Moral Language Use by U.S. Political Elites

Sze-Yuh Nina Wang and Yoel Inbar*

Department of Psychology, University of Toronto, Toronto, Ontario, M1C 1A4, Canada.

* To whom correspondence should be addressed. Email: yoel.inbar@utoronto.ca

Abstract

We use a distributed language model to examine the moral language used by U.S. political elites. In Study 1, we analyze 687,360 Twitter messages (“tweets”) posted by accounts belonging to Democratic and Republican members of Congress from 2016-2018. In Study 2, we analyze 2,630,688 speeches given on the floor of the House and Senate from 1981-2017. We find that partisan differences in moral language use shift over time as the parties gain or lose political power. Overall, lower political power was associated with greater use of moral language for both Democrats and Republicans. On Twitter, Democrats used more moral language in the period after Donald Trump won the 2016 presidential election. In Congressional transcripts, both Democrats and Republicans used more of most kinds of moral language when they were in the minority.

Moral Language Use by U.S. Political Elites

Political elites use language to shift the public debate and persuade and motivate voters (Lakoff, 2004). Moral language is particularly powerful, as voters respond especially strongly to moral language from elites (Clifford & Jerit, 2013). Once people connect an issue to their moral values, they are more likely to act to be sure their side prevails, be it by persuasion, voting, protesting, or—in extreme cases—violence (Skitka & Mullen, 2002). Indeed, morality is so strong a motivator that it can sometimes justify any means to a moral end (Fiske & Rai, 2014; Skitka, 2010).

Despite the significant consequences of moralization, the use of moral language by political elites is poorly understood. Research has largely focused on ideological differences in moral language use using the theoretical framework of Moral Foundations Theory (Haidt & Graham, 2007), which posits five moral domains (“foundations”) said to constitute the basic building blocks of morality across cultures. The “individualizing” foundations of *harm* and *fairness* concern individual rights and well-being, whereas the “binding” foundations of *loyalty*, *authority*, and *purity* concern adherence to norms that maintain group cohesion. In the U.S. and elsewhere, both liberal and conservative members of the public rate the individualizing foundations as morally relevant, although liberals endorse them somewhat more strongly. In contrast, the binding foundations are endorsed much more strongly by conservatives than liberals (Graham, Haidt, & Nosek, 2009).

For elite communication, however, there is not a consistent pattern of ideological differences in moral language use. For example, in U.S. Senate debates on abortion from 1989-2006, Republicans used more purity-related language *and* more fairness-related language than

did Democrats (Sagi & Dehghani, 2014). In four months of Twitter posts from 2014, more liberal members of Congress used more moral language related to every foundation except authority, which was referenced more by conservative legislators (Sterling & Jost, 2017). However, in New York Times Op-Ed articles on stem-cell research (published 1999-2010), harm-related language was used more by writers taking the liberal position (i.e., supporting stem-cell research) and purity-related language, though rare overall, was used exclusively by writers taking the conservative position (Clifford & Jerit, 2013). Finally, in political texts (including State of the Union addresses, speeches from the House and Senate floors, and party platforms), associations between ideology and moral language use were inconsistent and dependent on the specific dictionaries used to assess moral language (Frimer, 2020).

These mixed results might arise for several reasons. First, many studies used dictionary-based word counting to assess moral language use (this was the case for every study described above besides Sagi & Dehghan, 2014). Especially when dictionary words are uncommon, word counting can produce unreliable results that are highly sensitive to the exact words included (Garten et al., 2017). Second, researchers often examined rhetoric around specific issues, which may introduce issue-specific differences (e.g., liberals may moralize some issues more than conservatives, and vice versa). Finally, no research has examined whether moral language use by elites changes over time, as features of the political environment change. One particularly important feature of the environment is political power, that is, whether a political party controls the executive and/or legislative branches of government. It may be that as parties gain and lose control of branches of government, their members change the moral language they use. If this is

the case, then research examining moral language use in different time periods would not necessarily be expected to show consistent results.

Political power and moral language

There are several reasons to think that lower power might lead to changes in moral language and, specifically, to greater moralizing. First, research on social identity (Tajfel & Turner, 1986) finds that threatened groups (in this case, the minority party) will be particularly motivated to defend their group identity (Branscombe & Wann, 1992). Moralizing is one way to do this: It can bolster a threatened group identity by elevating ingroup values to sacred principles and casting outgroup members as irredeemably wicked (Fiske & Rai, 2014; Haidt & Graham, 2012). Second, research on integrative complexity (Tetlock, 1986) has found that minority party members generally show lower integrative complexity in their public statements than do majority party members (Suedfeld, 2010). Although integrative complexity and morality are distinct, they share some characteristics: Low integrative complexity is characterized by black-and-white thinking and a rejection of compromise, just as moralizing is (Baron & Spranca, 1997; Tetlock, Lerner, Elson, & Green, 2000). Finally, parties have a strategic incentive to moralize when they need to mobilize their supporters to regain a majority, because moral language is particularly effective at increasing message transmission (Brady, Wills, Burkart, Jost, & Van Bavel, 2019; Brady, Wills, Jost, Tucker, & Van Bavel, 2017) and motivating supporters (Ryan, 2014; Skitka & Bauman, 2008).

The current research

In the current research, we examined 1) whether U.S. political elites show ideological differences in moral language use; and 2) whether these differences vary over time—in

particular, whether they change as parties gain and lose power. In Study 1, we examined the complete set of Twitter posts made by all accounts belonging to U.S. members of Congress in the two-year period beginning January 2016. In Study 2, we examined all speeches made on the floor of the House and Senate in the 97th-114th Congresses (1981-2017). In each corpus, we assessed how much Democrats and Republicans used language relevant to the moral foundations proposed by moral foundations theory (*harm, fairness, loyalty, authority, and purity*), and how this varied with shifts in political power.

Study 1

In Study 1, we analyzed the public tweets posted by accounts belonging to members of the U.S. Congress (i.e., the House of Representatives and the Senate). Nearly every member of Congress has a Twitter account, and many have a large number of followers (users who have subscribed to see another user's tweets). In our data set, for example, the most-followed account (Bernie Sanders) had over 7 million followers, but even the median account had over 20,000. Consequently, Twitter is a powerful tool for politicians to rally supporters and communicate with the public, and tweets from elected representatives are a natural way to examine the moral language used by political elites.

Methods

Data and code availability. All R scripts and links to data are available at <https://github.com/yoelinbar/moral-language>. Due to copyright restrictions, the full tweet text is not included. Full text for any tweet can be retrieved using Twitter's API and the tweet GUID in the data file.

Tweet collection and cleaning. We obtained the official Twitter accounts for members of Congresses 114 and 115 (from <https://gwu-libraries.github.io/sfm-ui/posts/2017-05-23-congress-seed-list>) and verified the list with web searches. Using a set of Python scripts and the “Tweepy” library, we used Twitter’s API to collect all original Tweets (i.e., omitting those marked by Twitter as retweets) posted by these accounts from January 1, 2016 to January 31, 2018, inclusive. We cleaned Tweets using a Python script that removed manual retweets (i.e., text following the strings “RT:” or “MT:”), URLs, and punctuation. We removed tweets from eight accounts that had posted fewer than 50 tweets each, because we thought they might cause estimation problems in our mixed-effects models. (All these removals are documented in the R script.) Dictionary similarity scores (see below) could not be computed for 2,701 cleaned tweets (these tweets were almost entirely non-words that had been hashtags, e.g., “VA10,” “WhyWeMarch,” or “FlyEaglesFly”). Our final sample comprised 687,360 Tweets from 578 unique accounts: 385,206 from Democratic members of Congress ($n = 261$) and 302,154 from Republicans ($n = 317$).

Moral language analysis. A major difficulty in examining naturally-occurring moral language use is that it is difficult to measure at scale. Human coders are considered the gold standard of natural language analysis, but human-coding hundreds of thousands of texts is infeasible. Automated text analysis by word counting (as implemented by, for example, the Linguistic Inquiry and Word Count [LIWC] software) is an attractive alternative to human coding but is limited in different ways. Most problematically, researchers must specify the entire dictionary of words or word stems that match a concept. If the dictionary coverage is too narrow, important parts of the concept will be omitted and text will not match when it should. But if the

coverage is too broad, text will match when it shouldn't. And because word-counting software does not know what words mean, only what they look like, including more terms can easily lead to false-positives (e.g., "happ*" matches "happy" and happiness," but also "happened" and "happenstance"). This problem is exacerbated for short texts, such as tweets (which were until recently limited to 140 characters and are currently limited to 280), where the large majority of texts might contain no dictionary words at all (Garten et al., 2017). Especially for dictionaries comprising mainly low-frequency words (as is the case for the moral foundations dictionaries), researchers' subjective judgments about whether a term should be included can have a large influence on results.

In the present research we avoid these limitations by using a text analysis technique built on a distributed language model. Distributed language models, which derive from a long tradition of research in computational linguistics, cognitive psychology, and computer science (Landauer & Dumais, 1997; Osgood, Suci, & Tannenbaum, 1957; Rogers & McClelland, 2004; Salton, Wong, & Yang, 1975), encode word meaning as a point in a many-dimensional space. Semantically similar words (e.g. "car" and "automobile") are close to each other in this space. Each word's location in space is described by a n -dimensional vector of real numbers (where n = the number of dimensions in the space), and mathematical operations on these vectors are meaningful. The method we use, called distributed dictionary representations (DDR; Garten et al., 2017) combines a small researcher-specified dictionary of terms representing a concept with a distributed language model. Averaging the normalized vectors for each word in the dictionary produces a single composite representation of the dictionary's meaning. In the same way, the words in a text (in this case, a tweet) can be combined into a composite representation. Finally,

the similarity between the text and the dictionary can be computed by cosine similarity (scores range from -1 to 1, with higher values reflecting greater similarity). Importantly, the text need not contain any of the dictionary words for similarity to be computed. For example, a text containing “car” might be considered highly similar to a dictionary consisting of “automobile,” “auto,” and “vehicle.” This means that dictionaries can be short but still cover a concept well, and texts can be short but still be considered similar to a dictionary. This eliminates the two major shortcomings of word-counting approaches. In validation studies comparing automated methods to human coders, this method was substantially more accurate than word-counting (Garten et al., 2017).

We used moral foundations dictionaries developed and validated by Garten et al. (2017) in combination with a publicly available distributed language representation (<https://code.google.com/p/word2vec/>) trained on approximately 100 million words from Google News Articles (Mikolov, Sutskever, Chen, Corrado, & Dean, 2013). We used the “ddr” Python module developed by Garten et al. (2017) to compute the semantic similarity between each tweet and short dictionaries representing positive and negative aspects (“virtue” and “vice”) of each of the moral foundations (see Table 1). For each tweet we therefore had 10 values representing the similarity between the tweet and the positive and negative aspects of each foundation.

Table 1

Central concerns for each of the five moral foundations, and the DDR seed words used to measure positive (“virtue”) and negative (“vice”) moral language pertaining to each foundation.

<u>Moral Foundation</u>	<u>Concerns</u>	<u>DDR seed words</u>
Harm	Kindness, compassion, protection of the helpless or innocent	Virtue: <i>kindness, compassion, nurture, empathy</i> Vice: <i>suffer, cruel, hurt, harm</i>
Fairness	Justice, rights, cooperation, reciprocity	Virtue: <i>fairness, equality, justice, rights</i> Vice: <i>cheat, fraud, unfair, injustice</i>
Loyalty	Patriotism, heroism, fidelity, self-sacrifice	Virtue: <i>loyal, solidarity, patriot, fidelity</i> Vice: <i>betray, treason, disloyal, traitor</i>
Authority	Respect, duty, leadership, magnanimity	Virtue: <i>authority, obey, respect, tradition</i> Vice: <i>subversion, disobey, disrespect, chaos</i>
Purity	Chastity, piety, sacredness, self-restraint	Virtue: <i>purity, sanctity, sacred, wholesome</i> Vice: <i>impurity, depravity, degradation, unnatural</i>

Results

Language model validation. In a validation study looking specifically at moral language use on Twitter, DDR with these dictionaries showed the highest agreement with human coders of any automated method tested (Garten et al., 2017). In particular, it substantially outperformed word counting using the full Moral Foundations Dictionary (Graham et al., 2009).¹ We conducted a similar validation comparing a subset of the tweets in our sample to ratings given by human coders, and found that DDR with these dictionaries showed substantial agreement with human ratings (mean F1 score = .865, see Supplemental Methods for full details). Table S7 in the Supplemental Material shows, separately for Democrats and Republicans, the tweets in our sample that DDR identified as expressing each foundation and aspect most strongly.

Analytic strategy. We used DDR to compute the similarity of each tweet to each foundation and aspect; these similarity scores were our primary dependent variables. In order to account for non-independence of observations (i.e., multiple tweets from the same username), we used mixed-effects models for most of our analyses. Except where otherwise noted, we used the *lme4* package (Bates, Mächler, Bolker, & Walker, 2015) with restricted maximum likelihood estimation to fit our models and the *lmerTest* package (Kuznetsova, Brockhoff, & Christensen, 2017) to test significance of model coefficients using Satterthwaite-approximated degrees of freedom. These analyses were conducted in R 3.3.2 (R Core Team, 2016).

We fit separate models for each foundation and aspect. Thus, for each analysis, we fit 10 separate models. To model non-independence between multiple tweets from the same account,

¹ In terms of F1, a commonly used measure of classification accuracy that takes into account both false-positives and false-negatives, DDR achieved a mean F1 = .496 whereas word counting achieved a mean F1 = .275 (Garten et al., 2017).

we first added random person intercepts in addition to the fixed effects. Examining residual ACF plots from these models showed that there was also substantial autocorrelation across days (i.e., model errors from day n were correlated with those from day $n + 1$). To account for this autocorrelation we used the *splines* function in R to model a third-order polynomial effect of time (i.e., days since January 1, 2016) that was allowed to vary randomly by individual (see Verbyla, Curtis, Kenward, & Welham, 1999). We also added random intercepts for each day. We initially attempted to fit these models to the full set of tweets for each foundation and aspect, but no model converged. We therefore simplified the data by computing per-day average similarity scores for each account; that is, for each foundation and aspect we treated the mean similarity score of all tweets posted on a given day by a given account as one observation. These models converged and showed acceptably low levels of autocorrelation in the residuals.

Our key tests were implemented by dummy variables for party, time period, or their interaction. Because multilevel models are designed to accommodate varying numbers of observations between groups (see Snijders & Bosker, 2012, p. 56), tests of these dummy variables are unbiased despite the fact that we have different numbers of observations between different individuals, parties, or time periods.

Partisan differences in moral language use. For these analyses, we created a dummy variable for party (1 = Democrat; 0 = Republican); statistical tests are of the dummy coefficient. Table S1 in the Supplemental Results shows tests of the effects of party on moral language use for each moral foundation and aspect. For every foundation and aspect, Democrats used moral language more than Republicans, all $ps < .001$. As a measure of effect size, we standardized moral language scores so that the dummy coefficients can be interpreted as the difference in

scores between Democrats and Republicans, expressed in standard deviations (Lorah, 2018). Coefficients ranged from $b = .19$ (*loyalty-virtue*) to $b = .39$ (*fairness-vice*). Table S2 and Figure S1 in the Supplemental Results show that this pattern is consistent when predicting moral language use from each individual's DW-Nominate score, a continuous measure of ideology based on Congressional voting records (Poole & Rosenthal, 1997).

Changes in political power. Republicans were the majority party in the House and Senate 2016-2018. However, the presidency changed from Democratic to Republican control following the 2016 presidential election. We therefore tested the association between political power and moral language use by adding a dummy variable to our mixed-effects models that coded whether each tweet was posted before or after the 2016 presidential election. Moral language use by both Democrats and Republicans was higher in the post-election time period (i.e., after November 8, 2016). However, this difference between time periods was larger for Democrats than Republicans ($ps < .001$) for every foundation and aspect except *harm-virtue* (where it was equally large for both parties). For Democrats, pre- vs. post-election differences ranged from $b = .17$ (*harm-virtue*) to $b = .45$ (*harm-vice* and *fairness-vice*), all $ps < .001$. In contrast, for Republicans the largest increase was $b = .22$ (*fairness-virtue*), and in several cases pre- vs. post-election differences were very small ($bs < .10$), though still statistically significant. Table S3 in the Supplemental Results shows statistical tests from mixed-effects models of the differences in moral language between time periods and their interactions with party.

We next tested whether moral language use showed a discontinuity specifically at election day (November 8, 2016) with an interrupted time series analyses (Wagner, Soumerai, Zhang, & Ross-Degnan, 2002). These analyses, which are reported in detail in the Supplemental

Results and depicted in Figure 2, show positive and significant discontinuities for Democrats for every foundation except *harm-virtue*. The pattern for Republicans was much more mixed. For some foundations/aspects, moral language use increased, for others it did not change significantly or even decreased.

Figure 1. Mean similarity scores for each moral language category for tweets by Republicans (red) and Democrats (blue). Within a moral language category, each point represents the average of all tweets from one member of Congress. Points are jittered for legibility. Overlaid boxes show the 25th (bottom line), 50th (middle line), 75th (top line) percentiles. Whiskers extend to 1.5 * inter-quartile range.

Figure 2. Mean similarity scores for each moral language category pre- and post-election, separately for Democrats and Republicans. Each point represents the mean similarity for tweets from one day. Green and orange lines are separate linear fits of the time trend in scores pre- and post-election, respectively. The vertical gap between lines shows the discontinuity at November 8, 2016.

Topic models. Finally, we conducted an exploratory analysis of whether different topics of discussion were associated with the use of particular kinds of moral language, and with political ideology. We did this using a Structural Topic Model (STM; Roberts et al., 2014), a variation of the popular Latent Dirichlet Allocation (LDA; Blei, Ng, & Jordan, 2003) that allows for the inclusion of additional document-level metadata which may alter either the prevalence or content of topics. We fit our structural topic model using the ‘stm’ package in R (Roberts, Stewart, & Tingley, in press). Our final model had 77 topics. We included political affiliation and a dummy-coded variable of whether the tweet was made before or after the 2016 Presidential election as topical prevalence covariates, along with the tweet’s similarity score for each of the 10 moral foundations/aspects. These covariates and their interactions were significant for almost all topics, and an inspection of the top words and tweets for each topic revealed that the topics were interpretable and reflected meaningful facets of political discussion (see Table S8 in the Supplemental Results). This suggests that, as expected, there are differences in the topics that are moralized, the topics that are discussed by Democrats and Republicans, and the topics that were discussed before and after the election. These differences are plotted in Figures S4-S15 in the Supplemental Results. Figure S3 shows the semantic coherence (a measure of topic quality) of these topics.

Discussion

Posts from Democratic lawmakers' accounts consistently used more moral language relevant to all moral foundations. Partisan differences increased after the election of Donald Trump due to a substantial increase in moral language use by Democrats in the post-election period. Thus, ideological differences in moral language use were inconsistent with moral foundations theory (which predicts that *authority*, *loyalty*, and *purity* should be referenced more by Republicans than Democrats), and variant over time. In Study 2, we examine the relationship between political power and moral language use in a different corpus over a substantially longer time period.

Study 2

In Study 2, we analyzed the United States Congressional Record, which contains transcripts of the debates and proceedings of the U.S. Congress. Congressional transcripts are commonly used to study the speech of legislators (e.g., Gentzkow, Shapiro, & Taddy, 2019; Thomas, Pang, & Lee, 2006). Unlike tweets, congressional transcripts are available over a much longer period of time. This allowed us to test the effects of shifting political power on moral language use as the parties gained and lost control of the presidency, House, and Senate.

Methods

We used a dataset of the United States Congressional Record that parses the text by speech, including metadata on speeches and their speakers (Gentzkow, Shapiro, & Taddy, 2019). We restricted our analyses to the 97th Congress (beginning January 3, 1981) onwards, because our hypotheses concern moral language use by the minority party. Between the 72nd Congress (beginning March 4, 1931) and the 97th, Democrats had nearly unbroken control over both the House and Senate. Thus, for these years our

analyses would have confounded minority status and party. Additionally, before the civil rights era in the 1960s, the ideological composition of the two parties was different enough that analyses of partisan differences would not be particularly meaningful across longer periods of time (Brady & Stewart, 1982).

The dataset we used ends with the 114th Congress (ending January 3, 2017) so all our analyses are of the 97th-114th Congresses (i.e., January 3, 1981-January 3, 2017).

Moral language analysis. As in Study 1, we used DDR (Garten et al., 2017) to compute the semantic similarity between each speech and the positive and negative aspects of each moral foundation. We again found good agreement (mean F1 score = .791) between human coders and the DDR scores (please see Supplemental Materials for details).

Results

Analytic Strategy. Our analytic strategy is similar to Study 1. We again used mixed-effects models to account for non-independence of observations. These models included random intercepts for speakers and for Congresses (in the modern era, a single Congress lasts two years and consists of two sessions in consecutive years). We again fit separate models for each foundation and aspect, resulting in 10 models for each analysis. Initial models failed to converge, so we computed the per-Congress average similarity scores for each speaker (in other words, we treated all speeches given by an individual in a single Congress as a single observation). After this change, every model converged and showed acceptably low levels of autocorrelation in the residuals.

Partisan differences in moral language use. As in Study 1, we created a dummy variable for party (1 = Democrat; 0 = Republican). Table S9 in the Supplemental Results shows tests from mixed-effect models of the effect of party affiliation on moral language

use for each moral foundation and aspect. As in Study 1, when there were partisan differences in moral language use, Democrats tended to use more of it. However, differences were much less consistent than in Study 1. Overall, Democrats used significantly more moral language in two categories, *harm-virtue* ($b = 0.15$) and *fairness-virtue* ($b = 0.17$), with no significant effects of political party in other categories. This suggests that the strong partisan differences we observed in Study 1 are specific to that political environment.

Effects of political power. To test the association between power and moral language use, we created dummy variables representing whether or not a speaker's political party held control of their chamber of Congress (1 = yes; 0 = no). We ran mixed models assessing the effects of affiliation and political control separately for the House and Senate. Results of these mixed models are shown in Tables S10 and S11 in the Supplemental Results. We find largely consistent patterns for the effects of political control. In the Senate, the minority party used more moral language for all foundations/aspects except *fairness-virtue* and *authority-vice* (where there were no significant effects), and *authority-virtue* (which was used significantly more by the majority party). In the House, the minority party used more moral language for all foundations/aspects except *purity-vice* (where there were no significant effects), *fairness-virtue* (which was used significantly more by the majority party), and *authority-virtue* (which was used significantly more by the majority party). These results are shown in Figures 3 and 4. To determine whether party affiliation moderated the effects of political control, we compared models that contained only main effects of political control and party to models that also included an interaction term between the two. Likelihood ratio tests comparing each set of models found no significant improvements in

model fit from the inclusion of an interaction term for any model, indicating that the effects of political power on moral language use was not moderated by party.

We also tested for cross-chamber effects, that is, whether party control of the *other* chamber of Congress affected a speaker's moral language use, controlling for minority status in the speaker's own chamber. We found a significant cross-chamber effect for *authority-virtue* in the Senate, with no other significant effects in the Senate. Cross-chamber effects were significant for five moral language categories in the House, but effect sizes were small (ranging from $b = 0.055$ to $b = 0.083$). Likewise, we tested whether the speaker's party controlled the presidency and found significant effects for four categories such that the minority party used more moral language related to *fairness-virtue*, *fairness-vice*, and *authority-virtue*, and the majority party used more moral language related to *loyalty-virtue*, again with small effect sizes ($b = 0.048$ to $b = 0.072$). Full results of these models are in the Supplemental Materials.

Overall, then, relationships between party control and moral language use were most consistent and strongest for control of the speaker's own chamber of Congress: when speakers were in the minority in their chamber, they tended to use moral language more.

Figure 3. Mean loadings for each moral language category for Democratic and Republican Senators in Congresses 97-114. Shading illustrates majority control of the Senate.

Figure 4. Mean loadings for each moral language category for Democratic and Republican Representatives in Congresses 97-114. Shading illustrates majority control of the House.

General Discussion

We examined ideological differences in moral language use, and whether these differences vary over time as political power shifts. In Study 1, which examined all Twitter posts from members of Congress from January 2016-2018, we found that Democrats used more moral language than Republicans across all moral foundations, including binding foundations, and that these differences increased after the election of Donald Trump. In Study 2, which examined 36 years of Congressional transcripts, we found that Democrats used two kinds of moral language more overall (*harm-virtue* and *fairness-virtue*) but also that partisan differences in the use of most moral language fluctuated over time, because minority-party legislators in a chamber used more moral language of most kinds.

Theoretical Implications

Assessing the moral concerns of laypeople via self-report consistently finds that conservative respondents see the binding moral foundations (*loyalty*, *authority* and *purity*) as much more morally-relevant than do liberal respondents (Graham, Haidt, & Nosek, 2009). Based on this, one would expect that conservative elites would reference the binding foundations more than liberal elites do. Some investigations of elite moral language use have found this pattern, at least in part (e.g., Clifford & Jerit, 2013; Graham, Haidt, & Nosek, 2009, Study 4), but others have not (e.g., Frimer, 2020; Sagi & Dehghani, 2014; Sterling & Jost, 2017).

Together with this previous research, our findings suggest that self-report ratings of relevance cannot be straightforwardly extrapolated to the moral language used by elites. However, our results suggest some promising possibilities for theoretical refinement. We found that moral language use changes over time as parties' political fortunes rise and fall, so it is

likely too simplistic to treat observed ideological differences in morality as static. For a better understanding of ideological differences in moral language, theorists need to acknowledge that any differences may shift over time, and that differences at a specific time may not be generalizable to other time periods. Ideally, theorists should model the effects of specific influences as they change over time, as we did here for political power. For example, liberals and conservatives might use different moral language in response to domestic crises, conflicts with foreign adversaries, or power struggles between different branches of government. Of course, it might also be that liberal and conservative moral language responds to these events similarly (as we found to be the case for political power). In either case, however, our theoretical understanding will be enhanced by treating moral language use as dynamic, rather than static.

Open Questions

In Study 2, we observed the strongest and most consistent power effects for party control of the legislator's own chamber. Effects for control of the presidency were much smaller and less consistent. This contrasts with Study 1, where we observed significant increases in the moral language used by Democrats across all categories after the 2016 election. This suggests that there may be something unique about the Trump presidency—perhaps the fact that Trump, himself, is especially polarizing (Gallup, 2018), or the fact that Clinton was widely expected to win—that increases moral language use beyond the level one would expect in a more normal political environment. It is also possible that the medium of communication—i.e., Twitter posts vs. Congressional debates and proceedings—may be responsible for some of these differences. Both of these questions are promising ones for future research.

Although minority party members generally used more moral language across most categories, this was not true for all types of moral language in both chambers. Although it was uncommon for the *majority* party to use more moral language, one exception in both the House and Senate was the *authority-virtue* category. This category includes concepts such as respect for authority figures and obedience; its greater use by the majority party therefore seems reasonable. However, future research should use topic modeling or similar approaches to better understand what is being discussed when moral language is used, and how this changes over time.

Future Directions

For our Twitter dataset, we conducted an exploratory analysis (reported in the results of Study 1) using structural topic models to determine whether different topics of discussion were associated with particular kinds of moral language. We found that different topics were indeed associated with different kinds of moral language for Democrats and Republicans (see also Sterling & Jost, 2017). Moreover, Democratic and Republican lawmakers differed in how likely they were to tweet about specific topics, and topic prevalence changed over time (see Figures S2-S14). These preliminary results suggest that tracing the rise and fall in prevalence of particular topics over time, and the different moral language they evoke from liberals vs. conservatives, is a promising avenue for understanding ideological differences in morality.

Author Contribution Statement

The authors jointly designed and performed the research, analyzed the data, and wrote the paper.

References

- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software*, 67.
- Blei, D.M., Ng, A.Y., & Jordan, M.I. (2003). Latent Dirichlet Allocation. *Journal of Machine Learning Research*, 3, 993-1022.
- Brady, W. J., Wills, J. A., Burkart, D., Jost, J. T., & Van Bavel, J. J. (2019). An ideological asymmetry in the diffusion of moralized content on social media among political leaders. *Journal of Experimental Psychology: General*, 148, 1802–1813.
- Brady, W. J., Wills, J. A., Jost, J. T., Tucker, J. A., & Van Bavel, J. J. (2017). Emotion shapes the diffusion of moralized content in social networks. *Proceedings of the National Academy of Sciences*, 114, 7313–7318.
- Branscombe, N. R., & Wann, D. L. (1992). Physiological arousal and reactions to outgroup members during competitions that implicate an important social identity. *Aggressive Behavior*, 18, 85-93.
- Clifford, S., & Jerit, J. (2013). How Words Do the Work of Politics: Moral Foundations Theory and the Debate over Stem Cell Research. *The Journal of Politics*, 75, 659–671.
- Descioli, P., & Kurzban, R. (2012). A Solution to the Mysteries of Morality. *Psychological Bulletin*, 1–21.
- Fiske, A. P., & Rai, T. S. (2015). *Virtuous Violence*. Cambridge University Press.
- Garten, J., Hoover, J., Johnson, K. M., Boghrati, R., Iskiwitch, C., & Dehghani, M. (2018). Dictionaries and distributions: Combining expert knowledge and large scale textual data content analysis. *Behavior Research Methods*, 50, 344-361

- Gentzkow, M., Shapiro, J.M., Taddy, M. (2019). Measuring group differences in high-dimensional choices: Method and application to congressional speech. *Econometrica*, 87(4), 1307-1340.
- Graham, J., Haidt, J., & Nosek, B. A. (2009). Liberals and conservatives rely on different sets of moral foundations. *Journal of Personality and Social Psychology*, 96, 1029–1046.
- Haidt, J. (2007). The New Synthesis in Moral Psychology. *Science*, 316(5827), 998–1002.
- Haidt, J., & Graham, J. (2007). When Morality Opposes Justice: Conservatives Have Moral Intuitions that Liberals may not Recognize. *Social Justice Research*, 20, 98–116.
- Haidt, Jonathan. (2013). *The Righteous Mind: Why Good People are Divided by Politics and Religion*. Vintage Books.
- Hofmann, W., Wisneski, D. C., Brandt, M. J., & Skitka, L. J. (2014). Morality in everyday life. *Science*, 345, 1340–1343.
- Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). lmerTest Package: Tests in Linear Mixed Effects Models. *Journal of Statistical Software*, 82.
- Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato’s problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, 104, 211–240.
- Lorah, J. (2018). Effect size measures for multilevel models: definition, interpretation, and TIMSS example. *Large-scale Assessments in Education*, 6(8), 1-11.
- McCarty, N., Poole, K. T., & Rosenthal, H. (2016). *Polarized America: The Dance of Ideology and Unequal Riches*. MIT Press.

- Mikolov, T., Sutskever, I., Chen, K., Corrado, G., & Dean, J. (2013). Distributed Representations of Words and Phrases and their Compositionality. *Proc. NIPS*, 1–9.
- Osgood, C. E., Suci, G. J., & Tannenbaum, P. H. (1957). *The Measurement of Meaning*. University of Illinois Press.
- Poole, K. T., & Rosenthal, H. (1997). *Congress: A Political-economic History of Roll Call Voting*. Oxford University Press.
- R Core Team. (2016). *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Roberts, M., Stewart, B., & Tingley, D. (in press). stm: R Package for Structural Topic Models. *Journal of Statistical Software*.
- Roberts, M., Stewart, B., Tingley, D., Lucas, C., Leder-Luis, J., Gadarian, S.K., ..., Rand, D.G. (2014). Structural topic models for open-ended survey responses. *American Journal of Political Science*, 58(4), 1064-82.
- Rogers, T. T., & McClelland, J. L. (2004). *Semantic Cognition: A Parallel Distributed Processing Approach*. MIT Press.
- Ryan, T.J. (2014). Reconsidering Moral Issues in Politics. *The Journal of Politics*, 76(2), 380-397.
- Sagi, E., & Dehghani, M. (2014). Measuring Moral Rhetoric in Text. *Social Science Computer Review*, 32, 132–144.
- Salton, G., Wong, A., & Yang, C.-S. (1975). A vector space model for automatic indexing. *Communications of the ACM*, 18, 613–620.

- Skitka, L.J., Bauman, C.W. (2008). Moral Conviction and Political Engagement. *Political Psychology*, 29(1). 29-54.
- Skitka, L. J., & Mullen, E. (2002). Understanding Judgments of Fairness in a Real-World Political Context: A Test of the Value Protection Model of Justice Reasoning. *Personality and Social Psychology Bulletin*, 28, 1419–1429.
- Snijders, T. A. B., & Bosker, R. J. (2012). *Multilevel analysis: An introduction to basic and advanced multilevel modeling* (2nd ed.). London: Sage.
- Tajfel, H., & Turner, J. C. C. (1986). The social identity theory of intergroup behavior. In S. Worchel & W. G. Austin (Eds.), *Psychology of Intergroup Relations* (pp. 7-24). Chicago: Nelson-Hall.
- Tetlock, P. E. (1986). A value pluralism model of ideological reasoning. *Journal of Personality and Social Psychology*, 50, 819–827.
- Thomas, M., Pang, B., & Lee, L. (2006). Get out the vote: Determining support or opposition from Congressional floor-debate transcripts. *Proceedings of EMNLP*, 327-335.
- Verbyla, A. P., Cullis, B. R., Kenward, M. G., & Welham, S. J. (1999). The analysis of designed experiments and longitudinal data by using smoothing splines. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 48, 269-311.
- Wagner, A. K., Soumerai, S. B., Zhang, F., & Ross-Degnan, D. (2002). Segmented regression analysis of interrupted time series studies in medication use research. *Journal of Clinical Pharmacy and Therapeutics*, 27, 299–309.

Supplemental Methods

Inter-rater Reliability

Study 1

To counteract the low base rate of use of moral language, we selected the top 100 loading tweets from each moral category for a total of 1000 tweets. To represent non-moral tweets, an additional 100 tweets were randomly selected from the tweets that had moral loadings that fell within one standard deviation of 0.

Four coders each coded these tweets on 11 dimensions: the five moral foundations with a Virtue and Vice aspect for each, and a non-moral category. Coders were instructed that a given tweet could be coded as expressing more than one moral category. After an initial training session, coders rated an initial subset of 183 tweets. Due to low inter-rater reliability, these ratings were discarded and further training sessions were conducted using the guidelines in the Moral Values Coding Guide (Hoover, Johnson, Dehghani, & Graham, 2017), resulting in 917 tweets that were used in the final validation. Inter-rater agreement was measured using Prevalence and Bias Adjusted Kappa (Byrt, Bishop, & Carlin, 1993; Sim & Wright, 2005). Inter-rater agreement was good, with a mean of 0.84 across dimensions ($SD = 0.06$).

Study 2

As in Study 1, we selected the top 100 loading speeches from each moral category for a total of 1000 speeches and randomly selected an additional 100 non-moral speeches.

Three trained coders each coded these tweets on the same 11 dimensions. Ratings for an initial subset of 119 speeches were discarded due to low inter-rater reliability. After further training, a total of 981 speeches were used in the final validation. Inter-rater agreement was good, with a mean Prevalence and Bias Adjusted Kappa of 0.60 across dimensions ($SD = 0.19$).

Validation

To evaluate the agreement between human coders (who were treated as the gold standard) and DDR, we used F1 scores. F1 scores are calculated as the harmonic mean of precision (1 – false positive rate) and recall (1 - false negative rate; Powers, 2011). All evaluations were based on 10-fold cross validations using the cosine similarities between each Tweet and each moral foundation concept representation as features for use in classification. Because of the unbalanced nature of the data (i.e., many more tweets are coded as not being a member of a moral category than being a member of that category), positive cases were upsampled using the Synthetic Minority Over-sampling Technique (SMOTE; Chawla, Boyer, Hall, & Kegelmeyer, 2002). Using the human-rated data as the gold standard, the F1 score in Study 1 averaged across coders and dimensions was 0.865, with a mean recall of 0.773 and a mean precision of 0.867. In Study 2, the F1 score averaged across coders and dimensions was 0.791, with a mean recall of 0.772 and a mean precision of 0.789. This compares favorably to the F1 score of .496 reported by Garten et al. (2018) in their validation of DDR moral foundations dictionaries. It also far exceeds the F1 score of .275 Garten et al. reported for word counting using the Moral Foundations Dictionary, i.e., the method used by Graham, Haidt, & Nosek (2009, Study 4).

Interrupted time series analyses

We tested whether moral language use showed a discontinuity specifically at election day (November 8, 2016) with an interrupted time series analyses (Wagner, Soumerai, Zhang, & Ross-Degnan, 2002). In order to account for autocorrelation between time points, we fit these models using the *gl*s function in the *nlme* package in R. Examining residual ACF plots showed that they had an AR(1) structure, so we specified an order-1 autoregressive process for the residuals. Models were specified as:

$$Y_d = \beta_0 + \beta_1 * day_d + \beta_2 * election_d + \beta_3 * days\ post\ election_d + e \quad (1)$$

Where: Y_d was the mean similarity of tweets on day d to a particular foundation/aspect, day was a continuous variable indicating days since January 1, 2016, $election$ was a dummy variable encoding whether day d was before or after the election (before = 0, after = 1), $days\ post\ election$ was a continuous variable coded 0 before the election and counting the number of days since the election afterwards, and e was an autoregressive error term. In these models, Y_d is the mean amount of moral language for a particular foundation/aspect on a given day d , β_1 is the slope pre-election, β_3 is the change in slope post-election, and β_2 is the discontinuity or “jump” in moral language at election day. We first fit separate models for Democrats and Republicans for each foundation/aspect; for these models β_2 tests the hypothesis that moral language increased at election day. We then fit a new series of models on the full data set that included a dummy for party as well as its interaction with election day:

$$Y_d = \beta_0 + \beta_1 * day_d + \beta_2 * election_d + \beta_3 * party + \beta_4 * election_d * party + \beta_5 * days\ post\ election_d + e \quad (2)$$

Here, β_4 tests the hypothesis that the discontinuity differs for Democrats and Republicans.

Among Democrats there was a significant increase in moral language use at election day for every foundation and aspect except *care-virtue*. The pattern for Republicans was more mixed—for some foundations/aspects, moral language use increased, for others it did not change significantly or even decreased. These results are shown in Figure 2. Statistical tests of the interrupted time series models for each foundation and aspect are shown in Table S4 in the Supplemental Results.

Topic Modeling

To examine whether moral language and political ideology were associated with different topics of discussion, we ran a Structural Topic Model (STM; Roberts et al., 2014), a variation of the popular Latent Dirichlet Allocation (LDA; Blei, Ng, & Jordan, 2003) that allows for the inclusion of additional document-level metadata which may alter either the prevalence or content of topics. We fit our structural topic model using the ‘stm’ package in R (Roberts, Stewart, & Tingley, in press). Using the Lee and Mimno algorithm (2014) to fit the number of topics, our final model had 77 topics. We included political affiliation and a dummy-coded variable of whether the tweet was made before or after the 2016 Presidential election as topical prevalence covariates, along with the DDR loadings for the 10 moral foundations/aspects. These covariates and their interactions were significant for almost all topics, and an inspection of the top words and tweets for each topic revealed that the topics were interpretable and reflected meaningful facets of political discussion (see Table S8 in the Supplemental Results). This suggests that there are differences in the topics that are moralized, the topics that are discussed by Democrats and

Republicans, and the topics that were discussed before and after the election. These differences are plotted in Figures S4-S15 in the Supplemental Results. Figure S3 shows the semantic coherence (a measure of topic quality) of these topics.

Retweets and Likes

Finally, we examined how moral language affected mass audiences by testing whether the amount of moral language a tweet expressed predicted how much it was retweeted (i.e., re-posted by members of the public to their own followers) and “liked” (which does not re-post a tweet but indicates public approval). Both variables were highly right-skewed with excess zeros. For example, 9% of tweets received no retweets, 90% received less than 100, but the most retweeted received 455,142. We therefore used zero-inflated negative binomial models, which are intended for count variables and model excess zeros as well as over-dispersion (CITE ECONOMETRICS BOOK). We fit separate models for each foundation/aspect using the “zinb” command in Stata 16.0, using cluster-robust standard errors to account for non-independence of tweets posted by the same account. We modeled each tweet’s retweet count as a function of its similarity score for that foundation/aspect and the posting account’s follower count (as of January 20, 2018). Zero-inflation was modeled as a function of follower count plus a constant (i.e., an intercept term). (We included follower count in the model because we thought it would affect both how much tweets were retweeted and their probability of not being retweeted at all.) As before, we first fit separate models for Democrats and Republicans and then a full model including a party dummy and the interaction between this dummy and similarity score. As Figure 2 shows, greater moral language was associated with more retweets for every foundation and aspect except *care-virtue* for Republicans. The association was significantly stronger for Democrats than Republicans for *care-vice* and *fairness-vice*, but no other interactions with party were significant (see Table S5 in the Supplemental Results). Examining likes rather than retweets showed a very similar pattern, although in this case no party interactions were significant (see Table S6)

Supplemental Results

Table S1. Differences between Democrats and Republicans for each moral language category. Positive coefficients reflect greater moral language use by Democrats. Statistics are from mixed linear models with random intercepts for authors; degrees of freedom are by Satterthwaite approximation.

Category	<i>b</i> (SE)	df	<i>t</i>	<i>p</i>
Care-virtue	0.23 (0.023)	558.33	9.69	<.001
Care-vice	0.37 (0.027)	550.33	13.46	<.001
Fairness-virtue	0.34 (0.023)	560.32	14.36	<.001
Fairness-vice	0.39 (0.027)	554.10	14.31	<.001
Loyalty-virtue	0.19 (0.023)	551.21	8.18	<.001

Loyalty-vice	0.30 (0.028)	555.22	10.92	<.001
Authority- virtue	0.24 (0.026)	545.61	9.25	<.001
Authority-vice	0.38 (0.027)	554.19	13.74	<.001
Sanctity-virtue	0.29 (0.024)	558.32	12.46	<.001
Sanctity-vice	0.37 (0.026)	564.55	14.66	<.001

Table S2. Relationship between individual's DW-Nominate score (Dimension 1) and moral language use for each foundation and aspect. Higher DW-Nominate scores indicate a more conservative voting record, so negative coefficients indicate greater moral language use by more liberal individuals. Statistics are from mixed linear models with random intercepts for authors; degrees of freedom are by Satterthwaite approximation.

Category	<i>b</i> (SE)	df	<i>t</i>	<i>p</i>
Care-virtue	-0.25 (0.026)	558.3	-9.68	< .001
Care-vice	-0.38 (0.031)	548.70	-12.37	< .001
Fairness-virtue	-0.36 (0.026)	556.5	-13.66	< .001
Fairness-vice	-0.41 (0.030)	551.2	-13.48	< .001
Loyalty-virtue	-0.21 (0.026)	549.3	-8.03	< .001
Loyalty-vice	-0.31 (0.031)	553.3	-10.11	< .001
Authority-virtue	-0.25 (0.029)	544.3	-8.59	< .001
Authority-vice	-0.39 (0.031)	551.3	-12.84	< .001
Purity-virtue	-0.31 (0.026)	555.8	-11.95	< .001
Purity-vice	-0.41 (0.028)	560.3	-14.60	< .001

Table S3. Effects of time period (0 = pre-election, 1 = post-election) for each moral language category. Effects are shown separately for Democrats and Republicans, followed by the interaction between party and time period for each category. Statistics are from mixed linear models with random intercepts for authors; degrees of freedom are by Satterthwaite approximation. Effect sizes are computed as recommended in (31).

	<u>Democrats</u>				<u>Republicans</u>				<u>Party x Time Period Interaction</u>			
	<i>b</i> (SE)	df	<i>t</i>	<i>p</i>	<i>b</i> (SE)	df	<i>t</i>	<i>p</i>	<i>b</i> (SE)	df	<i>t</i>	<i>p</i>
Care-virtue	0.17 (0.023)	870.0	7.093	<.001	0.17 (0.027)	932.2	6.33	<.001	0.0047 (0.015)	1663	.31	0.76
Care-vice	0.45 (0.024)	780.5	18.55	<.001	0.17 (0.022)	885.6	7.74	<.001	0.23 (0.016)	2330	14.54	<.001
Fairness-virtue	0.29 (0.027)	867.7	11.15	<.001	0.22 (0.026)	927.2	8.43	<.001	0.064 (0.016)	2099	4.11	<.001
Fairness-vice	0.45 (0.025)	849.7	17.56	<.001	0.15 (0.022)	686.2	6.86	<.001	0.27 (0.016)	2303	17.12	<.001
Loyalty-virtue	0.27 (0.024)	893.7	9.28	<.001	0.18 (0.028)	913.5	6.36	<.001	0.032 (0.015)	1561	2.13	.034
Loyalty-vice	0.43 (0.023)	861.3	18.56	<.001	0.078 (0.023)	941.1	3.37	<.001	0.30 (0.015)	1900	19.69	<.001
Authority-virtue	0.32 (0.023)	819.1	14.01	<.001	0.21 (0.025)	938.5	8.21	<.001	0.076 (0.016)	2123	4.81	<.001
Authority-vice	0.39 (0.024)	851.8	16.08	<.001	0.093 (0.024)	938.4	3.95	<.001	0.23 (0.015)	1954	15.08	<.001
Purity-virtue	0.28 (0.023)	858	11.94	<.001	0.14 (0.028)	936.2	5.16	<.001	0.10 (0.015)	1742	6.80	<.001
Purity-vice	0.27 (0.024)	859.5	11.29	<.001	0.067 (0.023)	930.9	2.90	.004	0.18 (0.015)	1929	11.70	<.001

Table S4. Tests of the election-day dummy variable from interrupted time series analyses for each moral language category. Positive coefficients indicate an increase in moral language at election day. Effects are shown separately for Democrats and Republicans, followed by the interaction between party and election day for each category.

	<u>Democrats</u>				<u>Republicans</u>				<u>Party x Election Day Interaction</u>			
	<i>b</i> (SE)	<i>t</i>	df	<i>p</i>	<i>b</i> (SE)	<i>t</i>	df	<i>p</i>	<i>b</i> (SE)	df	<i>t</i>	<i>p</i>
Care-virtue	0.0002 (.003)	.88	758	.38	-.003 (.003)	-.82	758	.41	.001 (.001)	1518	.99	.32
Care-vice	0.014 (.003)	4.12	758	<.001	-.006 (.003)	-1.80	758	.07	.020 (.001)	1518	17.80	<.001
Fairness-virtue	0.015 (.004)	3.86	758	<.001	.008 (.004)	1.79	758	.07	.009 (.001)	1518	7.35	<.001
Fairness-vice	0.013 (.004)	3.53	758	<.001	-0.006 (.003)	-1.80	758	.08	.019 (.001)	1518	19.56	<.001
Loyalty-virtue	0.010 (.003)	3.09	758	.002	.008 (.004)	2.37	758	.02	.004 (.001)	1518	3.77	<.001
Loyalty-vice	0.023 (.003)	6.91	758	<.001	.010 (.003)	3.14	758	.002	.021 (.001)	1518	20.11	<.001
Authority-virtue	0.012 (.003)	4.41	758	<.001	.007 (.003)	2.12	758	.03	.009 (.001)	1518	9.67	<.001
Authority-vice	0.015 (.003)	4.68	758	<.001	.005 (.003)	1.36	758	.18	.019 (.001)	1518	19.24	<.001
Purity-virtue	0.010 (.002)	3.79	758	<.001	.005 (.003)	1.47	758	.14	.010 (.001)	1518	9.76	<.001
Purity-vice	.006 (.003)	2.33	758	.02	-.003 (.003)	-.95	758	.34	.013 (.001)	1518	14.25	<.001

Table S5. Association between similarity to each moral language category and number of retweets for each tweet. Effects are shown separately for Democrats and Republicans, followed by the interaction between party and similarity for each category. Statistics are from zero-inflated negative binomial models with cluster-robust standard errors (fit using *zinb* in Stata 16.0) that include number of followers as a covariate.

	<u>Democrats</u>			<u>Republicans</u>			<u>Party x Similarity Interaction</u>		
	<i>b</i> (SE)	<i>Z</i>	<i>p</i>	<i>b</i> (SE)	<i>Z</i>	<i>p</i>	<i>b</i> (SE)	<i>Z</i>	<i>p</i>
Care-virtue	1.60 (.50)	3.18	.001	.69 (1.18)	.58	.56	1.00 (1.30)	.77	.44
Care-vice	7.00 (.69)	10.15	<.001	3.67 (1.55)	2.36	.02	3.45 (1.67)	2.07	.04
Fairness-virtue	5.18 (.58)	8.89	<.001	4.85 (1.40)	3.47	.001	.45 (1.51)	.30	.77
Fairness-vice	9.25 (.63)	14.78	<.001	6.32 (1.45)	4.36	<.001	3.05 (1.56)	1.95	.05
Loyalty-virtue	4.24 (.48)	8.84	<.001	4.56 (.87)	5.24	<.001	-.26 (1.02)	-.26	.80
Loyalty-vice	9.84 (.63)	15.47	<.001	9.50 (.68)	13.95	<.001	.47 (.89)	.53	.60
Authority- virtue	5.73 (.65)	8.83	<.001	3.72 (1.51)	2.47	.01	2.14 (1.63)	1.30	.19
Authority-vice	10.07 (.72)	13.91	<.001	8.73 (.93)	9.38	<.001	1.47 (1.14)	1.29	.20
Purity-virtue	6.65 (.64)	10.35	<.001	6.12 (.96)	6.34	<.001	.64 (1.15)	.56	.58
Purity-vice	7.85 (.65)	12.09	<.001	9.09	10.23	<.001	-1.15 (1.08)	-1.06	.29

Table S6. Association between similarity to each moral category and number of “likes” for each tweet. Effects are shown separately for Democrats and Republicans, followed by the interaction between party and similarity for each category. Statistics are from zero-inflated negative binomial models with cluster-robust standard errors (fit using *zinb* in Stata 16.0) that include number of followers as a covariate.

	<u>Democrats</u>			<u>Republicans</u>			<u>Party x Similarity Interaction</u>		
	<i>b</i> (SE)	<i>Z</i>	<i>p</i>	<i>b</i> (SE)	<i>Z</i>	<i>p</i>	<i>b</i> (SE)	<i>Z</i>	<i>p</i>
Care-virtue	2.59 (.48)	5.41	<.001	2.54 (.62)	4.12	.001	.09 (.78)	.11	.91
Care-vice	6.54 (.66)	9.85	<.001	4.98 (.81)	6.13	<.001	1.63 (.99)	1.64	.10
Fairness-virtue	5.01 (.55)	9.04	<.001	5.62 (.69)	8.11	<.001	-.56 (.87)	-.65	.52
Fairness-vice	8.37 (.59)	14.14	<.001	6.76 (.79)	8.58	<.001	1.67 (.94)	1.77	.08
Loyalty-virtue	5.42 (.43)	12.49	<.001	5.57 (.53)	10.58	<.001	-.13 (.70)	-.19	.85
Loyalty-vice	9.87 (.59)	16.75	<.001	9.09 (.61)	14.81	<.001	.88 (.81)	1.09	.28
Authority- virtue	6.01 (.63)	9.47	<.001	5.17 (.68)	7.62	<.001	.90 (.92)	.97	.33
Authority-vice	9.69 (.69)	13.98	<.001	8.35 (.76)	11.02	<.001	1.43 (.98)	1.45	.15
Purity-virtue	7.11 (.60)	11.81	<.001	6.56 (.72)	9.12	<.001	.62 (.92)	.67	.50
Purity-vice	7.52 (.61)	12.42	<.001	7.26 (.90)	8.09	<.001	.34 (1.06)	.32	.75

Table S7. Tweets scored by DDR as most similar to each category for Democrats (top) and Republicans (bottom). Raw (i.e., uncleaned) tweets are displayed; punctuation was removed for analysis.

<u>Democrats</u>		
Care	<i>Virtue</i>	May our empathy be courageous, our kindness contagious, and our love of others unyielding and tenacious.
	<i>Vice</i>	#HR732 makes it harder for victims of environmental harm to get restitution - while rewarding those who caused the harm. #VoteNO
Fairness	<i>Virtue</i>	On #HumanRightsDay, we celebrate our common humanity, and the universal values of justice, fairness, and equality.
	<i>Vice</i>	I stand in opposition to this unfair, immoral #GOPBudget.
Loyalty	<i>Virtue</i>	Our brave transgender servicemembers deserve honor and praise, not this President's fake news. A patriot is a patriot.
	<i>Vice</i>	Aiders and abettors to obstruction of Justice and treason! Shame!
Authority	<i>Virtue</i>	Welcome news as our independent judiciary steps in to uphold the rule of law & compel Trump to obey Constitution.
	<i>Vice</i>	Consistent appeasement of hate, intolerance and lawlessness, #disgraceful #unfit #unAmerican
Purity	<i>Virtue</i>	"Never forget" is a moral imperative, as sacred and as holy as a commandment – an 11th commandment for the Jewish people. #YomHashoah
	<i>Vice</i>	Cruelty, corruption, betrayal, hypocrisy, snobbery, violence. #GameofThrones or #GameofTrumps?
<u>Republicans</u>		
Care	<i>Virtue</i>	@RAWTheTerror01 thank you for that expression of kindness and rationality
	<i>Vice</i>	Communism in articles like @Yuxin_Gao's hurts those who suffered + is harmful to those still suffering from brutal Chinese/Cuban regimes 3/3
Fairness	<i>Virtue</i>	The law not only binds us together, it embodies the virtues we cherish: justice, mercy, equality, and fairness.
	<i>Vice</i>	Today, the Senate dealt a huge blow to criminals who use spoofing scams to deceive, intimidate, & defraud victims
Loyalty	<i>Virtue</i>	Thanking @StandwthRand for his Courage and willingness to run. He is a leader of liberty causes and a patriot whom I proudly call friend.
	<i>Vice</i>	#Castro stooge @citaconsilvio #SilvioRodriguez lying abt #AnaBelenMontes treason. Shameless campaign 2 free her.
Authority	<i>Virtue</i>	Judge Gorsuch's respect for the Constitution, commitment to the rule of law, and devotion to branch integrity will serve our nation well.
	<i>Vice</i>	@wearepepito @briansikma if Islamic radicalization not reaction 2 Assad's repression so stupid, refute it instead of a personal attack.
Purity	<i>Virtue</i>	Great meeting w/pastors on our shared commitment to protecting unborn life, religious liberty & traditional values.
	<i>Vice</i>	No such dignity exists in the atrocities ISIS is committing around the world with such subhuman savagery.

Table S8. Examples of topics extracted from structural topic model.

Topic	Top Words
Tax Reform	job, reform, creat, worker, taxreform, econom, fair
Mueller Investigation	trump, russia, investig, elect, russian, administr, independ
Trump	potus, realdonaldtrump, border, interest, promis, wall, order
Mass Shootings	live, victim, attack, lost, thought, famili, prayer
SCOTUS (Supreme Court)	justic, scotus, court, judg, nomin, nomine, senat
Military	militari, secur, forc, threat, defens, war, resolut
Legislation	bill, hous, act, pass, legisl, bipartisan, introduc
Climate Change	chang, energi, water, climat, epa, clean, economi
Education	student, learn, educ, school, opportun, colleg, young
Gun Violence	congress, take, action, gun, violenc, enough, demand
Voting	make, want, sure, made, voic, place, everyon
Opioid Crisis	must, opioid, crisi, continu, effort, address, fight
Healthcare	million, trumpcar, will, cost, coverag, medicaid, insur
Tech	leader, futur, innov, develop, annual, tech, san
Military Service	honor, serv, servic, today, hero, men, sacrific
GOP Tax Plan	tax, cut, famili, pay, class, goptaxscam, middl
Law	law, rule, enforc, feder, sign, must, decis
Immigration/Refugees	end, must, immigr, dreamer, daca, solut, thousand
Business	busi, local, small, grow, economi, feder, compani
Veterans	veteran, provid, servic, vet, benefit, joe, ensur

Table S9. Differences between Democrats and Republicans for each moral language category. Positive coefficients reflect greater moral language use by Democrats. Statistics are from mixed linear models with random intercepts for speakers and congresses; degrees of freedom are by Satterthwaite approximation.

Category	<i>b</i> (SE)	df	<i>t</i>	<i>p</i>
Care-virtue	0.15 (0.037)	1716	3.96	<.001***
Care-vice	0.069 (0.036)	1666	1.89	0.059
Fairness-virtue	0.17 (0.037)	1728.41	4.72	<.001***
Fairness-vice	0.044 (0.036)	1676	1.22	0.22
Loyalty-virtue	0.087 (0.039)	1986	2.23	0.026
Loyalty-vice	-0.094 (0.037)	1848.73	-2.51	0.012
Authority-virtue	0.032 (0.035)	1594	0.90	0.37
Authority-vice	-0.011 (0.037)	1670.64	-0.30	0.76
Sanctity-virtue	0.059 (0.037)	1822	1.57	0.12
Sanctity-vice	0.067 (0.036)	1883	1.87	0.062

Table S10. Effects of political control in the Senate for each moral language category. Positive coefficients for Party reflect greater moral language use by Democrats. Negative coefficients for Match reflect greater moral language use when there is a mismatch between a speaker's party and the party in control of the Senate. Statistics are from mixed linear models with random intercepts for speakers and congresses; degrees of freedom are by Satterthwaite approximation.

Category		<i>b</i> (SE)	df	<i>t</i>	<i>p</i>
Care-virtue	Party	0.15 (0.09)	311.44	1.67	0.09
	Match	-0.19 (0.03)	1515.39	-6.05	<.001***
Care-vice	Party	0.06 (0.07)	217.98	0.83	0.41
	Match	-0.29 (0.03)	1495.95	-9.17	<.001***
Fairness-virtue	Party	0.11 (0.0096)	340.4	1.20	9.23
	Match	0.0050 (0.0029)	1477	1.75	0.08
Fairness-vice	Party	-0.026 (0.088)	288.60	-0.29	0.77
	Match	-0.18 (0.029)	1466.90	-6.29	<.001***
Loyalty-virtue	Party	0.10 (0.096)	316.23	1.07	0.28
	Match	-0.21 (0.033)	1507.94	-6.33	<.001***
Loyalty-vice	Party	-0.08 (0.09)	259.71	-0.89	0.37
	Match	-0.16 (0.035)	1505.86	-4.65	<.001***
Authority-virtue	Party	-0.0028 (0.11)	243.9	-0.026	0.98
	Match	0.15 (0.041)	147.6	3.74	<.001***
Authority-vice	Party	-0.084 (0.099)	258.22	-0.85	0.40
	Match	-0.088 (0.035)	1455.29	2.51	0.01
Sanctity-virtue	Party	0.12 (0.078)	285	1.56	0.12
	Match	-0.188 (0.035)	1603.31	-5.31	<.001***
Sanctity-vice	Party	0.076 (0.076)	320.86	1.01	0.31
	Match	-0.20 (0.036)	1621.28	-5.61	<.001***

Table S11. Effects of political control in the House for each moral language category. Positive coefficients for Party reflect greater moral language use by Democrats. Negative coefficients for Match reflect greater moral language use when there is a mismatch between a speaker's party and the party in control of the Senate. Statistics are from mixed linear models with random intercepts for speakers and congresses; degrees of freedom are by Satterthwaite approximation.

Category		<i>b</i> (SE)	df	<i>t</i>	<i>p</i>
Care-virtue	Party	0.13 (0.039)	1446.80	3.31	<.001 ***
	Match	-0.039 (0.019)	7747.43	-1.98	0.048*
Care-vice	Party	0.057 (0.040)	1460.24	1.43	0.15
	Match	-0.17 (0.021)	7858.61	-8.07	<.001***
Fairness-virtue	Party	0.18 (0.040)	1529.05	4.50	<.001***
	Match	0.099 (0.020)	7796	4.81	<.001***
Fairness-vice	Party	-0.046 (0.039)	1473.70	1.16	0.24
	Match	-0.15 (0.021)	7852.12	-7.37	<.001***
Loyalty-virtue	Party	0.08 (0.035)	1577.08	2.41	0.021
	Match	-0.13 (0.018)	7756.05	-7.60	<.001***
Loyalty-vice	Party	-0.099 (0.039)	1608.20	-2.60	0.01
	Match	-0.21 (0.019)	7762.55	-10.84	<.001***
Authority-virtue	Party	0.020 (0.036)	1459.52	0.56	0.57
	Match	0.17 (0.020)	7935.55	8.55	<.001***
Authority-vice	Party	-0.0098(0.038)	1497	-0.25	0.80
	Match	-0.071(0.02)	7833	-3.51	<.001***
Sanctity-virtue	Party	0.032 (0.038)	1482.32	0.85	0.40
	Match	-0.040 (0.020)	7813.46	-0.20	0.04*
Sanctity-vice	Party	0.058 (0.037)	1574.96	1.56	0.12
	Match	-0.034 (0.020)	7837.89	-1.73	0.08

Table S12. Effects of political control across chambers in the Senate for each moral language category. Positive coefficients for Party reflect greater moral language use by Democrats. Negative coefficients for Match reflect greater moral language use when there is a mismatch between a speaker's party and the party in control of the Senate (own chamber) and House (cross chamber). Statistics are from mixed linear models with random intercepts for speakers and congresses; degrees of freedom are by Satterthwaite approximation.

Category		<i>b</i> (SE)	df	<i>t</i>	<i>p</i>
Care-virtue	Cross Chamber	-0.069 (0.037)	1677.80	-1.85	0.06
	Own Chamber	-0.16 (0.034)	1558.60	-4.85	<.001***
	Party	0.16 (0.091)	311.57	1.72	0.086
Care-vice	Cross Chamber	-0.088 (0.037)	1758.66	-2.38	0.018
	Own Chamber	-0.26 (0.034)	1564.26	-7.58	<.001***
	Party	0.067 (0.07)	218.07	0.90	0.37
Fairness-virtue	Cross Chamber	0.066 (0.034)	1609	1.91	0.056
	Own Chamber	0.027 (0.031)	1511	0.86	0.39
	Party	0.11 (0.096)	3432	1.16	0.25
Fairness-vice	Cross Chamber	-0.058 (0.035)	1634.75	-1.68	0.09
	Own Chamber	-0.16 (0.031)	1510.47	-5.13	<.001***
	Party	-0.022 (0.088)	288.09	-0.26	0.80
Loyalty-virtue	Cross Chamber	-0.024 (0.039)	1668.65	-0.64	0.52
	Own Chamber	-0.20 (0.035)	1550.43	-5.58	<.001***
	Party	0.10 (0.10)	316.60	1.09	0.28
Loyalty-vice	Cross Chamber	0.038 (0.042)	1717.94	0.91	0.36
	Own Chamber	-0.18 (0.038)	1562.25	-4.64	<.001***
	Party	-0.083)	260.46	-0.91	0.36
Authority-virtue	Cross Chamber	0.14 (0.049)	1688.37	2.95	0.0033**
	Own Chamber	0.10 (0.045)	1532.46	2.32	0.021
	Party	-0.012 (0.11)	248.47	-0.11	0.91
Authority-vice	Cross Chamber	0.060 (0.042)	1650.03	1.45	0.15
	Own Chamber	-0.11 (0.038)	1506.57	-2.88	0.0040**
	Party	-0.088 (0.10)	260.27	-0.88	0.38
Sanctity-virtue	Cross Chamber	-0.047 (0.041)	1794.31	-1.15	0.25
	Own Chamber	-0.17 (0.038)	1657.58	-4.49	<.001***
	Party	0.12 (0.078)	285.19	1.60	0.11
Sanctity-vice	Cross Chamber	-0.072 (0.041)	1799	-1.75	0.080
	Own Chamber	-0.17 (0.038)	1672.51	-4.54	<.001***
	Party	0.08 (0.07)	321.32	1.07	0.28

Table S13. Effects of political control across chambers in the House for each moral language category. Positive coefficients for Party reflect greater moral language use by Democrats. Negative coefficients for Match reflect greater moral language use when there is a mismatch between a speaker's party and the party in control of the House (own chamber) and Senate (cross chamber). Statistics are from mixed linear models with random intercepts for speakers and congresses; degrees of freedom are by Satterthwaite approximation.

Category		<i>b</i> (SE)	df	<i>t</i>	<i>p</i>
Care-virtue	Cross Chamber	-0.020 (0.019)	7114.25	-1.06	0.29
	Own Chamber	-0.031 (0.021)	7848.17	-1.46	0.14
	Party	0.12 (0.039)	1477.89	3.17	0.0015 **
Care-vice	Cross Chamber	-0.078 (0.021)	7239.54	-3.80	<.001 ***
	Own Chamber	-0.14 (0.023)	7929.83	-6.21	<.001 ***
	Party	0.039 (0.040)	1498.63	0.98	0.33
Fairness-virtue	Cross Chamber	-0.026 (0.020)	7202.75	-1.33	0.18
	Own Chamber	0.11 (0.022)	7882.80	4.98	<.001 ***
	Party	0.17 (0.040)	1562.99	4.32	<.001 ***
Fairness-vice	Cross Chamber	-0.083 (0.020)	7238.25	-4.12	<.001 ***
	Own Chamber	-0.12 (0.022)	7924.88	-5.34	<.001 ***
	Party	0.027 (0.040)	1511.99	0.68	0.50
Loyalty-virtue	Cross Chamber	-0.015 (0.017)	7185.84	-0.85	0.39
	Own Chamber	-0.13 (0.019)	7850.05	-6.79	<.001 ***
	Party	0.078 (0.035)	1610.38	2.20	0.028
Loyalty-vice	Cross Chamber	-0.055 (0.019)	7205.45	-2.97	0.003 **
	Own Chamber	-0.19 (0.021)	7855.28	-9.07	<.001 ***
	Party	-0.11 (0.039)	1642.89	-2.87	0.004 **
Authority-virtue	Cross Chamber	-0.074 (0.020)	7334	-3.72	<.001 ***
	Own Chamber	0.20 (0.022)	7971	9.32	<.001 ***
	Party	0.0033 (0.036)	1502	0.091	0.93
Authority-vice	Cross Chamber	-0.074 (0.019)	7227.87	-3.78	<.001 ***
	Own Chamber	-0.042 (0.022)	7911.18	-1.95	0.051
	Party	-0.026 (0.039)	1533.72	-0.69	0.49
Sanctity-virtue	Cross Chamber	-0.043 (0.019)	7201.36	-2.28	0.023
	Own Chamber	-0.023 (0.021)	7897.58	-1.09	0.27
	Party	0.022 (0.038)	1517.73	0.58	0.56
Sanctity-vice	Cross Chamber	-0.042 (0.019)	7255.29	2.20	0.026
	Own Chamber	-0.017 (0.021)	7910.16	-2.23	0.41
	Party	0.048 (0.038)	1612.857	1.29	0.20

Table S14. Effects of political control of the White House for each moral language category. Negative coefficients for Match reflect greater moral language use when there is a mismatch between a speaker's party and the party in control of the White House. Statistics are from mixed linear models with random intercepts for speakers and congresses; degrees of freedom are by Satterthwaite approximation.

Category	<i>b</i> (SE)	df	<i>t</i>	<i>p</i>
Care-virtue	0.0058 (0.016)	9111	0.35	0.72
Care-vice	-0.040 (0.017)	9325.80	-2.27	0.023
Fairness-virtue	-0.048 (0.017)	9219.43	-2.85	0.0043 **
Fairness-vice	-0.061 (0.017)	9296.62	-3.55	0.0038 **
Loyalty-virtue	0.056 (0.016)	8978	3.63	<.001 ***
Loyalty-vice	0.033 (0.017)	9206	1.97	0.05
Authority-virtue	-0.072 (0.018)	9440.81	-4.01	<.001***
Authority-vice	-0.032 (0.017)	9223.41	-1.86	0.063
Sanctity-virtue	0.0079 (0.017)	9202	0.47	0.64
Sanctity-vice	0.0058 (0.017)	9295	0.34	0.73

Table S15. Effects of single-party governance (control of the Senate, House, and presidency) for each moral language category. Statistics are from mixed linear models with random intercepts for speakers and congresses; degrees of freedom are by Satterthwaite approximation.

Category	<i>b</i> (SE)	df	<i>t</i>	<i>p</i>
Care-virtue	0.12 (0.074)	15.43	1.59	0.13
Care-vice	-0.053 (0.085)	15.97	-0.62	0.54
Fairness-virtue	0.011 (0.083)	15.99	0.13	0.90
Fairness-vice	-0.059 (0.097)	16	-0.60	0.55
Loyalty-virtue	0.051 (0.081)	16	0.63	0.54
Loyalty-vice	0.075 (0.099)	16.02	0.75	0.46
Authority-virtue	-0.029 (0.096)	16.12	-0.30	0.77
Authority-vice	-0.033 (0.10)	15.91	-0.32	0.75
Sanctity-virtue	0.0031 (0.065)	15.89	0.048	0.96
Sanctity-vice	-0.22 (0.20)	15.79	-1.10	0.28

Table S16. Correlations between dictionary similarity scores for each foundation and aspect, Study 1.

[illegible]

Figure S1. Relationship between individual's DW-Nominate score (Dimension 1) and moral language use for each foundation and aspect. Each point represents the average of all tweets from one individual. Lines are simple linear fits; shaded areas are 95% confidence intervals.

Figure S2. Semantic coherence and exclusivity measures for topic model. Each point represents a topic within the model.

Figure S3. Mean difference in topic proportions by political affiliation. Negative values indicate a topic discussed more by conservatives. Positive values indicate a topic discussed more by liberals.

Figure S4. Mean difference in topic proportions before and after the 2016 Presidential election. Negative values indicate a topic discussed more prior to election; positive values indicate a topic discussed more after the election.

Figure S5. Mean difference in topic proportions at one standard deviation above and one standard deviation below the mean loading for the moral aspect of Care-virtue. Higher values indicate a topic discussed more in tweets that loaded highly on Care-virtue.

Figure S6. Mean difference in topic proportions at one standard deviation above and one standard deviation below the mean loading for the moral aspect of Care-vice. Higher values indicate a topic discussed more in tweets that loaded highly on Care-vice.

Figure S7. Mean difference in topic proportions at one standard deviation above and one standard deviation below the mean loading for the moral aspect of Fairness-virtue. Higher values indicate a topic discussed more in tweets that loaded highly on Fairness-virtue.

Figure S8. Mean difference in topic proportions at one standard deviation above and one standard deviation below the mean loading for the moral aspect of Fairness-vice. Higher values indicate a topic discussed more in tweets that loaded highly on Fairness-vice.

Figure S9. Mean difference in topic proportions at one standard deviation above and one standard deviation below the mean loading for the moral aspect of Loyalty-virtue. Higher values indicate a topic discussed more in tweets that loaded highly on Loyalty-virtue.

Figure S10. Mean difference in topic proportions at one standard deviation above and one standard deviation below the mean loading for the moral aspect of Loyalty-vice. Higher values indicate a topic discussed more in tweets that loaded highly on Loyalty-vice.

Figure S11. Mean difference in topic proportions at one standard deviation above and one standard deviation below the mean loading for the moral aspect of Authority-virtue. Higher values indicate a topic discussed more in tweets that loaded highly on Authority-virtue.

Figure S12. Mean difference in topic proportions at one standard deviation above and one standard deviation below the mean loading for the moral aspect of Authority-vice. Higher values indicate a topic discussed more in tweets that loaded highly on Authority-vice.

Figure S13. Mean difference in topic proportions at one standard deviation above and one standard deviation below the mean loading for the moral aspect of Purity-virtue. Higher values indicate a topic discussed more in tweets that loaded highly on Purity-virtue.

Figure S14. Mean difference in topic proportions at one standard deviation above and one standard deviation below the mean loading for the moral aspect of Purity-vice. Higher values indicate a topic discussed more in tweets that loaded highly on Purity-vice.

Figure S15. Predicted number of retweets by similarity score for each moral language category separately for Democrats (blue) and Republicans (red). Predicted values were generated by the Stata 16.0 *margins* command based on zero-inflated negative binomial models at the observed median number of followers across all accounts (30,547). Shaded regions are 95% confidence intervals.

References

- Byrt, T., Bishop, J., Carlin, J.B. (1993). Bias, prevalence and kappa. *Journal of Clinical Epidemiology*, 46(5), 423-429.
- Chalwa, N.V., Bowyer, K.W., Hall, L.O., Kegelmeyer, W.P. (2002). SMOTE: Synthetic Minority Over-sampling Technique. *Journal of Artificial Intelligence Research*, 16, 321-57.
- Fleiss, J. L. (1971). Measuring nominal scale agreement among many raters. *Psychological Bulletin*, 76(5), 378–82.
- Garten, J., Hoover, J., Johnson, K. M., Boghrati, R., Iskiwitch, C., & Dehghani, M. (2018). Dictionaries and distributions: Combining expert knowledge and large scale textual data content analysis. *Behavior Research Methods*, 344-361.
- Graham, J., Haidt, J., & Nosek, B. A. (2009). Liberals and conservatives rely on different sets of moral foundations. *Journal of Personality and Social Psychology*, 96, 1029-1046.
- Hoover, J., Johnson-Grey, K. M., Dehghani, M., & Graham, J. (2017). Moral Values Coding Guide. <https://doi.org/10.31234/osf.io/5dmgi>
- Powers, D.M.W. (2011). Evaluation: From precision, recall, and F-measure to ROC, informedness, markedness, & correlation. *Journal of Machine Learning Technologies*, 2(1), 37-63.
- Sim, J., Wright, C.C. (2005). The kappa statistics in reliability studies: use, interpretation, and sample size requirements. *Physical Therapy*, 85(3), 257-268.