A comment on Everett et al. (2020): No evidence for the effectiveness of moral messages on public health behavioural intentions during the COVID-19 pandemic

Farid Anvari

University of Southern Denmark

This is a review of (version 4 of) a preprint titled, “The effectiveness of moral messages on public health behavioral intentions during the COVID-19 pandemic” (Everett et al., 2020). It is commendable that the authors preregistered their study and have made their preregistration, study materials, and data available for scrutiny along with the posting of their preprint. In this spirit, I will evaluate the preprint as presented. Furthermore, after having sent the authors my review on the first version of their preprint, they responded by not only updating their preprint based on the criticisms I presented, but they were extremely amicable and invited me to be a coauthor in their ongoing project in the hopes of producing as high quality research as possible. This second version of my review addresses remaining points of concern that I have on their updated preprint. In reviewing the preprint, after a quick aside about the exploratory analyses, I will mainly focus on the preregistered hypotheses and analyses as they pertain to (intentions) of behavioural change because I consider these the most important aspects of the paper.

First, some comments on the exploratory analyses. In exploratory analyses, where the number of statistical tests conducted are not (and often cannot) be reliably tracked, Type 1 errors (the likelihood of inferring that a true effect exists, when there is no true effect) are not controlled and are likely to be inflated. This applies to all the analyses and conclusions in the preprint that are based on ANCOVAs, using demographic variables as covariates, and the section titled “Results: Utilitarianism”. The preregistration mentions that demographic factors would be

used in exploratory analyses, but as moderators of the experimental effects and to examine their direct relationships with the dependent variables, not as covariates.

Furthermore, and notwithstanding inflated Type 1 error rates, many of the exploratory results are statistically nonsignificant, given the preregistered alpha level of .05—the preregistration states the use of “conventional statistical significance of $p < .05$ ”. With nonsignificant results we can conclude very little, and researchers have warned that when testing against a zero-point null hypothesis, as in the study being reviewed, the only justified conclusion based on a nonsignificant result is that we cannot be confident about the direction of the effect (e.g., Cohen, 1994; Tukey 1991). Taken together, I propose (as the authors themselves note) that the results and conclusions derived from the exploratory analyses should be treated with caution and verified in preregistered tests on new data before being acted upon.

Now I begin my review of the sections of the preprint examining the preregistered hypotheses pertaining to (intentions) of behavioural change. Preregistration does not automatically make a study better or worse, but, rather, it functions to allow “researchers to evaluate the severity with which a claim is tested” (Lakens 2019, p. 5). In this review, I evaluate how well the preregistered hypotheses have been tested, the extent with which we can claim that the preregistered hypotheses have been supported, and whether the conclusions presented in the preprint are supported by the results. First, I will summarize the relevant parts of the preregistration relating to the independent and dependent variables. Next, I present the corresponding results from the preprint followed by my critique.

Farid Anvari, Strategic Organization Design group, Department of Marketing and Management, University of Southern Denmark.

Correspondence concerning this paper should be addressed to Farid Anvari. Email: faridanvari.phd@gmail.com

The Preregistered Variables and Analyses

The preregistration (see here <https://osf.io/wy93v>) gives details about the variables of interest. There are two independent (manipulated) variables: message source (citizen vs leader); and message type (deontological, virtue ethics, utilitarianism, and non-moral). The first three categories of message type are moral justifications. An important limitation of the study, as alluded to by the authors and worth pointing out here, is that only a single item was used for each ethical framework. This means that we have very little information about the generalizability of the results. Because the generalizability of these findings will be extremely important if they are to be applied in practice, a replication study would benefit from increasing the set of messages for each ethical framework to see if the results are generalizable. Ideally, researchers would include a large enough set of stimuli for each message type so that stimuli can be modelled as random factors (see Yarkoni, 2019). In the updated preprint, the authors have acknowledged this limitation.

The measured, or dependent, variables in the preregistration, and of relevance to the preregistered hypotheses, are first-order behavioural intentions (4 items), second-order beliefs (4 items), and person perception (or impressions of messenger; 2 items). Note, however, that the preregistration does not make clear whether the dependent variables would be examined as composite measures (e.g., behavioural intentions measured as a mean of the 4 items) or whether each item would be examined separately. This is one area of vagueness in the preregistration that, in my view, reduces the severity with which the hypotheses were tested, due to leaving room for flexibility in the analyses. Nevertheless, I will treat how the authors presented the results in their paper as what was intended at the preregistration stage: that each item would serve as a dependent variable itself. Thus, there are a total of 10 dependent variables.

The preregistration is sparse on details with respect to the analyses that would be conducted. It merely states that the authors would conduct “a 2x4 between-subjects ANOVA with planned contrasts carried out via independent samples *t*-tests”. I, and I presume most readers would, take this to mean that

for each preregistered dependent variable the authors planned to conduct a 2 (messenger: citizen vs leader) by 4 (message type: deontological vs virtue vs utilitarian vs non-moral) ANOVA. Indeed, this is what they did. Critically, the preregistration states that the authors would “use conventional statistical significance of $p < .05$, and two-tailed tests”.

The Results: Behavioural Intentions and Second Order Beliefs

The first preregistered hypothesis predicted a stronger effect of moral justifications (or moral messages), than non-moral messages, on behavioural intentions. Moreover, the preregistration states that utilitarian justifications will be less effective than deontological or virtue-based justifications. The authors conducted 2 by 4 ANOVAs on each of the 4 items measuring behavioural intentions. Message type had a significant main effect on only 1 measure of behavioural intentions (intention to share the message on social media) with a *p*-value reported as .01 (more precisely it is .0082). Planned contrasts showed that the deontological message resulted in higher intentions to share the message than non-moral messaging ($p = .01$) and virtue-based messaging ($p = .02$).

The second preregistered hypothesis makes the same prediction as the hypothesis about behavioural intentions except that the dependent variable is second-order beliefs (beliefs about what other people would do). Based on 2 by 4 ANOVAs on each of the 4 items measuring second-order beliefs, the authors report a marginally significant main effect of message type on two items (second-order beliefs about handwashing, $p = .07$, and second-order beliefs about avoiding social gatherings, $p = .08$).

Given that 4 ANOVAs were conducted to draw inferences about behavioural intentions, and 4 ANOVAs to draw inferences about second-order beliefs about other people's behaviours, there should have been corrections for multiple comparisons. If we assume that behavioural intentions and second-order beliefs are separate hypotheses and thus involving separate families of tests, then, using the Holm-Bonferroni correction (Holm, 1979; Lakens, 2016), the corrected alpha level for the ANOVA with the lowest *p* value testing each of the hypotheses would be $(.05/4) .0125$, the next lowest *p* value

would be tested against $(.05/3) .0167$, and so on. The significant main effect on the behavioural intention to share the message is still significant at this alpha level. However, the main effects of message type on second-order beliefs become further marginal; more accurately, they remain statistically nonsignificant (which means we cannot reject an effect of zero or, even, some effects in the opposite direction).

Notably, the authors expected ceiling effects due to social desirability for the behavioural intention measures (as stated in both the preprint and the pre-registration). In addition to the possibility that they would be more sensitive to the manipulations, second-order beliefs were measured to address the issue of ceiling effects and socially desirable responding to the behavioural intention items. Indeed, the authors state, “Because we expected potential ceiling effects for behavioral intentions due to the strong social desirability of these behaviors, we also measured participants’ second-order beliefs about the behavioral intentions of others, which we expected would be less sensitive to social desirability effects.” (p. 4). The reasoning implies that the second-order beliefs would be used to draw inferences about which message type would be best for changing people’s intentions—that second-order beliefs and behavioural intentions would both be used to draw inferences about the effectiveness of moral messaging. In other words, it does not matter what people believe about others *per se* but rather what this belief reveals about people’s own (potentially hidden) intentions. Thus, the hypotheses on behavioural intentions and second-order beliefs do not represent a separate family of hypotheses but, rather, tap into the same question: Which ethical framework is most effective for influencing public health behavioural intentions?

Therefore, the 4 items measuring behavioural intentions and the 4 items measuring second-order beliefs can be considered to be of the same family of tests examining which moral framings are most effective for influencing people’s behaviours. This results in a total of 8 dependent variables corresponding to 8 (ANOVA) tests. Applying the Holm-Bonferroni approach to correct for multiple comparisons means that the lowest p -value from the results of these 8 tests should be tested against an alpha of $(.05/8) .00625$, the second lowest p -value against $(.05/7) .0071$, the third lowest p -value against $(.05/6)$

$.0083$, and so on. Based on these corrected alpha levels, even the main effect of message type on the behavioural intention item of sharing the message on social media ($p = .0082$) becomes statistically nonsignificant. These corrections might feel overly strict, but when we are comparing many independent tests and are willing to take any effect as support for a claim, this level of strictness is required to not fool ourselves about the existence of an effect. In the updated (version 4 of the) preprint the authors consider this and refer the reader to the Supplementary Note (Everett et al., 2020, p. 5) where they present their disagreement and argue for applying Holm-Bonferroni corrections for 4 comparisons (Everett et al., 2020, p. 22). They conclude this section of the note with, “the robustness of our findings are better addressed by pre-registered replication studies rather than debates about the appropriate level of stringency for multiple comparisons correction” (p. 22). I tend to agree.

The Conclusions: Behavioural Intentions and Second Order Beliefs

When testing against a null hypothesis of zero, as done in the study being reviewed and as done in psychology more generally, the only conclusion we can draw from a statistically nonsignificant result is that we are unsure about the direction of the effect (Cohen, 1994; Tukey, 1991). However, even though the authors have noted that some hypotheses do not survive corrections for multiple comparisons, the conclusions in the preprint go somewhat beyond what can be justified.

In the discussion section, the authors note that “this study tested which [ethical framework] might be most effective” (p. 12) for changing people’s behaviour to be in line with advice from public health experts. In the following two paragraphs, the authors then report how, relative to non-moral messaging, deontological messaging increased behavioural intentions for sharing the public health message on social media (confirming the preregistered hypothesis), and that exploratory analyses controlling for demographic variables showed that deontological messaging showed “trend-level effects” for behavioural intentions and second-order beliefs for handwashing and second-order beliefs about avoiding social gatherings. Critically, as described earlier,

after correcting for multiple comparisons none of the effects on behavioural intentions or second-order beliefs are statistically significant (which the authors now note). Therefore, the conclusions about the effects of message type on behavioural intentions and second-order beliefs, both with respect to the preregistered hypotheses and exploratory analyses, should go no further than “we cannot be sure about the direction of the effect”.

The authors notably conclude that “Perhaps most importantly, our effects appear to be small and in many cases do not pass conventional levels of statistical significance, which means estimated effect sizes are unreliable and need to be confirmed with future studies” (p. 14). This amendment following my initial review is laudable. Although I believe in justifying alpha levels on a per study basis (see Lakens et al., 2018), the authors should have specified what they considered an acceptable alpha level before, not after, observing the p -values. Moreover, especially when selectively reporting significant results (as the authors do, given that the discussion makes no note of the preregistered hypotheses which were not supported), the importance of controlling erroneous interpretations becomes a non-trivial concern. In other words, a more precise wording would be, “our preregistered hypotheses were statistically not significant based on our preregistered alpha level”.

Furthermore, implicit in the conclusion that I have quoted at the beginning of the previous paragraph, is the assumption that small effects would in some way stack to have a meaningful impact on a population level. This assumption is not certain to hold. The authors could improve on all the above analyses, including the exploratory ones, by prespecifying (in their future research) which effects would likely matter in practice (ideally including arguments for why and how they might accumulate), which would also make nonsignificant results more interpretable through equivalence tests (e.g., Lakens, 2017; Lakens et al., 2018; Rogers et al., 1993).

In sum, after applying Holm-Bonferroni corrections to the (prespecified and preregistered) alpha levels to control for multiple comparisons, none of the preregistered hypotheses for behavioural intentions or second-order beliefs were supported. The main effects of message type for every ANOVA are statistically nonsignificant, despite the reasonably

large sample size. (I recommend that the authors report the sample size per condition, the descriptive statistics, and the effect sizes for the preregistered analyses.)

Conclusion of this Review

Based on my review of the preregistration, the results in the preprint, the conclusions in the preprint, and results following corrections for multiple comparisons, I conclude that the severity with which the hypotheses were tested is lacking and that the conclusions presented in the preprint are stronger than can be supported by the data. I agree with the authors that replication is needed because many of the effects should not be expected to be reliable. Researchers and policymakers should treat the conclusions in the preprint with caution, at least until more severe tests and larger replications have been conducted. Indeed, by implementing interventions based on such results, we may suffer opportunity costs, such as that of using other types of (more effective) messaging, but which would remain unexplored if people act on the conclusions presented in the preprint.

Notes and Acknowledgements

Prior to publishing this comment on PsyArXiv, I emailed the authors with an initial version. The authors responded very kindly, taking many of my comments on board, and provided feedback which resulted in some relatively minor changes to my arguments. In the interests of ensuring that the highest possible quality of science is published to address the current crisis, the authors and I are discussing the potentiality of working together, following an invitation from the authors for me to join their project. Following further (email) discussions with the authors, they laudably updated their preprint (now at version 4) and I subsequently have updated my review (version 2).

Finally, I would like to thank Amy Orben and Daniël Lakens for reading through a draft of this review and providing insightful feedback.

Yarkoni, T. (2019, November 22). The Generalizability Crisis.
<https://doi.org/10.31234/osf.io/jqw35>

References

- Cohen, J. (1994). The earth is round ($p < .05$). *American Psychologist* 49(12), 997-1003. <https://doi.org/10.1037/0003-066x.50.12.1103>
- Everett, J. A. C., Colombatto, C., Chituc, V., Brady, W. J., & Crockett, M. (2020, March 20). The effectiveness of moral messages on public health behavioral intentions during the COVID-19 pandemic. <https://doi.org/10.31234/osf.io/9yqs8>
- Holm, S. 1979. A simple sequential rejective multiple test procedure. *Scandinavian Journal of Statistics* (6), 65-70.
- Lakens, D. (2016, January 1). Error Control in Exploratory ANOVA's: The How and the Why. *The 20% Statistician*. Retrieved from <http://daniellakens.blogspot.com/2016/01/error-control-in-exploratory-anovas-how.html>
- Lakens, D. (2017). Equivalence tests: a practical primer for t tests, correlations, and meta-analyses. *Social psychological and personality science*, 8(4), 355-362. <https://doi.org/10.1177/1948550617697177>
- Lakens, D. (2019, November 18). The Value of Preregistration for Psychological Science: A Conceptual Analysis. <https://doi.org/10.31234/osf.io/jbh4w>
- Lakens, D., Adolphi, F. G., Albers, C. J., Anvari, F., Apps, M. A., Argamon, S. E., ... & Buchanan, E. M. (2018). Justify your alpha. *Nature Human Behaviour*, 2(3), 168-171. <https://doi.org/10.1038/s41562-018-0311-x>
- Lakens, D., Scheel, A. M., & Isager, P. M. (2018). Equivalence testing for psychological research: A tutorial. *Advances in Methods and Practices in Psychological Science*, 1(2), 259-269. <https://doi.org/10.1177/2515245918770963>
- Rogers, J. L., Howard, K. I., & Vessey, J. T. (1993). Using significance tests to evaluate equivalence between two experimental groups. *Psychological Bulletin*, 113(3), 553. <https://doi.org/10.1037/0033-2909.113.3.553>
- Tukey, J. W. (1991). The philosophy of multiple comparisons. *Statistical Science* 6(1), 100-116. <https://doi.org/10.1214/ss/1177011945>