When protests turn violent: The roles of moralization and moral convergence

Marlon Mooijman*

Department of Management and Organizations, Kellogg School of Management,
Northwestern University

Joe Hoover*

Brain and Creativity Institute, Psychology Department, University of Southern
California

Ying Lin & Heng Ji

Department of Computer Science, Rensselaer Polytechnic Institute

Morteza Dehghani

Brain and Creativity Institute, Psychology Department, Department of Computer
Science, University of Southern California

Author Note

*Both authors contributed equally to this work. This research was sponsored by the Army Research Lab. The content of this publication does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred. Correspondence regarding this article should be addressed to Marlon Mooijman (marlon.mooijman@kellogg.northwestern.edu), Joe Hoover (jehoover@usc.edu), or Morteza Dehghani (mdehghan@usc.edu).

Abstract

We propose that the risk of violence at protests can be estimated as a function of individual moralization and perceived moral convergence. Using data from the 2015 Baltimore protests, we find that not only did the rate of moral rhetoric on social media increase on days with violent protests, but also that the hourly frequency of morally relevant tweets predicted the future rates of arrest during protests, suggesting an association between moralization and protest violence. To understand the structure of this association, we ran a series of controlled behavioral experiments demonstrating that people are more likely to endorse violent protest for a given issue when they moralize the issue; however, this effect is moderated by the degree people believe others share their values. We discuss how online social networks may contribute to inflations of protest violence.

Keywords: Moral Psychology, Protests, Police Arrests, Natural Language Processing, Social Media

When protests turn violent: The roles of moralization and moral convergence

Protest is widely seen as an important component of democratic societies. It enables constituents to express grievances, communicate directly with the public and representatives, and promote change in accordance with their beliefs. While protests associated with popular platforms often attract large numbers of attendees, they are frequently peaceful events, even when they target controversial issues. Influential theories on social movements suggest that people engage in peaceful protests for a plethora of reasons, including rational deliberations, identification with a political cause, and feelings of relative deprivation (Olson, 1965; Van Stekelenburg & Klandermans, 2013; Walker & Pettigrew, 1984).

However, protests can also quickly erupt into violence. For instance, in the United States, protesters clashed violently with police in Ferguson, Missouri after the killing of Michael Brown (Davey & Bosman, 2014). More recently, far-right protesters clashed violently with counter-protesters in Charlottesville, Virginia in response to the proposed removal of a confederate statue (Heim, 2017). Recent polls suggest that the acceptability of using violence at protests is rising among the US public and college students in particular (Rampell, 2017). How do we understand, and predict, the rising acceptability of violent protests and the increasing emergence of actual violent behavior at protests?

While many factors likely contribute to a protest's risk for violence, we propose that it can be in part understood as a function of two key risk factors: (a) the degree to which people see protest as a moral issue and (b) the rate of perceived moral convergence – the degree to which participants believe that others share these moral attitudes. Further, we suggest that reliable measurements of these risk factors (moralization and moral convergence) can be obtained from online social network activity; we find not only that violent protest is preceded by an increase in online moral rhetoric, but also that hourly signals of online moral rhetoric predict future hourly arrest rates during violent protests.

Our hypotheses are grounded in the observation that protests are typically preceded by extensive discussions on Twitter and Facebook about moral topics such as societal unfairness and injustice (Manjoo, 2017). Social media platforms, in other words, have become important tools for people to express their moral disapproval with social and political developments such as government corruption, the killing of unarmed citizens by police, and the removal of culturally meaningful symbols and statues. Due to the scale of social networks, messages that contain references to moral terms such as injustice and unfairness are likely to spread to thousands if not millions of others and reflect the moral sentiments of a given population (Brady, Wills, Jost, Tucker, & Van Bavel, 2017). Thus, social media discourse materially encodes signals of moralization and moral outrage (Crockett, 2017).

Importantly, such signals are often not what one might refer to as mere rhetoric, as moral sentiments can provide the foundation for violence (Fiske & Rai, 2014). For instance, the perceived moral obligation to oppose a political candidate, discipline one's child, or prevent a criminal from reoffending often justifies the use of force and violence (Darley, 2009; Zaal, Laar, Ståhl, Ellemers, & Derks, 2011). Suicide bombers kill themselves and others in the name of a divine authority that commands obligation (Atran & Ginges, 2012; Fiske & Rai, 2014). When attitudes towards a social or political issue are seen as a reflection of moral beliefs, people are more willing to use violence to protect their moral beliefs and achieve desired ends (Skitka & Morgan, 2014). This association between violence and moral beliefs provides the foundation for the first risk factor in our theory: violence at protests is associated with protest-relevant moral rhetoric on online social networks, because this rhetoric reflects moralization and moralization is a risk factor for violence (Skitka & Morgan, 2014).

However, while moralization is a risk factor for violence, we propose that in isolation it is not always sufficient for violent protest. Rather, in our view, protest violence is contingent on both moralization and the belief that others share one's moral attitudes, a phenomenon which we refer to as perceived moral convergence. When people encounter others who share their moral attitudes, those attitudes are validated

and reinforced and, as moral beliefs become more intransigent, the likelihood of advocating or enacting violence to achieve desired moral ends (e.g., topple a corrupt government, alter policing practices, stop the removal of a statue, or defend the purity of one's race) may increase (Boothby, Clark, & Bargh, 2014).

Accordingly, we propose that the risk of violent protest is not simply a function of moralization, but also the perception of moral convergence, which we believe can be influenced by social media dynamics. People do not only rely on social media to signal their moral sentiments; they also use social media to gauge the moral sentiments held by others. Thus, as morally-relevant messages proliferate, impressions of convergence may become more robust and . Notably, such impressions can be dramatically biased given the tendency for online social networks to function as digital echo chambers (Barberá, Jost, Nagler, Tucker, & Bonneau, 2015). Nonetheless, we suggest that the rate of moral rhetoric in online social networks encodes both a signal of moralization and moral convergence.

To test these hypotheses, we collected tweets relevant to the 2015 Baltimore protests. These protests, which were motivated proximally by the eventual death of Freddie Gray, extended across multiple weeks and included both periods of peace and violence, allowing us to examine the links between moral rhetoric and violent protests. Second, across three behavioral experiments, we directly test whether the interaction between moralization and perceived moral convergence explains support for violent protest.

Study 1

To test the hypothesized association between moralization and violent protest, we conducted an observational study of the relationship between online moral rhetoric and real-world indicators of violent protest during the 2015 Baltimore protests. We first investigated whether expressions of online moral rhetoric increase on days with violent protests, relative to days without. We then conducted two fine-grained hour-level analyses investigating the association between rates of online moral rhetoric and rates of

arrests in the Baltimore area as reported by the Baltimore police.

Data Collection. Approximately 19 million tweets posted during the Baltimore protests (04/12/2015 to 05/08/2015) were purchased from Gnip.com. These tweets were filtered based on geolocation information, constrained to cities where protests related to the death of Freddie Gray occurred. Compared to focusing on a specific list of hashtags, filtering by geo-location allows us to focus more holistically on the nation-wide social stream rather than on some small portion associated with experimenter-determined hashtags.

Moral Rhetoric Classification. To evaluate the association between online moral rhetoric and violent protest, it was first necessary to predict the moral sentiment of the tweets in the Baltimore corpus. We accomplished this using a Long Short-Term Memory neural network trained on a set of 5000 tweets labeled by expert annotators (Ying, Hoover, Dehghani, Mooijman, & Ji, 2017). We first developed a coding manual (Hoover, Johnson, Dehghani, & Graham, 2017) and trained three human annotators to code the tweets for moral content based on the Moral Foundations Theory (Graham, Haidt, & Nosek, 2009; Haidt, Graham, & Joseph, 2009; Haidt, Joseph, et al., 2007). Specifically, each coder was asked to annotate each tweet depending on whether it was related to any of the five Moral Foundations or not. Agreement was measured using Prevalence and Bias adjusted Kappa (PABAK) (Byrt, Bishop, & Carlin, 1993; Sim & Wright, 2005), an extension of Cohen’s Kappa robust to unbalanced data sets. The calculated agreement was high for all dimensions (for the moral dimensions averaged across coder pairs, $M = 0.723$, $SD = 0.168$). The agreement for moral/non-moral categories was 0.636.

The annotated tweets were then used to train a deep neural network-based model to automatically predict moral values involved in Twitter posts (Ying et al., 2017). This model consists of three layers, an embedding (lookup) layer, a recurrent neural network (RNN) with long short-term memory (LSTM) (Hochreiter & Schmidhuber, 1997) and an output layer. The first layer converts words in an input tweet to a sequence of pre-trained word embeddings, which are low-dimensional dense vectors that represent

semantic meanings of words. After that, the LSTM layer processes these embeddings in succession and outputs a fixed-sized vector which encodes critical information for moral value prediction. Compared to vanilla RNNs, LSTM is capable of storing inputs over long sequences to model long-term dependencies. Other dense features, such as a vector representing percentage of words that match each category in the Moral Foundation Dictionary (Graham et al., 2009), are concatenated with the LSTM feature vector. On top of the model, a fully connected layer with Softmax activation function is added to perform binary classification. It takes as input the concatenated feature vector and predicts whether a tweet reflects a moral concern or not. We trained a separate model for each moral foundation and combined all results. Previous research has found that the moral content labels predicted by this model are generally not distinguishable from labels generated by humans (Ying et al., 2017). Our model achieved a cross-validated accuracy of 89.01% ($F_1 = 87.96$) compared to 75.15% ($F_1 = 72.17$) of MFD on the tweets used for training.

Study 1A

To evaluate the association between online moral sentiment and violent protest we first investigated whether more moral tweets were posted on days with violent protests. To do this, we estimated the daily count of moral tweets as a function of whether a violent protest occurred or not.

Data. Using the LSTM Tweet Categorization Model (see above), we generated binary labels indicating the presence of each Moral Foundation domain and aggregated these labels to form a binary moral vs. non-moral label for each tweet in our corpus. Both hourly and daily counts of moral tweets were then generated from these binary labels. Then, for each day in the range of dates included in the twitter corpus, media coverage of the timeline of the Baltimore protests was used to label each day for whether it had no protest ($n = 12$), a peaceful protest ($n = 11$), or a violent protest ($n = 4$).

Design and Analysis. Because the dependent variable (daily count of moral tweets) is a count variable and overdispersed ($M = 8188.45$, $s^2 = 23,065,122$), we used

negative binomial regression. In total, three models were estimated: a null, intercept-only model (Model 1); a model with a sufficient serial autocorrelation structure (Model 2); and a model with both an autocorrelation component and dummy variables indicating protest type (Model 3). Unless otherwise stated, all models in studies 1A and 1B were estimated in the *R* statistical computing environment version 3.3.1 (R Core Team, 2013) using the package *tscount* version 1.4.0 (Liboschik, Fokianos, & Fried, 2015).

To test our hypothesis, we first used Model 1 to conduct an intervention test (Fokianos & Fried, 2010), which tests the null hypothesis that a specified intervention has no effect on the data generation process (DGP). In our case, the alternative hypothesis assumes that days with violent protests statistically intervene on the rate of moral tweets. Functionally, this test treats intervention time-points (i.e. violent protest days) as potential outliers and estimates the likelihood that they were generated from the baseline DGP. If the intervention time-points appear to have come from a different DGP, this counts as evidence for an intervention effect. We follow this intervention test with Model 3, which directly estimates the effect of violent protest days on the rate of moral tweets.

Results. First, a baseline model (Model 1) including only the intercept and dispersion parameter was estimated. Because the dependent variable is a time-series, residual autocorrelation was both expected and observed in this model, Box-Pierce $\chi^2 = 10.26, p = 0.001$. Subsequent examination of autocorrelation and partial autocorrelation plots indicated an AR(1) structure and the baseline model was refit with a AR(1) component (Model 2). Examination of the Model 1 residuals indicated that the AR(1) sufficiently accounted for the residual autocorrelation observed in Model 1, Box-Pierce $\chi^2 = 10.26, p = 0.58$.

To evaluate the association between daily moral tweet counts and days with violent protests, we first conducted an intervention test (Fokianos & Fried, 2010). This indicated an effect of day-level violent protest on the daily rate of moral tweets, such that the rate of tweets on days with violent protests was substantially different from the

rate of moral tweets on other days, $\chi^2 = 11.48, p = 0.02$. We then estimated a third model (Model 3) by modifying Model 2 to include two dummy coded factors reflecting whether, for a given day, no protest occurred, a peaceful protest occurred, or a violent protest occurred. Convergent with our hypothesis and the intervention test, the parameters estimated in this model indicate a positive association between the occurrence of a violent protest and the daily rate of moral tweets, $b = 0.63, SE = 0.2, 95\%CI = [0.24, 1.02]$. That is, on days with violent protests, the log expected counts of moral tweets is expected to increase by 0.63. In terms of incidence ratios, this means that the rate of moral tweets on days with violent protests is 1.88 times that of days with no protests, holding the other variables in the model constant. No such association was observed for peaceful protest days, $b = 0.11, SE = 0.14, 95\%CI = [-0.17, 0.39]$ (Figure 1).

Discussion. This analysis provides evidence for the operational hypothesis that days with violent protests have higher rates of moral tweets. That is, if a moral protest occurs on a given day, this model indicates that we should expect the rate of moral tweets to also increase on that day. This is a necessary condition for our hypothesis; however, while it is consistent with our hypothesis, it does not provide direct confirmatory evidence. For example, the rate of moral tweets might increase on violent protest days because people tweet about the violent protests after they occur. In the next study, we address this ambiguity by conducting a more fine-grained time-series analysis. Specifically, we investigate whether the rate of moral tweets predicts the rate of arrests during protests.

Study 1B

Study 1A provided evidence for the hypothesis that the rate of moral tweets increases on days with violent protests. However, as discussed above, we are primarily interested in whether the rate of moral tweets predicts violent protest. To do address this question, we tested the hypothesis that the hourly rate of moral tweets predicts protest violence. However, because hourly estimates of violence during the Baltimore

protests do not exist, it was not possible to rely on direct measurements of violence. To overcome this issue, we used hourly arrest rates in the Baltimore area as an indicator of protest violence.

Data. Hourly counts of moral tweets were calculated using the labeled Twitter data described in study 1A. Hourly arrest counts for the Baltimore area were obtained from the Open Baltimore data portal at <https://data.baltimorecity.gov/>.

Design and Analysis. Both the hourly rate of moral tweets and the hourly rate of arrests are time-series (see Figure 2A for arrest and moral tweet count time series). Thus, whether the former predicts the latter cannot be determined with conventional regression models because shared temporal trends (such as a seasonal pattern at the hour level) can induce spurious correlations between otherwise independent time-series (Chatfield, 2016). One widely-used method for overcoming this problem is the Granger causality test (Granger, 1969). This test assumes that a dependent variable Y is at least partially a function of some set of endogenous factors and it tests whether an independent variable X provides additional predictive information. In practice, Granger causality is tested by regressing Y at $t + 1$ both on lagged values of Y and X and evaluating whether the lagged X components improve the model. If lagged X values do not predict Y at $t + 1$ above and beyond the degree to which lagged Y values predict Y at $t + 1$, then Granger non-causality is assumed. However, if the lagged X values contribute additional predictive information to the model, then Granger causality is assumed.

To evaluate the relationship between the hourly rates of moral tweets and arrests, we first test for Granger causal relationship between moral tweets and arrests on days with peaceful or violent protests using the Toda and Yamamoto (Toda & Yamamoto, 1995) procedure. Because it may also be the case that the rate of arrests Granger causes the rate of moral tweets, we also test the reverse model. That is, while it may be that the rate of moral tweets is an indicator of protest violence, it may also be the case that protest violence granger causes the rate of moral tweets.

Following these tests of Granger causality, we then estimate a dynamic time-series

model in order to estimate the degree to which the lagged rate of moral tweets is associated with the rate of arrests. Because we do not have a hypothesis about the specific temporal association between moral tweets and arrests, to estimate this association we average the hourly rates of moral tweets over a four hour window and use this four-hour rolling mean to predict the rate of arrests at $t + 1$. That is, we estimate the rate of arrests at $t + 1$ as a function of the average rate of moral tweets across $t - 3$, $t - 2$, $t - 1$, and t . Lastly, we estimate two versions of this model. In the first version, the effect of moral tweets on arrest rate is modeled as a fixed effect. That is, it is assumed that this effect does not vary across days. In the second and final version of the model, we relax this assumption, as well as the assumption of fixed intercepts across days.

Thus, in this final model, we include all of the terms included in the former model, but also allow the intercept and the slope of the moral tweets variable to vary across days. In other terms, we estimate a hierarchical model with so-called random intercepts and slopes. By relaxing the fixed slope assumption, this model better reflects the uncertainty of the slope estimates (Gelman & Hill, 2007). This model is estimated using Bayesian estimation via the *rstanarm* R package version 2.15.3 (Carpenter et al., 2016) with default priors. In all of these models, hourly moral tweet counts ($M = 406.5, SD = 418.92$) was log-transformed and standardized so that its scale was comparable to that of hourly arrest counts ($M = 3.10, SD = 3.57$). Further, as in Study 1A, hourly arrest count is a count variable and cannot be expected to be normally distributed. To allow for the possibility of overdispersion, both null Poisson and Negative Binomial models are compared.

Results. A Vuong model comparison test (Vuong, 1989) indicated that hourly arrest count was better fit by the null Negative Binomial model, compared to a Poisson model. Per the Toda and Yamamoto procedure, the maximum order of integration p across tweet counts and arrest counts ($p = 1$) was identified across both moral tweet counts and arrest counts using augmented Dickey-Fuller (Banerjee, Dolado, Galbraith, Hendry, et al., 1993) and Kwiatkowski-Phillips-Schmidt-Shin (Kwiatkowski, Phillips,

Schmidt, & Shin, 1992) tests for stationarity. Autocorrelation and partial autocorrelation plots and Box-Pierce (Box & Pierce, 1970) tests were then used to identify the maximum number of lags m ($m = 2$). Examination of the residual correlation plots also revealed seasonal correlation at the hourly level. Finally, per the Toda and Yamamoto procedure, a final model with $m + p = 2 + 1 = 3$ lags for both number of arrests and number of moral tweets was estimated. This model also included a 24 hour arrest count lag to account for hourly seasonality. Finally, the Wald χ^2 test of the null hypothesis of Granger non-causality, that the slopes of the m moral tweet lags are not distinguishable from zero, was rejected, $\chi^2(2) = 17.5, p = 0.0002$. That is, the standardized log rate of moral tweets Granger causes the rate of arrests. To determine whether the count of arrests Granger causes the standardized log count of moral tweets, the transformed moral tweets variable was entered in a linear regression with the same predictors as the previous model. The Wald χ^2 test of Granger non-causality was also rejected for this model, $\chi^2(2) = 6.4, p = 0.04$. Accordingly, these analyses indicate an asymmetric, bi-directional granger causal relationship, such that the rate of moral tweets predicts the future rate of arrests and the rate of arrests predicts the future rate of moral tweets.

To estimate the magnitude of the effect of lagged moral tweet counts on future arrest rates, we estimated two final Negative Binomial $AR(2)(1)[24]$ models. The first model was estimated using the *tscount* package and the effect of standardized log count of moral tweets was treated as fixed. In this model, the expected effect of the moral tweets variable ($b = 0.22, SE = 0.09, 95\%CI = [0.04, 0.39]$) indicated that the rate of moral tweets across a four hour window predicts the rate of arrests in the next hour. The second version of this model, which was estimated using Bayesian estimation and which allowed the intercept and moral tweets effect to vary across days, revealed substantial between-days variation for the intercept ($SD = 0.20$) and the effect of the moral tweets variable ($SD = 0.16$). However, the estimate of the fixed effect of the moral tweets variable was comparable to the previous model ($b = 0.24, SD = 0.10, 95\%HPDI = [0.07, 0.44]$). Thus, even after accounting for arrest

rates during previous hours, variable baseline arrest rates across days, and potential variation in the effect of moral tweets on arrest rates, these models suggest that for a one unit increase of the average standardized log count of moral tweets over a four hour time span, the log count of arrests is expected to increase by 0.24. In terms of incidence ratios, this means that the rate of arrests increases by 1.27 for every one unit increase in the moral tweets variable (see Figure 2B for observed and predicted hourly arrest counts).

Discussion. This analysis indicates that the rate of moral tweets predicts the rate of arrests. We find evidence for this both via tests of Granger causality and direct estimates of the association between lagged moral tweets and number of arrests. Specifically, the Granger causality analysis found that moral tweet rates predicted future arrest rates above and beyond current arrest rates. Further, we found evidence for our directional hypothesis; as the rate of moral tweets increases, the expected future rate of arrests also increases. Thus, even after accounting for arrest rates during previous hours, variable baseline arrest rates across days, and potential variation in the effect of moral tweets on arrest rates, our model indicates a relationship between the rate of moral tweets and the future rate of arrests. In other words, this model suggests that observing expressions of moral sentiment in social media can help predict when future protests will take on the sort of characteristics that lead to higher rates of arrests.

While these results correspond to our general hypothesis, they do not address the question of what, exactly, higher rates of moral sentiment indicate. As discussed above, an increase in moral sentiment could simply reflect an increase in moralization; however, given that people rely on social media to track popular opinions, it may also be the case that an increase in moral sentiment also indexes perceived convergence. Accordingly, we followed this study with three experimental studies, which allowed us to conduct more precise hypothesis tests and to disentangle the potential effects of moralization and moral convergence.

Study 2

In Study 2, we tested the degree to which moralizing a protest predicts the acceptability of using violence at this protest. Specifically, we used the violent protests between the far-right and counter-protesters in Charlottesville, Virginia in August 2017 as an example. We tested to what extent protesting the far-right was seen as a moral issue and to what extent this moralization predicted the perceived acceptability of using violence against the far-right.

Method

Participants. Two hundred and seventy five participants (157 men; $M_{age} = 33.82$ years, $SD_{age} = 9.61$) were recruited from Mechanical Turk (Mturk; Buhrmester, Kwang, & Gosling, 2011). A power analysis indicated that our sample size provided approximately 90% power to detect a small to medium effect size ($\beta \approx .25$). All reported power analyses were conducted using G*Power 6 (Faul, Erdfelder, Lang, & Buchner, 2007).

Design and Procedure. We introduced participants with a short description of the Unite the Right rally in Charlottesville, Virginia, United States, where far-right protesters and counter-protesters clashed. Specifically, participants read the following: “The Unite the Right rally (also known as the Charlottesville rally) was a far-right rally in Charlottesville, Virginia, United States, from August 11–12, 2017. The rally occurred amidst the backdrop of controversy generated by the removal of several Confederate monuments”. Participants were informed that we were interested in their opinion about the counter-protesters that protested the far-right protesters. Participants indicated on a four item scale (1 = disagree completely, 5 = agree completely) to what extent they thought that protesting the far right is a moral issue (i.e., whether protesting the far right was: a reflection of their moral beliefs and convictions, connected to their beliefs about moral right and wrong, based on moral principle, and a moral stance; taken from Skitka and Morgan (2014), Cronbach’s $\alpha = .94$).

Participants also indicated (1 = disagree completely, 7 = agree completely) to

what extent they agreed with the following statements: it is acceptable to use violence against far-right protesters, the use of violence against far-right protesters is justified, violence against far-right protesters is acceptable if its means fewer future protests from the far-right, using force during a protest is wrong even if its leads to positive change, using force during a protest against the far-right is immoral even if its leads to positive change (the last two items were reverse-coded; $\alpha = .91$).

Results and Discussion

A regression analysis in which violence was regressed on moralization showed that, as expected, the moralization scale was positively associated with the acceptability of using violence at protests ($\beta = .19, t[275] = 3.22, p = .001$). These findings suggest that moralization is indeed associated with the increased acceptability of violent protests.

Study 3

Study 3 aimed to replicate Study 2 while also disentangling the potential effects of moralization and moral convergence on violence acceptability.

Method

Participants. Two hundred and five participants (132 men; $M_{age} = 35.04$ years, $SD_{age} = 11.61$) were recruited from Mechanical Turk and randomly assigned to two conditions: high versus low moral convergence. A power analysis indicated that our sample size provided approximately 80% power to detect a small to medium effect size ($\beta \approx .25$).

Design and Procedure. Participants were confronted with the same excerpt and moralization questionnaire as in Experiment 1 ($\alpha = .93$). In addition, in the moral convergence [non-convergence] condition, participants were told that based on their responses, “the majority of [few] people in the United States share your particular moral values. Other people in the United States think about this protest in a similar [different] manner compared to you”. Similar to Experiment 1, participants then

indicated on six items (1 = disagree completely, 7 = agree completely) to what extent they considered using violence against far-right protesters acceptable, $\alpha = .89$).

Results and Discussion

Regression analysis was used to test the interactive effects of moralization and moral convergence on the perceived acceptability of violence. For the first step, moralization (standardized) and moral convergence were added. For the second step, the interaction between moralization and convergence was added. Results demonstrated that moralization was positively associated with violence

($\beta = .15, t[204] = 2.22, p = .028$) and that moral convergence was (overall) unrelated to violence ($\beta = .02, t[204] = 0.25, p = .80$). Crucially, we observed a significant interaction effect between moralization and moral convergence ($\beta = .44, t[204] = 1.97, p = .050$).

Moralization predicted the acceptability of violence at protests when moral convergence with others was high ($\beta = .27, t = 2.82, p = .006$) but not when moral convergence with others was low ($\beta = .01, t = 0.05, p = .96$).

Similar to Study 2, findings from Study 3 indicate that participants found violence more acceptable when they moralized a protest. In addition, Study 3 shows this effect to be moderated by the degree to which people believed others shared their moralized attitudes. Moralization predicted violence only when participants perceived that they shared their moralized attitudes with others.

Study 4

In Study 4 we aimed to directly replicate the findings from Study 3 while also measuring attitude certainty as a possible explanation for why moralization only predicts the acceptability of violence under conditions of moral convergence.

Method

Participants. Two hundred and eight participants (157 men; $M_{age} = 36.25$ years, $SD_{age} = 12.22$) were recruited from Mechanical Turk. A power analysis indicated

that our sample size provided approximately 80% power to detect a small to medium effect size ($\beta \approx .25$).

Design and Procedure. Participants were confronted with the same excerpt and moralization questionnaire as in Experiment 1 and 2 ($\alpha = .78$), and the same convergence manipulation as in Experiment 2. Similar to the previous experiments, participants then indicated on six items (1 = disagree completely, 7 = agree completely) to what extent they considered using violence against far-right protesters acceptable, $\alpha = .89$). Participants also indicated to what extent they were certain about their attitude towards the protest (1 = extremely uncertain, 7 = extremely certain).

Results and Discussion

We conducted similar analyses in Experiment 2. Results demonstrated that moralization was positively associated with violence certainty ($\beta = .15, t[208] = 2.19, p = .029$) and that moral convergence was (overall) unrelated to violence ($\beta = .07, t[208] = 0.98, p = .33$). Crucially, we observed a significant interaction effect between moralization and moral convergence ($\beta = .20, t[208] = 2.92, p = .004$). Moralization predicted the acceptability of violence at protests when moral convergence with others was high ($\beta = .35, t = 3.65, p < .001$) but not when moral convergence with others was low ($\beta = .05, t = 0.05, p = .62$).

In addition, a 5,000 resamples bootstrapping analysis (Hayes, Preacher, & Myers, 2011) demonstrates that attitude certainty (95%CI = [0.34, 0.42]) mediates this interaction effect. Specifically, the effect of moralization on violence in the moral convergence condition was mediated by an increase in attitude certainty (95%CI = [0.28, 0.99]).

Studies 2, 3, and 4 confirm that the moralization of a protest can increase the acceptability of using violence at this protest. In addition, this only occurred when participants perceived to share their moralized attitudes with others and increased their attitude certainty. Overall, the findings from these three behavioral experiments are consistent with the findings from the Baltimore protests such that the moralization of

protests and the moral convergence of these moralized attitudes drives the acceptability of using violence at protests.

Discussion

Taken together, the findings from 18 million tweets posted during the 2015 Baltimore protests and three behavioral experiments suggest that violence at protests can be understood as a function of two key risk factors: the degree to which people see protest as a moral issue and the rate of perceived moral convergence. A rise in violence at protests may thus reflect the increasing moralization and polarization of political issues in online echo chambers. The risk of violent protest, in other words, may not be simply a function of moralization, but also the perception that others agree with one's moral position, which can be strongly influenced by social media dynamics. Indeed, previous research has shown that people do not just rely on social media to signal their moral sentiment but they also use social media to gauge the moral sentiments held by others (Barberá et al., 2015; Dehghani et al., 2016).

As morally-relevant messages proliferate in social networks, impressions of convergence may become more robust, as suggested by the findings reported in the current manuscript. This may increase the degree to which people overcome their objections to using violence aimed at perceived opponents (Zaal et al., 2011). Notably, impressions of moral convergence on online social networks can be dramatically biased given the tendency for online social networks to function as digital echo chambers (Barberá et al., 2015). It is estimated that seven out of ten Americans are connected to some online social network (Pew Research Center, 2017) and that political polarization is at an all-time high (Doherty, 2014), implying that online social networks are currently playing a significant role in shaping, and perhaps causing, our attitudes surrounding the use of violence aimed at ideological opponents at protests.

These findings come at a time when recent polls suggest that the acceptability of using violence at protests is rising among the US public and college students in particular (Rampell, 2017). In the last few years, protesters in the US have clashed

violently with police and counter-protesters on numerous occasions, whether this was due to the killing of young men by police, protests by the far-right, or the invitation of controversial speakers on college campuses. The findings reported in the current manuscript shed some light onto these social development while also providing suggestions for counteracting the increasing acceptability of violence. Increased support for using violence at protests occurred only when people perceived to share moralized attitudes with other but not when people did not perceive to share moralized attitudes with others. This implies that decreasing the moralization of attitudes and diluting the perception that others agree with one's moral position may attenuate the rise of the acceptability of violence. Although people may try to "sort" themselves into networks with high levels of moral convergence, our findings imply that ways of combating this may be effective at decreasing the acceptability of using violence at protests.

However, while we propose that the combination of moral outrage and perceived moral convergence are prerequisites for violent protest, they are by no means sufficient conditions for violent protest. In addition to these factors, whether a protest will become violent or not might also depend on a range of other contextual factors, such as the propensity for violence among the protesting population (i.e., base-rate levels of violent inclinations), likelihood of instigation from involved non-protesting agents, and the specific nature of the issues being protested (e.g., a march explicitly aimed at being peaceful). Nonetheless, for any given configuration of contextual factors, the current work suggests that the risk of violence increases as a function of moralization and perceived convergence of moralization. Indeed, our work extends previous research on (sacred) values and violence by suggesting that strong moral convictions can increase the acceptability of violence (Atran & Ginges, 2012; Dehghani et al., 2010) and that (a) this happens primarily when convergence is high, and (b) the link between moralization and violence can be tracked, and is shaped, by online social networks.

As a consequence, our findings have far-reaching practical implications, as a key decision-making problem for government officials is to understand which protests will turn violent and how resources should be allocated to prevent protests from spiraling

out of control (Hvistendahl, 2016). Although more research should be done to replicate and extend the findings reported in the current manuscript, we believe that our findings suggest that online social networks and moral psychology can be of some help: Given the increasing importance of social networks in our daily lives, the moral language used on online social networks can be directly linked to violent protests, implying that online social networks can also be used by policy makers to track and predict the emergence of violence at protests.

References

- Atran, S., & Ginges, J. (2012). Religious and sacred imperatives in human conflict. *Science*, 336(6083), 855–857.
- Banerjee, A., Dolado, J. J., Galbraith, J. W., Hendry, D., et al. (1993). Co-integration, error correction, and the econometric analysis of non-stationary data. *OUP Catalogue*.
- Barberá, P., Jost, J. T., Nagler, J., Tucker, J. A., & Bonneau, R. (2015). Tweeting from left to right: Is online political communication more than an echo chamber? *Psychological science*, 26(10), 1531–1542.
- Boothby, E. J., Clark, M. S., & Bargh, J. A. (2014). Shared experiences are amplified. *Psychological science*, 25(12), 2209–2216.
- Box, G. E., & Pierce, D. A. (1970). Distribution of residual autocorrelations in autoregressive-integrated moving average time series models. *Journal of the American statistical Association*, 65(332), 1509–1526.
- 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(28), 7313–7318.
- Buhrmester, M., Kwang, T., & Gosling, S. D. (2011). Amazon’s mechanical turk: A new source of inexpensive, yet high-quality, data? *Perspectives on psychological science*, 6(1), 3–5.
- Byrt, T., Bishop, J., & Carlin, J. B. (1993). Bias, prevalence and kappa. *Journal of clinical epidemiology*, 46(5), 423–429.
- Carpenter, B., Gelman, A., Hoffman, M., Lee, D., Goodrich, B., Betancourt, M., ... Riddell, A. (2016). Stan: A probabilistic programming language. *Journal of Statistical Software*, 20, 1–37.
- Chatfield, C. (2016). *The analysis of time series: an introduction*. CRC press.
- Crockett, M. (2017). Moral outrage in the digital age. *Nature Human Behaviour*.
- Darley, J. M. (2009). Morality in the law: The psychological foundations of citizens’ desires to punish transgressions. *Annual Review of Law and Social Science*, 5,

1–23.

- Davey, M., & Bosman, J. (2014). Protests flare after ferguson police officer is not indicted. *The New York Times*. Retrieved from <https://www.nytimes.com/2014/11/25/us/ferguson-darren-wilson-shooting-michael-brown-grand-jury.html>
- Dehghani, M., Atran, S., Iliev, R., Sachdeva, S., Medin, D., & Ginges, J. (2010). Sacred values and conflict over iran's nuclear program. *Judgment and Decision Making*, 5(7), 540.
- Dehghani, M., Johnson, K., Hoover, J., Sagi, E., Garten, J., Parmar, N. J., . . . Graham, J. (2016). Purity homophily in social networks. *Journal of Experimental Psychology: General*, 145(3), 366.
- Doherty, C. (2014). 7 things to know about polarization in america. *Pew Research Center*. Retrieved from <http://www.pewresearch.org/fact-tank/2014/06/12/7-things-to-know-about-polarization-in-america/>
- Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G* power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior research methods*, 39(2), 175–191.
- Fiske, A. P., & Rai, T. S. (2014). *Virtuous violence: Hurting and killing to create, sustain, end, and honor social relationships*. Cambridge University Press.
- Fokianos, K., & Fried, R. (2010). Interventions in ingarch processes. *Journal of Time Series Analysis*, 31(3), 210–225.
- Gelman, A., & Hill, J. (2007). *Data analysis using regression and multilevelhierarchical models* (Vol. 1). Cambridge University Press New York, NY, USA.
- 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(5), 1029.
- Granger, C. W. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: Journal of the Econometric Society*, 424–438.

- Haidt, J., Graham, J., & Joseph, C. (2009). Above and below left-right: Ideological narratives and moral foundations. *Psychological Inquiry*, 20(2-3), 110–119.
- Haidt, J., Joseph, C., et al. (2007). The moral mind: How five sets of innate intuitions guide the development of many culture-specific virtues, and perhaps even modules. *The innate mind*, 3, 367–391.
- Hayes, A. F., Preacher, K. J., & Myers, T. A. (2011). Mediation and the estimation of indirect effects in political communication research. *Sourcebook for political communication research: Methods, measures, and analytical techniques*, 23, 434–465.
- Heim, J. (2017). Recounting a day of rage, hate, violence and death. *The Washington Post*. Retrieved from https://www.washingtonpost.com/graphics/2017/local/charlottesville-timeline/?utm_term=.b960304937fe
- Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural computation*, 9(8), 1735–1780.
- Hoover, J., Johnson, K. M., Dehghani, M., & Graham, J. (2017). Moral values coding guide. *PsyArXiv*. Retrieved from psyarxiv.com/5dmgj
- Hvistendahl, M. (2016). Can ‘predictive policing’ prevent crime before it happens. *Science Magazine*. Retrieved from www.sciencemag.org/news/2016/09/can-predictive-policing-prevent-crime-it-happens
- Kwiatkowski, D., Phillips, P. C., Schmidt, P., & Shin, Y. (1992). Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root? *Journal of econometrics*, 54(1-3), 159–178.
- Liboschik, T., Fokianos, K., & Fried, R. (2015). *tscount: An r package for analysis of count time series following generalized linear models*. Universitätsbibliothek Dortmund.
- Manjoo, F. (2017). The alt-majority: How social networks empowered mass protests against trump. *The New York Times*. Retrieved from <https://www.nytimes.com/2017/01/30/technology/>

- `donald-trump-social-networks-protests.html?mcubz=0`
- Olson, M. (1965). *Logic of collective action: Public goods and the theory of groups*. Harvard University Press.
- Pew Research Center. (2017). Social media fact sheet. *Pew Research Center*. Retrieved from <http://www.pewinternet.org/fact-sheet/social-media/>
- R Core Team. (2013). R: A language and environment for statistical computing [Computer software manual]. Vienna, Austria. Retrieved from <http://www.R-project.org/>
- Rampell, C. (2017). Protests flare after ferguson police officer is not indicted. *The Washington Post*. Retrieved from https://www.washingtonpost.com/opinions/a-chilling-study-shows-how-hostile-college-students-are-toward-free-speech/2017/09/18/cbb1a234-9ca8-11e7-9083-fbfdddf6804c2_story.html
- Sim, J., & Wright, C. C. (2005). The kappa statistic in reliability studies: use, interpretation, and sample size requirements. *Physical therapy*, 85(3), 257–268.
- Skitka, L. J., & Morgan, G. S. (2014). The social and political implications of moral conviction. *Political Psychology*, 35(S1), 95–110.
- Toda, H. Y., & Yamamoto, T. (1995). Statistical inference in vector autoregressions with possibly integrated processes. *Journal of econometrics*, 66(1), 225–250.
- Van Stekelenburg, J., & Klandermans, B. (2013). The social psychology of protest. *Current Sociology*, 61(5-6), 886–905.
- Vuong, Q. H. (1989). Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica: Journal of the Econometric Society*, 307–333.
- Walker, I., & Pettigrew, T. F. (1984). Relative deprivation theory: An overview and conceptual critique. *British Journal of Social Psychology*, 23(4), 301–310.
- Ying, L., Hoover, J., Dehghani, M., Mooijman, M., & Ji, H. (2017). Acquiring background knowledge to improve moral value prediction. *ArXiv*. Retrieved from arxiv.org/abs/1709.05467
- Zaal, M. P., Laar, C. V., Ståhl, T., Ellemers, N., & Derks, B. (2011). By any means necessary: The effects of regulatory focus and moral conviction on hostile and

benevolent forms of collective action. *British Journal of Social Psychology*, 50(4), 670–689.

Figure 1. Expected and observed daily moral tweet counts by protest type. The shaded band around the expected trend line indicates the 95% HPD interval.

Figure 2. A: Z-Scored time-series of smoothed hourly counts of arrests and moral tweets. B: Observed and predicted arrest counts. The gray band indicates the 95% HPDI predictive interval.

Figure 3. Acceptability of violence as a function of moralization and moral convergence.