How Real is Moral Contagion in Online Social Networks?

Jason W. Burton (jburto03@mail.bbk.ac.uk)
Nicole Cruz (ncruzd01@mail.bbk.ac.uk)
Ulrike Hahn (u.hahn@bbk.ac.uk)

Department of Psychological Sciences, Birkbeck, University of London
Malet Street, London, WC1E 7HX, UK

Abstract

People increasingly turn to online social networks for information and debate. This means that the structures and properties of these networks, and the information they propagate, play crucial roles in the development of social beliefs, attitudes, and morals. Recently, research has shown that the presence of specific language drives the diffusion of moral messages, regardless of the informational quality, in a phenomenon dubbed *moral contagion* (Brady et al., 2017). Due to the widespread attention and implications of such findings for science and society, we investigate the presence of moral contagion across six sets of data that capture the communications of naturally-occurring networks on Twitter. Across a large corpus of diverse tweets (n = 525,229), we find moral contagion to be an inconsistent and often absent phenomenon that does not effectively predict message diffusion. The implications and reasons for this finding are discussed.

Keywords: moral contagion; social networks; social influence; computational social science; Twitter

Introduction

The advent of internet-based communications has permitted global connections at previously unfathomable speeds and volumes. While these connections present immense opportunities via global knowledge sharing and peer-to-peer collaboration, the magnitude of our interconnectedness also presents new liabilities, such as new forms of political persuasion.

The spread of psychological and behavioral phenomena is often likened to that of a pathogen moving from node to node, individual to individual, as a result of repeated exposure. Whereas concepts of “peer effects” and interpersonal influence have existed in psychology and sociology domains for quite some time (e.g., Allport, 1920; Redl, 1949), the formalization of *social contagion theory* has done well to shed light on the impacts that social networks have on everyday life. In a series of seminal studies, Christakis and Fowler (2007, 2008; Fowler & Christakis, 2009, 2010) utilized mass longitudinal datasets and network statistics to show that everything from obesity and smoking to happiness and cooperative behavior can cascade and cluster across social networks. From these findings, the development of collective behaviors, norms, and ideologies is understood to be a product of not only the aggregation of individuals, but also the topology of how individuals are arranged. For example, the proximity and volume of interpersonal ties of groups in social spaces, be it digital or not, increases the probability of both social information and behavior being transmitted amongst them. Simultaneously, this reinforcement of social homogeneity makes it difficult for intergroup connections to be made (i.e., echo chambers). As such, social contagion theory provides a lens through which the diffusion of information and the creation of collective intelligence can be examined.

Moral Contagion

In an interesting application of social contagion theory, Brady, Wills, Jost, Tucker, and Van Bavel (2017) present their conceptualization of *moral contagion*, which directly applies the process of social transmission to information diffusion. Extant literature suggests that morality is a powerful force in human reasoning and rationalization, with studies showing that one’s moral beliefs are the foundation for one’s ideology and political views (Graham, Haidt, & Nosek, 2009; Haidt, 2001). But how does an issue become moralized in the first place? Adopting the social intuitionist approach, Brady et al. (2017) explain that moral beliefs are less a product of private, individual reasoning and more the result of interpersonal processes and cultural norms (Haidt, 2001). What’s more, they elaborate that the communication of moral ideas is tied to the use of emotion in social transmission. In other words, emotions, which serve as demonstrated contagions in social networks (e.g., Coviello et al., 2014; Ferrara & Yang, 2015; Kramer, Guillory, & Hancock, 2014), are highly associated with moral judgements and may serve as a segue to moralizing debates that would be otherwise nonmoral (Brady et al., 2017). In their analysis of a large corpus of Twitter communications, they find that not only does emotion drive the diffusion of moral content through social networks, but that the mere presence of moral-emotional words in a tweet increases its transmission by approximately 20% (Brady et al., 2017). Within the contemporary context of ideological polarization, the finding that moral-emotional language diffuses at such a high rate is concerning. As people increasingly rely on their online networks as news sources, blending spaces of socialization with information, the tendency of moral-emotional language to diffuse across networks means that feelings of outrage or disgust might be weaponized as tools of persuasion. Of course, there is a time and place for moralization, but the claims of Brady et al. (2017) suggest that their emotionally-driven moral contagion is highly impactful across domains, going so far as to say that "it seems likely that politicians, community leaders, and organizers of social movements express moral emotions…in an effort to increase message exposure and to influence perceived norms within social networks" (p. 7316).

Because of the implications for society's ability to effectively reason and debate with contentious issues, the present study seeks to explore the prevalence of moral contagion across diverse, naturally-occurring social networks. More specifically, we aim to put the conclusions drawn by Brady et al. (2017) to the test by recreating their methodology and assessing whether moral-emotional language does in fact predict the diffusion of moral information regardless of quality or "truthiness."

Method

To investigate the presence of moral and emotional contagion in online social networks, an adaptation of Brady et al.'s (2017) methodology was employed. Specifically, we use the R programming language to recreate the main analysis strategy from Brady et al. (2017), reproduce their findings with their cleaned aggregated data, and then apply the analyses to five unique Twitter datasets that capture naturally-occurring social networks. Datasets and R scripts are made available at https://osf.io/943zm/.

Datasets

A total of six datasets were analyzed in this study. Four pre-existing datasets were obtained via the Open Science Framework (OSF) and Google's dataset search engine, which hosts links to a wide range of open data repositories. One dataset (*#MuellerReport*) was self-collected by connecting to the Twitter REST API with the `rtweet` package in R. While no specific dataset or topic was initially targeted, certain criteria were employed. To be considered for this study, datasets had to contain Twitter data (i.e., tweet messages and retweet counts), contain a significant number of messages written in English, and relate to a polarizing or morally-charged real-world issue, event, or social movement. Datasets were further narrowed by collapsing repeated messages into a single observation (to generate a composite diffusion count that combined the raw retweet count with the number of times the message appeared in the dataset) and removing non-English messages. Since the found datasets did not include language identifying metadata, the `textcat` package in R was employed to extract English tweets in these instances.

Brady et al. (2017) First and foremost, the present study drew directly from the recent study of moral contagion in social networks by Brady et al. (2017). Their data, which is generously shared on a public OSF project page, was thus crucial to the present study for both inspiration and corroboration. The data collected by Brady et al. (2017) focused on topical political issues in the United States: gun control, same-sex marriage, and climate change. Using the Twitter API and sets of topic-related filter words (e.g., *guns*, *gun control*, and *NRA* for the gun control topic), tweets and metadata were extracted between 30 October and 15 December 2015.

***#MeToo* Tweets** A second dataset comprised of Twitter messages containing the *#metoo* hashtag was obtained from the data.world repository. The raw dataset (n = 393,135) was extracted with the Twitter API between 29 November and 25 December 2017, little more than a month after the *#metoo* hashtag first appeared online in coordination with the "Me Too movement" (Turner, 2018). The "Me Too movement" is a movement against sexual harassment and assault. It was ignited by Hollywood sexual abuse allegations and has since become an international phenomenon garnering widespread media attention, support, and critique.

***#WomensMarch* Tweets** A third dataset with tweets containing the *#womensmarch* hashtag was also obtained from the data.world repository. Using the Twitter API, 15,000 messages were collected that referenced the pro-women's rights, and effectively anti-Trump, protest that took place in the wake of the presidential inauguration on 21 January 2017 (Adhokshaja, 2017). The Women's March has since become a worldwide movement with annual marches in late January to non-violently protest for women's reproductive rights, LGBTQ rights, immigration and healthcare reform, as well as racial, gender, and religious equality.

Post-Brexit Tweets A fourth dataset containing unfiltered tweets and metadata from the morning that Brexit was announced was obtained from the Mendeley Data repository. This unfiltered dataset (n = 17,998) was collected with NCapture from QSR, and employed a tight temporal parameter so as to capture the global public's reaction to the political event (Parker, 2017). Brexit refers to the result of the 2016 EU Referendum in the United Kingdom, and this dataset includes Twitter responses from across the globe.

Viral 2016 US Election Tweets A fifth dataset (n = 9,001) containing viral tweets (those with 1,000+ retweets) from the 2016 US Presidential Election was obtained from the Zenodo repository. The set of tweets was collected with the Twitter API and extracted messages that contained specific hashtags (*#MyVote2016*, *#ElectionDay*, and *#electionnight*) and/or user handles (*@realDonaldTrump* and *@HillaryClinton*) (Amador, Oehmichen, & Molina-Solana, 2017). This dataset was of special interest as it contained many "fake news" messages as coded by the curators, which one would might expect to use especially morally- and emotionally-charged language to garner extra attention.

***#MuellerReport* Tweets** A sixth dataset (n = 229,046) was collected by using the *#muellerreport* hashtag to retrieve tweets from the Twitter API created between 23 and 25 March 2019 — the weekend during which US Attorney General William Barr released his summary of Special Counsel Robert Mueller's investigation into Donald Trump's 2016 presidential campaign. This corpus was of special interest because the Mueller Report has been (and at the time of writing, *still is*) a major source of controversy. While originally a non-polarized issue, the public opinion has divided overtime (Thomson-DeVeaux, 2019), meaning that moral-emotion might have played a part in moralizing conversations on Twitter.

Procedure and Analysis

All datasets were wrangled with R. Tweets were preprocessed with the `tm` and `dplyr` packages, and then a simple dictionary-based approach was employed to quantify the use of specific rhetoric in each message. To do so, the same three dictionaries used and validated by Brady et al. (2017) were used. One dictionary contains distinctly moral words and stems (n = 316; e.g., *fair*, *racism*, *family*), one contains distinctly emotional words and stems (n = 819; e.g., *panic*, *fear*, *heartwarming*), and one contains moral-emotional words and stems (n = 72; e.g., *shame*, *victimize*, *disgust*) that appeared in both of the original moral and emotional dictionaries (i.e., a subset of the moral and emotional dictionaries that was extracted to form the third unique dictionary). Through this categorization, "moral emotions" are considered distinct from "nonmoral emotions" because they are linked to triggers and functions specific to moral contexts, making them especially relevant to political debate (Haidt, 2003; Brady et al., 2017). For instance, outrage and disgust are often considered prototypical moral emotions because their expression can be elicited by perceiving a moral transgression, the breaking of some social axiom that threatens the collective order (e.g., infringement of human rights). In contrast, sadness is a nonmoral emotion because it can be triggered by nonmoral cues (e.g., the death of a loved one). The presence of these categorized words (moral, emotional, and moral-emotional) in each tweet was counted, so that each observation was coded with a discrete word count for each dictionary.

To accurately assess the degree to which each tweet in a given dataset diffused across the social media platform, the present study utilized a collapse-and-count scheme similar to that of Brady et al. (2017). Essentially, there are two measures of diffusion that can be calculated in an observational Twitter dataset: the retweet count displayed in collected metadata and the number of times a message appears in the dataset itself. Thus, the present study quantified diffusion by counting the presence of identical messages in each dataset, adding this count to the message's actual retweet count recorded in the metadata, and then collapsing repeated messages into a single observation.

To measure contagion effects, a negative binomial regression model was used. This model accounts for the overdispersion of data and effectively models count variables (i.e., discrete word counts and diffusion counts). In an effort to maintain consistency with Brady et al. (2017), incidence rate ratios (IRRs) were used as the ultimate indicator of the existence and magnitude of contagion effects. The `MASS` and `lmtest` packages were used for the main analysis.

Results

For the main analysis, negative binomial regression models with maximum likelihood estimation were fit onto each dataset to follow in line with the methodology of Brady et al. (2017), and to allow for a consistent measurement of moral, emotional, and moral-emotional contagion. The presence of contagion was determined by exponentiating the regression coefficients to generate incidence rate ratios (IRR) for each language dictionary, which were then used to plot diffusion prediction lines (Figure 1).

Brady et al. (2017) Across the corpus of 313,002 tweets, there was an average of 0.23 moral-emotional, 0.36 moral, and 0.69 emotional words per tweet. The main analysis indicated that moral contagion did indeed exist in the data collected by Brady et al. (2017). For distinctly moral language, there was a slight main effect (IRR = 1.02, $p < 0.05$, 95% CI = 1.01, 1.04), and the same went for distinctly emotional language (IRR = 1.03, $p < 0.001$, 95% CI = 1.01, 1.04). Crucially, the presence of moral-emotional language appeared to have a strong effect on message diffusion (IRR = 1.21, $p < 0.001$, 95% CI = 1.19, 1.24). This result corroborates the statistics reported in Brady et al. (2017) to show that the use of moral-emotional language in tweets increases the likelihood of getting retweeted or otherwise shared among individuals in the social network platform by up to 21%. We also performed likelihood ratio tests to assess the statistical model's goodness of fit against nested univariate models that used *only* moral, emotional, *or* moral-emotional language as a predictor of diffusion (Table 1). These tests confirmed that the multivariate negative binomial regression model was effective for predicting message diffusion in the dataset.

***#MeToo* Tweets** After preprocessing the dataset, 151,572 unique tweets remained for analysis with an average of 0.21 moral-emotional, 0.30 moral, and 1.03 emotional words per tweet. The negative binomial regression model displayed a small but significant effect of distinctly moral language (IRR = 1.05, $p < 0.001$, 95% CI = 1.02, 1.09), as well as a significant effect of emotional contagion (IRR = 1.13, $p < 0.001$, 95% CI = 1.11, 1.15) such that emotional language increased a message's diffusion by 13%. Curiously, while both moral and emotional language had a significant relationship with increased diffusion, moral-emotional language was significantly associated with message diffusion in a negative direction (IRR = 0.89, $p < 0.001$, 95% CI = 0.79, 0.85). Likelihood ratio tests confirmed that the multivariate model was the best fit for the dataset (Table 1).

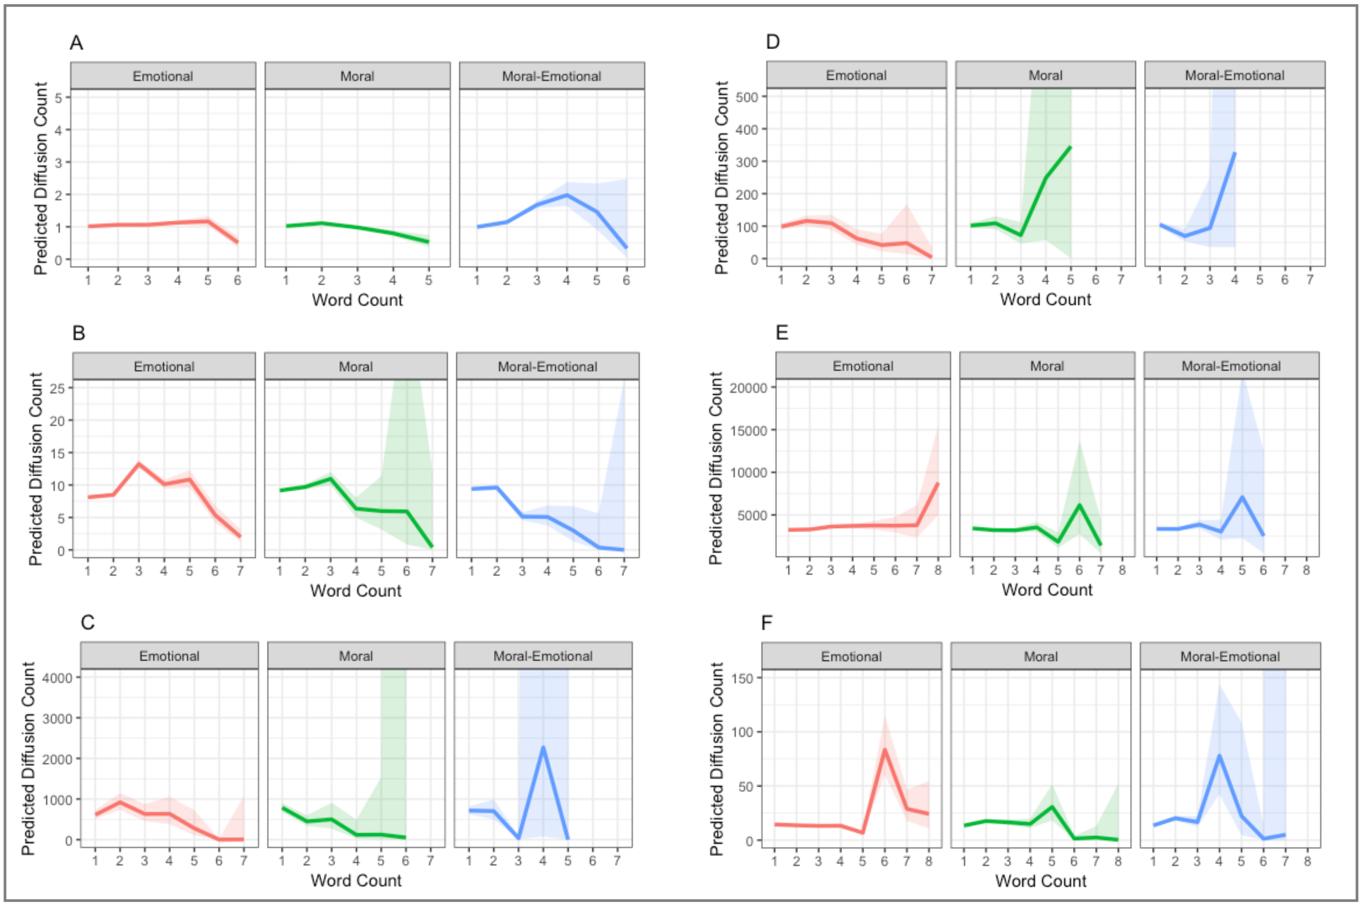

Figure 1: Predicted message diffusion trends as determined by multivariate negative binomial regression models. A = Brady et al. (2017) aggregate dataset; B = *#MeToo* dataset; C = *#WomensMarch* dataset; D = Post-Brexit dataset; E = Viral 2016 US election dataset; F = *#MuellerReport* dataset. 95% CIs are represented with the shaded areas. Note that scales vary widely as a result of the range of diffusion counts present in each dataset.

Post-Brexit Tweets Of the 7,124 analyzable tweets from the morning of the Brexit announcement, there was an average of 0.08 moral-emotional, 0.20 moral, and 0.69 emotional words per tweet. The regression model showed that there was no significant effect of moral language (IRR = 1.05, *n.s.*, 95% CI = 0.91, 1.22), emotional language IRR = 0.95, *n.s.,* 95% CI = 0.88, 0.1.04), or moral-emotional language (IRR = 0.86, *n.s.*, 95% CI = 0.71, 1.07) on message diffusion. While this goes against the claims of Brady et al. (2017), the large confidence intervals and low levels of moral, emotional, and moral-emotional language present in the dataset make the findings here generally inconclusive. In fact, a likelihood ratio test demonstrated that the main multivariate model was not suited to this corpus of tweets. Univariate models, where an isolated word dictionary alone is used to predict diffusion rates instead of the combined three, were slightly better at explaining the data (Table 1).

***#WomensMarch* Tweets** The 3,783 analyzable tweets, messages pertaining to the Women's March movement had an average of 0.17 moral-emotional, 0.31 moral, and 0.86 emotional words per tweet. Upon fitting the negative binomial regression model to the data, it was found that there was no effect of distinctly emotional (IRR = 0.98, *n.s.*, 95% CI = 0.86, 1.13) or moral-emotional language on diffusion (IRR = 0.95, *n.s.*, 95% CI = 0.72, 1.28). However, there was a significant negative effect of distinctly moral language on diffusion (IRR = 0.67, $p < 0.001$, 95% CI = 0.55, 0.82). This finding also contradicts that of Brady et al. (2017) and suggests that both emotional and moral contagion effects are domain specific. Likelihood ratio tests also indicated that a univariate model with *only* moral language was a better predictor of diffusion within the dataset (Table 1).

Viral 2016 US Election Tweets The 8,243 analyzable viral tweets from the 2016 US Presidential Election were found to have an average of 0.17 moral-emotional, 0.35 moral, and 0.98 emotional words per tweet. Analysis showed that there was indeed a small effect of emotional contagion (IRR = 1.05, $p < 0.001$, 95% CI = 1.03, 1.07), whereas the association with distinctly moral language decreased message diffusion (IRR = 0.96, $p < 0.01$, 95% CI = 0.93, 0.99). These findings set up what would have been an ideal case for emotion to drive the diffusion of moral content, however the regression model showed that there was no significant relationship between moral-emotional language and message diffusion (IRR = 1.02, *n.s.*, 95% CI = 0.98, 1.06), despite hinting at an association in the expected positive direction. Nevertheless, likelihood ratio tests indicated that the multivariate model was the most appropriate predictor of the dataset, having outperformed each of the possible nested univariate models (Table 1).

***#MuellerReport* Tweets** In 41,505 unique analyzable tweets from the *#MuellerReport* corpus, an average of 0.18 moral-emotional, 0.47 moral, and 1.25 emotional words per message was found—the highest level of distinctly moral and distinctly emotional language of all datasets. Interestingly, a textbook moral contagion effect as per Brady et al. (2017) was found here. The negative binomial regression model showed that there was a significant effect of emotional contagion (IRR = 1.07, $p < 0.001$, 95% CI = 1.04, 1.09), and that the association between distinctly moral language and message diffusion was not statistically significant (IRR = 1.05, *n.s.*, 95% CI = 1.00, 1.11). And most importantly, there was a significant relationship between moral-emotional language and diffusion (IRR = 1.33, $p < 0.001$, 95% CI = 1.23, 1.46). This effect is even stronger than that of Brady et al. (2017), suggesting that the presence of moral-emotional language can increase a message's diffusion by 33%. Likelihood ratio tests supported the main multivariate model as the best explanation of the dataset (Table 1).

Aggregated Data Finally, we sought to rule out that the observed differences to Brady et al. (2017) were due to non-content differences in the samples, such as sample size. On top of analyzing each individual dataset, an aggregate dataset was compiled from the *#MeToo*, *#WomensMarch*, Post-Brexit, Viral 2016 US Election, and *#MuellerReport* datasets. This was done in an effort to present an analysis of a novel corpus that is similar in size to that addressed by Brady et al. (2017). However, it should be noted that statistics are skewed toward the *#MeToo* dataset as it is significantly larger than the others, comprising 71% of the aggregated data. Like the Brady et al. (2017) dataset, which captured discourse around multiple topics, this compilation of Twitter data also reaches a diverse range of contentious topics, as well as, in theory, a diverse range of individual Twitter users. This aggregation of five datasets into a single corpus (n = 212,227) displayed and average of 0.20 moral-emotional, 0.33 moral, and 1.05 emotional words per tweet. Analysis here showed that neither moral-emotional language (IRR = 0.90, $p < 0.001$, 95% CI = 0.87, 0.94), nor moral language (IRR = 1.00, *n.s.,* 95% CI = 0.97, 1.03), nor emotional language (IRR = 0.99, *n.s.,* 95% CI = 0.97, 1.00) predicted an *increase* in message diffusion. However, moral-emotional language was the only key variable that displayed a significant association with diffusion. This finding is reiterated by likelihood ratio tests, which showed that the multivariate model was slightly outperformed by a nested univariate model that used moral-emotional language only (Table 1). Further analysis with larger datasets, and examinations of specific moral-emotions (e.g., positive vs. negative affect; high- versus low-arousal) is planned for future studies in order to explore possible explanations.

Table 3: Likelihood ratio test statistics of deviance for goodness of fit. Significance indicates that the multivariate model (composed of moral-emotional language, distinctly moral language, *and* distinctly emotional language predicting diffusion) is the better fit for the dataset than the respective univariate model (composed of moral-emotional language, distinctly moral language, *or* distinctly emotional language predicting diffusion). The "Aggregate" column refers to the combined data from the *#MeToo*, *#WomensMarch*, Post-Brexit, Viral 2016 US Election, and *#MuellerReport* datasets. Significance codes: '***' < 0.001 '**' < 0.01 '*' < 0.05.

Univariate model	Dataset						
	Brady et al. (2017)	#MeToo	#Womens March	Post-Brexit	Viral US 2016 Election	#Mueller Report	Aggregate
Moral-emotional language only	30.07***	191.50***	13.06**	1.50	43.70***	27.73***	2.56
Distinctly moral language only	446.74***	212.59***	0.20	3.52	38.92***	86.72***	30.63***
Distinctly emotional language only	432.70***	43.62***	14.07**	2.12	9.62**	62.47***	26.85***

Discussion

Our results suggest that moral contagion driven by moral-emotional language is not as general a phenomenon as Brady et al. (2017) propose. In fact, the statistical models displayed no noteworthy effects of moral contagion in four of the six observational datasets analyzed. While the significant results of the likelihood ratio tests (Table 1) effectively link the use of moral, emotional, and moral-emotional language with information diffusion in most cases, the domain specificity of certain contagion effects in our results spurs a series of conceptual and methodological considerations.

Invoking morality in reasoning is known to harden existing belief structures, delegitimize authority, and, in extreme cases, dehumanize opposing perspectives (Ben-Nun Bloom & Levitan, 2011; Crockett, 2017). While morality can of course be a force for good—providing shared identities and guiding ethical behavior—the introduction of *unnecessary* morality and its emotional underpinnings can jeopardize rational debate. It is for this reason that moral justifications carry weight in some domains but not others. For example, loading an argument with moral-emotional language might be an effective strategy in discourse pertaining to human rights, yet that same strategy is likely to be penalized in an argument over mathematics. Sentiments about where morality is appropriate may be changing, and this may very well be a factor driving ideological polarization. But it seems unlikely a priori that moral language will be viewed the same in all domains. Our results are in keeping with such considerations.

There are also a number of methodological issues that potentially restrict the generality of studies such as this. Perhaps most conspicuous is the inability to parse true *causal* contagion from network homophily. The observational data used here and in Brady et al. (2017) fails to distinguish actual contagion (where exposure to a "contagious" condition has a causal effect on an individual's shift from state A to state B) from manifested homogeneity (where individuals with similar characteristics act in similar ways, irrespective of conditional exposure). It could be argued that the act of retweeting or sharing a message is a behavioral metric because it requires some motivated action. However, Brady et al. (2017) note that where moral contagion has been documented, it has been "bounded by [ideological] group membership" (Brady et al., 2017, p. 7313). This makes it important for future research to heed the substantial body of literature concerning the homophily-contagion problem (e.g., Aral, Muchnik, & Sundararajan, 2009; Shalizi & Thomas, 2011). Plus, Dehghani et al. (2016) specifically show that expressions of moral purity can predict the distance between users on Twitter, which further suggests that moral contagion may simply be an inadvertent measure of moral homophily. Along similar lines, the measurement of diffusion is also an imperfect operationalization of social influence. While collapsing repeated messages into a single observation ensures the language of a single repeated message does not skew analysis, it effectively penalizes unconventional retweeting (e.g., paraphrasing a message's content rather than clicking "retweet"), and is prone to overlook retweet chains (e.g., retweets of retweets) that might indicate true virality of a message (Brady et al., 2017). Crucially though, this imperfection applies to every dataset in the present study such that it cannot explain the discrepancies between datasets.

Needless to say, the use of social media analytics for investigations of the broader human condition has limitations with respect to external validity and representativeness (Tufekci, 2014). It is entirely possible that findings from studies conducted solely in the Twitterverse are in fact unique to the Twitterverse. Plus, the nature of Twitter metadata and correlational analyses like regression modelling mean that network agent variables (those pertaining to qualities of individual nodes/people) and structural variables like network topology are easily conflated. It may still be that human beings are susceptible to moral-emotionally-framed messages (Brady et al., 2017), but that unseen confounds, especially differences in network topology, can undermine contagion effects. Though the reverse is also possible, namely that contingent effects of topology may masquerade as a preference for moral-emotional language. Either way, the findings presented here demonstrate the need for a close partnership between descriptive accounts of "big data" analytics and controlled experimentation in order to draw confident conclusions about social rationality in the digitalized age.

Conclusion

Human reasoning is rarely, if ever, fully autonomous. We depend on our social environments for information and corroboration, and as these environments undergo digitalization, understanding how their evolution translates into new modes of influence is imperative for safeguarding spaces of rational debate. With high-profile papers (e.g., Brady et al., 2017) already pointing out concerning dynamics like moral contagion in real-world social networks, the present paper adds to this line of inquiry by illustrating the inconsistencies of such findings and offering theoretical and methodological explanations. Importantly, the results here indicate that given the diversity of naturally-occurring social networks, predicting the diffusion of information requires investigations of not only properties of the information itself, but also the domain specific topology of the networks through which it travels. Despite the limitations of current computational social science research, it is safe to say that exploring digital discourse can provide valuable insight into the state of human reasoning and argumentation in a time that has been labelled "post-truth."

References

Adhokshaja, P. (2017). #Inauguration and #WomensMarch. *Data.World*, (Version: a1689ca5). Retrieved from https://data.world/adhokshaja/inauguration-and-

womensmarch
Allport, F. H. (1920). The influence of the group upon association and thought. *Journal of Experimental Psychology*, *3*(3), 159–182. https://doi.org/http://dx.doi.org/10.1037/h0067891
Amador, J., Oehmichen, A., & Molina-Solana, M. (2017). Fakenews on 2016 US elections viral tweets (November 2016 - March 2017). *Zenodo*. https://doi.org/). http://dx.doi.org/10.5281/zenodo.1048826
Aral, S., Muchnik, L., & Sundararajan, A. (2009). Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks. *Proceedings of the National Academy of Sciences*, *106*(51), 21544–21549. https://doi.org/10.1073/pnas.0908800106
Ben-Nun Bloom, P., & Levitan, L. C. (2011). We're Closer than I Thought: Social Network Heterogeneity, Morality, and Political Persuasion. *Political Psychology*, *32*(4), 643–665. https://doi.org/10.1111/j.1467-9221.2011.00826.x
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. https://doi.org/10.1073/pnas.1618923114
Christakis, N. A., & Fowler, J. H. (2007). The Spread of Obesity in a Large Social Network over 32 Years. *New England Journal of Medicine*, *357*(4), 370–379. https://doi.org/10.1056/NEJMsa066082
Christakis, N. A., & Fowler, J. H. (2008). The Collective Dynamics of Smoking in a Large Social Network. *New England Journal of Medicine*, *358*(21), 2249–2258. https://doi.org/10.1056/NEJMsa0706154
Coviello, L., Sohn, Y., Kramer, A. D. I., Marlow, C., Franceschetti, M., Christakis, N. A., & Fowler, J. H. (2014). Detecting Emotional Contagion in Massive Social Networks, *9*(3), 1–6. https://doi.org/10.1371/journal.pone.0090315
Crockett, M. J. (2017). Moral outrage in the digital age. *Nature Human Behaviour*, 1–3. https://doi.org/10.1038/s41562-017-0213-3
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. https://doi.org/http://dx.doi.org/10.1037/xge0000139 This
Ferrara, E., & Yang, Z. (2015). Measuring emotional contagion in social media. *PLoS ONE*, *10*(11), 1–14. https://doi.org/10.1371/journal.pone.0142390
Fowler, J. H., & Christakis, N. A. (2009). Dynamic spread of happiness in a large social network: Longitudinal analysis over 20 years in the Framingham Heart Study. *BMJ (Online)*, *338*(7685), 23–26. https://doi.org/10.1136/bmj.a2338
Fowler, J. H., & Christakis, N. A. (2010). Cooperative Behavior Cascades in Human Social Networks. *Proceedings of the Human Factors and Ergonomics Society*, *107*(12), 5334–5338. https://doi.org/10.1073/pnas.0913149107
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), 1029–1046. https://doi.org/10.1037/a0015141
Haidt, J. (2001). The Emotional Dog and its Rational Tail: A Social Intuitionist Approach to Moral Judgment. *Psychological Review*, *108*, 814–834. https://doi.org/10.1016/j.bbrc.2005.09.115
Haidt, J. (2003). The Moral Emotions. In R. J. Davidson, K. R. Scherer, & H. H. Goldsmith (Eds.), *Handbook of affective sciences* (pp. 852–870). Oxford: Oxford University Press.
Kramer, A. D. I., Guillory, J. E., & Hancock, J. T. (2014). Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences*, *111*(24), 8788–8790. https://doi.org/10.1073/pnas.1412469111
Parker, C. (2017). Brexit Tweets from the morning of it's announcement. *Mendeley Data*, *v2*. https://doi.org/10.17632/x9wkrghz23.2
Redl, F. (1949). The phenomenon of contagion and "shock effect" in group therapy. In *Searchlights on delinquency* (K.R. Eissl, pp. 315–328). New York: International Universities Press.
Shalizi, C. R., & Thomas, A. C. (2011). Homophily and Contagion Are Generically Confounded in Observational Social Network Studies. *Sociological Methods & Research*, *40*(2), 211–239. https://doi.org/10.1177/0049124111404820
Thomson-DeVeaux, A. (2019). The Politics Surrounding Mueller Have Changed A Lot SInce He Started.
Tufekci, Z. (2014). Big Questions for Social Media Big Data: Respresentativeness, Validity and Other Methodological Pitfalls. In *Proceedings of the Eighth International AAAI Conference on Weblogs and Social Media Big* (pp. 505–514). https://doi.org/10.1016/0022-5193(78)90170-4
Turner, A. (2018). 390,000 #MeToo Tweets. *Data.World*, (Version: 2249b175). Retrieved from https://data.world/balexturner/390-000-metoo-tweets