# Slanted images: Measuring nonverbal media bias

Boxell, Levi

Stanford University

17 September 2018

# Slanted images:
# Measuring nonverbal media bias

Levi Boxell, *Stanford University*[*]

September 2018

**Abstract**

I build a dataset of over one million images used on the front page of websites around the 2016 election period. I then use machine-learning tools to detect the faces of politicians across the images and measure the nonverbal emotional content expressed by each politician. Combining this with data on the partisan composition of each website's users, I show that websites portray politicians that align with the partisan preferences of their users with more positive emotions. I also find that nonverbal coverage by Republican-leaning websites was not consistent over the 2016 election, but became more favorable towards Donald Trump after he clinched the Republican nomination.

[*]E-mail: lboxell@stanford.edu. I thank Matthew Gentzkow and Jesse M. Shapiro for their comments and suggestions. Funding was generously provided by the National Science Foundation (grant number: DGE-1656518) and the Institute for Humane Studies.

# 1 Introduction

Nonverbal information is more memorable and more persuasive than verbal information (Sullivan and Masters 1988; Graber 1990; Graber 1996). Furthermore, much of the media consumed today is nonverbal—political coverage is watched on television, images of politicians are posted alongside newspaper stories whether online or in print, and social media is littered with attention grabbing photos on shared posts. Despite this, the literature on media bias has focused almost exclusively on textual or verbal measures of media bias (Gentzkow and Shapiro 2010; Martin and Yurukoglu 2017).

To better understand the degree to which nonverbal bias is present across online media firms, I build a dataset of over one million images with which I use facial recognition tools to extract nearly 80,000 faces of 61 different politicians from 92 websites around the 2016 election cycle. For each face, I use the Microsoft Emotion API to characterize the face on eight different emotional categories: happiness, anger, fear, surprise, disgust, contempt, sadness, and neutral. The displayed facial expression of a politician on a website is a good indicator of nonverbal media slant as facial expressions are readily altered by choosing different images and have been show to influence opinions. For example, video excerpts of candidates displaying happy or reassuring facial expressions are more influential in shaping respondents' attitudes towards politicians than verbal information such as party affiliation and ideological beliefs (Sullivan and Masters 1988). Other studies reach similar conclusions regarding the ability to manipulate voter's preferences via the way in which a candidate is portrayed visually (Barrett and Barrington 2005b; Rosenberg and McCafferty 1987; Rosenberg et al. 1991).

I then combine these images with the partisanship score from the 2016 Berkman-Klein Report (Faris et al. 2017). The Berkman-Klein partisanship score is the relative frequency with which Trump supporters versus Clinton supporters shared links from each website on Twitter. I use this index to show that websites with a higher Republican user composition portray Republican politicians with more positive emotions and less negative emotions relative to Democratic politicians.

I also show that media firms vary their degree of bias over the election cycle. Republican-leaning websites gave more favorable nonverbal coverage towards Trump after he clinched the Republican nomination. Before Trump clinched the nomination, Republican-leaning websites portrayed Trump in an increasingly negative fashion. On the other hand, Democrat-leaning websites have a peaked portrayal around the period in which he clinched the nomination, but quickly taper off to previous levels.

With regards to nonverbal slant, previous work has suggested the presence of nonverbal biases in the media and has found corroborating evidence in small-scale manual codings (Kepplinger 1982; Moriarty and Garramone 1986; Waldman and Devitt 1998; Barrett and Barrington 2005a; Coleman and Banning 2006). I implement the first automated examination of nonverbal slant that is scalable and applicable across domains. The size of the data and the measurement approach also allows for the examination of trends in slant across the election cycle, a topic which we have limited knowledge of. The measure of nonverbal slant can also be used to measure differences across any identifiable groups, not just political parties. Furthermore, with the increasing capabilities of computer vision, this work highlights the new opportunities for using images as data to complement previous work on text as data.

# 2 Data

## 2.1 Website partisanship

The baseline set of websites comes from Appendix 3 in Faris et al. (2017). They use 2016 media link and social media sharing data along with information on whether the Twitter users retweeted Trump or Clinton to construct a measure of the partisan composition of each website, subsequently denoted as the "partisanship score." The partisanship score measures the relative frequency of Twitter shares made by Trump or Clinton supporters for each website on scale from -1 to 1. Positive partisanship scores indicate a higher frequency of shares by Trump supporters relative to Clinton supporters. Their Appendix 3 lists the partisanship score for 115 websites that fall in the top 100 Twitter shares, Facebook shares, or media links along with the associated number of shares or links. These 115 websites form the base sample of websites.[1]

## 2.2 Website images

To build the dataset of politician images, I scrape the archived version of the front page of each website from the Internet Archive's Wayback Machine and download all images from this page. I scrape each day between September 2015 and April 2017 separately.[2] This yields over one

[1]Mediaite.com is dropped.

[2]I attempt to scrape the noon archive of each website, but accept the default re-direct from the Wayback Machine to alternative archives on the same day.

million scraped images.

The Wayback Machine's choice of whether or not to archive a version of a website is not random, but is likely a function of the popularity of the website along with their overall archiving capabilities at a given point in time. Websites can also explicitly restrict the archiving. Furthermore, idiosyncratic issues with the scraping or the original archiving may have prevented some images from being downloaded for certain websites.

## 2.3 Identifying politicians

I select the set of politicians by identifying, for each year in 2008, 2012, and 2016, the main Republican and Democratic presidential candidates, the vice presidents selected by each nominee, and the main congressional leaders for each party.[3] I then construct a small dataset of labeled images for each politician.

The scraped images are not labeled and may contain multiple faces. To identify the faces of politicians across the images, I first use Matlab's eye detector to filter images that are likely to contain a face.[4] This reduces the number of images to roughly 350,000.

Using Microsoft's Face API and the manually labeled images, I use a facial recognition machine learning algorithm to identify politician faces in the unlabeled images. The algorithm first searches for faces in the image and then, for each face, assigns a match a confidence score between zero and one for a given politician. Whether or not the API detects a face (or matches a face to a politician) depends on several aspects of the photo, such as the resolution, the size of the face, the angle of the face, and whether there are any obstructions to the face (e.g., a hand). The baseline results restrict the faces to those with a confidence score of at least 0.5.

## 2.4 Estimating emotions

For each face, I use Microsoft's Emotion API to estimate the emotional content of the image on eight different dimensions: happiness, anger, fear, surprise, disgust, contempt, sadness, and neutral. The level of emotion in each category roughly sums to one.

One important measure is the difference between positive emotions (happiness) and the sum of the negative emotions (anger, sadness, contempt, disgust, surprise, and fear), which

[3] See Online Appendix for details.

[4] I also restrict to images that I was able to load into Matlab and that were at least 1 KB in size, which is an API size restriction later.

indicates the relative favorability towards an individual in a given image. The Online Appendix contains a histogram of the relative favorability measure along with the positive and negative emotion scores. The relative favorability scores cluster at -1, 0, and 1, indicating negative, neutral, and positive emotions respectively. This is to be expected as the API gives the likelihood of each emotion being displayed in the image and will assign a value of 1 for images with clear emotional expressions. As such, the relative favorability and emotion scores should be interpreted as the (relative) likelihood of containing a given emotion.

# 3 Results

To measure whether a website's nonverbal coverage of politicians is correlated with the partisan composition of its users, I estimate the following equation via OLS:

$$y_{ijt} = X_{ijt}\beta + c_j 1(i \in R)\gamma + e_{ijt},$$

where $y_{ijt}$ is the score for a given emotion category for the $t$th instance of politician $i$ on website $j$, $X_{ijt}$ is a set of controls that includes website and politician indicators with corresponding parameter vector $\beta$, $c_j$ is the partisanship score for website $j$, $1(i \in R)$ is an indicator for whether politician $i$ is in the Republican party, $\gamma$ is a scalar parameter, and $e_{ijt}$ is the error term.

Table 1 gives the main results from this regression for each of the eight emotion categories. Since treatment varies at the website-level, standard errors are clustered at the same level throughout. For the happiness emotion, $\hat{\gamma}$ is positive and statistically significant at conventional levels. For the neutral emotion, $\hat{\gamma}$ is statistically indistinguishable from zero. For the negative emotions, $\hat{\gamma}$ is negative throughout and statistically significant in most cases. When using the sum of the negative emotion scores as the outcome $y_{ijt}$, $\hat{\gamma}$ is negative with roughly the same magnitude and statistical significance as the happiness results.

These results corroborate the use of the measure of relative favorability which is defined as the happiness emotion score minus the sum of the negative emotion scores. Using relative favorability as an outcome, column (10) shows that going from equal partisanship to a completely Republican partisanship score increases the relative favorability of Republican politicians by 0.16, which is nearly half of the average non-neutral emotion score as reported in the Online Appendix. The Online Appendix shows that these results are robust to restricting the data to higher levels of match confidence, using alternative specifications, and restricting the data to certain websites or politicians.

Figure 1 plots the average relative favorability towards Republicans minus the average relative favorability towards Democrats against the Berkman Klein partisanship score for each website. It also repeats this exercise for Donald Trump and Hillary Clinton. These website-aggregated results are consistent with the findings in table 1. There is a strong, positive correlation between the partisanship scores and Republican-leaning slant. The Online Appendix ranks select websites according their Republican-Democrat slant and and Trump-Clinton slant. The Daily Kos, PoliticusUSA, and CNN provided some of the most Democrat-slanted visual coverages according to both measures. On the other hand, The Daily Caller, InfoWars, and The Gateway Pundit provided some of the most Republican-slanted visual coverages according to both measures.

The Online Appendix shows that partisan websites are also more likely to cover the opposing party. Overall, these results are consistent with previous findings that partisan sources are more likely to cover political scandals of the opposing party (Puglisi and Snyder 2011), and tend to express bias by criticizing the opposing party (Budak et al. 2016).

## 3.1 Trends

The previous analysis presents the correlation between the partisan composition of users and the nonverbal slant of websites averaged over the entire time period. Does the degree of slant change over the election cycle? Figure 2 plots the relative favorability towards Trump and Clinton for Republican-leaning and Democrat-leaning websites separately across the election, aggregating across images displayed in a given day.

Focusing on Republican-leaning websites' portrayal of Trump, figure 2 shows an increase in favorable coverage after Trump clinched the Republican nomination.[5] Before Trump clinched the nomination, Republican-leaning websites tended to portray Trump in an increasingly negative fashion. On the other hand, Democrat-leaning websites have a peaked portrayal around the period in which he clinched the nomination, but quickly taper off to previous levels. This gives some indication that Republican-leaning websites shifted how they covered Trump in response to him becoming the presumptive party nominee.

Examining the variability in coverage between Trump and Clinton, one sees that Trump's coverage by Democrat-leaning websites is relatively stable, whereas the same websites' cover-

[5]Defined to be May 3, 2016 when Ted Cruz suspended his campaign. The corresponding date for Hillary Clinton is defined to be June 6, 2016.

age of Clinton exhibits substantial swings and is actually trending downwards in the run up to the election. There is also a large divergence in coverage of Hillary Clinton during the post-election period between Republican- and Democrat-leaning websites, with the former giving less favorable coverage overtime.

The Online Appendix restricts attention to Trump and Clinton and reproduces the main correlation between nonverbal slant and the partisan composition of website users for three periods: primaries, post-primaries, and post-election. The findings suggest slant increases over the election cycle and does not taper after the election. The increase in favorability appears to be driven by positive coverage during the post-primary period, but negative coverage during the post-election period.

## 3.2 Reference points

Figure 3 plots the average relative favorability among websites with negative partisanship scores against websites with positive partisanship scores for politicians with more than 200 images. Movements to the left of the 45 degree line indicate more positive portrayal by Democrat-leaning websites relative to Republican-leaning websites, and vice-versa for movements to the right. Figure 3 shows strikingly different baseline levels of emotion for Donald Trump and Hillary Clinton, with Donald Trump being portrayed with more negative emotions overall. Conceptually, slant must be measured relative to a reference point. For example, movements along the 45 degree line could be due to either politician-specific differences in emotions or systematic media bias against certain politicians.

To examine this, I measure the average emotion portrayed in the 2016 Presidential Debates, Google images, and politician headshots for both Trump and Clinton. Figure 4 compares the nonverbal coverage of Clinton against the nonverbal coverage of Trump for these three reference points along with the equivalent measure for each website. Across all three reference points, the average emotional expression for Trump is more negative than that of Clinton—suggesting that politician-specific differences in emotional expressions are likely driving much of the movement along the 45 degree line in figure 3 for Trump and Clinton. Figure 4 also shows that nonverbal coverage of Trump in the media was slightly more negative than what a viewer of the debates would have observed, whereas the nonverbal coverage of Clinton in the media was more favorable relative to the debates. The Online Appendix shows that this still holds when restricting to images on websites the day after each debate—a period in which firms are likely selecting images from the universe of debate clips.

# 4 Discussion

Most research on media bias has focused on textual measures of media bias. While verbal media coverage dominated the media diets of consumers a century ago, media consumption today is highly nonverbal. In fact, presidential candidates on broadcast news networks are primarily shown visually with commentary voiced over, rather than being heard directly (Bucy and Grabe 2007).

If the endogenous consumption of nonverbally biased media impacts the political beliefs and feelings of consumers as exogenous consumption of such media has been shown to do, then the rising amount of nonverbal information in political news sources may play an important role in explaining contemporary (affective) political polarization in the United States (Iyengar et al. 2012). This is consistent with the role of cable news in driving political polarization (Martin and Yurukoglu 2017). While Boxell et al. (2017; 2018) argue that the role of the internet in driving recent trends in political polarization or the outcome of the 2016 election is limited relative to what has often been suggested, increasing trends towards visual (rather than verbal) information may still be problematic in the media ecosystem more generally. The nonverbal slant measure provides an important tool for examining this question, and future research should continue to examine this relationship.

# References


Barrett, Andrew W. and Lowell W. Barrington. 2005a. Bias in newspaper photograph selection. *Political Research Quarterly.* 58(4):609–618.

Barrett, Andrew W. and Lowell W. Barrington. 2005b. Is a picture worth a thousand words? Newspaper photographs and voter evaluations of political candidates. *The International Journal of Press/Politics*. 10(4):98–113.

Boxell, Levi, Matthew Gentzkow, and Jesse M. Shapiro. 2018. A note on internet use and the 2016 US presidential election outcome. *PLOS One.* 13(7): e0199571.

Boxell, Levi, Matthew Gentzkow, and Jesse M. Shapiro. 2017. Greater internet use is not associated with faster growth in political polarization among US demographic groups. *Proceedings of the National Academy of Sciences.* 114(40): 10612–10617.

Bucy, Erik P. and Maria Elizabeth Grabe. 2007. Taking television seriously: A sound and image bite analysis of presidential campaign coverage, 1992–2004. *Journal of Communication*. 57(4):652–675.

Budak, Ceren, Sharad Goel, and Justin M. Rao. 2016. Fair and balanced? Quantifying media bias through crowdsourced content analysis. *Public Opinion Quarterly.* 80(S1): 250–271.

Coleman, Renita and Stephen Banning. 2006. Network TV news' affective framing of the presidential candidates: Evidence for a second-level agenda-setting effect through visual framing. *Journalism & Mass Communication Quarterly*. 83(2):313–328.

Faris, Robert M., Hal Roberts, Bruce Etling, Nikki Bourassa, Ethan Zuckerman, and Yochai Benkler. 2017. Partisanship, propaganda, and disinformation: Online media and the 2016 U.S. presidential election. *Berkman Klein Center for Internet & Society Research Paper.*

Gentzkow, Matthew and Jesse M. Shapiro. 2010. What drives media slant? Evidence from U.S. daily newspapers. *Econometrica*. 78(1):35–71.

Graber, Doris A. 1990. Seeing is remembering: How visuals contribute to learning from television news. *Journal of Communication*. 40(3):134–156.

Graber, Doris A. 1996. Say it with pictures. *The ANNALS of the American Academy of Political and Social Science*. 546(1):85–96.

Iyengar, Shanto, Gaurav Sood, and Ypthach Lelkes. 2012. Affect, not ideology: A social identity perspective on polarization. *Public Opinion Quarterly*. 76(3):405–431.

Kepplinger, Hans Mathias. 1982. Visual biases in television campaign coverage. *Communication Research*. 9(3):432–446.

Martin, Gregory J. and Ali Yurukoglu. 2017. Bias in cable news: Persuasion and polarization. *American Economic Review.* 107(9):2565–2599.

Moriarty, Sandra E. and Gina M. Garramone. 1986. A study of newsmagazine photographs of the 1984 presidential campaign. *Journalism Quarterly*. 63(4):728–734.

Puglisi, Riccardo and James M. Snyder, Jr. 2011. Newspaper coverage of political scandals. *Journal of Politics.* 73(3): 931–950.

Rosenberg, Shawn W. and Patrick McCafferty. 1987. The image and the vote manipulating voter's preferences. *Public Opinion Quarterly*. 51(1):31–47.

Rosenberg, Shawn W., Shulamit Kahn, Thuy Tran, and Minh-Thu Le. 1991. Creating a political image: Shaping appearance and manipulating the vote. *Political Behavior*. 13(4):345–367.

Sullivan, Denis G. and Roger D. Masters. 1988. "Happy Warriors": Leaders' facial displays, viewers' emotions, and political support. *American Journal of Political Science*. 32(2):345–368.

Waldman, Paul and James Devitt. 1998. Newspaper photographs and the 1996 presidential election: The question of bias. *Journalism & Mass Communication Quarterly*. 75(2):302–311.

Figure 1: Nonverbal Slant and Segregation of Twitter Shares

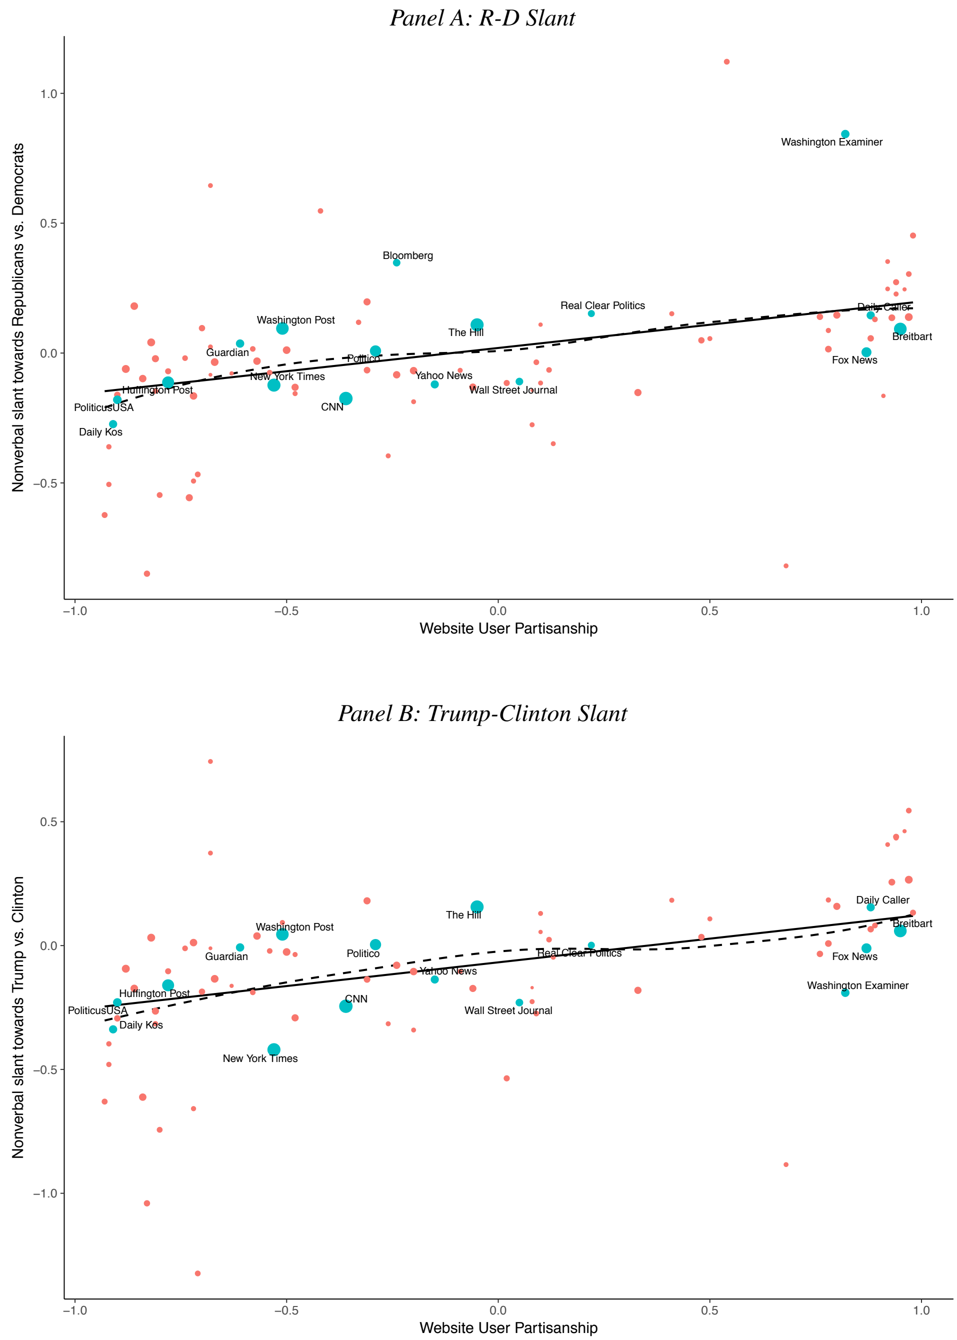


Notes: Panel A plots the average relative favorability towards Republicans (demeaned across all websites) minus the average relative favorability towards Democrats (demeaned across all websites) against the Berkman Klein partisanship score for each website. The black solid line is the linear best fit weighted by shares on Twitter; the black dashed line is a loess fit weighted by shares on Twitter. The labelled blue dots indicate websites with at least 10,000 Twitter shares or 2,000 media inlinks in the Berkman Klein data. Panel B plots average relative favorability towards Donald Trump (demeaned across all websites) minus the average relative favorability towards Hillary Clinton (demeaned across all websites) against the Berkman Klein partisanship score for each website.

Figure 2: Favorability towards Trump and Clinton across 2016

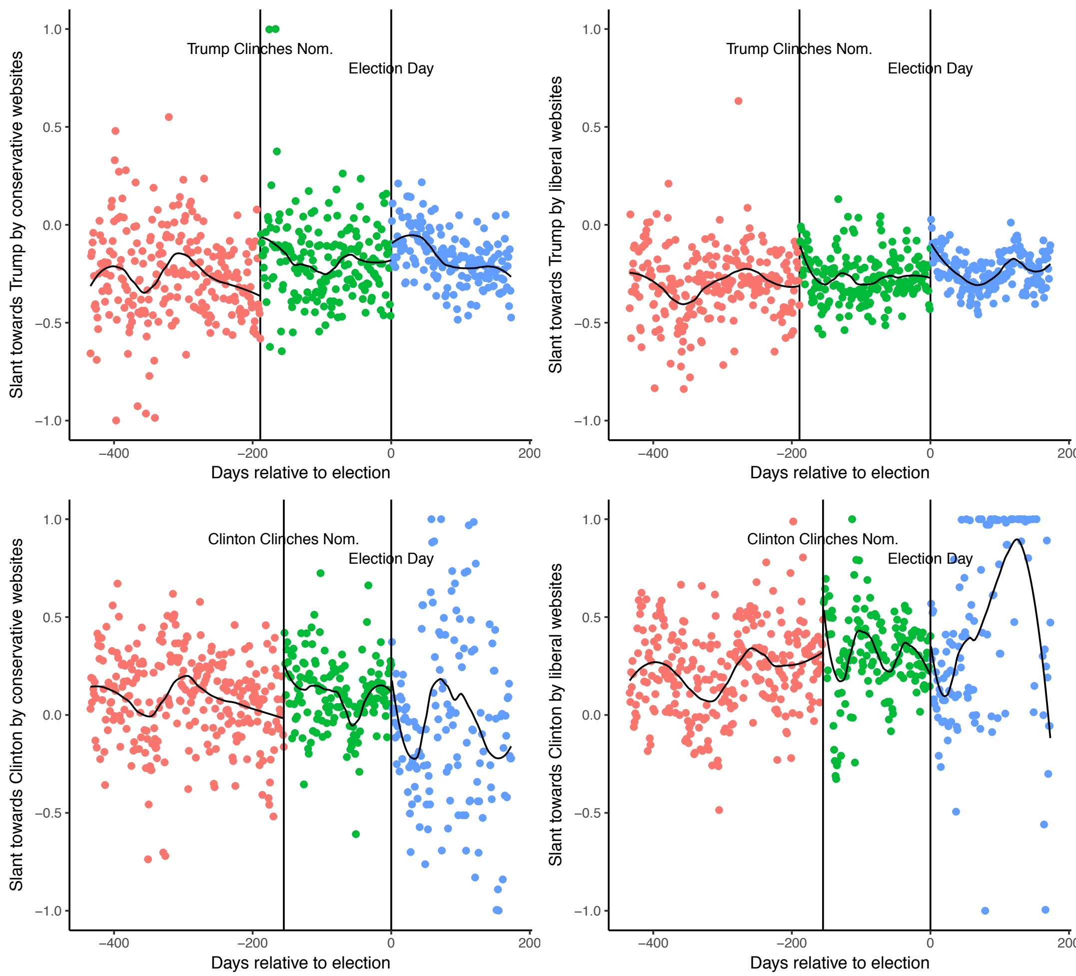


Notes: The top left (right) figure plots the average relative favorability towards Donald Trump across all images for websites with a partisanship score greater (less) than 0 for each day in 2016. The bottom left (right) figure plots the average relative favorability towards Hillary Clinton across all images for websites with a partisanship score greater (less) than 0 for each day in 2016. Loess smoothing lines are fit on each side of the day the candidate clinched their party's nomination and election day.

Figure 3: Favorability of Politicians

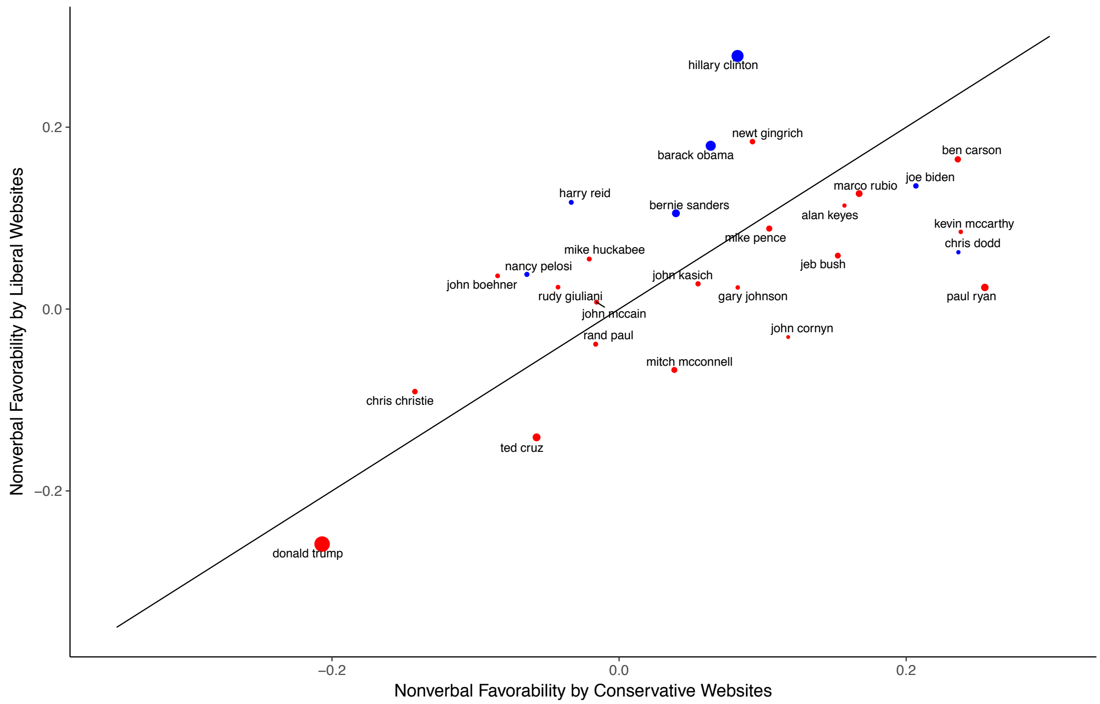


Notes: Figure plots the average relative favorability among websites with a partisanship score less than 0 against the average relative favorability among websites with a partisanship score greater than 0 for each politician with more than 200 images. Blue dots indicate Democrats; Red dots indicate Republicans. The solid line indicates the 45 degree line.

Figure 4: Trump and Clinton Favorability with Respect to Reference Points

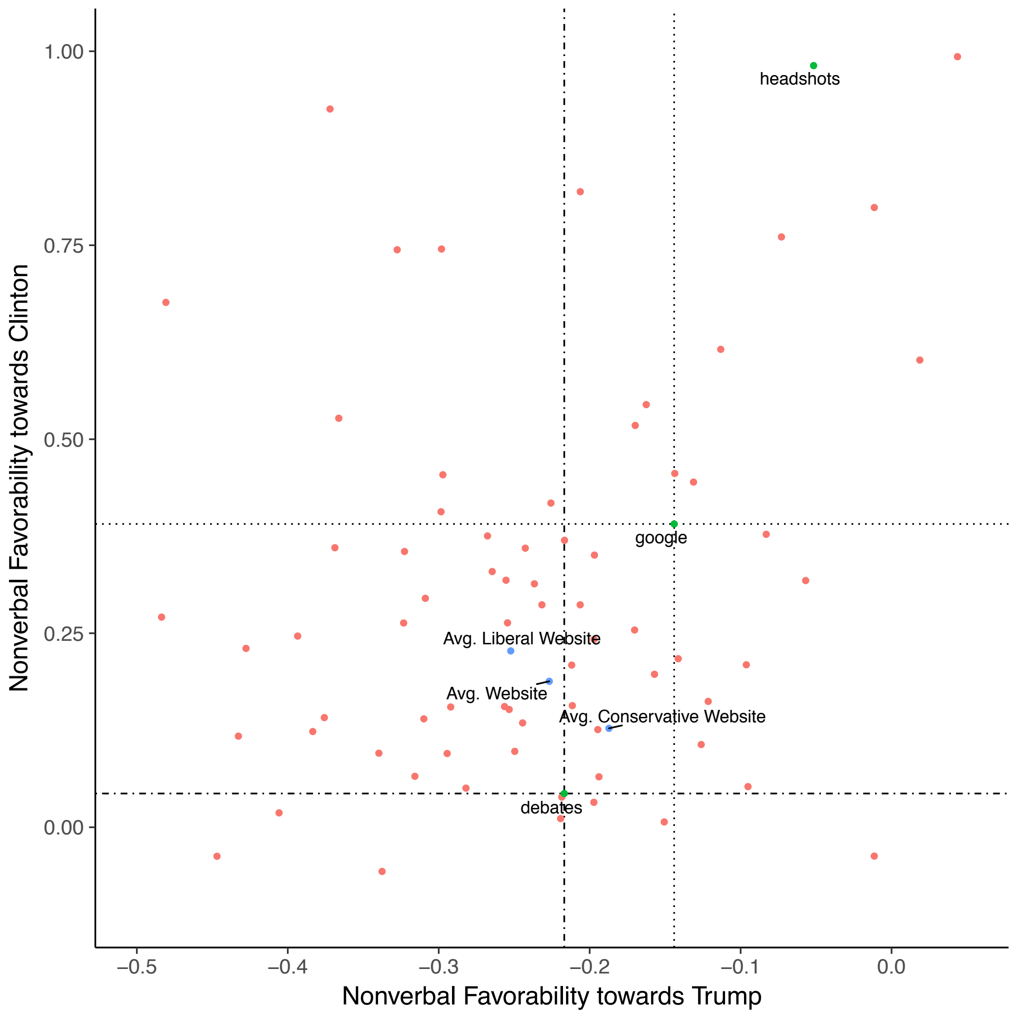


Notes: Figure plots the average relative favorability towards Clinton against the average relative favorability towards Trump for each website. The headshots, google, and debate points indicate the average relative favorability across politican controlled headshots, the first several hundred images on google images, and the 2016 presidential debates. The other labeled points are the average relative favorability towards a candidate averaged across either all websites, websites with a positive partisanship score, or websites with a negative partisanship score and weighted by the number of Twitter shares. The x and y-axis of the plot are artificially restricted to zoom in on the majority of the websites.

Table 1: Partisanship of users and politician emotions

| | Dependent Variable: Emotion Score | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) | (10) |
| | Neutral | Happiness | Anger | Sadness | Contempt | Disgust | Surprise | Fear | Negative | Relative Fav. |
| Partisanship Score x Republican Politician | 0.0206 | 0.0702 | -0.0208 | -0.0230 | -0.0030 | -0.0019 | -0.0342 | -0.0080 | -0.0908 | 0.1610 |
| | (0.0147) | (0.0166) | (0.0053) | (0.0111) | (0.0018) | (0.0008) | (0.0056) | (0.0023) | (0.0117) | (0.0247) |
| Website F.E. | Y | Y | Y | Y | Y | Y | Y | Y | Y | Y |
| Politician F.E. | Y | Y | Y | Y | Y | Y | Y | Y | Y | Y |
| Observations | 79761 | 79761 | 79761 | 79761 | 79761 | 79761 | 79761 | 79761 | 79761 | 79761 |
| Clusters | 92 | 92 | 92 | 92 | 92 | 92 | 92 | 92 | 92 | 92 |

Notes: Table shows the results from OLS regressions where the emotion score for a politician's face is the dependent variable and with politician and website fixed effects. 'Partisanship Score x Republican Politician' is the interaction between the Berkman Klein partisanship score with an indicator for whether the politician is a Republican. The emotion that is used as the dependent variable for each regression is noted in the column header. 'Negative' denotes the sum of the anger, sadness, contempt, disgust, surprise, and fear scores. 'Relative Fav.' denotes the happiness score minus the negative score. Standard errors clustered by website are in parentheses.

# Online Appendix

## Slanted images:
## Measuring nonverbal media bias

Levi Boxell, *Stanford University*[*]

September 2018

[*]E-mail: lboxell@stanford.edu. I thank Matthew Gentzkow and Jesse M. Shapiro for their comments and suggestions. Funding was generously provided by the National Science Foundation (grant number: DGE-1656518) and the Institute for Humane Studies.

## List of Figures



## List of Tables



# 1 Politician Sample

The main Republican and Democratic presidential candidates are defined as follows for each year. For 2008, the sample includes candidates that withdrew after the primaries started. For 2012, the sample includes candidates that appeared on at least three primary ballots for Republicans and candidates that captured at least one percent of the primary vote for Democrats. For 2016, the sample includes candidates that withdrew after the primaries started for Republicans and candidates that were on at least 6 state ballots and invited to a forum or debate for Democrats.

The main congressional leaders include the President of the Senate, the President pro tempore, the Speaker of the House, the majority and minority leader of the House and Senate, and the majority and minority whip of the House and Senate.

Figure 1: Histogram of emotions

*Panel A: Relative Favorability*

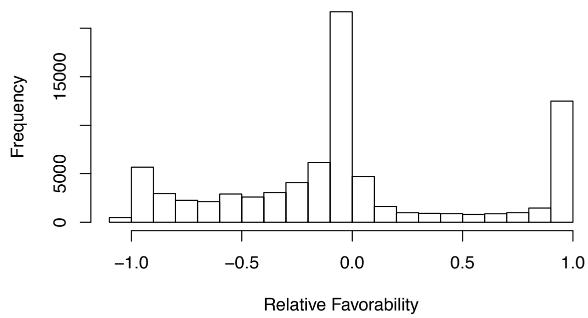


*Panel B: Happiness*

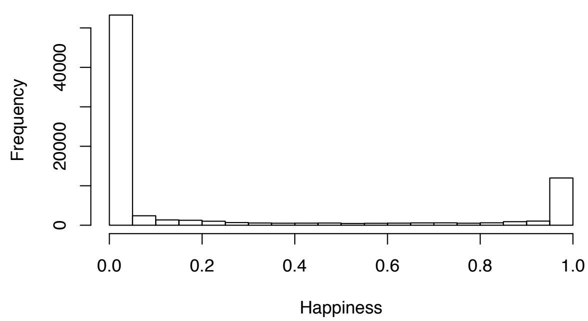


*Panel C: Negative Emotions*

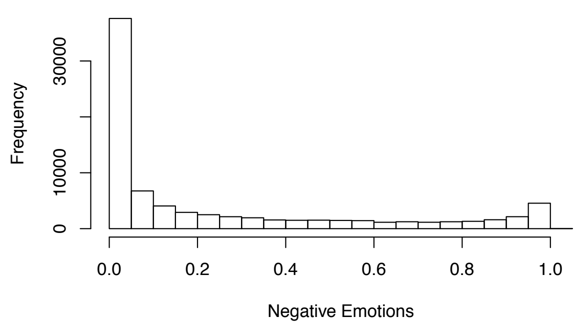


Notes: Each panel plots a histogram of the emotion scores for images with at least 0.5 match confidence.

Figure 2: Nonverbal Slant and Segregation of Twitter Shares

*Panel A: Relative Favorability Towards Republicans*

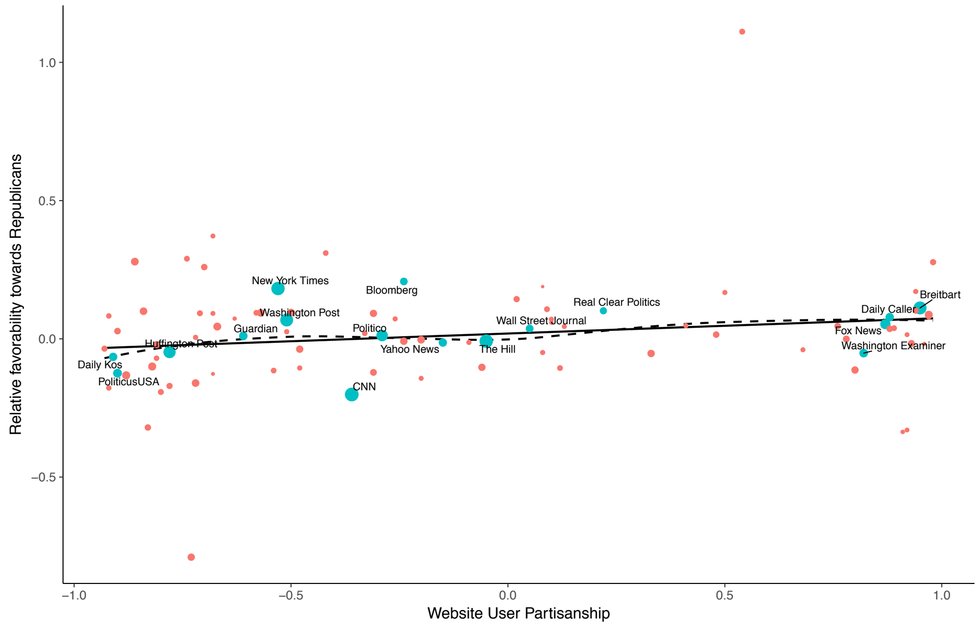


*Panel B: Relative Favorability Towards Democrats*

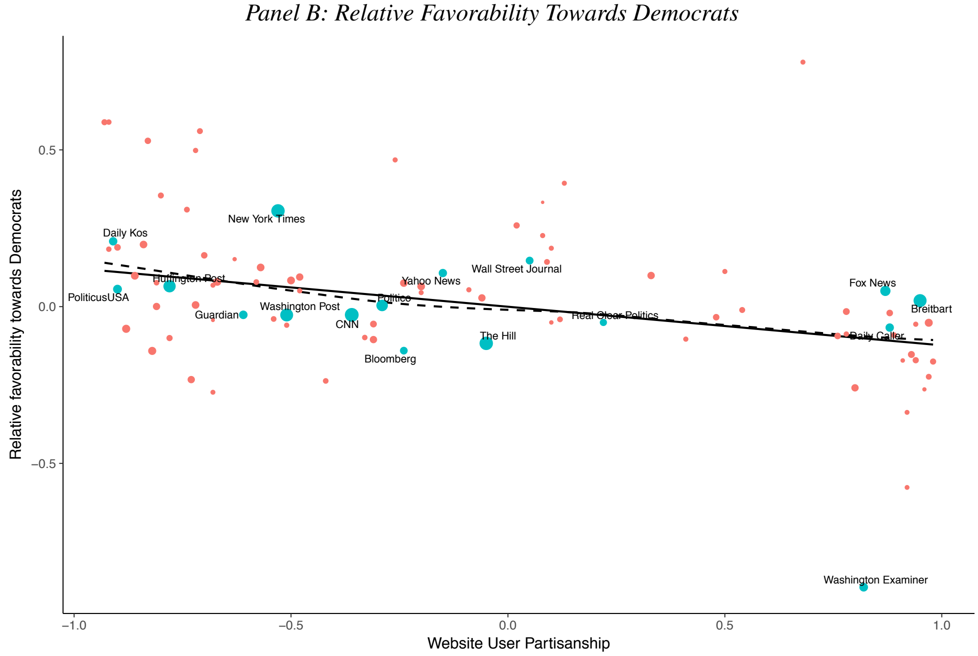


Notes: Panel A plots average relative favorability towards Republicans (demeaned across all websites). Panel B plots the average relative favorability towards Democrats (demeaned across all websites). The black solid line is the linear best fit weighted by shares on Twitter; the black dashed line is a loess fit weighted by shares on Twitter. The labelled blue dots indicate websites with at least 10,000 Twitter shares or 2,000 media inlinks in the Berkman Klein data.

Figure 3: Nonverbal Slant and Segregation of Twitter Shares, Trump-Clinton

*Panel A: Relative Favorability Towards Trump*

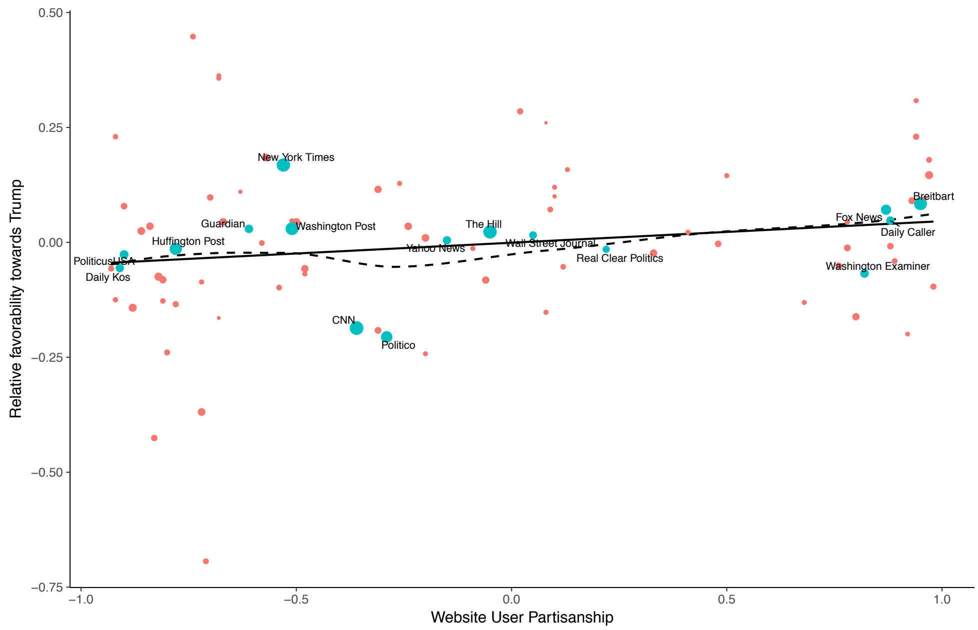


*Panel B: Relative Favorability Towards Clinton*

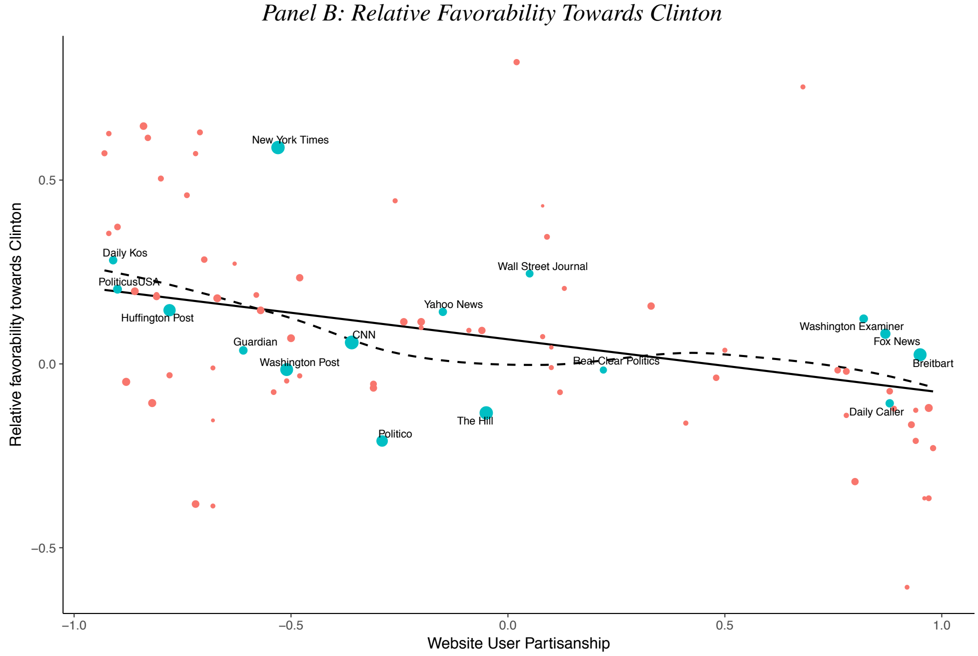


Notes: Panel A plots average relative favorability towards Donald Trump (demeaned across all websites). Panel B plots the average relative favorability towards Hillary Clinton (demeaned across all websites). The black solid line is the linear best fit weighted by shares on Twitter; the black dashed line is a loess fit weighted by shares on Twitter. The labelled blue dots indicate websites with at least 10,000 Twitter shares or 2,000 media inlinks in the Berkman Klein data.

Figure 4: Trump and Clinton Favorability, Day After Debates

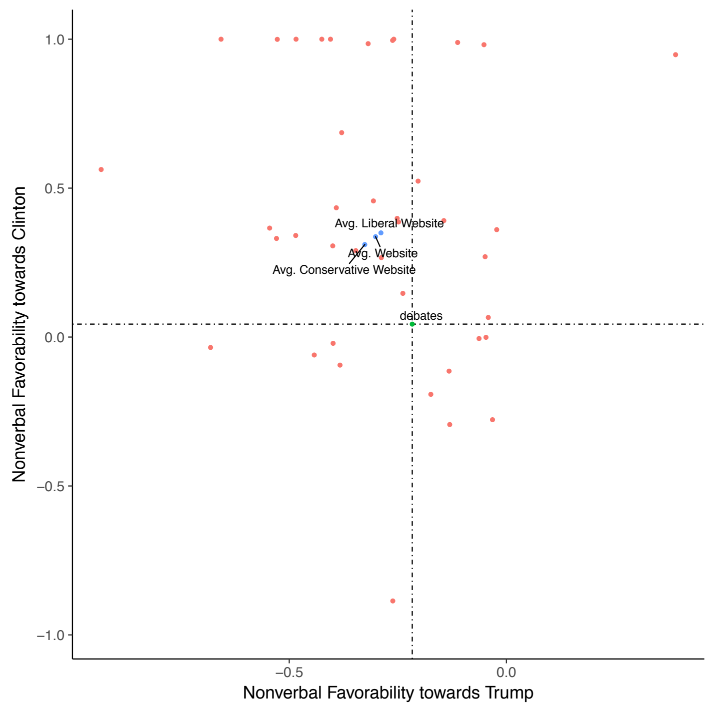


Notes: Figure plots the average relative favorability towards Clinton against the average relative favorability towards Trump for each website after restricting observations to the subsequent day after one of the 2016 presidential debates. The debate point indicate the average relative favorability across the 2016 presidential debates. The other labeled points are the average relative favorability towards a candidate averaged across either all websites, websites with a positive partisanship score, or websites with a negative partisanship score and weighted by the number of Twitter shares.

Table 1: Number of images by politician

| Politician | Total |
|---|---|
| Alan Keyes | 242 |
| Barack Obama | 8336 |
| Ben Carson | 1471 |
| Bernie Sanders | 3557 |
| Carly Fiorina | 676 |
| Chris Christie | 803 |
| Chris Dodd | 252 |
| Donald Trump | 28658 |
| Gary Johnson | 290 |
| Harry Reid | 452 |
| Hillary Clinton | 14348 |
| Jeb Bush | 1054 |
| Joe Biden | 739 |
| John Boehner | 370 |
| John Cornyn | 207 |
| John Kasich | 622 |
| John McCain | 444 |
| Kevin McCarthy | 291 |
| Marco Rubio | 2063 |
| Mike Huckabee | 458 |
| Mike Pence | 1276 |
| Mitch McConnell | 1143 |
| Mitt Romney | 357 |
| Nancy Pelosi | 554 |
| Newt Gingrich | 769 |
| Paul Ryan | 2960 |
| Rand Paul | 482 |
| Richard Durbin | 129 |
| Rick Santorum | 160 |
| Robert Byrd | 145 |
| Rudy Giuliani | 360 |
| Sarah Palin | 314 |
| Steny Hoyer | 256 |
| Ted Cruz | 3596 |
| Tim Kaine | 538 |
| Total | 79761 |

Notes: Table shows the number of images in the baseline sample for each politician. Only politicians with at least 125 total images are included. The ‘Total’ row includes all politicians in the sample including those not mentioned explicitly in the table.

Table 2: Number of images by website

| Website | Image Count | Website | Image Count |
|---|---|---|---|
| ABC News | 413 | Alternet | 1608 |
| Bipartisan Report | 365 | BizPacReview | 259 |
| Breitbart | 3989 | Business Insider | 874 |
| Buzzfeed | 381 | CBS News | 184 |
| CNBC | 140 | CNN | 204 |
| Conservative Tribune | 1414 | CSPAN | 334 |
| Daily Caller | 4726 | Daily Kos | 416 |
| Daily Newsbin | 524 | EndingtheFed | 216 |
| FactCheck.org | 1274 | FiveThirtyEight | 476 |
| Fox News | 629 | Free Beacon | 1474 |
| Gateway Pundit | 663 | Gawker | 201 |
| Guardian | 365 | Huffington Post | 2155 |
| IBTimes | 680 | InfoWars | 4029 |
| Inquisitr | 150 | Judicial Watch | 192 |
| Media Matters | 3262 | Mother Jones | 318 |
| MSNBC | 1646 | NBC News | 5692 |
| Newsweek | 1318 | Observer | 710 |
| Occupy Democrats | 1271 | opensecrets.org | 354 |
| PBS NewsHour | 615 | People | 407 |
| Political Insider | 173 | Politico | 420 |
| PoliticusUSA | 4800 | Raw Story | 3151 |
| Real Clear Politics | 1442 | RedState | 2076 |
| Reuters | 257 | Right Scoop | 2733 |
| RT | 476 | Salon | 409 |
| sanders.senate.gov | 328 | Talking Points Memo | 2331 |
| tedcruz.org | 222 | The Federalist | 1419 |
| The Hill | 3233 | The Intercept | 190 |
| The Nation | 847 | The Onion | 1268 |
| The Week | 3285 | Time | 164 |
| townhall.com | 1346 | US News & World Report | 176 |
| US Uncut | 877 | Vanity Fair | 161 |
| Vox | 194 | Wall Street Journal | 431 |
| Washington Post | 524 | Washington Times | 1096 |
| Yahoo News | 459 | Zero Hedge | 332 |
| | | Total | 79761 |

Notes: Table shows the number of images for each website after restricting to images with at least 50 percent match confidence. Only website with at least 125 images are shown. The total row is the sum across all websites.

Table 3: Average emotion by politician

| Politician | Neutral | Happiness | Anger | Sadness | Contempt | Disgust | Surprise | Fear |
|---|---|---|---|---|---|---|---|---|
| Alan Keyes | 0.826 | 0.160 | 0.003 | 0.004 | 0.004 | 0.001 | 0.002 | 0.000 |
| Barack Obama | 0.696 | 0.215 | 0.013 | 0.048 | 0.010 | 0.002 | 0.015 | 0.001 |
| Ben Carson | 0.642 | 0.273 | 0.002 | 0.045 | 0.025 | 0.001 | 0.008 | 0.003 |
| Bernie Sanders | 0.595 | 0.248 | 0.087 | 0.014 | 0.021 | 0.002 | 0.031 | 0.002 |
| Carly Fiorina | 0.490 | 0.438 | 0.025 | 0.005 | 0.006 | 0.002 | 0.033 | 0.001 |
| Chris Christie | 0.698 | 0.096 | 0.085 | 0.009 | 0.005 | 0.010 | 0.091 | 0.004 |
| Chris Dodd | 0.719 | 0.198 | 0.058 | 0.012 | 0.010 | 0.000 | 0.002 | 0.000 |
| Donald Trump | 0.438 | 0.160 | 0.140 | 0.161 | 0.016 | 0.015 | 0.067 | 0.004 |
| Gary Johnson | 0.578 | 0.240 | 0.045 | 0.005 | 0.014 | 0.002 | 0.112 | 0.003 |
| Harry Reid | 0.798 | 0.126 | 0.017 | 0.015 | 0.008 | 0.002 | 0.034 | 0.000 |
| Hillary Clinton | 0.422 | 0.375 | 0.020 | 0.031 | 0.010 | 0.001 | 0.119 | 0.022 |
| Jeb Bush | 0.688 | 0.201 | 0.014 | 0.010 | 0.005 | 0.001 | 0.080 | 0.002 |
| Joe Biden | 0.569 | 0.295 | 0.049 | 0.063 | 0.004 | 0.003 | 0.015 | 0.001 |
| John Boehner | 0.710 | 0.138 | 0.007 | 0.116 | 0.012 | 0.007 | 0.009 | 0.001 |
| John Cornyn | 0.868 | 0.075 | 0.003 | 0.040 | 0.001 | 0.000 | 0.012 | 0.000 |
| John Kasich | 0.746 | 0.146 | 0.022 | 0.021 | 0.012 | 0.001 | 0.050 | 0.001 |
| John McCain | 0.754 | 0.122 | 0.026 | 0.048 | 0.009 | 0.003 | 0.035 | 0.002 |
| Kevin McCarthy | 0.687 | 0.228 | 0.017 | 0.027 | 0.002 | 0.002 | 0.036 | 0.001 |
| Marco Rubio | 0.678 | 0.232 | 0.009 | 0.022 | 0.004 | 0.000 | 0.053 | 0.001 |
| Mike Huckabee | 0.633 | 0.199 | 0.075 | 0.009 | 0.007 | 0.004 | 0.071 | 0.002 |
| Mike Pence | 0.665 | 0.214 | 0.028 | 0.065 | 0.008 | 0.003 | 0.016 | 0.001 |
| Mitch McConnell | 0.636 | 0.157 | 0.003 | 0.111 | 0.002 | 0.001 | 0.084 | 0.006 |
| Mitt Romney | 0.467 | 0.361 | 0.040 | 0.050 | 0.012 | 0.005 | 0.061 | 0.004 |
| Nancy Pelosi | 0.401 | 0.296 | 0.012 | 0.028 | 0.004 | 0.002 | 0.238 | 0.020 |
| Newt Gingrich | 0.786 | 0.170 | 0.028 | 0.009 | 0.003 | 0.001 | 0.002 | 0.000 |
| Paul Ryan | 0.608 | 0.240 | 0.004 | 0.099 | 0.026 | 0.000 | 0.022 | 0.001 |
| Rand Paul | 0.813 | 0.081 | 0.007 | 0.008 | 0.004 | 0.001 | 0.085 | 0.000 |
| Richard Durbin | 0.818 | 0.149 | 0.002 | 0.016 | 0.002 | 0.000 | 0.012 | 0.000 |
| Rick Santorum | 0.339 | 0.510 | 0.021 | 0.075 | 0.012 | 0.003 | 0.038 | 0.002 |
| Robert Byrd | 0.753 | 0.179 | 0.039 | 0.017 | 0.007 | 0.001 | 0.005 | 0.000 |
| Rudy Giuliani | 0.647 | 0.178 | 0.024 | 0.019 | 0.008 | 0.009 | 0.110 | 0.006 |
| Sarah Palin | 0.342 | 0.452 | 0.020 | 0.020 | 0.017 | 0.007 | 0.137 | 0.005 |
| Steny Hoyer | 0.754 | 0.174 | 0.018 | 0.020 | 0.012 | 0.001 | 0.020 | 0.000 |
| Ted Cruz | 0.469 | 0.216 | 0.006 | 0.276 | 0.008 | 0.011 | 0.010 | 0.004 |
| Tim Kaine | 0.354 | 0.525 | 0.032 | 0.023 | 0.006 | 0.002 | 0.056 | 0.001 |
| Total | 0.646 | 0.235 | 0.027 | 0.032 | 0.008 | 0.003 | 0.047 | 0.003 |

Notes: Table shows the average emotion values for images in the baseline sample for each politician. Only politicians with at least 125 total images are included. The ‘Total’ row includes all politicians in the sample including those not mentioned explicitly in the table and is the average across politicians’ averages.

Table 4: Website Slant

| Website | Overall Slant | SD | Website | Trump-Clinton Slant | SD |
|---|---|---|---|---|---|
| Daily Kos | -0.274 | (0.081) | New York Times | -0.420 | (0.270) |
| PoliticusUSA | -0.180 | (0.021) | Daily Kos | -0.338 | (0.145) |
| CNN | -0.175 | (0.110) | Time | -0.292 | (0.150) |
| RT | -0.152 | (0.075) | Talking Points Memo | -0.265 | (0.070) |
| Time | -0.132 | (0.109) | CNN | -0.245 | (0.153) |
| Business Insider | -0.131 | (0.062) | Wall Street Journal | -0.230 | (0.098) |
| New York Times | -0.123 | (0.201) | PoliticusUSA | -0.230 | (0.045) |
| Yahoo News | -0.121 | (0.078) | RT | -0.181 | (0.109) |
| Huffington Post | -0.113 | (0.035) | Business Insider | -0.173 | (0.096) |
| Wall Street Journal | -0.110 | (0.077) | Mother Jones | -0.173 | (0.192) |
| ABC News | -0.084 | (0.078) | Huffington Post | -0.160 | (0.053) |
| Reuters | -0.068 | (0.112) | Yahoo News | -0.137 | (0.108) |
| Raw Story | -0.061 | (0.028) | MSNBC | -0.134 | (0.059) |
| MSNBC | -0.035 | (0.038) | Reuters | -0.105 | (0.157) |
| Buzzfeed | -0.031 | (0.090) | Raw Story | -0.093 | (0.046) |
| Talking Points Memo | -0.022 | (0.044) | ABC News | -0.079 | (0.120) |
| Fox News | 0.003 | (0.066) | NBC News | -0.026 | (0.031) |
| Politico | 0.008 | (0.071) | Fox News | -0.011 | (0.092) |
| NBC News | 0.011 | (0.021) | Guardian | -0.007 | (0.141) |
| New York Post | 0.015 | (0.198) | Real Clear Politics | 0.002 | (0.067) |
| Guardian | 0.037 | (0.084) | Politico | 0.004 | (0.096) |
| Salon | 0.041 | (0.089) | New York Post | 0.008 | (0.259) |
| Breitbart | 0.093 | (0.026) | Salon | 0.032 | (0.128) |
| Washington Post | 0.095 | (0.071) | Buzzfeed | 0.039 | (0.139) |
| The Hill | 0.109 | (0.022) | Washington Post | 0.045 | (0.110) |
| InfoWars | 0.136 | (0.025) | Breitbart | 0.059 | (0.041) |
| Gateway Pundit | 0.139 | (0.064) | Daily Caller | 0.155 | (0.035) |
| Daily Caller | 0.145 | (0.024) | The Hill | 0.156 | (0.037) |
| Right Scoop | 0.146 | (0.043) | Right Scoop | 0.158 | (0.068) |
| Real Clear Politics | 0.152 | (0.037) | InfoWars | 0.256 | (0.031) |
| Mother Jones | 0.181 | (0.106) | Gateway Pundit | 0.266 | (0.080) |

Notes: Table reports, on the left side, the average relative favorability towards Republicans (demeaned across all websites) minus the average relative favorability towards Democrats (demeaned across all websites). The right side of the table reports the same differential measure of favorability, but for Donald Trump and Hillary Clinton. The websites ranked have more than 5,000 Twitter shares or 2,000 media inlinks in the Berkman Klein data and more than 5 images of both Trump and Clinton. The standard deviation estimates are constructed by taking the standard deviation of the demeaned average relative favorability for each website-partisan group separately, dividing by the square root of the number of images in each website-partisan group, and summing across both partisan groups. The Republican-leaning estimate for Mother Jones is driven by a disproportionate number of positive Carly Fiorina images—dropping Carly Fiorina or restricting attention to Donald Trump and Hillary Clinton removes this discrepancy.

Table 5: Partisanship of users and politician emotions, robustness

| | Dependent Variable: Relative Favorability | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) | (10) | (11) |
| Partisanship Score x | | | | | | | | | | | |
| Republican Politician | 0.1675 | 0.1651 | 0.1395 | 0.1610 | 0.1610 | 0.0879 | 0.1806 | 0.3010 | 0.1934 | 0.0768 | 0.2105 |
| | (0.0246) | (0.0264) | (0.0548) | (0.0237) | (0.0059) | (0.0123) | (0.0288) | (0.0609) | (0.0695) | (0.0205) | (0.0296) |
| Website F.E. | Y | Y | Y | Y | Y | Y | Y | Y | Y | Y | Y |
| Politician F.E. | Y | Y | Y | Y | Y | Y | Y | Y | Y | Y | Y |
| Observations | 69714 | 48787 | 13977 | 79761 | 79761 | 79761 | 79761 | 48098 | 31663 | 33596 | 43006 |
| Clusters | 92 | 92 | 89 | 61 | — | 92 | 92 | 55 | 37 | 40 | 90 |

Notes: Table shows the results from OLS regressions where the emotion score for a politician's face is the dependent variable and with politician and website fixed effects unless otherwise specified. 'Partisanship Score x Republican Politician' is the interaction between the Berkman Klein partisanship score with an indicator for whether the politician is a Republican. The emotion that is used as the dependent variable for each regression is noted in the column header. 'Negative' denotes the sum of the anger, sadness, contempt, disgust, surprise, and fear scores. 'Relative Fav.' denotes the happiness score minus the negative score. Columns (1)-(3) restrict to images with a match confidence of at least 0.6, 0.7, and 0.8 respectively. Columns (4) and (5) use standard errors clustered by politician and robust standard errors respectively. Column (6) uses the log of the relative favorability measure after shifting it to be positive. Column (7) uses a tobit estimator with bounds of -1 and 1 after constraining relative favorability values to fall within this range. Columns (8)–(10) restrict observations to websites with negative partisanship scores, positive partisanship scores, and at least 1000 media inlinks respectively. Column (11) restricts images to Donald Trump and Hillary Clinton. Standard errors clustered by website are in parentheses unless otherwise specified.

Table 6: Partisanship of users and politician party

| | Dependent Variable: Indicator for Republican Politician | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) |
| | OLS | Logit | OLS | Logit | OLS | Logit | OLS | Logit |
| Partisanship Score | -0.0824 | -0.3518 | -0.0837 | -0.3593 | -0.0462 | -0.1984 | -0.0472 | -0.2042 |
| | (0.0284) | (0.1212) | (0.0290) | (0.1254) | (0.0329) | (0.1416) | (0.0340) | (0.1479) |
| log(Twitter Shares) | | | 0.0112 | 0.0517 | | | 0.0123 | 0.0545 |
| | | | (0.0208) | (0.0925) | | | (0.0247) | (0.1089) |
| log(Media Inlinks) | | | -0.0124 | -0.0590 | | | -0.0163 | -0.0733 |
| | | | (0.0286) | (0.1290) | | | (0.0323) | (0.1449) |
| Observations | 79761 | 79761 | 79761 | 79761 | 48787 | 48787 | 48787 | 48787 |
| Clusters | 92 | 92 | 92 | 92 | 92 | 92 | 92 | 92 |

Notes: Table shows the results from OLS and logit regressions where an indicator for whether the politician is a Republican in the image is the dependent variable. ‘Partisanship Score’ is the measure of partisanship for the website from the Berkman Klein data, ‘log(Twitter Shares)’ is the log of the number of twitter shares in the Berkman Klein data, and ‘log(Media Inlinks)’ is the log of the number of media inlinks reported by the Berkman Klein data. Columns (1)-(4) restrict to images with at least 0.5 match confidence. Columns (5)-(8) restrict to images with at least 0.7 match confidence. Standard errors clustered by website are in parentheses.

Table 7: Slant over the election cycle

| | Primaries | | | Post-Primaries | | | Post-Election | | |
|---|---|---|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) |
| | Happiness | Negative | Relative Fav. | Happiness | Negative | Relative Fav. | Happiness | Negative | Relative Fav. |
| Partisanship Score x | | | | | | | | | |
| Republican Politician | 0.0446 | -0.0906 | 0.1351 | 0.1143 | -0.0814 | 0.1958 | 0.0784 | -0.1084 | 0.1868 |
| | (0.0182) | (0.0133) | (0.0277) | (0.0260) | (0.0097) | (0.0298) | (0.0195) | (0.0154) | (0.0276) |
| Website F.E. | Y | Y | Y | Y | Y | Y | Y | Y | Y |
| Politician F.E. | Y | Y | Y | Y | Y | Y | Y | Y | Y |
| Observations | 35704 | 35704 | 35704 | 22391 | 22391 | 22391 | 21666 | 21666 | 21666 |
| Clusters | 85 | 85 | 85 | 79 | 79 | 79 | 71 | 71 | 71 |

Notes: Table shows the results from OLS regressions where the emotion score for a politician's face is the dependent variable and with politician and website fixed effects. 'Share Conservative x Republican Politician' is the interaction between the share of website visitors that identify as conservatives with an indicator for whether the politician is a Republican. The emotion that is used as the dependent variable for each regression is noted in the column header. 'Negative' denotes the sum of the anger, sadness, contempt, disgust, surprise, and fear scores. 'Relative Fav.' denotes the happiness score minus the negative score. Columns (1)–(3) restrict data to images before Trump (for Republicans) or Clinton (for Democrats) clinched the nomination. Columns (4)–(6) restrict data to images after the nomination had been clinched, but before the election. Columns (7)–(9) restrict data to the post-election period. Standard errors clustered by website are in parentheses.

Table 8: Slant over the election cycle, Trump-Clinton

| | Primaries | | | Post-Primaries | | | Post-Election | | |
|---|---|---|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) |
| | Happiness | Negative | Relative Fav. | Happiness | Negative | Relative Fav. | Happiness | Negative | Relative Fav. |
| Partisanship Score x Republican Politician | 0.0540 | -0.1245 | 0.1785 | 0.1385 | -0.0834 | 0.2219 | 0.1059 | -0.1791 | 0.2850 |
| | (0.0257) | (0.0150) | (0.0310) | (0.0337) | (0.0117) | (0.0384) | (0.0358) | (0.0369) | (0.0578) |
| Website F.E. | Y | Y | Y | Y | Y | Y | Y | Y | Y |
| Politician F.E. | Y | Y | Y | Y | Y | Y | Y | Y | Y |
| Observations | 14362 | 14362 | 14362 | 15333 | 15333 | 15333 | 13311 | 13311 | 13311 |
| Clusters | 82 | 82 | 82 | 74 | 74 | 74 | 67 | 67 | 67 |

Notes: Table shows the results from OLS regressions where the emotion score for a politician's face is the dependent variable and with politician and website fixed effects. 'Share Conservative x Republican Politician' is the interaction between the share of website visitors that identify as conservatives with an indicator for whether the politician is a Republican. The emotion that is used as the dependent variable for each regression is noted in the column header. 'Negative' denotes the sum of the anger, sadness, contempt, disgust, surprise, and fear scores. 'Relative Fav.' denotes the happiness score minus the negative score. Columns (1)–(3) restrict data to images before Trump (for Trump) or Clinton (for Clinton) clinched the nomination. Columns (4)–(6) restrict data to images after the nomination had been clinched, but before the election. Columns (7)–(9) restrict data to the post-election period. Data is restricted to Donald Trump and Hillary Clinton. Standard errors clustered by website are in parentheses.