Aimed Emotions in American Presidential Politics

Andrew Trexler[a]*

Published in the

Journal of Information Technology & Politics

[a]*Sanford School of Public Policy, Duke University, Durham, NC, USA*

*Corresponding author. Please contact at andrew.trexler@duke.edu

The Version of Record of this manuscript has been published and is available in the *Journal of Information Technology & Politics* (August 17th, 2023).
https://www.tandfonline.com/10.1080/19331681.2023.2248111

Aimed Emotions in American Presidential Politics

Emotional appeals are essential tools for political candidates to motivate supporters, donors, and voters. Prior research has demonstrated the distinct behavioral consequences of discrete emotions, such as anger, anxiety, and enthusiasm. Do political candidates take advantage of these distinctions in their communication strategies? In this paper, I use supervised machine learning to classify emotional content in debate transcripts and contemporaneous tweets of American presidential candidates in the 2016 and 2020 elections, and show that candidates preference different emotional appeals in each communication medium. I argue that this behavior enables candidates to reap strategic benefits from two dissimilar audiences simultaneously.

Keywords: emotion; political communication; campaign strategy; social media; affective intelligence; machine learning

In June of 2019, flanked by nine other candidates vying for the Democratic nomination for president, then-candidate Joe Biden was asked his first debate question of the election cycle: could he clarify his position on income inequality? His response appealed to worries about people being left behind: "Too many people who are in the middle class and poor have had the bottom fall out." His Twitter account fired off a tweet on the same subject soon after, which said nothing of the worries that he had expressed on stage. Instead, he appealed to anti-elite anger and followed with a populist, feel-good sentiment: "This country wasn't built by Wall Street bankers and CEOs and hedge fund managers," he wrote on Twitter, "It was built by you. By ordinary people doing extraordinary things." Later that same evening, his rival Kamala Harris uttered what many pundits considered the most powerful line of the night, expressing disappointment in Biden for his previous support of segregationist politicians. "It was hurtful to hear you talk about [their reputations]. .

.You also worked with them to oppose bussing," she said, calmly but forcefully, "And you know, there was a little girl in California who was part of the second class to integrate her public schools and she was bussed to school every day, and that girl was me." Like Biden, she echoed this same message on Twitter minutes later, but again the tone was different—positive and celebratory rather than critical. Her tweet showed only a picture of a young, school-aged Kamala Harris along with the simple line, "There was a little girl in California who was bussed to school. That girl was me."

Politicians are, by profession, strategic about what they say and to whom. In a campaign context, candidates endeavor to say the right things to the right people—and emotional appeals are part and parcel of that effort. Emotions in politics have received substantial scholarly attention in recent decades, generating rich theories of emotion's effect on political behavior. Accordingly, political communication scholars have increasingly sought to uncover patterns in the use of emotional content—as distinct from substantive content—in political advertising (Brader 2006; Ridout and Searles 2011), campaign emails (Hassell et al. 2022), party manifestos (Crabtree et al. 2020), and social media (Berger and Milkman 2012; Gelman, Wilson, and Petrarca 2021; Webster 2020). Many contemporary studies examine which types of emotional appeals candidates use, changes in strategy over the course of a campaign, and how these messages are received and acted upon by voters. This paper is concerned primarily with the *medium* through which political candidates channel their chosen emotional appeals, and the relevant *audience* reached through that medium. Are politicians conveying different emotional messages to different segments of the electorate by strategically choosing channels of communication?

I contribute to the growing literature on political emotions by demonstrating that American presidential candidates differentiate their emotional messages across divergent media spaces. To do so, I take advantage of differences between the messages that the candidates convey during nationally-televised debates and the public messages (tweets) that they simultaneously publish on the popular social media platform Twitter. As the examples from Biden and Harris in their first debate together show, while each candidate is pursuing a single overarching campaign strategy at a given moment in the race, they are simultaneously communicating two messages with very different emotional content.

In the following section, I discuss theoretical expectations drawn from Affective Intelligence Theory (AIT) about the types of emotional appeals that are likely to appear in each medium. I then test these expectations using transcripts from nationally-televised primary and general election debates in the 2016 and 2020 presidential election cycles, combined with the tweets that each candidate in each debate published around the same time. I find that presidential candidates exhibit strategic behavior in partial accordance with the expectations I develop from AIT: they share comparatively more enthusiastic messages on social media, but favor the debate stage for both anger and anxiety appeals.

Theoretical Expectations

The theoretical approach I adopt is informed primarily by Affective Intelligence Theory (Marcus and MacKuen 1993; Marcus, Neuman, and MacKuen 2000), which has become a prominent theory of emotional or "affective" influences on political behavior since its formal conception thirty years ago. In its current form (Marcus et al. 2019), AIT acknowledges three primary emotion states that develop from two separate subconscious appraisal processes, which feed into conscious thought by shaping perceptions and evaluations of stimuli (Siegel et al. 2018). The first appraisal process, known as the "disposition system," assesses familiar incoming stimuli and assigns appropriate learned, habitual routines stored in procedural memory. Feelings of enthusiasm result from matches between goal-supporting stimuli and available habits, encouraging action via those habits. Feelings of anger, in contrast, result from matches between goal-opposing stimuli and available habits, again encouraging action via those habits to overcome the threat. Enthusiasm and anger are thus both responses to *familiar* stimuli and provide routines from memory that define the behavioral response. The primary difference between these two emotions is that enthusiasm arises from things we like (because they support our goals), and anger from things we do not (because they obstruct our goals), but in both cases the stimulus is recognizable enough that memory can suggest a response.

The second appraisal process is known as the "surveillance system." In contrast to the disposition system, the surveillance system monitors for *unfamiliar* or *uncertain* stimuli, acting as a sort of fire alarm for stimuli for which no suitable routine from memory is available. Feelings of anxiety (often also called "fear") result from this mismatch between environmental stimuli and available habits, discouraging the use of habitual routines, and instead encouraging conscious cognition to develop a considered response to the novel potential threat. Anger and anxiety, both considered negative emotions, thus arise from separate subconscious appraisal processes and produce divergent behavioral outcomes.

Researchers have amassed substantial evidence to support AIT's predictions for political behavior. Scholars have found, for example, that anger and enthusiasm both encourage political participation (Groenendyk and Banks 2014; Jones, Hoffman, and Young 2013; Valentino et al. 2011), that anxiety encourages information-seeking (Albertson and Gadarian 2015; Brader 2006; Brader, Valentino, and Suhay 2008; MacKuen et al. 2010), and that the two negative emotions of anger and anxiety produce dramatically different behavioral responses to political stimuli (Huddy, Feldman, and Cassese 2007; Lerner et al. 2003; Marcus et al. 2019).

These differences can be enormously useful for political candidates if the right emotions can be engendered in the right audience. In the contemporary American context, the voting public is politically polarized (Iyengar et al. 2019) and sorted into two opposing camps defined by multiple dimensions of social identity (Levendusky 2009; Mason 2018). Appealing to fellow partisans with enthusiasm and anger, which encourage reliance on routine (i.e., partisan) habits, can be an effective means for candidates to energize one's political base and boost turnout (Groenendyk and Banks 2014; Webster 2020). Eliciting anxiety in voters can instead serve to dislodge an opponent's supporters from their existing attachments and encourage political learning (Brader 2006), especially when focused on wedge issues (Hillygus and Shields 2008) where the mismatch between the voter's and preferred candidate's positions is unexpected (Johnston, Lavine, and Woodson 2015; Redlawsk, Civettini, and Emmerson 2010).

Sending the wrong emotional message to the wrong audience can harm a candidate's chances. The advertising literature has long grappled with the concept of "voter backlash" (Dowling and Wichowsky 2015; Lau et al. 1999; Ridout and Searles 2011; but see Krupnikov 2011; Mattes and Redlawsk 2015), which describes declining support for a candidate that pursues an advertising strategy that is perceived as too negative. Riling up the opposing party—such as by enthusiastically supporting a policy they dislike, or castigating a policy they treasure—motivates them to support their party's candidate, reducing one's own chances of victory. Raising fears and anxieties among supporters may actually induce paralysis and indecision, weakening turnout (Valentino et al. 2011). For example, raising doubts about election integrity may generate anxiety among party-line supporters—the most likely audience to both receive and accept this message (Zaller 1992)—potentially encouraging them to simply stay home on election day. Candidates therefore ought to be strategic about which emotions they aim to elicit in which audiences.

One means of doing so is to consider the characteristics of the typical audience for a given medium of communication, and to strategically vary the emotional content of messages conveyed through that medium—much like campaigns use digital targeting tools to strategically direct advertising messages to different audiences (Fowler et al. 2021). A candidate doing a tour of the cable networks' Sunday talk shows might consider, for example, both the typical viewer of Sunday talk shows in general, as well as the partisan lean of each network's audience. A Democrat appearing on *Fox News Sunday* should recognize that many viewers are reliable Republicans, and therefore avoid expressing anger about Republican policies—a tactic which could instead be quite fruitful when appearing on NBC's *Meet the Press*, which has a more liberal viewership.

In this paper, I focus on the contemporary American context, yet the same strategic consideration of *medium* and likely *audience* when manipulating emotional appeals is broadly applicable to electoral politics, at least in contexts where political communication has diversified and fragmented. That is, so long as politicians can direct appeals through distinct media and expect to reach different audiences, they likely have incentives to make different strategic emotional appeals in each. Here, I examine emotional appeals communicated through two media: nationally-televised candidate debates and the social media platform Twitter. Each medium has unique characteristics that could inform the strategic use of emotional messages. Both media are public-facing and are ostensibly directed at the general American public. Yet this public orientation belies key differences. Televised debates (in principle) constitute rare opportunities for political learning because the candidates are forced to contend with each other's ideas in direct dialogue, and are therefore attractive to voters who hope to gain information helpful for deciding which candidate to support (Benoit 2011). Indeed, many voters report the debates as a crucial period for deciding whom to support (Benoit and Hansen 2004; Bystrom et al. 1991; but see Prior 2012). Debating candidates can certainly expect some of their own strong supporters to be following along as well, or even hosting “watch parties”; a strong showing can be helpful for shoring up these voters’ loyalty. But, crucially, debaters can also expect that a large portion of the audience will be supporters of *other* candidates. Through a combination of pizazz and persuasive rhetoric, debaters can hope to alter voters’ impressions (Kenski and Jamieson 2011) and win some of their competitors’ supporters to their side—especially those who are only weakly attached to another candidate.

Twitter provides a much more skewed audience. The population of American Twitter users is younger, wealthier, and more ideologically extreme than the US general population (among other differences; see Barberá 2015), and these disparities are even more pronounced among those active in political discussions on Twitter (Barberá and Rivero 2015). Because of the platform’s directed network structure, however, each politician on the platform is communicating primarily with a narrow and peculiar slice of Twitter users: their “followers.” These, unsurprisingly, are dominated by co-partisans and likely supporters.[1] And while every candidate hopes to broaden their reach beyond direct followers through retweets and coverage by the traditional media (McGregor 2019), the network structure of the platform ensures that the bulk of the audience for most messages is likely to be sympathetic. On Twitter, then, a candidate can worry much less about upsetting out-partisans, but does need to worry about raising wedge issues (or otherwise generating anxiety) among their own base of support.

With these audience-oriented differences in mind, I expect that anger (Hypothesis 1a) and enthusiasm (Hypothesis 1b) appeals should be more prevalent on Twitter than on the debate stage.

[1] Indeed, this fact is essential to research that estimates the ideological leanings of the “Twitter public” by analyzing the platform’s directed network structure under a homophily assumption. Such mappings have strong face validity and convergent validity with other measures of ideology for political figures, such as Congress (Barberá 2015; King, Orlando, and Sparks 2016).

AIT argues that both anger and enthusiasm reinforce existing preferences, including partisan attachments as learned routines in memory. Because the audience on Twitter skews heavily towards existing supporters, anger and enthusiasm appeals are likely more effective there than when directed at undecided or leaning voters watching a debate. Note that we should still expect some anger and enthusiasm appeals on the debate stage—some die-hard supporters are no doubt watching as well—but these should appear less frequently in the debate messages than on Twitter. In contrast to the effects of anger and enthusiasm, AIT argues that anxiety discourages reliance on existing preferences, and instead encourages information-seeking. Because this emotional response is much more useful for persuading the undecided or weakly-attached voters more commonly found in debate audiences, I expect that anxiety appeals are more prevalent on the stage than on Twitter (Hypothesis 2). Again, this expectation does not mean that anxiety appeals should be absent on Twitter—not every follower is a supporter—but these should occur infrequently relative to what candidates say in front of the cameras.

Data & Methods

To test the above hypotheses, I analyze emotional appeals deployed by leading US presidential candidates of both parties in the 2016 and 2020 election cycles. Specifically, I analyze the emotional appeals present in these candidates' statements during all nationally-televised candidate debates, and in the messages that each candidate publishes to Twitter around the same time as their appearance on stage. In combination, these data allow me to evaluate differences in emotional content in messages from the same candidates at the same time across two public media environments. Previous work has explored evolutions in emotive strategy over the course of a campaign (e.g., Ridout and Searles 2011) or variation across candidates (e.g., Brader 2006;

Webster 2020). However, the communication strategy of a given candidate is likely to change over time, particularly for candidates that become their party's nominee and compete in the general election. Each candidate can also reasonably be expected to pursue a unique communication strategy, depending on their polling position in the race, personal background, issues priorities, and similar considerations. Instead, I take advantage of individual candidates communicating simultaneously through separate media with disparate audiences. The comparison here is more direct: the candidates are the same, the points in time are the same, and the stakes of the race are the same—only the medium, and thus the expected audience, differs.

The analyses below center on data from the 38 US presidential candidates who participated in at least one nationally-televised candidate debate during the 2016 or 2020 presidential election cycles.[2] As candidates for one of the world's most powerful elected offices, these individuals collectively spend billions of dollars on their campaigns. They simultaneously contest their political messages across a wide variety of media spaces while endeavoring to assemble broad coalitions, yet their messages are also broadly public and often subject to intense scrutiny by the press. As such, these candidates should be particularly strategic in how they communicate—including their use of emotional appeals—providing an ideal test case for distinguishing appeal strategies by medium.[3]

The televised debates under study consist of 12 Republican primary debates, nine Democratic primary debates, and three general election debates from the 2016 cycle, as well as 13 Democratic primary debates and two general election debates from the 2020 cycle. A list of the

[2] I do not include participants in the "junior varsity" second-tier debates held during the Republican primary season for the 2016 election, except when they appear in the "primary" first-tier debates.
[3] Still, one might still expect candidates in other contexts to exhibit similar behavior: while they may not have access to the same resources or media spaces, they might also face less public scrutiny and feel greater latitude to express different emotions through different channels.

debates and the participants is provided in Tables A1 and A2 in the Appendix. I collected transcripts from each of these debates, relying on the American Presidency Project archive at the University of California Santa Barbara for the 2016 cycle transcripts and various major news organizations (*The New York Times*, *The Washington Post*, CNN, CBS) for the 2020 cycle transcripts. The transcripts were lightly edited to ensure consistency in transcription across the set. Each statement spoken by each candidate was then separated into individual sentences, which form the units of analysis for the debate messages. A sentence can be as short as a single word—especially during verbal clashes when candidates speak over one another—but are approximately 15 words long on average, while (typically) containing a complete idea. The sentence is thus comparable with the average tweet (23 words) in terms of the information conveyed per unit of analysis. The dataset includes 43,739 sentences from candidate statements in total.

To obtain tweets from these same candidates, I collected all tweets from the primary individual Twitter account[4] for each candidate published within three days of each debate in which they participated. That is, for the Republican primary debate that occurred on 2015 August 6th, I collected tweets from the ten participating candidates over the course of the week from August 3rd through August 9th. Tweets are typically very short messages and had a maximum limit of 280 characters;[5] in this dataset, the average length is 23 words, slightly longer than the average sentence spoken in debate. Individual tweets are thus the unit of analysis in that medium, and the dataset comprises 29,972 total tweets published by the respective participating candidates in the one-week

[4] By "primary individual" account, I mean the account each candidate uses primarily as a personal account for campaigning, as distinguished from "official" accounts tied to their existing office (e.g., Senate office) and from "team" accounts that communicate on behalf of the campaign rather than the individual candidate. For example, for Elizabeth Warren I collected tweets from @ewarren, but not @SenWarren or @TeamWarren. A list of the accounts is provided in Appendix Table A3. Tweets for all candidates except Donald Trump were collected using the academictwitteR package in R on 2021 August 10th. Trump's tweets were collected previously in 2019 and 2020 via the twitteR package in R, and merged with the rest of the data.

[5] During the 2016 election cycle, the limit was 140 characters; Twitter doubled this limit to 280 characters in 2017.

period surrounding a given debate. The combined dataset thus contains 73,711 total messages from US presidential candidates across both media in the 2016 and 2020 cycles.

To evaluate the emotional content of each message, I adopt a machine learning approach using RoBERTa (Robustly optimized Bidrectional Encoder Representations from Transformers; see Liu et al. 2019) neural net classifiers. In brief, RoBERTa is a pre-trained general-use neural net recently developed by the natural language processing community, and is suitable for a variety of machine learning functions, including the classification of text. A major advantage of RoBERTa is that it can consider the context in which a given word occurs (i.e., what comes before and after the word) to evaluate its meaning, in contrast to many other common forms of text analysis which rely on the so-called "bag of words" assumption and treat each word as independent (e.g., topic modeling or dictionary-based lexicons).[6] A second major advantage is that RoBERTa is a form of "supervised" machine learning, meaning that a classification model can be further trained with a small hand-coded dataset to improve its performance on a specific downstream classification task. This training process allows the researcher to adapt RoBERTa to the specific context in which the data should be understood, and to the specific identification task of a particular emotion, by providing the neural net with comparison information on human evaluations of the same material.

For the main analyses that follow, I trained six RoBERTa classifiers: three classifiers to classify debate sentences on the presence of anger, anxiety, and enthusiasm (respectively), and three comparable classifiers to classify tweets on the same dimensions. To do so, I drew a random sample of 4,000 debate sentences and 4,000 tweets from the census datasets, and trained three research assistants to hand-code these messages as containing (1) or not containing (0) an appeal

[6] Indeed, dictionary-based methods for emotion classification are highly variable and exhibit poor inter-reliability and low convergent validity (Soroka, Young, and Balmas 2015); this is one of several reasons why much recent work on text analysis in political science favors supervised learning over dictionary approaches (e.g., Barberá et al. 2021; Grimmer, Roberts, and Stewart 2022).

to each of the three respective emotions. These coders were not aware of the hypotheses or research question, and were only provided the author, date, medium, and text of each message during the coding process.[7] A subset of the debate sentences and tweets (1,000 each) were coded independently by all three coders for the purpose of evaluating reliability;[8] the remaining 3,000 messages from each medium were each coded by a single coder. The inter-coder reliability scores were all in an acceptable if generally moderate range.[9]

These hand-coded data were then used to train the emotion classifiers (that is, the coded debate sentences were used to train the three debate classifiers, and the coded tweets to train the three tweet classifiers).[10] All classifiers were evaluated via out-of-sample validation; performance across the classifiers ranged from 78 to 86 percent accurate in the out-of-sample tests (performance statistics for each classifier are reported in Appendix Table A4).[11] All messages from each respective medium were then classified by the three trained classifiers for that medium.

---

[7] To train the research assistants, I provided and reviewed brief materials on anger, anxiety, and enthusiasm as defined under AIT, worked with the assistants to code several example messages, and provided two practice datasets for each assistant to code individually. After each practice dataset was completed, we met as a group to discuss differences in coding choices. Given the wide variety of topics discussed in these messages, no formal coding scheme was developed. After training, the assistants proceeded to code their respective datasets individually and independently; other assistants provided clarification or advice only in the extremely rare event that it was requested by a coder.

[8] For the use of these messages as training data, an emotional appeal was coded as present (1) if at least two of the three coders identified it as present, and 0 otherwise.

[9] For debate statements, the anger measure had 66 percent complete agreement (Cronbach's $\alpha = 0.61$), the anxiety measure 61 percent complete agreement ($\alpha = 0.66$), and the enthusiasm measure 76 percent complete agreement ($\alpha = 0.82$). For the tweets, the anger measure had 74 percent complete agreement ($\alpha = 0.74$), the anxiety measure 65 percent complete agreement ($\alpha = 0.76$), and the enthusiasm measure 51 percent complete agreement ($\alpha = 0.68$). Emotion is difficult to classify, even for human coders (see also Ridout and Searles 2011); the acceptable but mediocre reliabilities are thus unsurprising. Nevertheless, the accuracy ratings from the trained RoBERTa classifiers that I employ here represent a significant increase in reliability over common alternative approaches (e.g., dictionary-based lexicon classification; see Soroka, Young, and Balmas 2015).

[10] Classifier training and classification was implemented in Python via the simpletransformers package, using the Google CoLab environment.

[11] The accuracy of these classifiers compares very favorably against standard dictionary-based approaches. For example, one common dictionary-based classifier anger and anxiety, known as LIWC (Pennebaker et al. 2015), only accurately identifies 6 percent (anxiety) to 19 percent (anger) of true positives (recall percentage). For the RoBERTa classifiers, true positive recall ranges from 59 to 73 percent for these two emotions (see Appendix Table A4).

Results

Are American presidential candidates communicating more anger and enthusiasm appeals on Twitter, and more anxiety appeals during debates? To establish general patterns in the data, I first examine simple proportions of emotional appeals present in each medium. Figure 1 shows the proportion of messages from each medium that are classified as an appeal to anger, anxiety, or enthusiasm. Candidates do, in fact, make regular use of emotional appeals—as we would expect, given the persuasion and motivation aims of candidates in the campaign context. Across both media domains, 75 percent of all messages are classified as an appeal to at least one of the three emotions, including 67 percent of debate sentences and 86 percent of tweets.[12] Figure 1 also provides initial evidence that candidates vary the types of emotional messages they convey across the two media. Enthusiasm appeals, in particular, appear much more frequently in messages on Twitter than in debate. In contrast, negative emotional appeals to both anger and anxiety seem to be present in the messages on both media at roughly the same rate, with anger showing up in debate messages at a slightly higher frequency than on Twitter.

[Figure 1]

By looking at simple proportions, however, one cannot be certain that these differences are not driven by differing strategies of particular candidates, who may vary in the extent to which they communicate in each medium across a long campaign season. I estimate several logistic regression models that use the presence (1) or absence (0) of an anger, anxiety, or enthusiasm appeal in a given message as the outcome variable, and a main independent variable as a binary variable that identifies a given message as a tweet (1) or a debate statement (0). In each model, I also control for a baseline effect of whether a given message has any emotional appeal (1) or not

[12] These baseline rates are similar to those found in political advertising (Brader 2006; Fowler and Ridout 2012; Ridout and Searles 2011).

(0); the coefficient on the tweet variable is thus the expected change in log-odds that a message includes an (e.g.) anger appeal if that message is a tweet, holding constant the chance that the message includes *any* emotional appeal. I do this in order to evaluate the likelihood that each medium favors a *particular* flavor of emotion over other possible choices. A baseline model (Model 1) includes only these two indicator variables on the right-hand side, and is reported in columns 1-3 of Table 1.

I control for several relevant campaign and candidate factors in subsequent models. To account for changes key differences in the post-nomination period, I include an indicator variable for whether each message occurred during the post-convention general election phase of the campaign (1) or during the primary (0). To control for systemic differences in the types of appeals that candidates might prefer in the 2020 versus 2016 cycle, I include an indicator variable for the 2020 cycle (1) or not (0). Some voters punish female candidates for displays of anger that are accepted among male candidates (Boussalis et al. 2021), with similar racialized perceptions held against Black candidates (Wingfield 2007)—and news coverage often presents candidates' emotional communication through a gendered and racialized lens (Harmer, Savigny, and Siow 2021), leading women and minority candidates face different sets of incentives for emotional appeals. I therefore include an indicator for candidate gender and an indicator for racial minority status. Model 2 includes all of these covariates as controls, and is reported in columns 4-6 of Table 1. In Model 3, reported in columns 7-9 of Table 1, I also include fixed effects for each debate event for further granularity in changes over time, and fixed effects for each candidate (for sake of brevity, the coefficients for these fixed effects are not shown). Because each candidate can be assumed to prefer their own unique campaign messaging strategy—which may include preferences to favor (or avoid) negative (or positive) emotional messages—these fixed effects account for baseline propensities for each candidate to deploy a given emotional appeal in any medium.[13]

The evidence presented in Table 1 provides strong support for H2, that anxiety appeals should be less commonly deployed on Twitter, and more commonly deployed on the debate stage. Across all three model specifications, the coefficient on the tweet indicator is negative and highly significant ($p < .001$), reflecting that tweets are less likely to contain an anxiety appeal than debate messages. Table 1 shows mixed support for H1, however. Consistent with expectations (H1b), enthusiasm appeals are more likely to appear in tweets than debate messages across all three models ($p < .001$). But contrary to expectations (H1a), anger appeals are *less* likely to appear on Twitter rather than on the debate stage ($p < .001$). These results suggest that both negative emotions, anger and anxiety, are viewed by candidates as strategically valuable in a similar rather than divergent fashion: as negative attacks. I return to this possibility in the Discussion.

[Table 1]

Presidential campaigns are long affairs, and are commonly divided into two major phases: the primary campaign in which members of the same party compete with each other for their party's nomination, and the general election campaign in which the two major party's nominees go head-to-head. Although these two phases qualitatively differ in a number of important ways, the dynamics of concern here regarding the use of emotional appeals remain quite similar. The differences in audience characteristics across the two media are largely unchanged, and candidates in both phases can strategically direct their emotional appeals to take advantage of those audience differences. In Model 4, reported in columns 1-3 of Table 2, I explore whether the two overarching

[13] Indeed, as expected by the literature on gendered politics (e.g., Boussalis et al. 2021), women running for president exhibit less anger and more enthusiasm than their male counterparts, as Table 1 shows. Further, qualitative differences between the candidates who entered the 2020 contest and their 2016 predecessors appear to account for much of the overall change in emotional appeals between the cycles, as the candidate fixed effects in Model 3 wipe out the significant coefficients on the cycle indicator variable in Model 2.

campaign phases bear out similar patterns in the types of emotional messages the candidates use in each medium. To do so, I include an interaction variable for whether a message is a tweet posted during the general election phase, as well as an interaction that accounts for the baseline chance of comprising any emotional appeal during the general election phase. The main independent variables of interest are thus the tweet variable and its interaction with the general election variable, the coefficients of which respectively reflect the expected change in log-odds of a message containing an (e.g.) anger appeal if the message is a tweet, and if the message is a tweet that was posted during the final months of the campaign. The results show that the same patterns hold across both major phases of the campaign, but the differences across the two media are even more pronounced for anxiety and enthusiasm appeals in the general election phase. Anxiety appeals are favored on the debate stage rather than Twitter, but especially so during the final weeks of the campaign. The opposite is true for enthusiasm appeals, which appear most frequently on Twitter in the closing weeks—as we might expect, as campaigns increasingly pursue turnout mobilization over persuasion efforts as election day nears.

[Table 2]

Finally, I examine differences between Democrats and Republicans in their use of emotional appeals. I subset the data to Democratic (Republican) candidates to estimate Model 5 (Model 6) using a specification equivalent to Model 3; that is, covariate controls for any emotional content, campaign phase, campaign cycle, gender and racial identity, event-level fixed effects, and candidate-level fixed effects. Comparing Model 5 (columns 4-6) with Model 6 (columns 7-9), the candidates of both parties exhibit similar strategic choices, but the differences across the two media are much more pronounced for Republicans. Additionally, Democrats appear to favor tweets for anxiety appeals. This result casts some doubt on the support for H2 shown in the other models, and suggests that Democrats and Republicans may have pursued appreciably different emotional communication strategies in 2016 and 2020.

Figure 2 summarizes the findings with respect to the key coefficient of interest, the tweet indicator coefficient, across all six models for each of the three emotion outcomes. As Figure 2 shows, the coefficients are remarkably consistent, and their support (or not) for the hypotheses is reasonably robust to variety of model specifications. The results thus provide ample evidence of strategic behavior among presidential candidates in their emotive messaging strategies, by advantageously varying the emotions they appeal to across divergent media according to the characteristics of the audience. However, anger and anxiety (at least for Republicans) appear to be treated somewhat analogously as negative emotions, rather than discrete emotions with differing strategic implications.

[Figure 2]

I conduct three robustness checks of the above results. First, I replicate Models 1-6 with a restriction of the tweet data to the day-of each debate or the day after (when many modified quotes from the debate are posted to Twitter), rather than up to three days before or after, as presented above. This check ensures that the results are not driven by the tweets that are sent several days before or after each debate—in other words, that the results do not hinge on the specification of what "around the same time" means. The results of these replications (reported in Tables A5 and A6 in the Appendix) are broadly similar to the main findings; the negative coefficient for Twitter on anxiety is no longer significant in Model 4, but the interaction with the general election remains strongly negative and significant, suggesting a particularly strong role for anxiety appeals in the general election debates.

Second, I train another set of RoBERTa-based classifiers, but trained only on the 1,000 debate sentences and 1,000 tweets that were hand-coded by all three research assistants (using the majority's coding when they disagreed) as the training data. While each classifier has less data to work with during the training process, the training data is likely more consistent because each message was independently coded by three human researchers. I then use these new classifiers to reclassify the entire dataset, and replicate Models 1-6 using these classifications (reported in Tables A7 and A8 in the Appendix). The results of this replication are again very similar to the main results. As in the previous replication, the tweet coefficient on anxiety in Model 4 is no longer significant (but remains negative and significant in the interaction with the general election variable). Additionally, the tweet coefficient on anger in Model 5 (Democrats) is weakly positive, with strong differences by gender and racial identity, highlighting the potential role for these factors in shaping emotional appeal strategies among Democrats in particular.

Finally, because maximum likelihood estimators can produce inconsistent estimates in fixed effects models, I replicate the main analyses of Models 1-6 with linear probability models instead of logistic regression. The results of this replication are again substantively identical to the main findings, and are reported in Tables A9 and A10 of the Appendix.

Discussion

Although students of politics have long recognized politicians' proclivity to say one thing in one room, and quite another thing in another room, I demonstrate here that leading major party candidates substantially vary their emotional appeals through two media that are visible to a general public audience *at the same time*. In particular, I have shown that candidates of both parties vying for the highest office in the land preference Twitter for deploying enthusiasm appeals, a medium on which most of the message recipients are likely to be sympathetic. In contrast, most candidates (though perhaps not most Democrats) favor the debate stage for anxiety appeals, where encouraging political learning might be most fruitful.

My expectations with respect to the use of anger appeals, however, did not fare well. In contrast to the expectation that anger, as an attachment-reinforcing emotion, should be favored on Twitter rather than the stage, I find that candidates of both parties consistently prefer debates as a medium for conveying messages of anger. Although this finding does not fit with the behavioral consequences of anger that has been consistently identified by AIT researchers, it does fit with previous research in suggesting that political campaign staffers may think primarily in terms of valence, viewing both anger and anxiety appeals as negative messages that serve a similar strategic purpose (Ridout and Searles 2011). Expressions of anger on the debate stage may thus be intended to provide novel information about other candidates to the debate audience, as a means of prompting political learning and information seeking that could prompt weakly-attached voters to abandon opposing candidates. For example, a candidate may choose to deploy an anger appeal to emphasize that their opponent supports a broadly unpopular policy (or opposes a popular one), demonstrating that their opponent violates many voters' policy expectations and encouraging such voters to view that opponent less favorably (Johnston, Lavine, and Woodson 2015). Appeals to anger may thus prompt *feelings* of anxiety among those in the audience who both support the target candidate and disagree on policy. Twitter may be a less appropriate medium for this type of behavior precisely because its audience is skewed towards supporters, and cross-pressured supporters of other candidates are comparatively difficult to reach on the platform. Future research could tease out these potential mismatches between the appealed emotion, the strategic purpose, and the experience emotion in the target audience.

Tailored emotional messages such as these can have broad consequences beyond the implications for electoral politics. Emotions are central to political attitudes and behaviors, and negative emotions in particular can be particularly harmful for democratic enterprise. Previous work has connected the experience of anger to the activation of racial animus (Banks 2014) and authoritarian preferences (Marcus et al. 2019), as well as damaged trust in government and diminished support for key democratic principles (Webster 2020). Anxiety encourages active thinking but often leads to biased searches that prioritize information relieving the feeling of anxiety (Albertson and Gadarian 2015; Brader, Valentino, and Suhay 2008; but see MacKuen et al. 2010) rather than prioritizing accurate information. Beyond racial animus, recent scholarship has shown that the effects of emotional states on political behavior differ across racial groups (Phoenix 2020). As political campaigns around the world grow increasingly sophisticated in their ability to micro-target their messages to specific audiences, they will find opportunities to manipulate not only campaign behavior, but also larger patterns in the relationship between citizen and state. The acrimony of the American two-party system perhaps exacerbates these concerns, but is not a prerequisite: strategic manipulation of emotional appeals (and its consequences) is a feature of less polarized multiparty systems as well (Boussalis et al. 2021; Marcus et al. 2019). Further scholarly attention to political elites' strategic emotional communication targeting distinct segments of the mass public, and the downstream consequences of those emotive messages, is vital for understanding political polarization and democratic backsliding—both current and future. This paper contributes to that work by demonstrating the emotive targeting phenomenon occurring at a broad level in the upper echelons of American campaign politics. Similar studies in additional contexts can begin to unpack the many applications of this phenomenon, and their implications.

Acknowledgements: The DeWitt Wallace Center for Media and Democracy at Duke University generously supported this project. I thank Gus Gress, Euwan Kim, and Abigail Phillips for their valuable research assistance. I thank Kevin Aslett, Carolyn Barnes, Nick Carnes, Amy Funck, Chris Johnston, Sunshine Hillygus, Marayna Martinez, Phil Napoli, Marcos Rangel, David Redlawsk, Ken Rogerson, Mallory SoRelle, and two anonymous reviewers for helpful comments on the manuscript. Early manuscripts for this paper were previously presented at the 2021 Annual Meeting of the American Political Science Association and the 2022 Comparative Political Communication workshop at Duke University.

Funding Details: This work was supported by the DeWitt Wallace Center for Media and Democracy at Duke University.

Disclosure Statement: The author declares no competing interests.

Biographical Note: Andrew Trexler is a doctoral candidate of public policy and political science at Duke University.

Data Availability Statement: Data and replication materials for this article can be accessin the Journal of Information Technology & Politics Dataverse: https://doi.org/10.7910/DVN/LHT5MB.

References

Albertson, Bethany, and Shana K. Gadarian. 2015. *Anxious Politics: Democratic Citizenship in a Threatening World*. Cambridge: Cambridge University Press.

Banks, Antoine J. 2014. *Anger and Racial Politics: The Emotional Foundation of Racial Attitudes in America*. New York: Cambridge University Press.

Barberá, Pablo. 2015. "Birds of the Same Feather Tweet Together: Bayesian Ideal Point Estimation Using Twitter Data." *Political Analysis* 23: 76-91.

Barberá, Pablo, Amber E. Boydstun, Suzanna Linn, Ryan McMahon, and Jonathan Nagler. 2021. "Automated Text Classification of News Articles: A Practical Guide." *Political Analysis* 29 (1): 19-42.

Barberá, Pablo, and Geraldo Rivero. 2015. "Understanding the Political Representativeness of Twitter Users." *Social Science Computer Review* 33 (6): 712-729.

Benoit, William L. 2011. *Political Election Debates: Informing Voters about Policy and Character*. Plymouth: Lexington Books.

Benoit, William L., and Glenn J. Hansen. 2004. "Presidential debate watching, issue knowledge, character evaluation, and vote choice." *Human Communication Research* 30 (1): 121-144.

Berger, Jonah, and Katherine L. Milkman. 2012. "What Makes Online Content Viral?" *Journal of Marketing Research* 49 (2): 192-205.

Brader, Ted. 2006. *Campaigning for Hearts and Minds: How Emotional Appeals in Political Ads Work.* Chicago: University of Chicago Press.

Brader, Ted, Nicholas A. Valentino, and Elizabeth Suhay. 2008. "What Triggers Public Opposition to Immigration? Anxiety, Group Cues, and Immigration Threat." *American Journal of Political Science* 52 (4): 959-978.

Boussalis, Constantine, Travis G. Coan, Mirya R. Holman, and Stefan Müller. 2021. "Gender, Candidate Emotional Expression, and Voter Reactions During Televised Debates." *American Political Science Review* 115 (4): 1242-1257.

Bystrom, D. G., Cynthia Roper, Robert Gobetz, Tom Massey, and Carol Beall. 1991. "The effects of a televised gubernatorial debate." *Political Communication Review* 16: 57-80.

Crabtree, Charles, Matt Golder, Thomas Gschwend, and Indriđi H. Indriđason. 2020. "It Is Not Only What You Say, It Is Also How You Say It: The Strategic Use of Campaign Sentiment." *Journal of Politics* 82 (3): 1044-1060.

Dowling, Conor M., and Amber Wichowsky. 2015. "Attacks without Consequence? Candidates, Parties, Groups, and the Changing Face of Negative Advertising." *American Journal of Political Science* 59 (1): 19-36.

Fowler, Erika Franklin, Michael M. Franz, Gregory J. Martin, Zachary Peskowitz, and Travis N. Ridout. 2021. "Political Advertising Online and Offline." *American Political Science Review* 115 (1): 130-149.

Fowler, Erika Franklin, and Travis N. Ridout. 2012. "Negative, Angry, and Ubiquitous: Political Advertising in 2012." *The Forum* 10 (4): 51-61.

Gelman, Jeremy, Steven Lloyd Wilson, and Constanza Sanhueza Petrarca. 2021. "Mixing Messages: How Candidates Vary in Their Use of Twitter." *Journal of Information Technology & Politics* 18 (1): 101-115.

Grimmer, Justin, Margaret E. Roberts, and Brandon M. Stewart. 2022. *Text as Data: A New Framework for Machine Learning and the Social Sciences*. Princeton: Princeton University Press.

Groenendyk, Eric, and Antoine J. Banks. 2014. "Emotional Rescue: How Affects Helps Partisans

Overcome Collective Action Problems." *Political Psychology* 35 (3): 359-378.

Hassell, Hans J.G., Christopher D. Johnston, Jessica Khan, and Edgar Cook. 2022. "The Nature and Impact of Emotional Content in Congressional Candidate Emails to Supporters." *Electoral Studies* 79: 102501.

Harmer, Emily, Heather Savigny, and Orly Siow. 2021. "Trump, Clinton, and the Gendering of Newspaper Discourse About the 2016 U.S. Presidential Election Debates." *Women's Studies in Communication* 44 (1): 81-101.

Hillygus, D. Sunshine, and Todd G. Shields. 2008. *The Persuadable Voter: Wedge Issues in Presidential Campaigns*. Princeton: Princeton University Press.

Huddy, Leonie, Stanley Feldman, and Erin Cassese. 2007. "On the Distinct Political Effects of Anxiety and Anger." In *The Affect Effect: Dynamics of Emotion in Political Thinking and Behavior*, edited by W. Russell Neuman, George E. Marcus, Ann N. Crigler and Michael Mackuen: 202-230. Chicago: University of Chicago Press.

Iyengar, Shanto, Yphtach Lelkes, Matthew Levendusky, Neil Malhotra, and Sean J. Westwood. 2019. "The Origins and Consequences of Affective Polarization in the United States." *Annual Review of Political Science* 22: 129–46.

Johnston, Christopher D., Howard Lavine, and Benjamin Woodson. 2015. "Emotion and Political Judgement: Expectancy Violation and Affective Intelligence." *Political Research Quarterly* 68 (3): 474-492.

Jones, Philip Edward, Lindsay H. Hoffman, and Dannagal G. Young. 2013. "Online Emotional Appeals and Political Participation: The Effect of Candidate Affect on Mass Behavior." *New Media & Society* 15 (7): 1132-1150.

Kenski, Kate, and Kathleen Hall Jamieson. 2011. "Presidential and Vice Presidential Debates in

2008: A Profile of Audience Composition." *American Behavioral Scientist* 55 (3): 307-324.

King, Aaron S., Frank J. Orlando, and David B. Sparks. 2016. "Ideological Extremity and Success in Primary Elections: Drawing Inferences from the Twitter Network." *Social Science Computer Review* 34 (4): 395-415.

Krupnikov, Yanna. 2011. "When Does Negativity Demobilize? Tracing the Conditional Effect of Negative Campaigning on Voter Turnout." *American Journal of Political Science* 55 (4): 796-812.

Lau, Richard R., Lee Sigelman, Caroline Heldman, and Paul Babbitt. 1999. "The Effects of Negative Political Advertisements: A Meta-Analytic Assessment." *American Political Science Review* 93 (4): 851-875.

Lerner, Jennifer S., Roxana M. Gonzalez, Deborah A. Small, and Baruch Fischhoff. 2003. "Effects of Fear and Anger on Perceived Risks of Terrorism: A National Field Experiment." *Psychological Science* 14 (2): 144-150.

Levendusky, Matthew. 2009. *The Partisan Sort: How Liberals Became Democrats and Conservatives Became Republicans*. Chicago: University of Chicago Press.

Liu, Yinhan, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. "RoBERTa: A Robustly Optimized BERT Pretraining Approach." arXiv:1907.11692.

MacKuen, Michael, Jennifer Wolak, Luke Keele, and George E. Marcus. 2010. "Civic Engagements: Resolute Partisanship or Reflective Deliberation." *American Journal of Political Science* 54 (2): 440-458.

Marcus, George E., and Michael B. MacKuen. 1993. "Anxiety, Enthusiasm, and the Vote: The

Emotional Underpinnings of Learning and Involvement during Presidential Campaigns." *American Political Science Review* 87 (3): 688-701.

Marcus, George E., W. Russell Neuman, and Michael MacKuen. 2000. *Affective Intelligence and Political Judgement.* Chicago: University of Chicago Press.

Marcus, George E., Nicholas A. Valentino, Pavlos Vasilopoulos, and Martial Foucault. 2019. "Applying the Theory of Affective Intelligence to Support for Authoritarian Policies and Parties." *Advances in Political Psychology* 40 (S1): 109-139.

Mason, Lilliana. 2018. *Uncivil Agreement: How Politics Became Our Identity*. Chicago: University of Chicago Press.

Mattes, Kyle, and David P. Redlawsk. 2015. *The Positive Case for Negative Campaigning*. Chicago: University of Chicago Press.

McGregor, Shannon. 2019. "Social Media as Public Opinion: How Journalists Use Social Media to Represent Public Opinion." *Journalism* 20 (8): 1070-1086.

Pennebaker, James W., Ryan L. Boyd, Kayla Jordan, and Kate Blackburn. 2015. *The Development and Psychometric Properties of LIWC2015*.

Prior, Markus. 2012. "Who Watches Presidential Debates? Measurement Problems in Campaign Effects Research." *Public Opinion Quarterly* 76 (2): 350-363.

Phoenix, Davin L. 2020. *The Anger Gap: How Race Shapes Emotion in Politics.* Cambridge: Cambridge University Press.

Redlawsk, David P., Andrew J. W. Civettini, and Karen M. Emmerson. 2010. "The Affective Tipping Point: Do Motivated Reasoners Ever 'Get It'?" *Political Psychology* 31 (4): 563-593.

Ridout, Travis, and Kathleen Searles. 2011. "It's My Campaign and I'll Cry if I Want to: How and When Campaigns Use Emotional Appeals.” *Political Psychology* 32 (3): 439-458.

Siegel, Erika H., Jolie B. Wormwood, Karen S. Quigley, and Lisa Feldman Barrett. 2018. “Seeing What You Feel: Affect Drives Visual Perception of Structurally Neutral Faces.” *Psychological Science* 29 (4): 496-503.

Soroka, Stuart, Lori Young, and Meital Balmas. 2015. "Bad News or Mad News: Sentiment Scoring of Negativity, Fear, and Anger in News Content." *Annals of the American Association of Political and Social Science* 659 (1): 108-21.

Valentino, Nicholas A., Ted Brader, Eric W. Groenendyk, Krysha Gregorowicz, and Vincent L. Hutchings. 2011. "Election Night’s Alright for Fighting: The Role of Emotions in Political Participation." *The Journal of Politics* 73 (1): 156-70.

Webster, Steven W. 2020. *American Rage: How Anger Shapes Our Politics.* Cambridge: Cambridge University Press.

Wingfield, Adia Harvey. 2007. “The Modern Mammy and the Angry Black Man: African American Professionals’ Experiences with Gendered Racism in the Workplace.” *Race, Gender, & Class* 14 (1-2): 196-212.

Zaller, John R. 1992. *The Nature and Origin of Mass Opinion*. New York: Cambridge University Press.

Figures and Tables

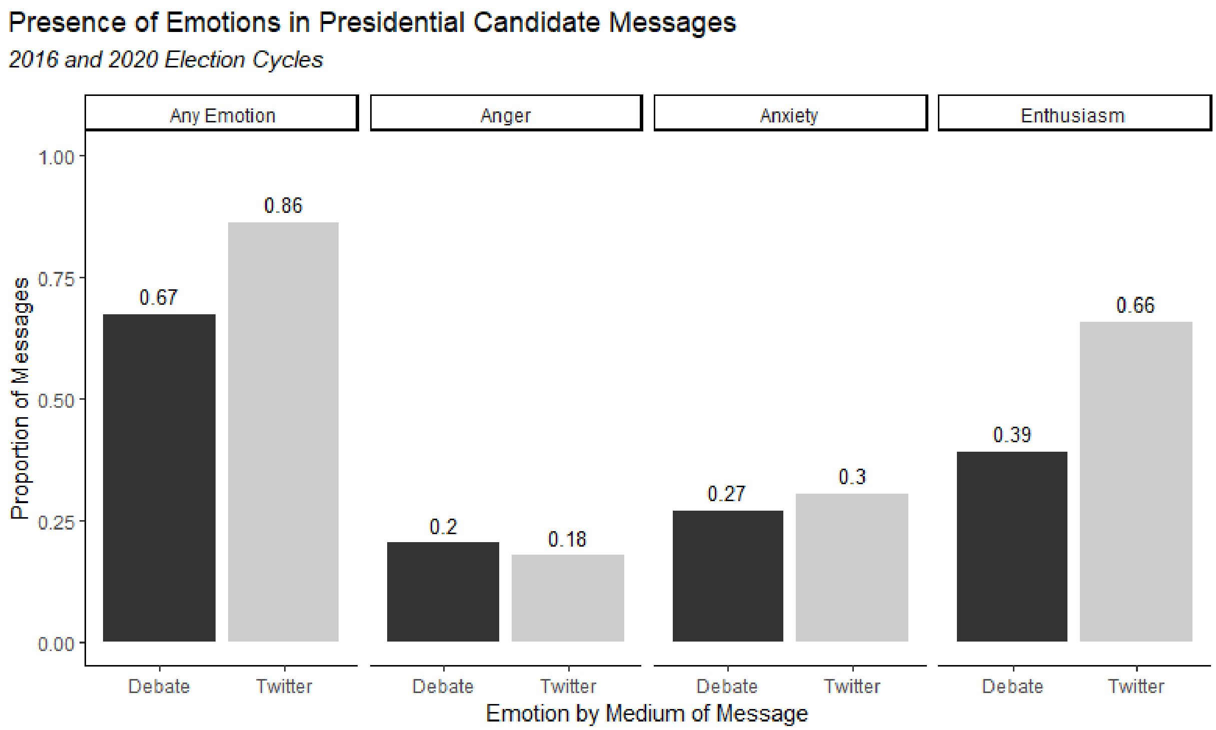

Figure 1. Data from 43,739 debate statements and 29,972 published tweets from presidential candidates in the 2016 and 2020 election cycles.

[Click here to return.]

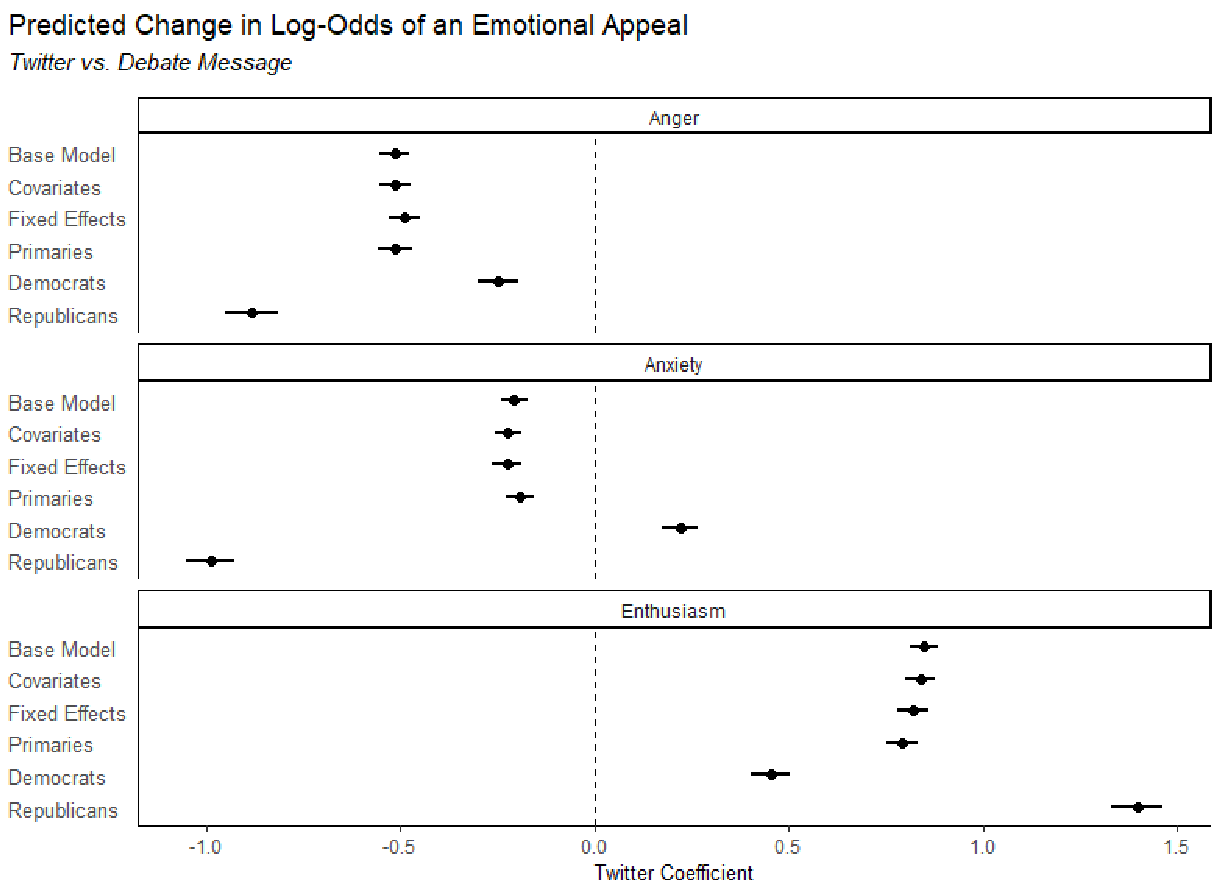

Figure 2. Data from 43,739 debate statements and 29,972 published tweets from presidential candidates in the 2016 and 2020 election cycles. Coefficients and 95 percent confidence intervals shown for primary independent variable only (Twitter binary); all models include additional control variables. See Table 1 and Table 2.

[Click here to return.]

Table 1: Emotional Appeals in Messages from US Presidential Candidates on Twitter versus Debates

	Model 1: Baseline			Model 2: Covariates			Model 3: Fixed Effects		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Dependent variable:	Anger	Anxiety	Enthusiasm	Anger	Anxiety	Enthusiasm	Anger	Anxiety	Enthusiasm
Tweet	−0.515***	−0.206***	0.847***	−0.512***	−0.223***	0.839***	−0.489***	−0.226***	0.820***
	(0.020)	(0.018)	(0.019)	(0.020)	(0.018)	(0.019)	(0.021)	(0.019)	(0.020)
Any Emotion	17.627	18.122	19.121	17.644	18.111	19.110	17.675	18.090	19.152
	(47.800)	(47.973)	(47.248)	(47.615)	(47.851)	(47.077)	(47.022)	(47.261)	(46.508)
General Election				0.378***	−0.026	−0.487***	0.735***	−0.217*	−0.410***
				(0.031)	(0.030)	(0.030)	(0.107)	(0.098)	(0.099)
2020 Cycle				0.193***	0.262***	−0.027	−0.104	0.163	−0.074
				(0.020)	(0.018)	(0.019)	(0.114)	(0.101)	(0.106)
Woman				−0.014	0.059**	0.133***	−0.561***	−0.933***	0.700***
				(0.024)	(0.021)	(0.023)	(0.091)	(0.087)	(0.090)
Nonwhite				−0.048	−0.040	−0.044	0.586***	0.997***	−0.806***
				(0.025)	(0.022)	(0.023)	(0.122)	(0.117)	(0.121)
Constant	−18.462	−18.521	−18.795	−18.596	−18.623	−18.736	−18.304	−17.897	−19.281
	(47.800)	(47.973)	(47.248)	(47.615)	(47.851)	(47.077)	(47.022)	(47.261)	(46.508)
Observations	73,711	73,711	73,711	73,711	73,711	73,711	73,711	73,711	73,711
Log Likelihood	−31,174.890	−36,600.700	−34,140.920	−31,051.360	−36,477.000	−33,991.520	−30,387.480	−35,788.930	−33,423.540
Akaike Inf. Crit.	62,355.770	73,207.400	68,287.840	62,116.720	72,968.000	67,997.040	60,926.960	71,729.850	66,999.090

*p<0.05; **p<0.01; ***p<0.001

Note: Each model is a logistic regression of a binary variable denoting the presence of the respective emotional appeal. The models show the likelihood that a tweet appeals to the respective emotion, relative to a message spoken in a candidate debate (the reference category), and holding constant a base probability of containing any emotional appeal at all. Model 3 includes fixed effects for event and candidate (coefficients and standard errors not shown). The observations comprise 43,739 debate sentences and 29,972 tweets from 38 leading US presidential candidates in the 2016 and 2020 election cycles.

[Click here to return.]

Table 2: Emotional Appeals in Messages from US Presidential Candidates on Twitter versus Debates

	Model 4: By Phase			Model 5: Democrats			Model 6: Republicans		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Dependent variable:	Anger	Anxiety	Enthusiasm	Anger	Anxiety	Enthusiasm	Anger	Anxiety	Enthusiasm
Tweet	−0.513***	−0.192***	0.791***	−0.248***	0.221***	0.453***	−0.882***	−0.989***	1.396***
	(0.022)	(0.019)	(0.020)	(0.027)	(0.024)	(0.025)	(0.035)	(0.032)	(0.033)
Any Emotion	17.568	18.083	19.205	17.612	18.145	19.195	18.759	19.063	20.195
	(51.672)	(51.541)	(51.183)	(66.821)	(67.161)	(66.641)	(108.597)	(108.857)	(105.427)
General Election	−0.021	−0.467	0.363	0.563***	−0.306**	−0.325**	−0.186	0.112	0.131
	(133.011)	(133.556)	(131.317)	(0.129)	(0.119)	(0.118)	(0.113)	(0.106)	(0.105)
2020 Cycle	0.089	0.278**	−0.289**	−0.116	0.242*	−0.094	0.565***	0.218	−0.335**
	(0.101)	(0.087)	(0.091)	(0.116)	(0.103)	(0.107)	(0.122)	(0.120)	(0.119)
Woman	−0.020	−0.036	0.113***	−0.525***	−0.875***	0.638***	0.040	−0.063	−0.214
	(0.026)	(0.023)	(0.025)	(0.091)	(0.087)	(0.089)	(0.170)	(0.152)	(0.162)
Nonwhite	−0.057*	0.061*	−0.054*	0.584***	1.001***	−0.791***	−0.211	0.068	−0.140
	(0.027)	(0.024)	(0.025)	(0.122)	(0.117)	(0.119)	(0.150)	(0.130)	(0.144)
Tweet*General	0.125	−0.217***	0.242***						
	(0.065)	(0.062)	(0.063)						
Any Emotion*General	0.502	0.100	−0.703						
	(133.011)	(133.556)	(131.317)						
Constant	−18.617	−18.357	−18.742	−18.271	−18.118	−19.231	−19.635	−19.159	−19.904
	(51.672)	(51.541)	(51.183)	(66.821)	(67.161)	(66.641)	(108.597)	(108.857)	(105.427)
Observations	73,711	73,711	73,711	42,372	42,372	42,372	31,339	31,339	31,339
Log Likelihood	−30,924.860	−36,152.980	−33,827.910	−18,658.930	−22,155.040	−20,389.710	−11,612.140	−13,143.740	−12,751.980
Akaike Inf. Crit.	61,935.720	72,391.970	67,741.820	37,423.860	44,416.090	40,885.430	23,282.280	26,345.470	25,561.960

*p<0.05; **p<0.01; ***p<0.001

Note: Each model is a logistic regression of a binary variable denoting the presence of the respective emotional appeal. The models show the likelihood that a tweet appeals to the respective emotion, relative to a message spoken in a candidate debate (the reference category), and holding constant a base probability containing any emotional appeal. All models include fixed effects for event, as well as candidate fixed effects for Models 5 and 6. Coefficients and standard errors are not shown. The observations comprise 43,739 debate sentencesand 29,972 tweets from 38 leading US presidential candidates in the 2016 and 2020 election cycles.

3

[Click here to return.]

Supplemental Online Material
for "Aimed Emotions in American Presidential Politics"

List of Contents

Table A1: 2016 Cycle Debates and Participants

	Republicans											Democrats				
Date	Bush	Carson	Christie	Cruz	Fiorina	Huckabee	Kasich	Paul	Rubio	Trump	Walker	Chafee	Clinton	O'Malley	Sanders	Webb
2015-08-06	✓	✓	✓	✓		✓	✓	✓	✓	✓	✓					
2015-09-16	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓					
2015-10-13												✓	✓	✓	✓	✓
2015-10-28	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓						
2015-11-10	✓	✓		✓	✓		✓	✓	✓	✓						
2015-11-14													✓	✓	✓	
2015-12-15	✓	✓	✓	✓	✓		✓	✓	✓	✓						
2015-12-19													✓	✓	✓	
2016-01-14	✓	✓	✓	✓			✓		✓	✓						
2016-01-17													✓	✓	✓	
2016-01-28	✓	✓	✓	✓			✓	✓	✓							
2016-02-04													✓		✓	
2016-02-06	✓	✓	✓	✓			✓		✓	✓						
2016-02-11													✓		✓	
2016-02-13	✓	✓		✓			✓		✓	✓						
2016-02-25		✓		✓			✓		✓	✓						
2016-03-03				✓			✓		✓	✓						
2016-03-06													✓		✓	
2016-03-09													✓		✓	
2016-03-10				✓			✓		✓	✓						
2016-04-14													✓		✓	
2016-09-26										✓			✓			
2016-10-09										✓			✓			
2016-10-19										✓			✓			

Note: Table indicates the participants in each of the nationally televised debates in the 2016 cycle included in the dataset.

Table A2: 2020 Cycle Debates and Participants

	Repub.	Democrats										
Date	Trump	Bennet	Biden	Bloomberg	Booker	Bullock	Buttigieg	Castro	De Blasio	Delaney	Gabbard	Gillibrand
2019-06-26					✓			✓	✓	✓	✓	
2019-06-27		✓	✓				✓					✓
2019-07-30						✓	✓			✓		
2019-07-31		✓	✓		✓			✓	✓		✓	✓
2019-09-12			✓		✓		✓	✓				
2019-10-15			✓		✓		✓	✓			✓	
2019-11-20			✓		✓		✓				✓	
2019-12-19			✓				✓					
2020-01-14			✓				✓					
2020-02-07			✓				✓					
2020-02-19			✓	✓			✓					
2020-02-25			✓	✓			✓					
2020-03-15			✓									
2020-09-29	✓		✓									
2020-10-22	✓		✓									

Note: Table indicates the participants in each of the nationally televised debates in the 2020 cycle included in the dataset.

Table A2 (cont.): 2020 Cycle Debates and Participants

	Democrats (cont.)											
Date	Harris	Hickenlooper	Inslee	Klobuchar	O'Rourke	Ryan	Sanders	Steyer	Swalwell	Warren	Williamson	Yang
2019-06-26			✓	✓	✓	✓				✓		
2019-06-27	✓	✓					✓		✓		✓	✓
2019-07-30		✓		✓	✓	✓	✓			✓	✓	
2019-07-31	✓		✓									✓
2019-09-12	✓			✓	✓		✓			✓		✓
2019-10-15	✓			✓	✓		✓	✓		✓		✓
2019-11-20	✓			✓			✓	✓		✓		✓
2019-12-19				✓			✓	✓		✓		✓
2020-01-14				✓			✓	✓		✓		
2020-02-07				✓			✓	✓		✓		✓
2020-02-19				✓			✓			✓		
2020-02-25				✓			✓	✓		✓		
2020-03-15							✓					
2020-09-29												
2020-10-22												

Note: Table indicates the participants in each of the nationally televised debates in the 2020 cycle included in the dataset.

Table A3: Twitter Data Source Handles by Presidential Candidate

Candidate	Twitter Handle	2016 Cycle	2020 Cycle
Michael Bennet	@MichaelBennet		✓
Joe Biden	@JoeBiden		✓
Mike Bloomberg	@MikeBloomberg		✓
Cory Booker	@CoryBooker		✓
Steve Bullock	GovernorBullock		✓
Jeb Bush	@JebBush	✓	
Pete Buttigieg	@PeteButtigieg		✓
Ben Carson	@RealBenCarson	✓	
Julián Castro	@JulianCastro		✓
Lincoln Chafee	@LincolnChafee	✓	
Chris Christie	@ChrisChristie	✓	
Hillary Clinton	@HillaryClinton	✓	
Ted Cruz	@TedCruz	✓	
Bill de Blasio	@BilldeBlasio		✓
John Delaney	@JohnDelaney		✓
Carly Fiorina	@CarlyFiorina	✓	
Tulsi Gabbard	@TulsiGabbard		✓
Kirsten Gillibrand	@SenGillibrand		✓
Kamala Harris	@KamalaHarris		✓
John Hickenlooper	@Hickenlooper		✓
Mike Huckabee	@GovMikeHuckabee	✓	
Jay Inslee	@JayInslee		✓
John Kasich	@JohnKasich	✓	
Amy Klobuchar	@amyklobuchar		✓
Martin O'Malley	@MartinOMalley	✓	
Beto O'Rourke	@BetoORourke		✓
Rand Paul	@RandPaul	✓	
Marco Rubio	@marcorubio	✓	
Tim Ryan	@TimRyan		✓
Bernie Sanders	@BernieSanders	✓	✓
Tom Steyer	@TomSteyer		✓
Eric Swalwell	@ericswalwell		✓
Donald Trump	@realDonaldTrump	✓	✓
Elizabeth Warren	@ewarren		✓
Marianne Williamson	@marwilliamson		✓
Andrew Yang	@AndrewYang		✓

Note: Table indicates source of Twitter data for each candidate.

Table A4: Performance Metrics for Trained RoBERTa Classifiers

	Debate Sentence Classifiers					
	Anger		Anxiety		Enthusiasm	
Metric	True Negative	True Positive	True Negative	True Positive	True Negative	True Positive
Precision	0.89	0.69	0.86	0.67	0.88	0.74
Recall	0.91	0.59	0.86	0.66	0.84	0.80
F1 Score	0.90	0.62	0.86	0.67	0.86	0.77
Overall Accuracy	0.84		0.80		0.82	

	Twitter Message Classifiers					
	Anger		Anxiety		Enthusiasm	
Metric	True Negative	True Positive	True Negative	True Positive	True Negative	True Positive
Precision	0.92	0.57	0.88	0.69	0.71	0.81
Recall	0.90	0.62	0.86	0.73	0.67	0.84
F1 Score	0.91	0.59	0.87	0.71	0.69	0.83
Overall Accuracy	0.86		0.82		0.78	

Note: Table shows performance on out-of-sample tests of the trained classifiers compared against human coders.

Table A5: Replication of Models 1 to 3 with Day-Of and Day-After Tweets Only

	Model 1: Baseline			Model 2: Covariates			Model 3: Candidate FE		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Dependent variable:	Anger	Anxiety	Enthusiasm	Anger	Anxiety	Enthusiasm	Anger	Anxiety	Enthusiasm
Tweet	−0.516***	−0.084***	0.809***	−0.521***	−0.097***	0.810***	−0.483***	−0.103***	0.773***
	(0.025)	(0.022)	(0.024)	(0.025)	(0.022)	(0.024)	(0.027)	(0.023)	(0.025)
Any Emotion	17.676	18.157	19.009	17.699	18.153	18.991	18.739	18.144	19.008
	(51.069)	(51.208)	(50.722)	(50.937)	(51.178)	(50.503)	(83.066)	(50.616)	(49.961)
General Election				0.386***	−0.021	−0.515***	0.818***	−0.080	−0.519***
				(0.034)	(0.032)	(0.032)	(0.114)	(0.106)	(0.106)
2020 Cycle				0.106***	0.110***	0.036	−0.159	0.174	−0.092
				(0.022)	(0.020)	(0.021)	(0.122)	(0.110)	(0.114)
Woman				−0.045	0.002	0.189***	−0.377***	−0.811***	0.550***
				(0.027)	(0.024)	(0.025)	(0.107)	(0.102)	(0.104)
Nonwhite				0.102***	0.117***	−0.182***	0.634***	1.164***	−0.857***
				(0.028)	(0.025)	(0.026)	(0.133)	(0.129)	(0.131)
Constant	−18.512	−18.556	−18.684	−18.637	−18.619	−18.626	−19.402	−18.109	−19.013
	(51.069)	(51.208)	(50.722)	(50.937)	(51.178)	(50.503)	(83.066)	(50.616)	(49.961)
Observations	58,537	58,537	58,537	58,537	58,537	58,537	58,537	58,537	58,537
Log Likelihood	−24,588.920	−28,392.230	−27,166.290	−24,514.260	−28,366.240	−26,992.430	−24,023.410	−27,895.130	−26,549.030
Akaike Inf. Crit.	49,183.840	56,790.450	54,338.590	49,042.510	56,746.470	53,998.870	48,198.820	55,942.260	53,250.070

*p<0.05; **p<0.01; ***p<0.001

Note: Replication conducted with 43,739 debate sentences and 14,798 tweets posted on the day of debate or the day after. See Table 1 for comparison.

Table A6: Replication of Models 4 to 6 with Day-Of and Day-After Tweets Only

	Model 4: By Phase			Model 5: Democrats			Model 6: Republicans		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Dependent variable:	Anger	Anxiety	Enthusiasm	Anger	Anxiety	Enthusiasm	Anger	Anxiety	Enthusiasm
Tweet	−0.499***	−0.040	0.729***	−0.196***	0.327***	0.384***	−0.973***	−0.833***	1.400***
	(0.028)	(0.024)	(0.026)	(0.033)	(0.029)	(0.031)	(0.046)	(0.040)	(0.043)
Any Emotion	17.637	18.136	19.079	17.664	18.154	19.123	18.829	19.159	19.954
	(55.686)	(55.625)	(55.285)	(70.751)	(70.974)	(70.588)	(117.918)	(118.167)	(115.386)
General Election	0.047	−0.321	0.215	0.657***	−0.101	−0.460***	0.037	0.022	−0.008
	(137.292)	(137.701)	(135.956)	(0.137)	(0.128)	(0.128)	(0.120)	(0.111)	(0.111)
2020 Cycle	0.066	0.288**	−0.363***	−0.183	0.212	−0.088	0.407**	0.052	−0.242
	(0.107)	(0.094)	(0.097)	(0.124)	(0.113)	(0.116)	(0.133)	(0.131)	(0.129)
Woman	−0.049	−0.065*	0.144***	−0.344**	−0.765***	0.501***	0.133	0.023	−0.268
	(0.029)	(0.026)	(0.028)	(0.106)	(0.102)	(0.103)	(0.192)	(0.169)	(0.180)
Nonwhite	0.113***	0.198***	−0.188***	0.628***	1.166***	−0.840***	−0.106	0.085	−0.179
	(0.030)	(0.027)	(0.028)	(0.132)	(0.129)	(0.130)	(0.177)	(0.151)	(0.165)
Tweet*General	0.081	−0.413***	0.353***						
	(0.080)	(0.076)	(0.078)						
Any Emotion*General	0.442	0.059	−0.632						
	(137.292)	(137.701)	(135.956)						
Constant	−18.665	−18.537	−18.502	−18.324	−18.196	−19.094	−19.799	−19.279	−19.613
	(55.686)	(55.625)	(55.285)	(70.751)	(70.974)	(70.588)	(117.918)	(118.167)	(115.386)
Observations	58,537	58,537	58,537	33,640	33,640	33,640	24,897	24,897	24,897
Log Likelihood	−24,404.550	−28,128.720	−26,826.760	−14,536.120	−16,928.790	−15,960.420	−9,384.303	−10,669.970	−10,388.340
Akaike Inf. Crit.	48,895.100	56,343.440	53,739.510	29,178.240	33,963.580	32,026.840	18,826.610	21,397.930	20,834.670

*p<0.05; **p<0.01; ***p<0.001

Note: Replication conducted with 43,739 debate sentences and 14,798 tweets posted on the day of debate or the day after. See Table 2 for comparison.

Table A7: Replication of Models 1 to 3 with Joint-Coding Trained Classifiers

	Model 1: Baseline			Model 2: Covariates			Model 3: Candidate FE		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Dependent variable:	Anger	Anxiety	Enthusiasm	Anger	Anxiety	Enthusiasm	Anger	Anxiety	Enthusiasm
Tweet	−0.325***	−0.044**	1.183***	−0.314***	−0.059***	1.168***	−0.273***	−0.045*	1.141***
	(0.019)	(0.017)	(0.017)	(0.019)	(0.017)	(0.017)	(0.020)	(0.018)	(0.018)
Any Emotion	1.318***	1.760***	1.653***	1.334***	1.744***	1.634***	1.355***	1.730***	1.626***
	(0.028)	(0.027)	(0.020)	(0.028)	(0.027)	(0.020)	(0.028)	(0.027)	(0.021)
General Election				0.391***	−0.090**	−0.486***	0.689***	−0.033	−0.452***
				(0.028)	(0.028)	(0.027)	(0.096)	(0.090)	(0.090)
2020 Cycle				0.195***	0.284***	0.086***	−0.097	−0.045	−0.027
				(0.018)	(0.017)	(0.017)	(0.106)	(0.094)	(0.095)
Woman				0.021	0.128***	0.296***	−0.643***	−0.896***	0.516***
				(0.022)	(0.020)	(0.021)	(0.088)	(0.082)	(0.084)
Nonwhite				−0.079***	−0.054*	0.065**	0.555***	1.000***	−0.214
				(0.024)	(0.021)	(0.021)	(0.119)	(0.111)	(0.116)
Constant	−2.228***	−2.253***	−1.577***	−2.370***	−2.368***	−1.615***	−2.068***	−1.597***	−1.901***
	(0.026)	(0.025)	(0.019)	(0.028)	(0.028)	(0.022)	(0.114)	(0.104)	(0.107)
Observations	73,711	73,711	73,711	73,711	73,711	73,711	73,711	73,711	73,711
Log Likelihood	−37,060.390	−42,298.650	−43,370.640	−36,889.970	−42,106.530	−43,067.390	−36,261.980	−41,444.110	−42,552.510
Akaike Inf. Crit.	74,126.790	84,603.290	86,747.290	73,793.950	84,227.060	86,148.780	72,675.950	83,040.230	85,257.020

*p<0.05; **p<0.01; ***p<0.001

Note: Replication conducted with values from RoBERTa classifiers trained on 2,000 messages coded by all three research assistants. See Table 1 for comparison.

Table A8: Replication of Models 4 to 6 with Joint-Code Trained Classifiers

	Model 4: By Phase			Model 5: Democrats			Model 6: Republicans		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Dependent variable:	Anger	Anxiety	Enthusiasm	Anger	Anxiety	Enthusiasm	Anger	Anxiety	Enthusiasm
Tweet	−0.308***	−0.017	1.107***	0.070**	0.383***	0.884***	−0.832***	−0.760***	1.472***
	(0.021)	(0.018)	(0.018)	(0.025)	(0.023)	(0.024)	(0.033)	(0.030)	(0.028)
Any Emotion	1.362***	1.710***	1.651***	1.327***	1.733***	1.733***	1.406***	1.745***	1.531***
	(0.031)	(0.029)	(0.022)	(0.039)	(0.037)	(0.029)	(0.041)	(0.041)	(0.031)
General Election	0.546***	−0.310**	−0.312**	0.455***	−0.187	−0.269*	−0.174	−0.062	0.200*
	(0.104)	(0.107)	(0.099)	(0.116)	(0.110)	(0.109)	(0.101)	(0.099)	(0.091)
2020 Cycle	0.166	0.125	−0.274***	−0.113	−0.009	0.015	0.401***	0.324**	−0.426***
	(0.093)	(0.081)	(0.080)	(0.109)	(0.097)	(0.097)	(0.108)	(0.110)	(0.105)
Woman	−0.012	0.036	0.221***	−0.597***	−0.850***	0.482***	0.508**	0.128	−0.189
	(0.024)	(0.022)	(0.022)	(0.088)	(0.083)	(0.083)	(0.163)	(0.140)	(0.134)
Nonwhite	−0.086***	0.025	0.083***	0.552***	1.007***	−0.212	0.135	0.170	−0.123
	(0.025)	(0.023)	(0.022)	(0.119)	(0.112)	(0.115)	(0.149)	(0.122)	(0.119)
Tweet*General	0.068	−0.352***	0.358***						
	(0.061)	(0.061)	(0.059)						
Any Emotion*General	−0.133	0.188*	−0.246***						
	(0.071)	(0.081)	(0.066)						
Constant	−2.504***	−2.095***	−1.495***	−2.093***	−1.714***	−1.983***	−2.513***	−1.982***	−1.512***
	(0.070)	(0.060)	(0.057)	(0.118)	(0.109)	(0.110)	(0.154)	(0.130)	(0.123)
Observations	73,711	73,711	73,711	42,372	42,372	42,372	31,339	31,339	31,339
Log Likelihood	−36,755.040	−41,801.520	−42,877.800	−21,709.420	−25,081.570	−24,761.830	−14,303.790	−15,876.590	−17,638.490
Akaike Inf. Crit.	73,596.070	83,689.030	85,841.600	43,524.830	50,269.150	49,629.660	28,665.580	31,811.180	35,334.980

*p<0.05; **p<0.01; ***p<0.001

Note: Replication conducted with values from RoBERTa classifiers trained on 2,000 messages coded by all three research assistants. See Table 2 for comparison.

Table A9: Replication of Models 1 to 3 with Linear Probability Model Estimation

	Model 1: Baseline			Model 2: Covariates			Model 3: Candidate FE		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Dependent variable:	Anger	Anxiety	Enthusiasm	Anger	Anxiety	Enthusiasm	Anger	Anxiety	Enthusiasm
Tweet	−0.078***	−0.039***	0.148***	−0.077***	−0.041***	0.146***	−0.070***	−0.040***	0.140***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Any Emotion	0.276***	0.388***	0.630***	0.277***	0.385***	0.626***	0.278***	0.379***	0.624***
	(0.003)	(0.004)	(0.004)	(0.003)	(0.004)	(0.004)	(0.003)	(0.004)	(0.004)
General Election				0.049***	−0.008	−0.072***	0.101***	−0.034*	−0.063***
				(0.004)	(0.005)	(0.005)	(0.015)	(0.016)	(0.016)
2020 Cycle				0.024***	0.046***	0.001	−0.009	0.024	−0.011
				(0.003)	(0.003)	(0.003)	(0.016)	(0.018)	(0.017)
Woman				−0.001	0.011**	0.021***	−0.094***	−0.176***	0.118***
				(0.003)	(0.004)	(0.004)	(0.014)	(0.016)	(0.015)
Nonwhite				−0.004	−0.006	−0.011**	0.104***	0.192***	−0.146***
				(0.003)	(0.004)	(0.004)	(0.019)	(0.022)	(0.021)
Constant	0.018***	0.009**	−0.033***	0.002	−0.008*	−0.024***	0.052**	0.132***	−0.111***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.004)	(0.004)	(0.018)	(0.020)	(0.019)
Observations	73,711	73,711	73,711	73,711	73,711	73,711	73,711	73,711	73,711
R^2	0.089	0.134	0.353	0.091	0.137	0.356	0.107	0.153	0.365
Adjusted R^2	0.089	0.134	0.353	0.091	0.137	0.356	0.106	0.152	0.364

*p<0.05; **p<0.01; ***p<0.001

Note: Replication conducted via linear probability models instead of logistic regression. See Table 1 for comparison.

Table A10: Replication of Models 4 to 6 with Linear Probability Model Estimation

	Model 4: By Phase			Model 5: Democrats			Model 6: Republicans		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Dependent variable:	Anger	Anxiety	Enthusiasm	Anger	Anxiety	Enthusiasm	Anger	Anxiety	Enthusiasm
Tweet	−0.075***	−0.035***	0.136***	−0.036***	0.047***	0.082***	−0.115***	−0.157***	0.216***
	(0.003)	(0.003)	(0.003)	(0.004)	(0.005)	(0.004)	(0.004)	(0.005)	(0.005)
Any Emotion	0.264***	0.379***	0.646***	0.272***	0.396***	0.645***	0.284***	0.358***	0.601***
	(0.004)	(0.004)	(0.004)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
General Election	0.003	−0.075***	0.051**	0.081***	−0.049*	−0.053**	−0.030*	0.009	0.027
	(0.015)	(0.017)	(0.016)	(0.019)	(0.021)	(0.020)	(0.014)	(0.015)	(0.015)
2020 Cycle	0.010	0.042**	−0.041**	−0.012	0.037	−0.011	0.078***	0.034*	−0.053**
	(0.014)	(0.015)	(0.015)	(0.017)	(0.019)	(0.018)	(0.016)	(0.017)	(0.017)
Woman	−0.003	−0.006	0.019***	−0.090***	−0.162***	0.112***	0.012	−0.002	−0.039
	(0.004)	(0.004)	(0.004)	(0.015)	(0.016)	(0.015)	(0.021)	(0.022)	(0.022)
Nonwhite	−0.006	0.010*	−0.012**	0.103***	0.189***	−0.145***	−0.023	0.015	−0.020
	(0.004)	(0.004)	(0.004)	(0.020)	(0.023)	(0.021)	(0.018)	(0.019)	(0.019)
Tweet*General	0.002	−0.045***	0.071***						
	(0.010)	(0.011)	(0.011)						
Any Emotion*General	0.101***	0.020	−0.161***						
	(0.010)	(0.011)	(0.010)						
Constant	−0.002	0.039***	−0.026*	0.052**	0.090***	−0.116***	−0.0003	0.067***	−0.014
	(0.010)	(0.011)	(0.010)	(0.019)	(0.021)	(0.020)	(0.019)	(0.020)	(0.020)
Observations	73,711	73,711	73,711	42,372	42,372	42,372	31,339	31,339	31,339
R^2	0.096	0.144	0.360	0.090	0.146	0.328	0.134	0.175	0.419
Adjusted R^2	0.095	0.144	0.360	0.089	0.145	0.327	0.134	0.174	0.418

*p<0.05; **p<0.01; ***p<0.001

Note: Replication conducted via linear probability models instead of logistic regression. See Table 2 for comparison.