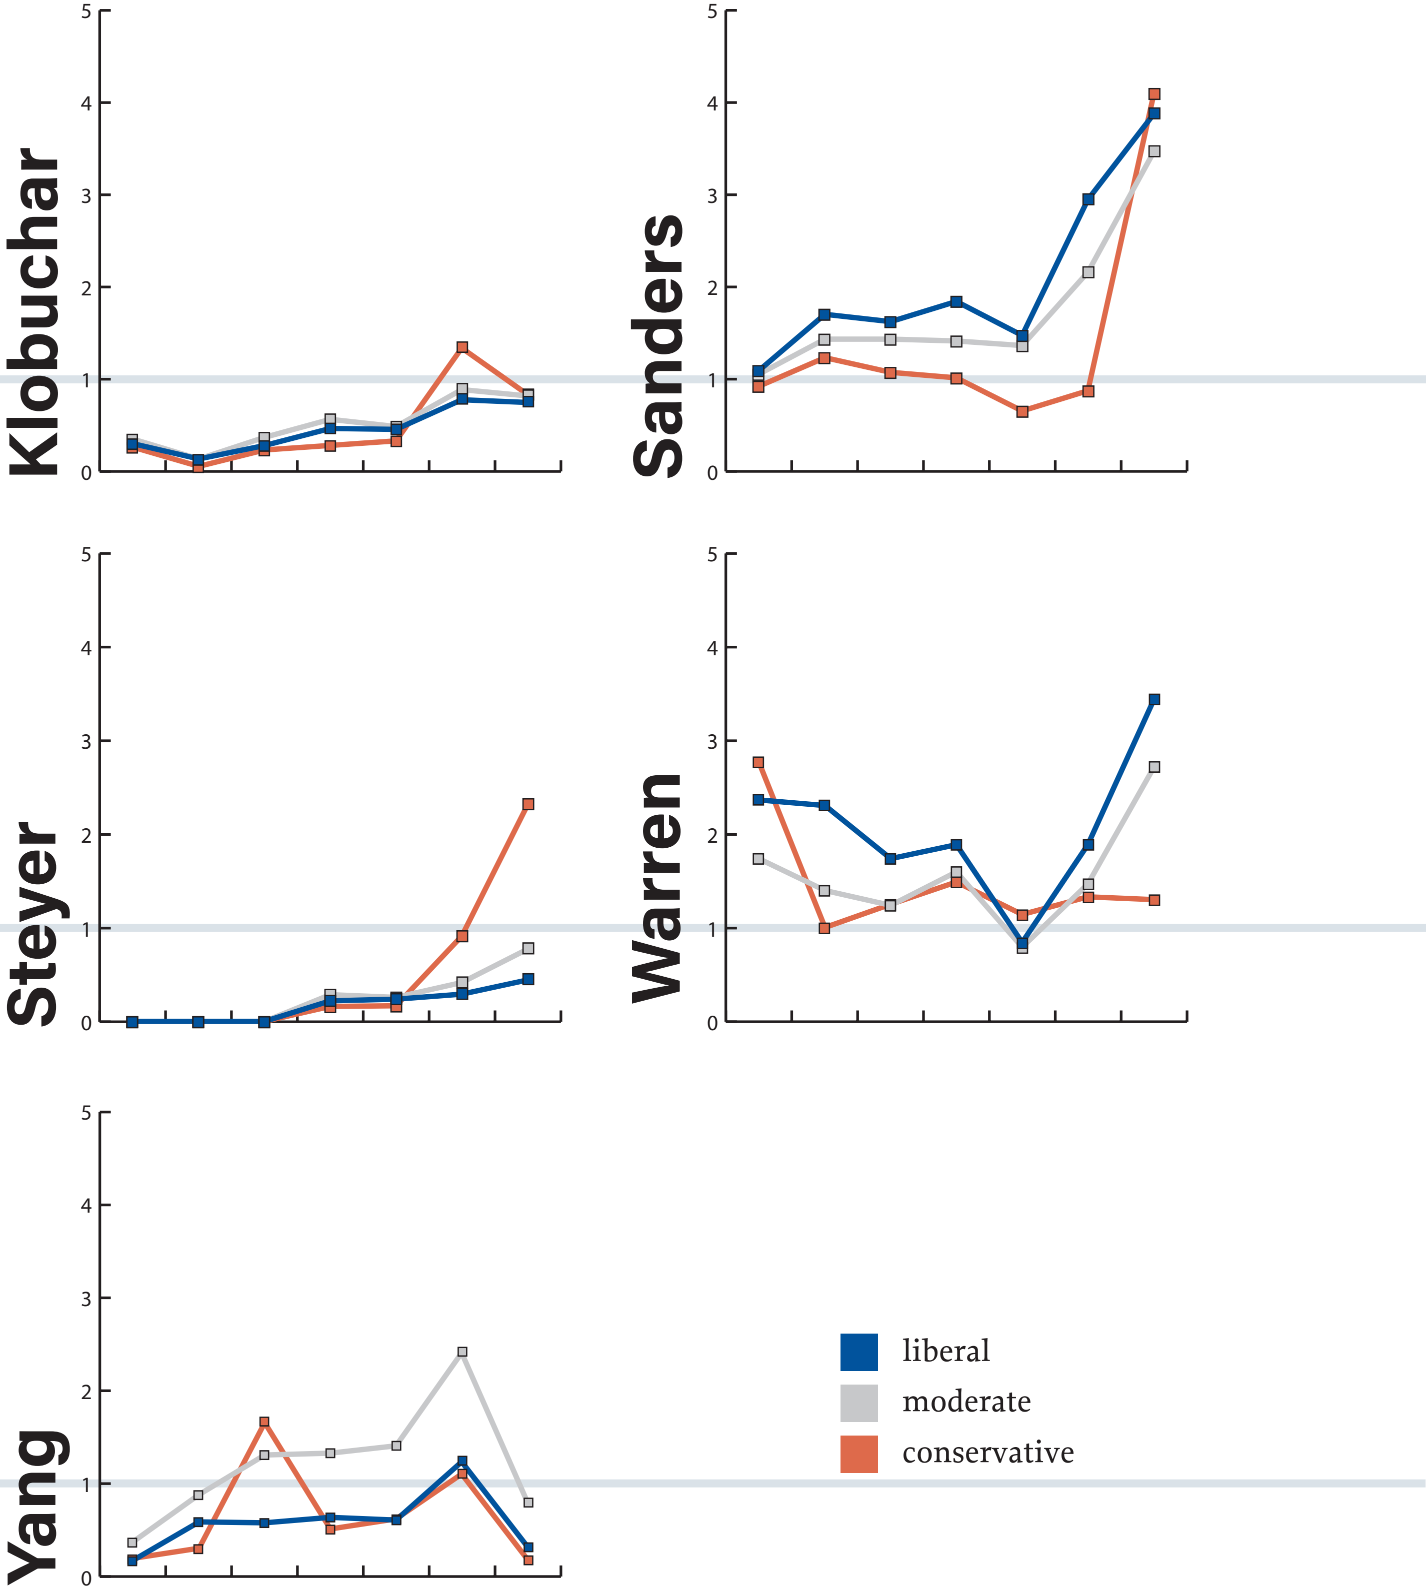

Figure 4: Candidate mention ratios by political affinity

Figure 5 shows the percentage of mentions per policy topic for each political affinity group, across the debates. Liberals, for example, paid more attention to civil rights, education, and the environment. Conservatives, on the other hand, paid more attention to the economy and immigration. Other topics, such as health, were equally important to all three groups.

These patterns are consistent with popular conceptions of how groups prioritize policy issues based on their ideological leaning, with some exceptions: Law and crime attracted more attention from liberals than conservatives, which goes against the grain of the "law and order conservative" trope. This is likely explained by the fact that many users tweeted about gun control — a popular issue for liberals — which also falls under the law and crime category.

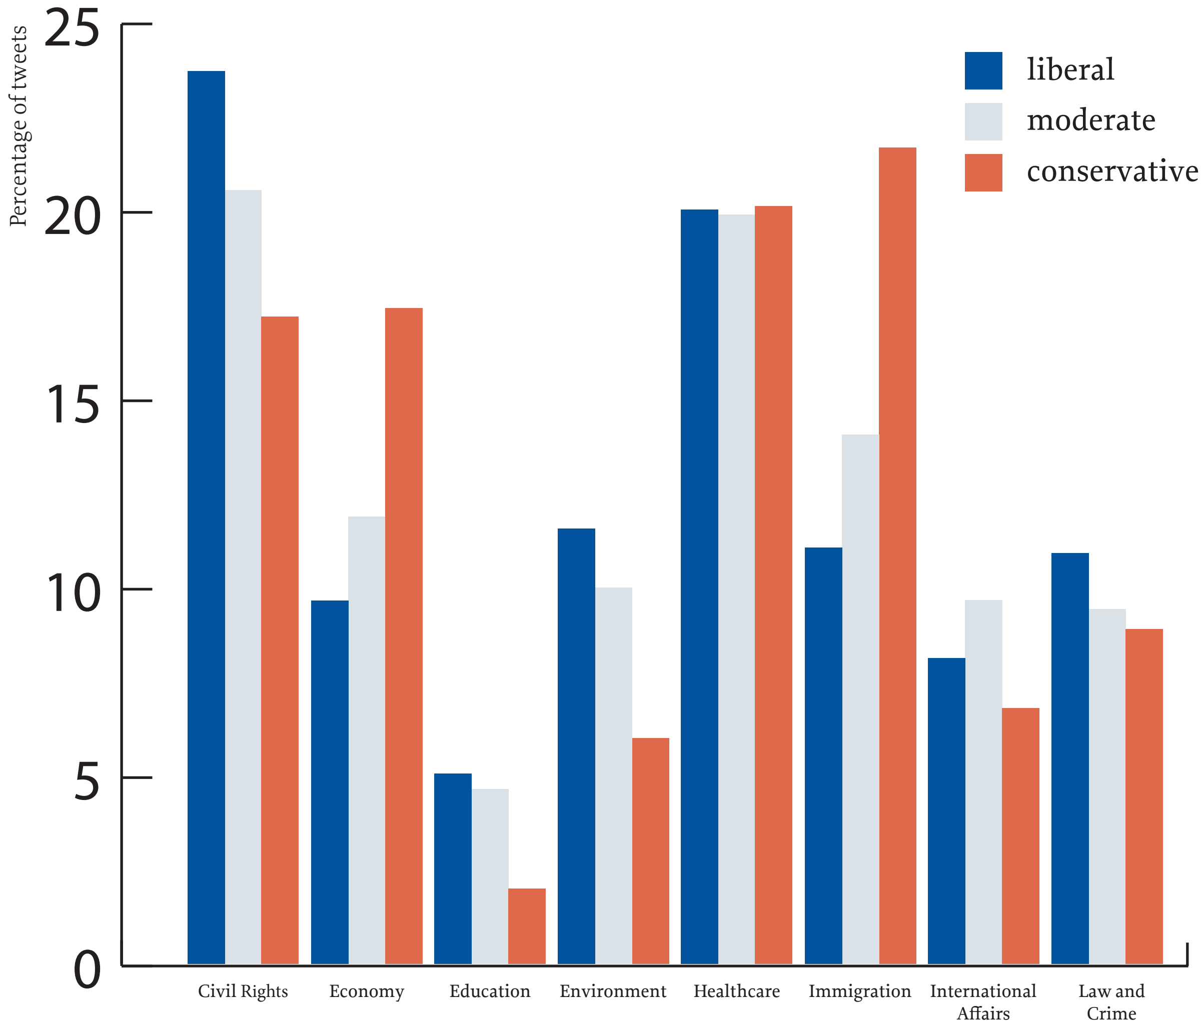

Figure 5: Percentage of tweets by topic and political affinity

Location

For the purposes of this report, we are interested in who and what are mentioned by users, and how this varies based on a user's estimated location. To estimate a user's location, we rely on self-reported location fields in their user profiles. To see further infromation on the location-parsing algorithm, see the appendix.

We are especially interested in the behavior of users in battleground states,[20] which 1) swing between Democratic and Republican candidates, with tight margins, in successive general elections, and 2) are crucial in determining the outcome of general presidential elections due to their shares of votes in the Electoral College.[21] To analyze differences among users across swing states, we focus on six: Arizona, Florida, Michigan, North Carolina, Pennsylvania, and Wisconsin.

Table 4 shows the percentage of candidate mentions in each of these swing states, as well as the total number of mentions from users classified as residents of one of these states. We see that Sanders is the most popular candidate across users in all swing states. These results are consistent with findings from Pew Research Center on the popularity of Sanders amongst Twitter users.[22] The table shows small-to-moderate levels of inter-state variation, though we did find variation among users who tweeted about Yang, who is more popular in North Carolina than in Wisconsin, and Buttigieg, who is more popular in Arizona than Wisconsin. Bloomberg, Booker, and Steyer are not very popular in any of the swing states we included in our analysis.

This could mean candidates receive higher mention proportions in states that neighbor or are in the same region as their home states. For example, Buttigieg, from the Midwestern state of Indiana, received the highest shares of mentions in Michigan and Wisconsin, which are also in the Midwest. The same pattern held true for Biden in Pennsylvania, which borders his home state of Delaware, and Klobuchar in Michigan and Wisconsin, which are close to her home state of Minnesota. However, this effect is moderate and not consistent across candidates. For example, Klobuchar receives a smaller share of mentions from Michigan than from Florida.

Overall, candidates' mention shares are very similar across swing states, and broadly in line with their mention shares for the entire United States. This indicates the candidates receive an approximately equal amount of attention on Twitter across battleground states.

Table 4: Percentage of candidate mentions by state

	AZ	FL	MI	NC	PA	WI
Joe Biden	13.3	13.9	13.0	12.1	13.4	13.1
Michael Bloomberg	0.2	0.2	0.2	0.1	0.2	0.1
Cory Booker	3.8	3.9	3.9	3.5	3.7	4.9
Pete Buttigieg	12.2	11.5	10.6	10.4	10.2	12.2
Tulsi Gabbard	7.3	6.2	6.2	6.4	6.5	5.4
Kamala Harris	11.4	11.7	11.4	12.3	11.7	10.5
Amy Klobuchar	4.3	4.5	4.7	4.0	4.7	6.2
Bernie Sanders	18.9	18.7	18.6	20.2	19.3	19.9
Tom Steyer	2.1	2.2	2.2	2.2	2.1	2.1
Elizabeth Warren	18.4	18.1	19.3	18.7	20.5	18.9
Andrew Yang	8.1	9.2	9.9	10.2	7.6	6.7
Total	28,010	84,341	31,098	24,876	45,225	18,372

Table 5 shows the share of topic mentions per swing state, which sheds light on state-level variation among policy issues discussed by users.

Regional variation among topic mention shares is somewhat more significant than it is among candidate mention shares. For example, 21.5% of tweets from users in North Carolina were about civil rights, compared with only 19.4% of tweets from users in Florida. One may assume this is due to North Carolina's larger African American population, but the explanation does not hold when comparing the state's civil rights mention shares to those in other states: Wisconsin — which has a comparatively small African American population — shows a larger share of tweets about civil rights than Florida, which has a higher population of African American residents. This inter-state variation in civil rights mention shares is not easily explained and could be caused by the variation in subtopics within the broader civil rights topic.

We also observed noticeable inter-state variation in tweets about immigration: States with the highest proportions of immigration-focused tweets are Arizona and Florida, which have higher immigrant populations than other states in our analysis.[23] This could mean immigration is more salient in these states than in those with lower immigrant populations. Overall, variation among topic mention shares ranges from very low to moderately low.

Table 5: Percentage of topic mentions by state

	AZ	FL	MI	NC	PA	WI
Civil Rights	19.8	19.4	20.4	21.5	20.4	20.7
Economy	11.9	13.5	11.9	11.7	11.9	12.1
Education	3.7	3.7	4.5	4.4	5.1	4.5
Environment	8.0	8.0	9.0	9.2	9.4	8.5
Healthcare	18.7	18.8	19.7	18.2	19.0	20.0
Immigration	13.6	13.9	12.2	11.9	11.9	10.5
International Affairs	13.6	13.9	12.2	11.9	11.9	10.5
Law and Crime	11.7	10.2	9.9	9.9	9.8	10.7
Total	12,450	34,030	11,793	9,127	17,607	6,905

How Candidates Connected With Voters

We compared the policy topics candidates discussed during the debates to those discussed by users who tweeted about the debates. We did so by selecting candidate responses that were classified under one of our eight policy topics of interest.

First, we estimated the proportion of all candidates' responses to questions about these topics as compared to the total number of their responses. We compared these proportions to the proportions of tweets about these issues across the debates (see Figure 6). The results suggest users pay more attention to issues related to law and crime (e.g., gun policy) than the amount of time candidates spent discussing these issues during the debates. This difference is most noticeable in the fifth and sixth debates. The proportion of tweets about education, on the other hand, is much smaller than the proportion of time candidates spent discussing education during the debates.

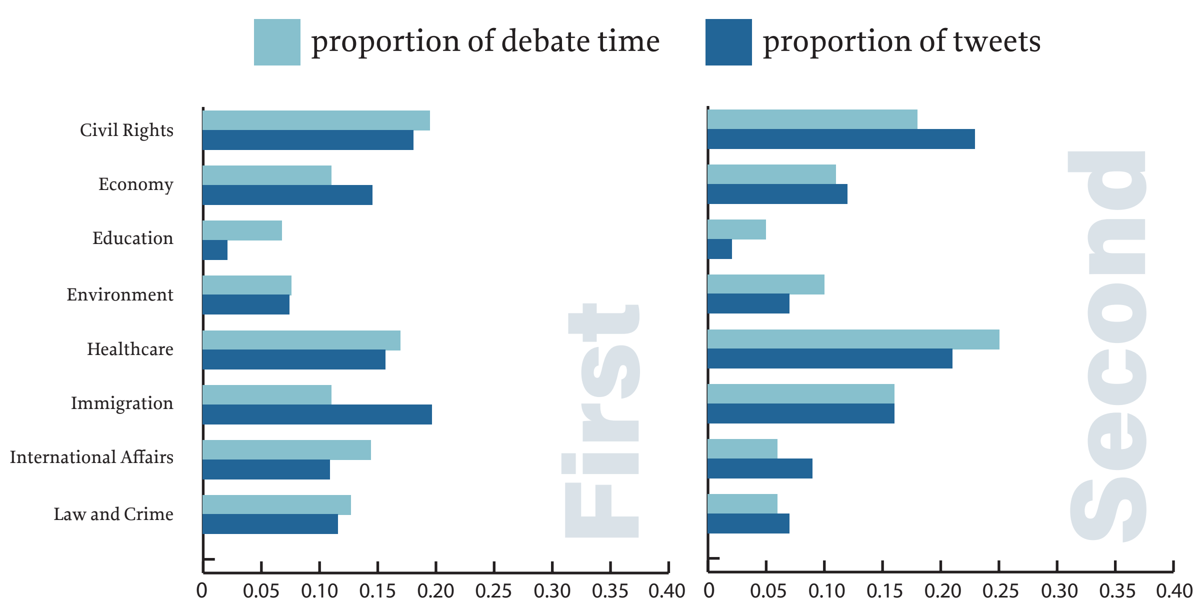

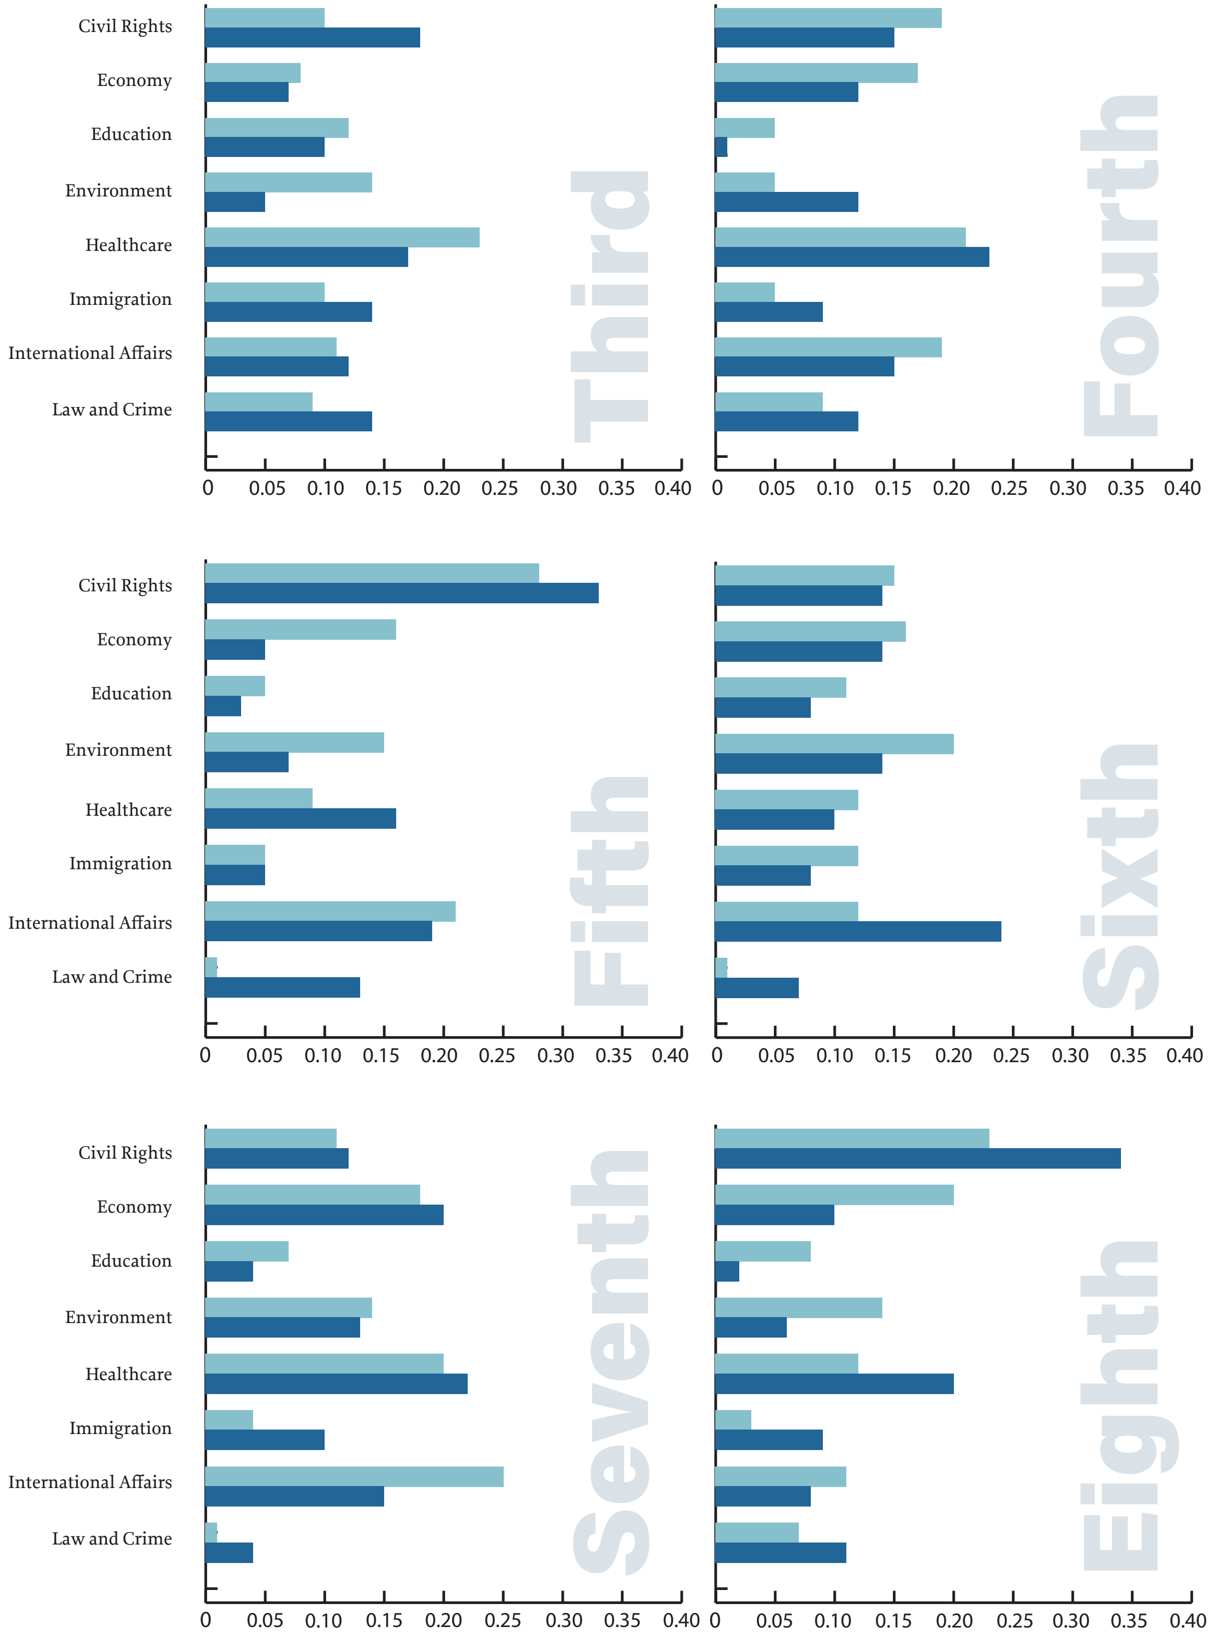

Figure 6: Proportion of tweets about topics compared to the proportion of time spent discussing these topics during the debates.

Second, we calculated the proportion of each candidate's responses to questions about these topics, compared with the total number of this candidate's responses.[24] We then examined the proportion of tweets and topics mentioning this candidate. In other words, we set out to determine whether the topics users tweet about when mentioning a candidate is a function of what the candidate talked about during the debates. Figure 7 shows the results of this analysis: When users mention a candidate, their tweets are likely to be about the issues this candidate talked about. However, there are several exceptions: Steyer frequently mentions the economy, the environment, and civil rights, but when users tweet about him, they tend to tweet mostly about environmental issues. Another exception is Yang, who also frequently mentions the environment, though users tend to respond more to his views on immigration.

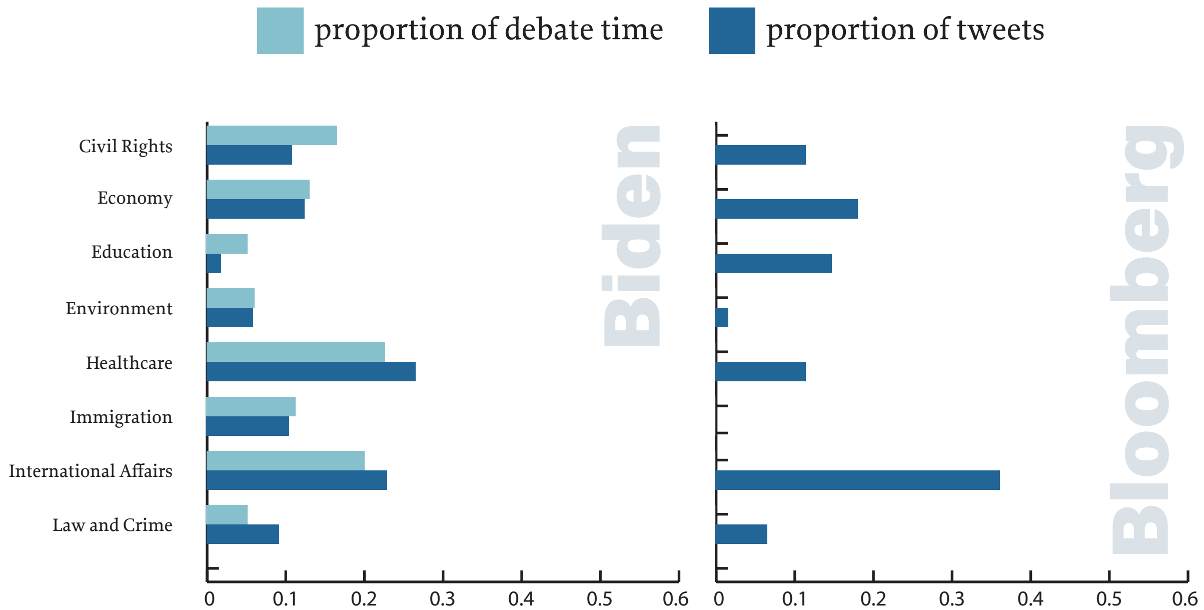


Civil Rights
Economy
Education
Environment
Healthcare
Immigration
International Affairs
Law and Crime
0 0.1 0.2 0.3 0.4 0.5 0.6
Booker
0 0.1 0.2 0.3 0.4 0.5 0.6
Buttigieg
Civil Rights
Economy
Education
Environment
Healthcare
Immigration
International Affairs
Law and Crime
0 0.1 0.2 0.3 0.4 0.5 0.6
Gabbard
0 0.1 0.2 0.3 0.4 0.5 0.6
Harris
Civil Rights
Economy
Education
Environment
Healthcare
Immigration
International Affairs
Law and Crime
0 0.1 0.2 0.3 0.4 0.5 0.6
Klobuchar
0 0.1 0.2 0.3 0.4 0.5 0.6
Sanders

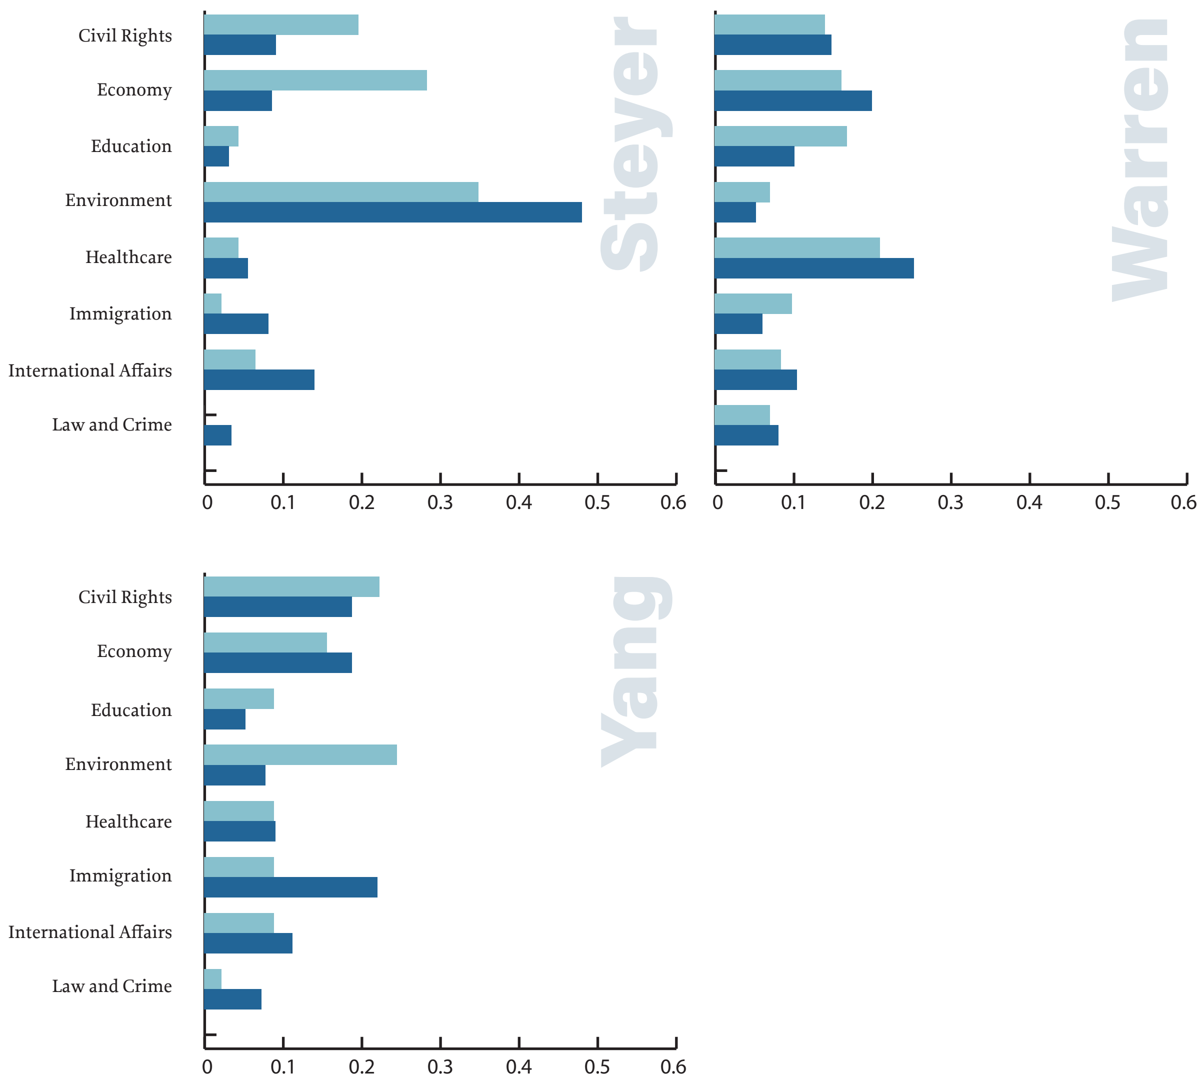

Figure 7: Proportion of tweets about topics compared to the proportion of time spent discussing these topics during the debates.

How Users Change Across Debates

A total of 1,597,362 unique Twitter users tweeted about the debates during our period of interest (debates 1–8). Of those users, 77,059 tweeted about at least six of the eight debates, which is 4.82% of the total users that tweeted about the debates. These users tweeted 4,392,166 of the 10,166,033 total tweets about the first eight debates, or 43.2%.

Of the users that tweeted during at least six of the eight debates, 1,541 tweeted about Booker more than any other candidate while he participated in the debates (debates 1-6). For users who tweeted the most about Booker, we calculated the percentage difference in tweets about other candidates before and after the date on which candidates' dropped out. We report these values in Figure 8.

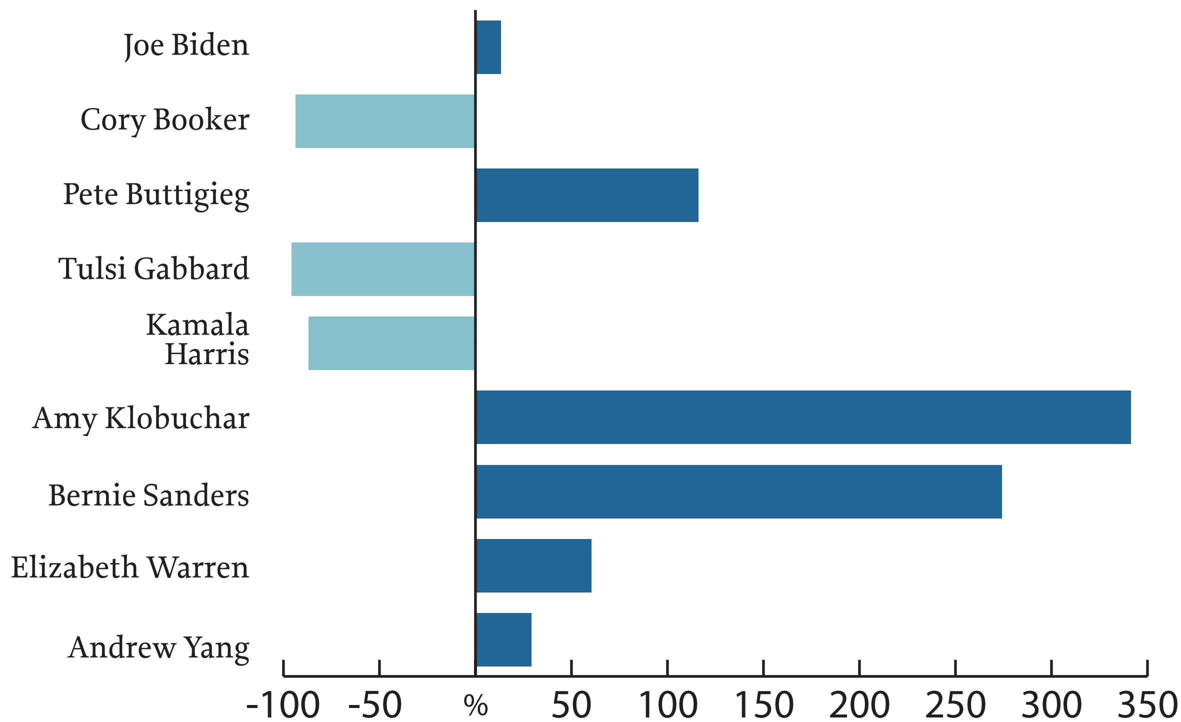

Figure 8: Percentage increase in mentions by users whose primary mentioned candidate was Cory Booker

Users who tweeted the most about Booker before he withdrew showed large increases in tweets about Klobuchar and Sanders, along with considerable increases in tweets about Yang, Biden, and Warren. They showed decreased mentions of Harris and Booker, who dropped out, and Gabbard, who is still a candidate but did not qualify for subsequent debates. Largely, users who tweeted the most about Booker matched the population at large. We see in Table 2 that over time, mentions of Sanders and Klobuchar increased for the overall population.

There were 10,673 users who tweeted about Harris the most across the debates she participated in (debates 1-5). We calculated the percentage change in tweet mentions in subsequent debates for these users.

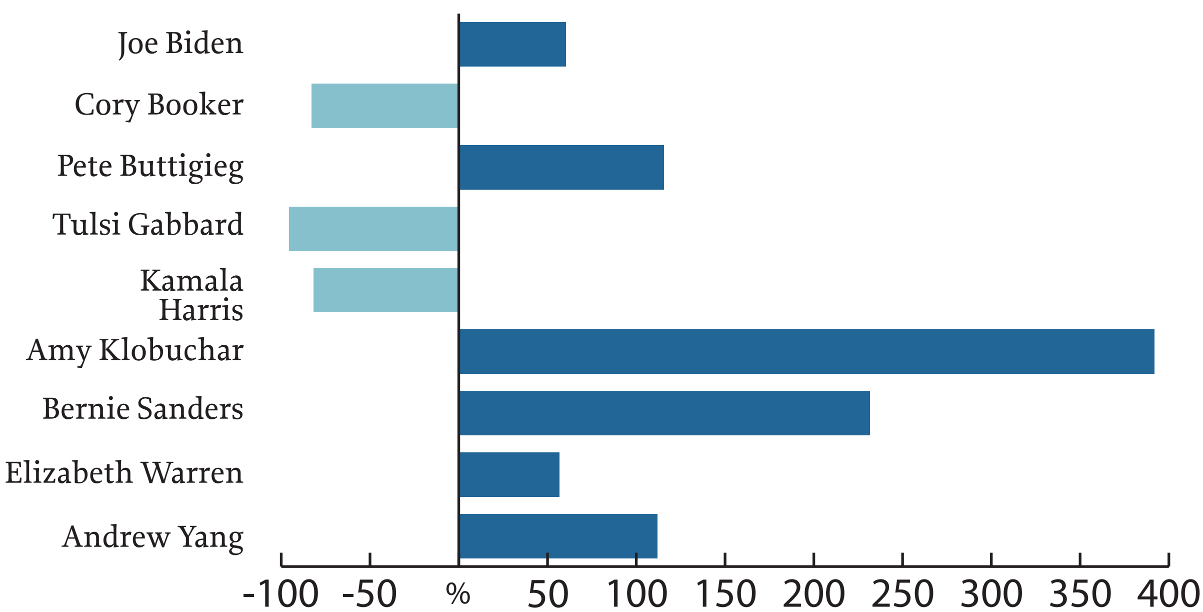

Figure 9: Percentage increase in mentions by users whose primary mentioned candidate was Kamala Harris

Like users who mentioned Booker the most, users who mentioned Harris the most began tweeting more about Klobuchar and Sanders and less about Harris, Booker, and Gabbard. The overall increase in tweets about Klobuchar was larger than that of users who tweeted about Booker.

Conclusions

As the most crowded Democratic primary field in U.S. history winnows to fewer candidates, we felt it was important to look back and understand how the debates drove the public conversation online.

Our analysis of a corpus of more than 11 million tweets and 1.7 million users shows what users most responded to in the debates and which candidates they responded to. By analyzing tweets, we can see how Twitter users change over time around primary debates.

Though the number of users who tweeted about the debates decreased over time, users who tweeted across the majority of the debates consistently contributed more content to the online conversation. And despite the lull in interest between the first and last few debates, we found interest rose again after primary season began.

We also found the policy issues users tweeted about remained fairly consistent in the population at large across debates.

We did see variation across users of different political affinities: Conservative users tweeted more about issues such as immigration and the economy, while liberal users tweeted more about issues such as civil rights, the environment, and healthcare.

Surprisingly, we saw little variation at the state level among swing states, both in topics and the candidates users mentioned. The small variation we saw was among the topic of immigration in Florida and Arizona — states with higher immigrant populations than the other swing states we analyzed: North Carolina, Michigan, Pennsylvania, and Wisconsin.

Our research suggests the debates had an impact on the online discussion: We find users respond strongly to topics that candidates talk about during the debates. This is consistent across the debates — even after candidates dropped out. And no candidate found a topic that generated significantly more online traction than expected based on the amount of time they spent discussing it during the debates.

When we separated tweets by topic into those mentioning specific candidates, we tested to see if some topics 'over-produced' for candidates (i.e., generating more reaction from the public than other topics). When Yang talked about immigration, for example, he generated more user activity per discussion-time than when he talked about the environment.

We found the candidates were not able to bring more attention to a key policy issue: Not one mentioned a topic that gained a significant amount of traction compared to the proportion of time they spent discussing it during the debates. When users mentioned individual candidates in tweets about policy issues, they did so in proportion to the time these same candidates spent talking about these same issues. However, there were some policy topics that did not gain as much traction online for particular candidates: For example, Steyer did not generate many tweets by discussing civil rights and economy, while Sanders and Klobuchar failed to gain momentum by discussing healthcare.

We believe this research is important because politicians tend to follow the discussions of public issues, and thus legislators are more likely to address the topics that dominate these public discussions.[25] Politicians following the public's lead on topics is encouraging for representative democracy.

Appendix

Data Collection

Tweet Data
Data for each of the debates was collected using the Twitter streaming API. We collected tweets containing the hashtags “#DemDebate,” “#DemocraticDebate,” “#Democrat,” “#2020Election,” “#Campaign2020,” “#2020Candidates,” “#POTUS2020,” “#DNC,” “#Debate,” “#FlipTheWhiteHouse,” “#DNCDebates,” and “#PresidentialDebate” for 24 hours after the start of each debate. We also included a hashtag specific to the debate number (e.g., “#DemDebate2”) for each debate after the first.

- Debate 1: 9pm ET, June 26, 2019 - 9pm ET, June 28, 2019
- Debate 2: 8pm ET, July 30, 2019 - 8pm ET, August 1, 2019
- Debate 3: 8pm ET, September 12, 2019 - 8pm ET, September 13, 2019
- Debate 4: 8pm ET, October 15, 2019 - 8pm ET, October 16, 2019
- Debate 5: 9pm ET, November 20, 2019 - 9pm ET November 21, 2019
- Debate 6: 8pm ET, December 19, 2019 - 8pm ET December 20, 2019
- Debate 7: 9pm ET, January 14, 2020 - 9pm ET, January 15, 2020
- Debate 8: 8pm ET, February 7, 2020 - 9pm ET, February 8, 2020
- Debate 9: 9pm ET, February 19, 2020 - 9pm ET, February 20, 2020

Debate Transcripts
Debate transcripts were collected for each of the debates following the broadcast of the debate.
- Debate 1, Night 1: The Washington Post[26]
- Debate 1, Night 2: The Washington Post[27]
- Debate 2, Night 1: The Washington Post[28]
- Debate 2, Night 2: The Washington Post[29]
- Debate 3: The Washington Post[30]
- Debate 4: The Washington Post[31]
- Debate 5: The Washington Post[32]
- Debate 6: The Washington Post[33]
- Debate 7: The Des Moines Register[34]
- Debate 8: rev.com[35]

Tweet-Level Data

Candidate Mentions
For each of the candidates, we tracked tweets that mentioned them using their Twitter handles (both campaign handles and personal handles), their last names, and first names if their first name was not considerably likely to be contained within other words.

For each tweet, if any of the terms for any of the candidates was contained within the tweet, the tweet counted as a mention for that candidate. It is possible for tweets to mention more than one candidate.

Table 7: Dictionary for Candidate Mentions Calculations

Candidate	Matching Terms
Michael Bennet	Bennet, MichaelBennet, SenatorBennet
Joe Biden	JoeBiden, Biden
Bill De Blasio	De Blasio, BilldeBlasio, NYCMayor
Michael Bloomberg	Michael Bloomberg, Bloomberg, Mike-Bloomberg
Cory Booker	Booker, CoryBooker, SenBooker
Pete Buttigieg	Pete, Buttigieg, Mayor Pete, PeteButtigieg
Julian Castro	Julian, Castro, JulianCastro
John Delaney	Delaney, JohnDelaney
Tulsi Gabbard	Tulsi, Gabbard, TulsiGabbard
Kirsten Gillibrand	Kirsten, Gillibrand, SenGillibrand, gillibrandny
Kamala Harris	Kamala, Harris, KamalaHarris, SenKamalaHarris
John Hickenlooper	Hickenlooper
Jay Inslee	Inslee, JayInslee, GovInslee
Amy Klobuchar	Klobuchar, AmyKlobuchar, SenAmyKlobuchar
Beto O'Rourke	Beto, O'Rourke, ORourke, BetoORourke
Tim Ryan	Ryan, RepTimRyan, TimRyan
Bernie Sanders	Bernie, Sanders, BernieSanders, SenSanders
Tom Steyer	Tom Steyer, Steyer, TomSteyer
Eric Swalwell	Swalwell, RepSwalwell, ericswalwell
Donald Trump	Donald Trump, realDonaldTrump, Trump
Elizabeth Warren	Elizabeth, Warren, ewarren, SenWarren
Marianne Williamson	Marianne, Williamson, marwilliamson, mari
Andrew Yang	Yang, AndrewYang

Topic Estimation

We trained a convolutional neural net (Kim, 2014)[36] to predict the presence of 21 topics on tweets: economy, civil rights, health, agriculture, labor, education, energy, immigration, transportation, law and crime, social welfare, housing, defense, science and technology, foreign trade, international affairs, government operations, public lands, partisan taunting, gun policy, and bureaucratic oversight (plus a "non policy issue" category for those tweets that were not about a particular policy topic). These topics are part of a well known and widely used topic classification in political science research, the Comparative Agendas Project.

To train this machine learning model, we combined and used 5 datasets of tweets that were manually labeled by trained research assistants.

- Debate tweets (N = 984): tweets mentioning at least one hashtags from the debate
- Media tweets (N = 6,219): tweets from U.S. media organizations (sent in 2018)
- Legislators tweets (N = 1,977): tweets from U.S. state legislators (sent in 2018)
- Followers tweets (N = 7,166): tweets from followers of U.S. state legislators (2018)
- Senators tweets (N = 58,630): tweets from U.S. senators (sent in 2013). These were coded by Russell (2018).[37]

To evaluate the performance of the algorithm, we split the data into a train, test, and validation set. First, each of the five datasets is divided into a train and test (following a 80/20 split) and then aggregated into a macro train and test sets. Then, the test set from the Debate tweets is also used as a validation set.

We trained the convolutional neural net for 50 epochs. Additional epochs improved the accuracy on the training set but not on the test nor the validation sets. The validation accuracy (so the accuracy at predicting manually labeled tweets from the Democratic debate) was around 72%, which is very high given the number of classes to predict (N = 22).

User Attributes

Location Parsing
To better understand the underlying trends of Twitter content related to the debates, it is necessary to classify individual-level user-locations. We are not interested in what users from Indonesia have to say about candidates, but we are highly interested in what users from the U.S. state of Indiana, who are potential voters, have to say. Besides excluding users unable to vote in the U.S., this also allows us to understand public opinion as expressed on Twitter in a geographically disaggregated way.

The goal of the Twitter user location-parsing algorithm applied in this research is to determine users' locations at 3 different levels: country, administrative area (e.g., U.S. or Indian state, Canadian province) and municipality, as well as latitude-longitude coordinates where available. However, this is not a trivial endeavor, as Twitter's user profile does not, by default, require users to enter a systematically parseable location. Rather, the profile offers a free-form, optional "location" text field where users are only restricted by a 30-character limit. Hence, many users choose either not to disclose any location information, or they enter something purposefully inaccurate, such as "Narnia".

However, approximately one-third to one-half of users provide an accurate, parsable location in their profiles, such as "Queens, New York", "Istanbul", or "Philippines". We classify these users' locations using the census-list-matching component of our geo-locating algorithm. We achieve this by extracting the text from their location field, performing some minimal text cleaning (e.g. removing non-alphanumeric characters) and using the GeoNames API to parse their locations into a 3-level (5 levels if lat-long is available) machine-readable data point. When GeoNames returns more than one result for a parsing query, we assign the top result to a given user. In general, this method has a classification accuracy of 92% (country-level) and 83% (administrative area level). For users who do not supply location information, those for whom this information is not parseable by GeoNames, we do not count them in our analysis.

Table 5 and Table 6 show the relative percentage of U.S. state-level mentions for candidates and pertinent topics over all debates, as well as the total number of mentions per state for each (far-right column).

Table 5: Candidate mention percentages by US state and Washington, D.C.

	J.B.	M.B.	C.B.	P.B.	T.G.	K.H.	A.K.	B.S.	T.S.	E.W.	A.Y.	Total
Alabama	15.7	0.4	4.4	11.8	5.9	12.9	4	15.8	2.2	19.4	7.4	12,691
Alaska	13.8	0.1	4.2	13.5	9.3	13.2	4.1	16	2	18.1	5.6	3,748
Arizona	13.3	0.2	3.8	12.2	7.3	11.4	4.3	18.9	2.1	18.4	8.1	28,010
Arkansas	13.4	0.2	3.1	13.2	5.7	10.1	4	21.6	2.3	18.7	7.5	9,900
California	12.1	0.2	3.6	11.1	6.9	13	4.2	18.9	2.3	18	9.8	2,217,14
Colorado	12.5	0.2	3.7	13.3	6.3	9.9	4.8	19.7	2.1	18.8	8.8	26,646
Connecticut	14.6	0.2	5.4	11.1	6.5	11.5	4.4	18.8	2.3	19.4	5.8	13,883
Delaware	17.8	0.1	3	12.9	5.6	9.9	3.5	14.2	1.1	22	9.8	977
Florida	13.9	0.2	3.9	11.5	6.2	11.7	4.5	18.7	2.2	18.1	9.2	84,341
Georgia	14.3	0.1	4.8	9.8	6.6	14.2	4.2	16.7	2	18	9.3	32,740
Hawaii	9.9	0.1	3.4	9.6	13.2	12.8	3.5	14.3	2.5	15.3	15.3	6,015
Idaho	12.2	0.2	3.5	20	7.9	10.5	4.1	14.9	2	16.5	8.1	4,444
Illinois	13.6	0.2	3.7	12.1	5.8	11.2	4.6	20.8	2.2	18.9	6.8	59,079
Indiana	11.7	0.2	3.2	24.5	6	8.5	3.8	15.6	2.1	15.1	9.3	19,105
Iowa	11.5	0.2	10	11.8	4.7	9.7	6.6	16.1	2.5	19.5	7.4	15,435
Kansas	12.5	0.1	3.2	12.4	5.7	9.4	5.8	20.7	2	19.2	9.2	6,249
Kentucky	14.7	0.2	4	12.2	6.4	12.1	4.3	16.5	2	19.4	8	11,233
Louisiana	14.3	0.1	4.4	9.7	5.9	13.3	4.2	19.4	2	18.2	8.5	12,249
Maine	12.8	0.1	4.5	11.5	5	9.7	3.8	23.1	2.5	16.1	11	5,970
Maryland	13.5	0.2	4.3	10.7	6.3	14.7	4.6	16.1	2.1	19.4	8.1	23,603
Massachusetts	11.9	0.2	3.2	11.7	4.9	8.8	4.7	17.9	2.2	27.4	7.1	44,430
Michigan	13	0.2	3.9	10.6	6.2	11.4	4.7	18.6	2.2	19.3	9.9	31,098
Minnesota	10.7	0.1	2.9	11.7	5.7	8.7	11.3	20.1	2	17.4	9.4	24,082
Mississippi	16.2	0.4	4.3	9.3	12.4	13.5	4.2	14.4	2.3	16.7	6.3	3,988
Missouri	12.7	0.2	3.7	14.3	5.8	12.6	4.1	17.1	2.3	17.9	9.4	17,694

	J.B.	M.B.	C.B.	P.B.	T.G.	K.H.	A.K.	B.S.	T.S.	E.W.	A.Y.	Total
Montana	14.1	0.2	3.8	11.7	5.5	10.7	4.2	23.7	2	19.6	4.6	3,699
Nebraska	11.6	0.2	3.1	23.7	4.7	9.1	7.5	17.7	1.8	14.7	6	5,750
Nevada	13.6	0.2	4.6	9.9	6.5	14.1	3.8	16.5	2.6	18.7	9.5	19,356
New Hampshire	11.1	0.2	9.5	13.9	6.3	8.4	7.9	15.2	2.1	18.3	7.2	8,803
New Jersey	13.5	0.2	6.4	9.7	6.3	11.8	4.3	21	2.1	16.8	7.9	33,284
New Mexico	10.5	0.1	3.4	15.9	6.5	12	4.5	20.1	1.8	19	6.3	7,599
New York	13.1	0.3	4.1	11.7	6	11.2	4.3	20.6	2.3	18.5	8.1	149,156
North Carolina	12.1	0.1	3.5	10.4	6.4	12.3	4	20.2	2.2	18.7	10.2	24,876
North Dakota	13.7	0	4.2	11.6	8.2	16.7	3.9	15.3	1.8	16.3	8.4	959
Ohio	14.4	0.2	4.6	10.8	6.3	12	4.3	17.6	2.2	18.3	9.4	37,665
Oklahoma	13.9	0.3	4	12.6	6.8	11.7	5	14.7	2.5	17.8	10.7	9,303
Oregon	11.6	0.1	3.1	10.3	6.3	10.5	4.5	21.4	2	20.5	9.7	29,393
Pennsylvania	13.4	0.2	3.7	10.2	6.5	11.7	4.7	19.3	2.1	20.5	7.6	45,225
Rhode Island	13.3	0.2	3.5	10.1	6.1	10.1	5.3	22.1	2.6	20.2	6.6	3,512
South Carolina	15.1	0.2	6.6	10.8	6.1	14.6	5.3	14.7	3.1	17.5	6.2	14,545
South Dakota	11.9	0.3	3	10.9	7.1	16.1	5.7	13.1	3.5	18.7	9.5	1,608
Tennessee	14.3	0.2	3.8	11.8	6.9	11.9	5.2	19.8	2.2	18.8	5.2	21,385
Texas	14.9	0.2	4.2	10.4	6.8	12.1	4.1	18.3	2.2	18	8.9	112,755
Utah	12.1	0.2	3.6	12.4	6.3	11.2	3.9	17.5	2.4	19.8	10.7	6,828
Vermont	10.4	0.1	1.9	9.2	5	11.1	4.1	31.8	2.5	19.5	4.5	3,945
Virginia	13.5	0.2	4.3	12.3	6.6	12.9	4.2	15	2.3	18.9	9.8	20,816
Washington	12.4	0.1	3.5	10.8	5.8	10.3	4.4	20.6	2.2	20.6	9.3	35,048
Washington, D.C.	15.1	0.2	5.8	12.8	5.8	11.6	5.5	15	2.6	19.2	6.4	67,361
West Virginia	14.2	0.2	4.5	14.1	6.6	13.7	5.3	14.5	2.1	18.8	6.2	14,340
Wisconsin	13.1	0.1	4.9	12.2	5.4	10.5	6.2	19.9	2.1	18.9	6.7	18,372
Wyoming	16	0.1	3.1	8.5	6.1	11.5	4.3	17.7	3.4	18.4	11	1,178

Table 6: Topic mention percentages for each state + DC

	Civil Rights	Economy	Education	Environment	Healthcare	Immigration	Intl. Affairs	Law and Crime	Total
Alabama	23	12.8	3.4	7	18.7	13.2	11.4	10.6	5550
Alaska	19.1	12.2	2.8	8	19.4	13.6	13.2	11.7	1636
Arizona	19.8	11.9	3.7	8	18.7	13.6	12.5	11.7	12450
Arkansas	19.8	12.7	4.3	8.6	18.5	12.5	12.6	11.1	4062
California	21.4	11	4.3	10.1	17.9	11.8	13.6	9.9	81018
Colorado	20.2	11.7	4.5	10	19.4	11.3	12.9	10	10838
Connecticut	21	11.3	4.4	9.4	18.1	11	12.5	12.2	5036
Delaware	21.9	7.6	5.5	6	19.6	12	13.1	14.4	383
Florida	19.4	13.5	3.7	8	18.8	13.9	12.6	10.2	34030
Georgia	23.1	12.1	4.2	8.3	17.7	13.2	11.7	9.7	13030
Hawaii	21	11.2	4	8.9	16.1	10.8	18.2	9.8	2019
Idaho	18.6	9.9	4	8.7	16.2	14.9	16.2	11.6	1613
Illinois	21.7	10.8	4.7	9.8	18.1	11.1	13.6	10.1	21664
Indiana	19.9	11.7	5	9	18	12.2	14.8	9.4	6991
Iowa	21	11	4.9	10	19.2	10.5	12.1	11.2	5517
Kansas	19	10.8	5.9	10.4	18.4	10.2	14	11.3	2253
Kentucky	19.4	13.2	4.5	7.8	18.5	13.2	12.8	10.4	4335
Louisiana	21.4	12.9	4	9.5	19.3	13.4	10.8	8.8	5075
Maine	21	11.2	3.7	10.3	20.9	11.7	12.5	8.6	2240
Maryland	24.5	10.2	5	8.5	17.4	11.6	11.8	10.9	8875
Massachusetts	21.4	10.3	5.1	10.7	17.9	10.4	14.2	10	16617
Michigan	20.4	11.9	4.5	9	19.7	12.2	12.5	9.9	11793
Minnesota	20.8	9.9	5.1	10.2	19.6	10.7	13.4	10.3	8720
Mississippi	20	15	3.7	6.1	15.6	14.6	16.3	8.6	1750
Missouri	21.6	12.2	4.3	8.6	18	11.1	12.8	11.5	7001

	Civil Rights	Economy	Education	Environment	Healthcare	Immigration	Intl. Affairs	Law and Crime	Total
Montana	18.3	11.4	4.1	9.9	22	12.4	12.6	9.4	1674
Nebraska	20.9	11.6	4.3	10.2	20.6	11.8	12	8.7	2277
Nevada	20.6	12.4	3.9	8	17.9	13.9	11.6	11.7	7992
New Hampshire	20	9.9	4.6	11.1	17.7	11.4	13.6	11.5	3169
New Jersey	20.6	11.8	3.9	8.3	18.6	12.6	13.3	10.9	12804
New Mexico	21.4	10.1	3.7	9.8	17.2	13.3	13.5	11	3147
New York	22.8	10.4	4.6	9.7	17.8	11.2	14	9.7	53187
North Carolina	21.5	11.7	4.4	9.2	18.2	11.9	13.2	9.9	9127
North Dakota	19.4	9.6	4.4	5.7	22.1	17.7	12	9.1	407
Ohio	20.3	11.9	4.6	8.6	19.2	12.9	12.8	9.8	14595
Oklahoma	20.5	13.1	4	7.4	19.1	13.6	12.2	10.1	3909
Oregon	20.3	11.8	4.8	10.9	19.3	10.2	13.1	9.7	10848
Pennsylvania	20.4	11.9	5.1	9.4	19	11.9	12.6	9.8	17607
Rhode Island	19.1	11.3	6.1	12.6	19.4	10.7	13.3	7.5	1444
South Carolina	20.9	13.1	3.5	9.3	17.8	13.9	11.9	9.7	6119
South Dakota	22.7	12.4	2.3	6.7	21.7	11.9	11.4	10.8	563
Tennessee	20.8	13.1	3.8	7.9	18.7	13.5	11.7	10.4	8860
Texas	20.2	12.4	3.6	7.8	17.4	15.1	12.2	11.4	47477
Utah	20	12	4.5	8.9	19.3	12.5	12.5	10.4	2842
Vermont	19.9	9.5	5.4	11.8	21.8	8.3	14.5	8.8	1410
Virginia	22.5	10.5	5	8.7	17	11.2	14.2	10.8	7933
Washington	20.1	10.5	4.7	15.3	17.6	9.8	12.9	9.1	14013
Washington, D.C.	22.8	9.9	5.2	10.5	16.9	10.7	14	9.9	29609
West Virginia	22.5	12.2	3.3	7.8	17.9	13.1	13.2	10.2	6040
Wisconsin	20.7	12.1	4.5	8.5	20	10.5	12.9	10.7	6905
Wyoming	18.2	16.8	4.6	8.9	18.6	11.6	12.1	9.1	570

Gender
We have 1,527,648 unique users across eight debates, 99% of whom have a potential first name (i.e., a display name with alphabetical characters). Of these, we were able to match 966,052 user names to a gender in the U.S. Social Security Administration baby name database. The distribution of gender is as follows:

Men: 57%
Women: 43%

What constitutes a first and last name: We define a first name as the first word in a user's display name, and a last name as the last word in a user's display name. Display names including only one word are assumed to be first names. Display names were pre-processed by removing non-alphabetical characters (e.g., emojis) and hashtags (e.g., #MAGA).

Political Affinity
Users who can be classified with an affinity must follow at least three politically salient Twitter accounts. We compiled a list of these accounts here. Affinity scores generated using this method are not naturally interpretable beyond greater than signifying a user is further to the right than score X, and less than meaning that a user is further to the left. However, the computed values are not on an interval scale, meaning that ranges between scores are not necessarily equally meaningful in every area of the scale, and indeed provide any interpretability at all in their own right.

Given this hard-to-interpret nature of the computed affinity scores, we use mass media outlets' known reference scores as a heuristic for defining cut-off points by which individual users can be divided into three groups: liberal, moderate and conservative. We classify users as liberal if their computed affinity score is equal to or less than (read: to the left of) the score of the Washington Post (-0.3396343). Moderate users have a score greater than that of the Washington Post, but lower than that of Fox News (0.8026926). Conservative users' scores are equal to or greater than Fox News's score. While these cutoff points are arbitrary to some degree — alternative cutoff points based on e.g., Members of Congress's computed score, or other media outlets' would also be feasible — this binning, which defines moderates widely and conservatives narrowly (rather than e.g., choosing more cut-off points and introducing intermediate categories, such as center-left or center-right) is particularly useful for the sample of users used in this research. Given the fact that the users underlying the analyses documented in this report are tweeting about candidates and topics related to the eminent center-left party in American politics, it is safe to assume the majority of tweeters will be more likely to be of a liberal persuasion than otherwise. Hence, identifying users who are definitely not liberal, even if defined rather generously, lets us draw more reliable inferences regarding the behaviour and preferences of non-typical users in our sample.

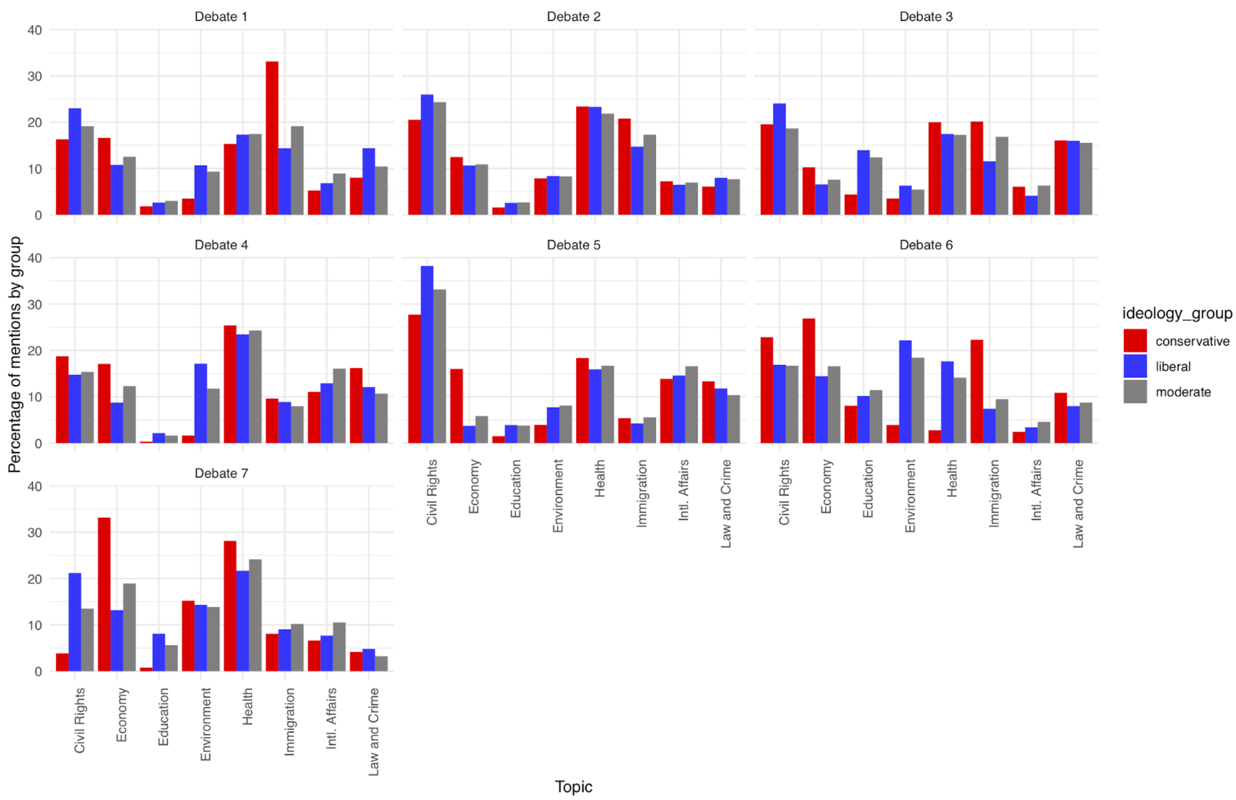

Figure 10: Topic mention percentages by political affinity, separated by debate

Endnotes

1 Lee, J.C. (2020, February 12). How the Democratic Presidential Field Has Narrowed. The New York Times. https://www.nytimes.com/interactive/2019/02/14/us/politics/2020-democratic-candidates-president.html

2 Fischer, S. (2020, February 20). Bloomberg Drives Record Primary Debate Ratings. Axios. https://www.axios.com/bloomberg-record-democratic-debate-ratings-bee460fa-1237-4748-a2e3-8900777e4ade.html

3 Benoit, W. L., McKinney, M. S., & Stephenson, M. T. (2002). Effects of Watching Primary Debates in the 2000 U.S. Presidential Campaign. Journal of Communication, (52)2, 316-331. https://onlinelibrary.wiley.com/doi/epdf/10.1111/j.1460-2466.2002.tb02547.x

4 Jennings, F.J., Coker, C.R., & McKinney, M.S. (2017). Tweeting Presidential Primary Debates: Debate Processing Through Motivated Twitter Instruction. American Behavioral Scientist, (61)4, 455-474. https://journals.sagepub.com/doi/abs/10.1177/0002764217704867

5 Dorsey, J. [@jack]. (2018, September 5). We believe many people use Twitter as a digital public square. They gather from all around the world to see what's happening, and have a conversation about what they see [Tweet]. Twitter. https://twitter.com/jack/status/1037399093084151808

6 McGregor, S.C. & Molyneux, L. (2018). Twitter's influence on news judgement: An experiment among journalists. Journalism, 00(0) 1-17. https://journals.sagepub.com/doi/abs/10.1177/1464884918802975

7 Perrin, A. & Anderson, M. (2019, April 10). Share of U.S. adults using social media, including Facebook, is mostly unchanged since 2018. Pew Research Center. https://www.pewresearch.org/fact-tank/2019/04/10/share-of-u-s-adults-using-social-media-including-facebook-is-mostly-unchanged-since-2018/

8 Wojcik, S. & Hughes, A. (2019, April 24). Sizing Up Twitter Users. Pew Research Center. https://www.pewresearch.org/internet/2019/04/24/sizing-up-twitter-users/

9 The Twitter Streaming API is an application programming interface that allows for streaming tweets in real time based on a set of keyword filters. Any tweet that matches one of the keyword filters is collected.

10 These hashtags were used by Democratic Party officials and candidates, and allowed us to collect as many relevant tweets as possible.

11 The keyword-based method counts mentions based on a dictionary of terms associated with each candidate. See the appendix for a full list of terms associated with each candidate.

12 We used a Convolutional Neural Network (CNN) model to classify the tweets. CNNs are a class of deep neural networks, that use a linear operation called convolution, often for image and text classification. To generate predictions, we trained the CNN on a corpus of tweets labeled according to 22 policy issue categories. Our training data included 983 tweets about the first Democratic debate (June 26-27, 2019); 6,154 tweets from followers of state legislators (timestamped in 2018); 1,887 tweets from state legislators (timestamped in 2018); 6,129 tweets from national and regional media accounts; and 57,962 tweets from senators (during the 113th congress).

13 Hearings. The Policy Agendas Project at the University of Texas at Austin, 2020. www.comparativeagendas.net.

14 Blevins, C., & Mullen, L. A. (2015). Jane, John ... Leslie? A Historical Method for Algorithmic Gender Prediction. Digital Humanities Quarterly, 9(3). http://www.digitalhumanities.org/dhq/vol/9/3/000223/000223.html

15 We implemented Blevin and Mullen's approach in the gender and genderdata libraries in the R software environment for statistical computing.

16 Barberá, P. (2015). Birds of the same feather tweet together: Bayesian ideal point estimation using twitter data. Political Analysis, 23(1), 76–91. https://doi.org/10.1093/pan/mpu011

17 Barberá, P., Jost, J. T., Nagler, J., Tucker, J. A., & Bonneau, R. (2015). Tweeting From Left to Right: Is Online Political Communication More Than an Echo Chamber? Psychological Science, 26(10), 1531–1542. https://doi.org/10.1177/0956797615594620

18 Users who did not Tweet during any of the previous debates

19 See appendix for the full dictionary of terms used to count candidate mentions.

20 2020 Presidential Election Interactive Map. (n.d.) Taegan Goddard's Electoral Vote Map. https://electoralvotemap.com/

21 The Electoral College consists of 538 electors. A majority of 270 electoral votes is required to elect the President of the United States. Each state has the same number of electoral votes as it does members of its congressional delegation. (See https://www.archives.gov/electoral-college/about.)

22 Smith, A., Hughes, A. Remy, E. & Shah, S. (2020, February 3). Democrats on Twitter more liberal, less focused on compromise than those not on the platform. Pew Research Center. https://www.pewresearch.org/fact-tank/2020/02/03/democrats-on-twitter-more-liberal-less-focused-on-compromise-than-those-not-on-the-platform/

23 Krogstad, J.M. & Keegan, M. (2014, May). 15 States with the highest share of immigrants in their population. Pew Research Center. https://www.pewresearch.org/fact-tank/2014/05/14/15-states-with-the-highest-share-of-immigrants-in-their-population/

24 We removed the responses that consisted of fewer than five words.

25 Barberá, P., Casas, A. Nagler, J. & Egan, P. (2019). Who Leads? Who Follows? Measuring Issue Attention and Agenda Setting by Legislators and the Mass Public Using Social Media Data. American Political Science Review, 113(4), 883-901. https://www.cambridge.org/core/journals/american-political-science-review/article/who-leads-who-follows-measuring-issue-attention-and-agenda-setting-by-legislators-and-the-mass-public-using-social-media-data/D855849CE288A241529E9EC2E4FBD3A8

26 Amber Phillips, E. S. (2019, June 27). Analysis | The first Democratic debate night transcript, annotated. Retrieved from https://www.washingtonpost.com/politics/2019/06/27/transcript-night-one-first-democratic-debate-annotated/

27 The Fix staff. (2019, June 28). Transcript: Night 2 of the first Democratic debate. Retrieved from https://www.washingtonpost.com/politics/2019/06/28/transcript-night-first-democratic-debate/

28 The Fix staff. (2019, July 31). Transcript: The first night of the second Democratic debate. Retrieved from https://www.washingtonpost.com/politics/2019/07/31/transcript-first-night-second-democratic-debate/

29 The Fix staff. (2019, August 1). Transcript: Night 2 of the second Democratic debate. Retrieved from https://www.washingtonpost.com/politics/2019/08/01/transcript-night-second-democratic-debate/

30 The Fix staff. (2019, September 13). Transcript: The third Democratic debate. Retrieved from https://www.washingtonpost.com/politics/2019/09/13/transcript-third-democratic-debate/

31 The Fix staff. (2019, October 16). The October Democratic debate transcript. Retrieved from https://www.washingtonpost.com/politics/2019/10/15/october-democratic-debate-transcript/

32 The Fix staff. (2019, November 21). Transcript: The November Democratic debate. Retrieved from https://www.washingtonpost.com/politics/2019/11/21/transcript-november-democratic-debate/

33 The Fix staff. (2019, December 20). Analysis | Transcript: The December Democratic debate. Retrieved from https://www.washingtonpost.com/politics/2019/12/20/transcript-december-democratic-debate/

34 Des Moines Register. (2020, January 17). Read the full transcript of Tuesday night's CNN/Des Moines Register debate. Retrieved from https://www.desmoinesregister.com/story/news/elections/presidential/caucus/2020/01/14/democratic-debate-transcript-what-the-candidates-said-quotes/4460789002/

35 New Hampshire Democratic Debate Transcript. (n.d.). Retrieved from https://www.rev.com/blog/transcripts/new-hampshire-democratic-debate-transcript

36 Kim, Y. (2014). Convolutional neural networks for sentence classification. arXiv preprint arXiv:1408.5882.

37 Russell, A. (2018). US Senators on Twitter: Asymmetric party rhetoric in 140 characters. American Politics Research, 46(4), 695-723.