***This paper has not yet been peer-reviewed*

***SI Appendix is available [here](#)*

**Estimating the effect size of moral contagion in online networks:
A pre-registered replication and meta-analysis**

William J. Brady^{1*} & Jay J. Van Bavel²

¹Yale University, Psychology Department

²New York University, Psychology Department

Author Note

*Correspondence concerning this article should be addressed to William Brady,
Department of Psychology, Yale University, william.brady@yale.edu; Jay Van Bavel
Department of Psychology, New York University, jay.vanbavel@nyu.edu.

Contributions

WJB and JJVB designed research; WJB collected and analyzed data with input from JJVB; WJB wrote the paper and JJVB and WJB both contributed to revisions.

Abstract

Over 4 billion people now use social media platforms. As our social lives become more entangled than ever before with online social networks, it is important to understand the dynamics of online information diffusion. This is particularly true for the political domain, as political elites, disinformation profiteers and social activists all utilize social media to gain influence by spreading information. Recent work found that emotional expressions related to the domain of morality (*moral emotion* expression) are associated with increased diffusion of political messages--a phenomenon we called '*moral contagion*'. Here, we perform a large, pre-registered direct replication ($N = 849,266$) of Brady et al. (2017), as well as a meta-analysis of all available data testing moral contagion (5 independent labs, 27 studies, $N = 4,821,006$). The estimate of moral contagion in the available population of studies is positive and significant ($IRR = 1.12$, 95% CI = [1.06, 1.19]), such that each message is 12% more likely to be shared for each additional moral-emotional word. The mean effect size of the large, pre-registered replication ($IRR = 1.15$) better estimated the effect size of the available population of studies than the original study ($IRR = 1.20$). These findings reinforce the importance of replication and producing a pre-registered analysis to generate accurate estimates of effect size for future studies.

Keywords: social media, morality, emotions, meta-analysis, replication

Estimating the effect size of moral contagion in online networks

In the digital age, more people than ever before communicate about politics, morality and news via online social networks (Shearer, 2018). Social media once promised to strengthen democracy and egalitarianism, yet we have also seen evidence suggesting platforms may contribute to the rise of political disinformation (Lazer et al., 2018), affective polarization (Finkel et al., 2020), hate speech and harassment (Silva et al., 2016), and organization of political violence (McEvoy, 2021). Examining the psychological processes that underlie these phenomena is key for understanding when and why online political conversations take these forms.

Recent work found that emotional expressions related to the domain of morality (*moral emotion* expression) are associated with increased diffusion of political messages in online social networks – a phenomenon we called '*moral contagion*'. The spread of moral emotions – particularly negative moral emotions like moral outrage – is linked to increased group identification and worsened intergroup relations (Brady et al., 2017; Brady & Van Bavel, 2021) as well as social pressure to express more negative emotions (Brady et al., 2020, 2021). Furthermore, moralized rhetoric online can spill over into offline violence (Mooijman et al., 2018). On the other hand, moral emotion expression online may also help raise awareness of moral issues and spark collective action (Brady & Crockett, 2019; Crockett, 2017; Spring et al., 2018). Given the importance of moral contagion in contentious political conversations online, there is a need to understand the extent to which the general phenomenon replicates in samples outside of the original sample used by Brady et al. (2017).

More generally, it is important to replicate empirical findings since several large replication projects have found that many original studies in psychology either did not replicate or overestimated the effect size of the phenomenon of interest (Klein et al., 2014, 2018; Open Science Collaboration, 2015). Relatedly, other work has questioned the generalizability of psychological findings (Judd et al., 2012; Yarkoni, 2019). This is especially likely to be the case in domains where the topics under investigation are sensitive to the social context (Van Bavel et al., 2016), such as political and moral psychology (Harris et al., 2021). In light of these considerations, there is a growing movement to conduct large pre-registered studies to determine replicability and provide more precise effect size estimates (Munafò et al., 2017). There are also now best practices for conducting replication studies to better estimate effect sizes (Brandt et al., 2014; Nature, 2018) and we sought to apply these principles to the phenomenon of moral contagion.

Current research

In the current paper, we estimated the effect size of moral contagion and assessed its generalizability using two approaches. First, we conducted a large, pre-registered direct replication of prior work ($N = 849,266$) spanning real social media conversations about 5 contentious political topics over a number of years in U.S. politics. Second, we searched the *PsychINFO* citation database and the Open Science Framework database (using the terms “moral contagion” and “moral-emotional language”) to perform a meta-analysis of all published and unpublished replications of the original Brady et al. (2017) studies that we were able to locate as of April, 2021. Our

meta-analysis includes 4,821,006 observations spanning 27 studies from 5 independent research labs.

The results of our pre-registered direct replication as well as the meta-analysis of existing datasets find that the estimate of moral contagion in the available population of studies is positive and significant, showcasing that it generally replicates outside of the original Brady et al. (2017) sample. We also find that the large, pre-registered replication better estimated the effect size of the available population of studies than the original study, and that heterogeneity exists in the effect size across studies. These findings highlight the importance of using replication and a pre-registered analysis plan when aiming to generate accurate estimates of effect size.

A pre-registered direct replication of Brady et al. (2017)

We first conducted a pre-registered, large-scale direct replication of the moral contagion effect ($N = 849,266$; see <https://osf.io/gd734>). We closely followed the guidelines and standards widely recommended by replicability advocates to increase the precision of the effect size estimate and provide a rigorous test of generalizability (Brandt et al., 2014; Nature, 2018). These standards include performing a direct replication, using a larger sample size than the original study, and utilizing pre-registration to guard against questionable research practices.

To perform the direct replication, we focused on five contentious political topics from recent years in U.S. politics that we expected to show evidence of moral contagion (Justice Kavanaugh Supreme Court hearings, the government shutdown over the border wall, the Trump administration's ban of transgendered individuals in the military,

the killing of George Floyd, and mask wearing during the COVID-19 pandemic). We then connected to Twitter's premium API with the goal of collecting 200,000 tweets from each of the 5 topics at specific dates (when the topic of interest was at its peak controversy in U.S. politics). The topics, their start dates, the data collection method, and statistical models were all pre-registered at the OSF link above before collection began (See **Table 1** for all topics and dates of the collection periods.) We note that for two of our datasets, 200,000 tweets could not be collected because there were not enough tweets available from Twitter's API at the exact date we chose before we collected the data.

Political Topic	Date	<i>N</i>	Description
Brett Kavanaugh Supreme Court Hearings	September 16, 2018	195,567	Hearings are held to address sexual assault allegations against Supreme Court nominee Brett Kavanaugh
President Trump Opts to Keep the Government Shutdown	January 16th, 2019	200,000	President Trump keeps the government shut down as leverage to build a wall at the U.S.-Mexico border
Transgender Ban in the Military Upheld by Supreme Court	January 23, 2019	53,649	Supreme court revives a ban on transgender individuals in the military
George Floyd Death	March 27, 2020	200,000	Protests erupt as a video is released showing George Floyd killed by a police officer

COVID-19 mask wearing	May 27, 2020	200,000	Conversations about wearing face masks two months after COVID-19 hit the U.S.
-----------------------	--------------	---------	---

Table 1. Data description for the pre-registered replication. Topics, start date for data collection, sample sizes (N), and description of topics used in the pre-registered replication of Brady et al. (2017).

To analyze each dataset, we followed the analysis pipeline of Brady et al. (2017) directly based on code from its OSF page: <https://osf.io/59uyz/>. This included preprocessing steps to prepare tweets for sentiment analysis and wrangling key covariates from meta-data. Our outcome variable was the retweet count for each tweet (measure of contagion), which was provided as metadata from Twitter's API. We also note that because we collected historical data from Twitter's Premium API, our retweet count variable is more accurate than previous studies, since collecting tweets in real-time does not account for retweets that occur outside the window of data collection.

To form our three word-category predictor variables, we computed the frequency of distinctly moral, distinctly emotional, and moral-emotional words present in each tweet to determine how these factors predicted retweet count. For each dataset, our statistical models included our three word-category predictor variables, as well as covariates known to affect retweet rate independent of message content (Brady et al., 2017; Stieglitz & Dang-Xuan, 2012), including number of Twitter followers the original message author had, whether media or a URL was attached to the message, and whether the message author was a "verified" Twitter user. We fit a negative-binomial model with maximum likelihood estimation to account for overdispersion in retweet counts (Hilbe, 2011).

Analyzing each dataset revealed a positive and significant association between moral-emotional language and message diffusion for 4 out of 5 datasets: Kavanaugh: IRR (Incident Rate Ratio) = 1.05; Trump wall: IRR = 1.07; Transgender Ban: IRR = 1.66; COVID-19: IRR = 1.12, all $ps < .001$. The George Floyd dataset was the exception with a negative association: IRR = 0.97, $p < .001$. The mean effect size across all datasets was IRR = 1.15. This effect size is slightly smaller than that estimated by Brady et al. (2017), where the mean effect size was IRR = 1.20. In order to determine whether the replication vs. the original studies better estimated the effect size of moral contagion in the broader population of studies, we next conducted a meta-analysis of all available published and unpublished data from our lab and other labs.

A meta-analysis of all available replications of Brady et al. (2017)

To provide a robust estimate of the effect size of moral contagion, we pulled all available published and unpublished data by searching *PsychINFO*'s database for articles using the keywords "moral contagion" and "moral-emotional language" and then selected papers that performed a replication of the main findings of the paper by reading the abstract. Next, we searched OSF's database of studies using the same terms and selected replications of Brady et al. (2017) for access to data, where required. Studies that were included either had their data available on OSF to reproduce or listed adequate information in their paper for us to perform the meta-analysis (i.e., the effect size and dispersion for each study was reported in results). Overall, we were able to collect data from 5 independent labs that have examined the association between moral-emotional words and retweet counts using the same method as Brady et al.

(2017). The final meta-analysis included 27 studies consisting of total $N = 4,821,006$, see **Table 2**.

Paper	N Studies	Total N (analyzed)	Publication
Ariztia (2020)	1	25,560	Unpublished
Brady et al. (2017)	3	313,003	<i>Proceedings of the National Academy of Sciences</i>
Brady et al. (2019)	2	286,255	<i>Journal of Experimental Psychology: General</i>
Pre-registered Replication	5	849,266	Current manuscript
Burton, Cruz & Hahn (2019)*	5	207,622	<i>Cognitive Science Society Proceedings</i>
Jordan (2019)	3	188,129	Unpublished
Rathje, Van Bavel & Van der Linden (2021)	8	2,951,322	Unpublished

Table 2. Studies included in meta-analysis of moral contagion. Table lists authors,

date of publication or posting of paper if unpublished, number of studies, the total number of observations across studies, and where the paper was published (or whether it was unpublished). *Burton et al. (2019) did not include covariates known to impact retweet count, and thus effect sizes were estimated using models that only contained the three word-category predictors.

The statistical model used for all studies in the meta-analysis was the same as Brady et al. (2017) and our replication study above, with the exception of one lab that did not collect covariates for their analysis (see **Table 2**). For those studies, we used the reported effect size and standard errors from the model that did not include covariates.

We performed the meta-analysis using a generic inverse-variance weighting method as implemented in the package *meta* (Schwarzer, 2007). As a more conservative estimate of the meta-analytic effect as well as to improve estimates of generalizability, we estimated a random effects model (Judd et al., 2012; Yarkoni, 2019) using the Hartung-Knapp-Sidik-Jonkman method (IntHout et al., 2014). We present the meta-analytic estimate of moral contagion in **Fig. 1**.

The meta-analytic estimate of moral contagion is positive and significant when all available data to date are taken into account. Results of the random effects model estimated an effect size of $IRR = 1.13$, 95% CI = [1.06, 1.20]. It is interesting to note that this meta-analytic estimate of effect size was smaller than Brady et al. (2017)'s original estimate ($IRR = 1.20$), but closer to our pre-registered replication study above ($IRR = 1.15$). Overall, however, the meta-analysis found that the results of Brady et al. (2017) generalize out of its original sample to various political topics, lab groups, and social media platforms.

In order to examine the datasets for evidence of publication bias, we conducted Egger's intercept test (Egger et al., 1997), which is a statistical test for asymmetry in the

relationship between study precision and study effect size. Egger's intercept test found that the pattern of results in our meta-analysis were not indicative of publication bias, intercept = 3.049, $p = .171$. Similarly, **Fig. 2** demonstrates that there is no correlation between precision of a study and the estimated effect size (potential publication bias would be indicated by lower standard errors being associated with greater effect sizes; Egger et al., 1997). Furthermore, a p-curve analysis of all studies in the meta-analysis revealed that evidential value for moral contagion was present, see SI Appendix, Fig. S1.

Fig. 1. Meta-analysis of studies examining moral contagion in online networks. A meta-analysis of the moral contagion phenomenon using a generic inverse-variance

weighting method with a random effects model. The effect size for moral contagion (the association between moral-emotional language and message diffusion) showed a significant, positive effect across 27 studies, $N = 4,821,006$. Effect size represents an incident rate ratio produced from negative binomial regression models. The dotted line represents an IRR of 1, which indicates a zero effect. Values greater than 1 represent a positive association between moral-emotional language and retweet counts, values less than 1 indicate a negative association.

Fig. 2. Plotting the relationship between study precision and effect size. Each observation represents a study, where the Y axis represents the standard error associated with the moral contagion estimate, and X axis represents the effect size in terms of the incident rate ratio (IRR). The red line highlights an IRR of 1.0, which indicates a zero effect. In the meta-analysis, lower study precision was not systematically related to effect size, which suggests that publication bias is not a contributor to the meta-analytic estimate.

Discussion

Our large, pre-registered replication and meta-analysis provides robust evidence of moral contagion in online social networks. In our replication, we found clear evidence that moral emotional is associated with increased sharing across multiple political topics. We did find one case where this pattern was not replicated: when people were discussing the killing of George Floyd (an event that precipitated massive international protests in support of Black Lives Matter). The pattern in our replication studies was also replicated in a meta-analysis of 27 studies that included over 4,821,006 messages. Together, these results provides strong evidence that the association between moral-emotional language and sharing replicates out of the original Brady et al. (2017) sample. Besides using a random effects model in our meta-analysis, the studies we combined also demonstrate the generalizability of moral contagion: a significant effect was found for various political topics, across multiple social media platforms (Twitter and Facebook), and among ordinary users as well as political elites.

Our replication and meta-analysis also provided evidence of heterogeneity in moral contagion as some studies produced negative coefficients or null findings. There are several reasons that can explain this heterogeneity, such as sampling error, real contextual features of the topic or social network that serve as moderators for moral contagion, or issues with the replication strategy and analytic decisions of the researchers. For example, Burton et al. (2019) did not collect covariates included in the models of Brady et al. (2017) which limits the ability to accurately estimate the effect size in these datasets. The heterogeneity we discovered highlights the importance of using pre-registration and identifying concrete boundary conditions for a phenomenon to

help guide future replications and extensions of the research (for an example applied to moral contagion, see Brady et al., 2020).

Our results demonstrate how using big data to test psychological hypotheses affords researchers the opportunity to conduct large-scale replications that help estimate effect size with better precision. Although our results underscore the general stability of the findings of Brady et al. (2017), it is still important to acknowledge several limitations with the methodological approach taken by Brady et al. (2017); most importantly the fact that word-count methods can produce noisy results and the data were purely correlational. These limitations point to the fact that meta-analyses following one method are not enough; there is a need to examine psychological phenomena with various methodological approaches *alongside* directly replicating the broader finding. For example, in the case of moral contagion, the finding now been replicated using advances in natural language processing that measure moral emotion expressions with much greater precision (Brady et al., 2021). Furthermore, carefully-controlled behavioral experiments have found that moral-emotional language causes increase in political message sharing (Brady & Van Bavel, 2021). Large-scale direct replications combined with multi-method approaches can help assess the robustness of psychological findings in the era of big data.

Conclusion

In sum, the current research reveals that the findings of Brady et al. (2017) replicate in multiple out-of-sample datasets and highlights the general robustness of the moral contagion phenomenon. Our estimate of the effect size from a large, pre-registered direct replication estimated the population effect size more precisely than the

original study, which reinforces the importance of replicating work following pre-registration standards to gain the most accurate expectations of effect sizes in future work. It appears that moral emotions play an important role in the diffusion of political messages across numerous words, topics, authors, and social media platforms.

Acknowledgements

We thank Molly Crockett for feedback on previous drafts.

References

- Ariztía, F. D. P. (2020). Contagio moral en redes sociales: cómo la emoción moral transmite mensajes y polariza diferentes tipos de conversaciones en Twitter.
- Brady, W. J., & Crockett, M. J. (2019). How Effective Is Online Outrage? *Trends in Cognitive Sciences*, 23(2), 79–80. <https://doi.org/10.1016/j.tics.2018.11.004>
- Brady, W. J., Crockett, M. J., & Van Bavel, J. J. (2020). The MAD Model of Moral Contagion: The Role of Motivation, Attention, and Design in the Spread of Moralized Content Online. *Perspectives on Psychological Science*, 15(4), 978–1010. <https://doi.org/10.1177/1745691620917336>
- Brady, W. J., McLoughlin, K., Doan, T. N., & Crockett, M. (2021). *How social learning amplifies moral outrage expression in online social networks*. PsyArXiv. <https://doi.org/10.31234/osf.io/gf7t5>

- Brady, W. J., & Van Bavel, J. J. (2021). *Social identity shapes antecedents and functional outcomes of moral emotion expression in online networks*. OSF Preprints. <https://doi.org/10.31219/osf.io/dgt6u>
- 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>
- Brandt, M. J., IJzerman, H., Dijksterhuis, A., Farach, F. J., Geller, J., Giner-Sorolla, R., Grange, J. A., Perugini, M., Spies, J. R., & van 't Veer, A. (2014). The Replication Recipe: What makes for a convincing replication? *Journal of Experimental Social Psychology*, 50, 217–224. <https://doi.org/10.1016/j.jesp.2013.10.005>
- Crockett, M. J. (2017). Moral outrage in the digital age. *Nature Human Behaviour*, 1(11), 769–771. <https://doi.org/10.1038/s41562-017-0213-3>
- Egger, M., Davey Smith, G., Schneider, M., & Minder, C. (1997). Bias in meta-analysis detected by a simple, graphical test. *BMJ (Clinical Research Ed.)*, 315(7109), 629–634. <https://doi.org/10.1136/bmj.315.7109.629>
- Finkel, E. J., Bail, C. A., Cikara, M., Ditto, P. H., Iyengar, S., Klar, S., Mason, L., McGrath, M. C., Nyhan, B., Rand, D. G., Skitka, L. J., Tucker, J. A., Bavel, J. J. V., Wang, C. S., & Druckman, J. N. (2020). Political sectarianism in America. *Science*, 370(6516), 533–536. <https://doi.org/10.1126/science.abe1715>
- Harris, E. A., Pärnamets, P., Brady, W. J., Robertson, C., & Bavel, J. J. V. (2021). *The role of generalizability in moral and political psychology*. OSF Preprints. <https://doi.org/10.31219/osf.io/nfjct>

Hilbe, J. M. (2011). *Negative Binomial Regression*. Cambridge University Press.

IntHout, J., Ioannidis, J. P., & Borm, G. F. (2014). The Hartung-Knapp-Sidik-Jonkman method for random effects meta-analysis is straightforward and considerably outperforms the standard DerSimonian-Laird method. *BMC Medical Research Methodology*, 14(1), 25. <https://doi.org/10.1186/1471-2288-14-25>

Judd, C. M., Westfall, J., & Kenny, D. A. (2012). Treating stimuli as a random factor in social psychology: A new and comprehensive solution to a pervasive but largely ignored problem. *Journal of Personality and Social Psychology*, 103(1), 54–69. <https://doi.org/10.1037/a0028347>

Klein, R., Ratliff, K., Vianello, M., Jr, R. A., Bahník, S., Bernstein, M., Bocian, K., Brandt, M., Brooks, B., Brumbaugh, C., Cemalcilar, Z., Chandler, J., Cheong, W., Davis, W., Devos, T., Eisner, M., Frankowska, N., Furrow, D., Galliani, E., ... Nosek, B. (2014). Data from Investigating Variation in Replicability: A “Many Labs” Replication Project. *Journal of Open Psychology Data*, 2(1), e4. <https://doi.org/10.5334/jopd.ad>

Klein, R., Vianello, M., Hasselman, F., Adams, B. G., Adams, R. B., Alper, S., Aveyard, M., Axt, J. R., Babalola, M. T., Bahník, Š., Batra, R., Berkics, M., Bernstein, M. J., Berry, D. R., Bialobrzeska, O., Binan, E. D., Bocian, K., Brandt, M. J., Busching, R., ... Nosek, B. A. (2018). Many Labs 2: Investigating Variation in Replicability Across Samples and Settings. *Advances in Methods and Practices in Psychological Science*, 1(4), 443–490. <https://doi.org/10.1177/2515245918810225>

- Lazer, D. M. J., Baum, M. A., Benkler, Y., Berinsky, A. J., Greenhill, K. M., Menczer, F., Metzger, M. J., Nyhan, B., Pennycook, G., Rothschild, D., Schudson, M., Sloman, S. A., Sunstein, C. R., Thorson, E. A., Watts, D. J., & Zittrain, J. L. (2018). The science of fake news. *Science*, 359(6380), 1094–1096.
<https://doi.org/10.1126/science.aao2998>
- McEvoy, J. (2021). *Capitol Attack Was Planned Openly Online For Weeks—Police Still Weren't Ready*. Forbes.
<https://www.forbes.com/sites/jemimamcevoy/2021/01/07/capitol-attack-was-planned-openly-online-for-weeks-police-still-werent-ready/>
- Mooijman, M., Hoover, J., Lin, Y., Ji, H., & Dehghani, M. (2018). Moralization in social networks and the emergence of violence during protests. *Nature Human Behaviour*, 2(6), 389–396. <https://doi.org/10.1038/s41562-018-0353-0>
- Munafò, M. R., Nosek, B. A., Bishop, D. V. M., Button, K. S., Chambers, C. D., Percie du Sert, N., Simonsohn, U., Wagenmakers, E.-J., Ware, J. J., & Ioannidis, J. P. A. (2017). A manifesto for reproducible science. *Nature Human Behaviour*, 1(1), 1–9. <https://doi.org/10.1038/s41562-016-0021>
- Nature. (2018). Learning from replication. *Nature Human Behaviour*, 2(9), 601–601.
<https://doi.org/10.1038/s41562-018-0441-1>
- Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Science*, 349(6251). <https://doi.org/10.1126/science.aac4716>
- Shearer, E. (2018). *Social media outpaces print newspapers in the U.S. as a news source*. Pew Research Center. <https://www.pewresearch.org/fact->

tank/2018/12/10/social-media-outpaces-print-newspapers-in-the-u-s-as-a-news-source/

Silva, L., Mondal, M., Correa, D., Benevenuto, F., & Weber, I. (2016). Analyzing the Targets of Hate in Online Social Media. *Proceedings of the International AAAI Conference on Web and Social Media*, 10(1), Article 1.

<https://ojs.aaai.org/index.php/ICWSM/article/view/14811>

Spring, V. L., Cameron, C. D., & Cikara, M. (2018). The Upside of Outrage. *Trends in Cognitive Sciences*, 22(12), 1067–1069.

<https://doi.org/10.1016/j.tics.2018.09.006>

Stieglitz, S., & Dang-Xuan, L. (2012). Political Communication and Influence through Microblogging—An Empirical Analysis of Sentiment in Twitter Messages and Retweet Behavior. *2012 45th Hawaii International Conference on System Sciences*, 3500–3509. <https://doi.org/10.1109/HICSS.2012.476>

Van Bavel, J. J., Mende-Siedlecki, P., Brady, W. J., & Reinero, D. A. (2016). Contextual sensitivity in scientific reproducibility. *Proceedings of the National Academy of Sciences*, 113(23), 6454–6459. <https://doi.org/10.1073/pnas.1521897113>

Yarkoni, T. (2019). *The Generalizability Crisis*. PsyArXiv.

<https://doi.org/10.31234/osf.io/jqw35>