1 Social reference cues can reduce misinformation sharing behaviour on social media

2 Christopher M. Jones^{1,2}, Daniel Diethei^{1,2}, Johannes Schöning^{2,3}, Rehana Shrestha^{1,2}, Tina
3 Jahnel^{1,2}, & Benjamin Schüz^{1,2}

4 ¹ University of Bremen, Germany

5 ² Leibniz ScienceCampus Digital Public Health Bremen, Germany

6 ³ University of St. Gallen, Switzerland

7 Author Note

8 Christopher Jones, Rehana Shrestha, Tina Jahnel and Benjamin Schüz, Institute for
9 Public Health and Nursing Research, University of Bremen, Bremen, Germany. Daniel
10 Diethei, Department for Human-Computer Interaction, University of Bremen, Bremen,
11 Germany. Johannes Schöning, School of Computer Science, University of St. Gallen,
12 St. Gallen, Switzerland.

13 Correspondence concerning this article should be addressed to Christopher M. Jones,
14 Grazer Str. 4, 28359 Bremen, Germany. E-mail: jones@uni-bremen.de

Abstract

Misinformation on social media is a key challenge to effective and timely public health responses. Existing mitigation measures include flagging misinformation or providing links to correct information but have not yet targeted social processes. Here, we examine whether providing balanced social reference cues in addition to flagging misinformation leads to reductions in sharing behavior and improvement of overall sharing quality. In 3 field experiments (N=824, N=322, and N=278) on Twitter, we show that highlighting which content others within the personal network share and, more importantly, not share combined with misinformation flags significantly and meaningfully reduces the amount of misinformation shared and improves sharing quality (Study 1-3). We show that this improvement could be driven by change in injunctive social norms (Study 2) but not social identity (Study 3). Social reference cues, combined with misinformation flags, are feasible and scalable means to effectively curb sharing misinformation on social media.

Keywords:

Word count: 3605

Social reference cues can reduce misinformation sharing behaviour on social media

Introduction

Misleading or fabricated health-related content on social media has developed into a fundamental challenge for public institutions and citizens alike. During the Covid-19 pandemic, an “epidemic of misinformation” (Zarocostas, 2020) has become a major impediment for pandemic control, in particular by spreading and multiplying misinformation on non-pharmaceutical interventions such as masking or social distancing, and by undermining vaccination efforts with fear and uncertainty (Imhoff & Lamberty, 2020). The fact-checking collective poynter.org reports more than 15,000 unique COVID-related misinformation items (Poynter, 2021), which are shared repeatedly or in slight variations to create a multitude of misleading social media posts. This is particularly concerning since misinformation content spreads faster and more expansively than content without false claims (Shahi, Dirkson, & Majchrzak, 2021; Vosoughi, Roy, & Aral, 2018). Ironically, most social media platforms’ architecture and algorithmic design facilitate sharing misinformation content (Charquero-Ballester, Walter, Nissen, & Bechmann, 2021) and the platforms’ own counter measures (e.g., warning labels, professional fact-checking) have proven not suited to be used at scale (Pennycook & Rand, 2021). Developing such effective measures to reduce the sharing of misleading content on social media is thus a paramount issue.

In line with the overall approach by platform providers such as Twitter, previous research has shown that sharing misinformation on social media can be due to knowledge or attention deficits (Pennycook et al., 2021), and that addressing these deficits can reduce the misinformation sharing. However, sharing misinformation (and in fact, any information) on social media is also the result of social processes - in particular, social reinforcements accruing through, for example, “likes” (Meshi, Tamir, & Heekeren, 2015). Features such as “likes” are key elements of social media platforms and facilitate user engagement. Besides a

social reinforcement, receiving and perceiving “likes” to messages on social media also provides the user with normative information about socially acceptable behavior within the group these messages are shared with (van Bavel et al., 2020). It is thus surprising that these core processes of social media itself, which are also evidence-based determinants of sharing behavior, have seldomly been targeted by interventions.

Sharing information in the digital space

So far, the main focus of interventions to reduce the sharing of misinformation has been on information processing. A key assumption underlying these interventions is that engaging with social media happens in an information-rich and attention-seeking environment as well as under time constraints, conditions which limit the cognitive and attentional resources available to judge the veracity and accuracy of social media posts (Pennycook & Rand, 2021). Accordingly, recent studies (Pennycook et al., 2021; Pennycook, McPhetres, Zhang, Lu, & Rand, 2020) show that if users’ attention is directed to the accuracy of social media content, they are more likely to identify misinformation as such and subsequently are less likely to share such content. This suggests that if social media users can be supported to not only correctly judge the accuracy of content but also consider it before sharing, the likelihood of sharing misinformation can be reduced. However, even when users’ attention is directed towards the accuracy of content they are about to share, this behavior still happens within the social media platform, an environment that is built to foster user engagement and sharing using social information (Charquero-Ballester et al., 2021). Social information such as social norms (e.g., information on content shared by other members of a user’s network) and social reinforcement (e.g., tallying of “likes” to own posts) are key elements of the user interface and functionality of most social media platforms, and it has been claimed that this social information might be more salient in guiding online behavior than the information content (Avram, Micallef, Patil, & Menczer, 2020; O’Connor & Weatherall, 2019). This guiding

role of social information is especially relevant when cognitive and attentional resources are scarce (Kozyreva, Lewandowsky, & Hertwig, 2020).

Social processes in sharing behavior

All information shared on social media is also social information, that is, information that shapes the perceptions of the choices, attitudes, and behaviors of users (Lerman, Yan, & Wu, 2016). Observing what others are doing and sharing on social media provides normative information - the quality, frequency and amount of behavior provides descriptive normative information; information and perceptions of whether others would approve or disapprove of behavior constitutes injunctive social norms (Conner & Norman, 2015; Schultz, Nolan, Cialdini, Goldstein, & Griskevicius, 2007). In particular injunctive norms resulting from engaging with personal networks or explicitly framed messaging on social media affect users' behavior (Bond et al., 2012), their reporting of potentially false information (Gimpel, Heger, Olenberger, & Utz, 2021) as well as their intention to share content (Andi & Akesson, 2021). Composing and sharing messages on social media however is also behavior used to signaling identity (Smaldino, 2022), that is, users may choose to express views in particular to facilitate association with others who might share similar views and values. Social identity theory (Tajfel, Turner, Austin, & Worchel, 1979) describes those processes related to in-group identification (the degree to which a person identifies with a particular social group) and inter-group bias (whether they favor this in-group over other out-groups). Sharing content on social media on particular topics can serve the function of increasing identification with a particular group and increase the difference to e.g., political out-groups (Rathje, Bavel, & van der Linden, 2021). Recent research on tweets by Democratic and Republican politicians on preventative measures (e.g., mask wearing) supports these assumptions as especially loyalty with the current party leadership was associated with lesser mask promotion (Shin et al., 2022). The key social information that is used in social media architecture is other users' engagement with

specific content, that is, whether they liked, commented or shared a particular post. Based on this information, the algorithms underlying the recommender systems then make it more (for posts with more engagement) or less (for posts with less engagement) likely that these posts will also appear in other users' timelines. This mechanism however is irrelevant of message content, i.e. whether users endorsed or opposed that content. While so far, these social cues have been used rather one-sidedly by platform providers to increase user engagement and sharing behavior, previous research suggests that targeting social processes could be an important additional building block to design interventions and environments that effectively reduce the sharing of misinformation on social media and empower its users - especially when cognitive and attentional resources are scarce (Kozyreva et al., 2020).

The present research

Here, we present a new open-source experimental paradigm to augment Twitter users' personal feed. We then use this paradigm to test the effects of social cues that provide users with more balanced information on what others in their personal network and on Twitter in general share and, most importantly, do *not* share. We additionally combine these cues with one of Twitter's standard misinformation labels (Figure 1). In two subsequent studies, we then examine the contribution of each social reference (1) own personal network, 2) all Twitter users) and of two hypothesized mechanisms behind the social cues, namely changes in social norms and social identity. We provide evidence that it is possible to implement a more balanced social reference within the general user interface of social media platforms - and, more importantly, that such social reference cues, together with the standard platform misinformation flags, can lead to reduced sharing of misinformation. Importantly, we examine these interventions within the natural environment of a widely-used social network platform (Twitter), thereby increasing the external validity of any potential effects and supporting its potential scalability.

Results

Study 1: Intervention

We begin by presenting the results examining the effects of social reference cues (personal network and all Twitter users) alone or combined with the standard misinformation flag on sharing of misinformation tweets. In short, participants were recruited for a study on a “new tweet recommender system” through an online research panel (Prolific) and were then randomly assigned to one of four experimental conditions:

- (1) Control (no flag or social cue),
- (2) Standard social media (Twitter) misinformation flag,
- (3) Social reference cues,
- (4) Combined (misinformation flag and social reference cues).

After providing initial consent, participants installed a custom browser extension for Chrome, logged into their Twitter account and were then asked to simply browse and interact with their personal feed as usual for 30 minutes on their computer. Participants in all groups had approximately 50% of their Twitter feed replaced with augmented content (pre-fabricated misinformation and control tweets; for details and examples see Methods). The remainder of participants’ “real” feed remained unchanged, so the other 50% were tweets from e.g. friends and others in their personal network. We recorded all interactions with the content, that is, what content participants saw, clicked on, liked, replied to or retweeted (Table 1), via the bespoke browser extension. After 30 minutes, participants then answered post-study questionnaires, were debriefed and provided informed consent again. Our final sample consisted of 824 individuals between 18 and 65 years ($M_{age} = 26.24$, $SD_{age} = 8.59$ years).

We find partial support for Hypothesis 1A as participants in the misinformation flag and the combined group shared less health-related misinformation than the control group. However, participants in the social cue and the misinformation flag group also shared less

control tweets (*H1B*) while only participants who saw the combined social cue and misinformation flag showed an improved ratio of shared mis- and control information (*H1C*). Thus, combining both interventions not only leads to a sizable and significant decrease in misinformation sharing (Table 2) but also improves overall sharing choices and quality. Participants in the combined intervention group shared on average only half of the amount of misinformation compared to the control group ($b = -0.66$, 95% CI $[-1.11, -0.20]$, $t = -2.86$, $p = .004$). As the thin tail of the model-based predictions distribution in Figure 2 demonstrates, presenting users with the combined intervention significantly reduced high sharing rates and thus “condensed” the distribution. This finding is especially important, as few users usually account for most of the shared information (Grinberg, Joseph, Friedland, Swire-Thompson, & Lazer, 2019). Reducing their activity might thus have additive effects on the amount shared overall. Beyond this intervention effect, participants’ age as well as their self-reported political orientation was significantly associated with the amount of shared misinformation. A one-point stronger right-wing political orientation was associated with an average 9% increase of sharing ($b = 0.09$, 95% CI $[0.02, 0.16]$, $t = 2.62$, $p = .009$), while one additional year of age was associated with an average 2% decrease of sharing ($b = -0.02$, 95% CI $[-0.04, 0.00]$, $t = -2.46$, $p = .014$).

Study 2 and 3: Social norms

Across study 2 and study 3, we examined whether the social reference users were provided with was important for the overall effect. Participants were shown different social reference points combined with the misinformation flag (reference points: (1) all Twitter users, (2) users in personal network only or (3) both references combined). We find that participants shared less misinformation and improved their overall sharing quality when shown a cue referencing their personal network or their personal network and all all Twitter users (*H2A*, *H2C*). The amount of shared control information was not significantly lower than the control group’s (please not the vast uncertainty around some estimates, Figure 3).

mistrust in science and evidence-based public health (Dryhurst et al., 2020; Freeman et al., 2020; Imhoff & Lamberty, 2020; Krause, Freiling, Beets, & Brossard, 2020). Together, our studies suggest that if social media users are provided with balanced social reference cues, i.e., information on the sharing and *not*-sharing of other individuals within their personal network additional to a misinformation flag, their perceived injunctive norms about the appropriateness of sharing behavior and their actual sharing behavior change. Importantly, the overall improvements of quality of information sharing were observed only if the social information referenced users' personal network and was combined with the misinformation flags currently employed within social networks; on its own, the social reference or misinformation flag had no consistent positive effects on sharing quality. Recent research on sharing behavior has highlighted the gap between which content users rate as trustworthy and which they actually share (Pennycook et al., 2021), as well as the importance of "attention economy" and a focus on the accuracy of message content in order to modify sharing behavior (Pennycook & Rand, 2021). Here, we provide first evidence that in addition to a focus on message content and on increasing accuracy in judging the veracity of online information, it is crucial to address the social determinants of sharing behavior. We argue that it is thus critical to understand misinformation sharing behavior as a result of the social processes and context users are embedded in and that has been created to foster user engagement by different social processes. While interaction behavior within the social media environment in itself is mainly based on social cues (such as recommender systems) and social reinforcement processes (such as likes), these processes have not yet been systematically addressed in interventions to modify sharing misinformation behavior.

Our findings in studies 2 and 3 show that the social reference cues in combination with the social network misinformation flags change users' subjective injunctive norms (perception of whether behaviors are deemed acceptable and appropriate within a given context). Changes in injunctive norms in turn were related to changes in sharing behavior, but descriptive norms, perceptions of the frequency of a specific behavior within a social

reference group, were not. This suggests that our social reference cue could provide a reference point for participants for which behaviors are socially acceptable within their networks. This is in line with previous findings as most social reinforcements that are relevant to users' sharing happen within their personal network. Thus, the approval of these important others is of key importance (Lewandowsky, Ecker, Seifert, Schwarz, & Cook, 2012). The social reference cue did not affect social identity in terms of in-group identification and intergroup bias, suggesting that these changes in norms were not due to a changing affiliation or identification with the users' social networks - or that our items (newly developed for this study) were not able to capture changes in social identity processes. Together, these findings highlight that normative perceptions play an important role in decision-making under load, that is, in conditions with limited time and large amounts of information to be judged before sharing. The fact that our studies replicate effects found in earlier studies increases our confidence in interpreting the intervention effects as generalizable. In particular, we found between-participant differences in sharing behavior such that there were few users who shared substantial amounts of misinformation (up to 26 tweets in 30 minutes), whereas many others did not engage in sharing behavior at all (Grinberg et al., 2019; Guess, Nagler, & Tucker, 2019). Similarly, we found that the general political orientation of participants influenced their sharing behavior such that in particular participants endorsing right-wing and authoritarian statements were more likely to share COVID-related misinformation (Guess et al., 2019; Lobato, Powell, Padilla, & Holbrook, 2020). This also points to the potential for exploring whether our findings replicate beyond the context of COVID-related misinformation. Our findings are to be considered in the light of some substantial limitations - first, as our paradigm had no access to and thus could not analyze the actual amount of misinformation shared within participants' personal networks, the social reference cues (i.e., the percentages reported) was a priori set by us. It remains to be seen whether basing the reference cues on actual shares has similar effects. Here, it will be especially important to test how lower rates of

non-shares affect users' behavior. Second, the significant results in study 2 fall just under the arbitrary alpha-level of .05. We thus only interpret them as the first indication of a potential mechanism and not sufficient evidence to rest major claims on. Third, the time frame within our paradigm was very short in relation to the usually vast amount of information read before sharing with also a usually lower percentage of fake news in users' feed (Grinberg et al., 2019). We thus added a very high amount of misinformation content into participants' feeds and it remains to be tested whether the intervention effects remain substantial over longer time frames and with (more realistic) lower density of misinformation.

On the other hand, our experimental paradigm has very high external validity; by embedding our social reference cues within the general platform of a social media network, we use the actual user interface that the general population is exposed to. This allows us to generalize our observed effects beyond those generated in laboratory settings (Mosleh, Pennycook, & Rand, 2021). At the same time, our paradigm allowed us to retain control over what participants were exposed to and thus provide participants with many different sources of misinformation - past research had only been able to analyze shared information from a given set of specific news sites or influential public persons. Additionally, and crucially, we were able to not only add misinformation but also control information and could thus extend our analyses to overall sharing quality. Especially in the light of several studies examining only sharing rates (or rather intentions thereof) without comparing them with all other sharing activities or / and a ratio thereof as an indicator of discernment, our paradigm and results add important points to the discourse. Our results suggest that it is possible to implement balanced social reference cues within the general user interface of social media platforms - and, more importantly, these social reference cues, together with the standard platform misinformation flags, can lead to reduced sharing of online misinformation and improved quality of overall sharing. In the context of major challenges to public health and public trust caused by excessive and strategically placed

misinformation in social media, such building blocks to effective mitigation measures have the potential to reduce the sharing of misinformation with all its associated negative consequences.

Materials & Methods

For the present studies, we created an open source browser extension to access and augment Twitter users' social media feed and an experimental paradigm to systematically examine the effects of social cues that provide users with more balanced information on what others in their personal network and on Twitter in general share and, most importantly, do *not* share. We additionally combined these social cues with one of Twitter's standard misinformation labels (Figure 1). In a pre-registered study, we tested whether social cues (alone or combined with Twitter's misinformation flag) 1) reduce the sharing of misleading content, 2) also reduce sharing of control stimuli, and 3) improve the ratio of control and misleading content shared by users as a measure of discernment. In two subsequent studies, we then examined different social references for the social cues and the contribution of two hypothesized mechanisms behind the social cues, namely changes in social norms and social identity. In total, we conducted three separate experimental studies sharing the same overall paradigm and similar procedures.

In particular, we test the following hypotheses:

Hypothesis 1A: Participants across all intervention groups share (sum of likes / retweets) less tweets containing health-related misinformation than participants in the control group

Hypothesis 1B: Participants across all intervention groups do not share (sum of likes / retweets) less tweets containing health-related control than participants in the control group

Hypothesis 1C: Participants across all intervention groups show a smaller ratio of

shared misinformation and control tweets than participants in the control group (indicating improved discernment)

Hypothesis 2A: Participants who see a social cue referencing their personal network, or their personal network and all Twitter users, share less tweets containing health-related misinformation than participants in the control group

Hypothesis 2B: Participants who see a social cue referencing their personal network, or their personal network and all Twitter users, do not share less tweets containing health-related control than participants in the control group

Hypothesis 2C: Participants who see a social cue referencing their personal network, or their personal network and all Twitter users, show a smaller ratio of shared misinformation and control tweets than participants in the control group (indicating improved discernment)

Hypothesis 3A: Participants in the social reference cues group and those in the combined group report less approval of sharing misleading information than participants in the control group (injunctive norms)

Hypothesis 3B: Participants in the social reference cues group and in the combined group perceive that others share less misleading information than participants in the control group (descriptive norms)

Hypothesis 4A: Participants in the social reference cues group and those in the combined group report stronger intergroup-bias than participants in the control group

Hypothesis 4B: Participants in the social reference cues group and those in the combined group will report stronger in-group identification than participants in the control group

Our research was approved by the University of St. Gallen Ethics Board (HSG-EC-20210816A). Participants provided informed consent before data collection. As

we used cover-up study aims, participants were debriefed about the true aims after data collection. All data were collected using the online crowd-sourcing platform Prolific (www.prolific.co) and German survey platform SoSci Survey (www.soscisurvey.de). Respondents were paid a fee deemed appropriate by Prolific, equivalent to £7.50/hour.

To determine the sample sizes, we conducted a priori power analyses. For study 1, the sample size to be able to detect small differences in the expected rate of events (between two negative binomial rates; e.g., $\mu_1 = 2$ and $\mu_2 = 3$) with a power of .90 and an α -level of .05, a total sample size of $N = 800$ was deemed sufficient ($N = 200$ per group). For studies 2 and 3, based on simulation studies, we targeted a minimum sample size of $N = 260$ to detect a growth condition of $d = 0.20$ with three repeated measurements, an α -level of .05 and power of .90 (Xitao & Xiaotao, 2005). This minimum sample size for the growth condition exceeds the required sample size to detect medium sized intercept and slope differences between groups with a power of .90 and an α -level of .05 (Fan, 2003).

Our *main outcome measures* across all studies were

- (1) the amount of misinformation tweets users shared,
- (2) the amount of control tweets users shared,
- (3) the ratio of misinformation and control tweets users shared (calculated as misinformation tweets minus control tweets shared).

Shared information includes all liked and retweeted misinformation / control tweets. Still, we also estimated all models for liked and retweeted misinformation tweets separately and provide results in the Supplementary Information (SI3, SI4). However, results do not vary substantially across these different outcomes.

In line with our pre-registration, we also provide results of a model with the amount of inhibited shares as outcome as Supplementary Information (SI5).

All data, code, and documentation can be found on the Open Science Framework

(OSF; <https://osf.io/49raq/>). We have created this manuscript with R and the papaja-package (Aust & Barth, 2020) to support reproducibility.

Divergence from pre-registration

We published the present manuscript as a pre-print and received valuable feedback from the research community. As a result, we now diverge from our pre-registration as we had to make several key adjustments to our analytic strategy after being made aware of several important methodological issues that could bias our results. Firstly, we learned that including the amount of misinformation participants' had seen in the model might introduce post-treatment bias (Montgomery, Nyhan, & Torres, 2018). Accordingly, we dropped the amount of misinformation as a predictor from all models. Secondly, we were made aware that only examining the amount of shared misinformation might lead to false claims about the overall effectiveness of the intervention. Thus, we have now added new analyses to also examine the effects on users' sharing of control stimuli as well as a changes to the ratio of misinformation and control information shared. Although we now diverge from our pre-registration substantially, the described changes are important to ensure the overall quality of the present study.

Study 1: Intervention

Participants. The crowdsourcing platform Prolific Academic was used to recruit participants. We invited 900 English-speaking participants with a minimum age of 18 years and with more than 100 posts submitted to Twitter in the last 12 months. We collected all the data for study 1 on September 3, 2021. As several participants dropped out before installing the browser extension, our final sample consisted of 824 individuals between 18 and 65 years ($M_{age} = 26.24$, $SD_{age} = 8.59$ years). Participants were either unemployed (and job seeking; $N = 240$, 29.13%), in full-time employment ($N = 233$, 28.28%) or in part-time employment ($N = 121$, 14.68%). 818 (99.27%) of the participants had at least an

undergraduate degree while only 0 participants (0.00%) reported having no formal education. Most participants reported the following countries as their current place of residence: in descending order, United Kingdom ($N = 132$, 16.02%), South Africa ($N = 123$), 14.93%), Portugal ($N = 122$, 14.81%), United States ($N = 101$, 12.26%). We did not drop any participants that provided full consent.

Material. We obtained demographic data (e.g., age, first language, employment status) from Prolific. Participants also completed pre-study self-report measures on education as well as on different potential mediators of the relationship between intervention and sharing behavior. We assessed *digital health-literacy* with four items asking their subjective ability to browse, search and find as well as assess the quality and trustworthiness of digital health-related content. We created these items in accordance with ability-oriented conceptions of digital literacy (Janssen et al., 2013) and the HLS EU-Q16 (Sørensen et al., 2013) which assesses the ability to find and process health-related information. Participants were asked to answer each question on a seven-point scale. The short scale exhibited good reliability with $\omega_u = 0.82$ representing the proportion of total-score variance due to a common single factor (Cronbach’s $\alpha = 0.80$).

Additionally, we asked participants to rate their *political orientation* on a continuum (ranging from 1-10): “Many people use the terms “left” and “right” when they want to describe different political views. Thinking of your own political views, where would you place these on this scale?”. The distribution of participants’ answers was right skewed ($M = 3.49$, $SD = 2.33$).

For our experimental procedure, we developed a study-specific *Chrome browser extension* to augment participants’ actual Twitter feed. This browser extension enabled us to add manipulated tweets to participants’ natural feed (e.g., content shared or posted by friends or other people users followed), record all manipulated content they viewed as well as all interactions with manipulated content. We stored all interaction data (tweet ID, timestamp, experimental condition) on a Google Firebase server. To protect participants

from spreading misleading information, the browser extension additionally intercepted all likes, replies, and retweets. We will share all code for this browser extension upon publication of this manuscript on GitHub (<https://github.com/dan91/tweet-recommender>). The extension is written in Javascript and published on the Google Chrome Web Store, but only accessible via a direct link to prevent dissemination outside of the study context. In total, we augmented users' Twitter feed with 50 manipulated tweets. Forty of the manipulated tweets contained false claims, ten contained short vaccine-positive, entertainment and sports-related information and served as *control stimuli*. As *misinformation stimuli*, we used actual tweets that either contained misinformation in the tweet itself or were sharing misinformation. Misinformation tweets were collected through the Google Fact Check API (<https://toolbox.google.com/factcheck/apis>). Through this API, we accessed claims that had been fact-checked by platforms such as politifact.com or Vera Files, who are mostly not-for-profit organizations. We included tweets with the keyword "COVID" that were rated as "false". A full list of the stimuli used can be found on the OSF (<https://osf.io/49raq/>).

On Twitter, we collected impressions (views), intended likes, intended shares and intended retweets of all pieces of information through the browser extension. Intended here means users clicked on like, share or retweet but did not confirm their action after possibly encountering a social reference.

For the social reference message, e.g., "In your personal Twitter network, 710 out of 740 people saw but did not like, reply or retweet this tweet.", we did not calculate the size of participants' actual Twitter network. Instead, the proportion of participants' personal Twitter network and total number of Twitter users not reacting to the tweet was calculated based on the number of retweets of the actual misinformation content. For the personal network the retweet count was multiplied with 10, for the complete Twitter network with 500. The proportions of seen and ignored tweets were then randomized between 95 and 99 percent. For example, if a misinformation tweet had been retweeted 10 times, the social

reference message could read "... 95 of the people in your personal network saw but did not share ... On twitter, 4950 other users saw but did not share...".

Our *main outcome measures* were

- (1) the amount of misinformation tweets users shared,
- (2) the amount of control tweets users shared,
- (3) the ratio of misinformation and control tweets users shared (calculated as misinformation tweets minus control tweets shared).

Procedure. To avoid expectation effects, participants were briefed to test a tweet recommendation system (Cheng et al., 2018). After providing consent, they then answered all pre-study questionnaires and were randomly assigned to one of four conditions:

- (1) Control (no flag or cues),
- (2) Standard social media (Twitter) misinformation flag,
- (3) Social reference cues,
- (4) Combined (misinformation flag and social reference cues).

After being redirected to Twitter, participants were instructed to browse and interact with their personal Twitter feed as usual for 30 minutes. During these 30 minutes, participants in all groups had approximately 50% of their Twitter feed replaced with misinformation and control tweets. All other content of participants' "real" feed remained unchanged. If participants had seen all added pieces of information, they would repeat in randomized order. This methodology is consistent with prior work (Geeng, Yee, & Roesner, 2020). After 30 minutes, participants were redirected again and debriefed about the actual study aims, and provided consent again. To ensure that misperceptions did not spread among participants after being exposed to misinformation (e.g., attaining an anti-vaccination attitude) we provided them with screenshots of all misinformation tweets

at the end of the study. Additionally, a list of all false claims with corresponding corrections from the fact-checking websites was presented. When participants decided to interact with the misinformation tweets, these interactions were intercepted to ensure they were not visible on their Twitter profiles.

Analysis. The *first outcome* variable was the count of shared misinformation tweets (sum of liked and retweeted misinformation tweets per participant). To account for the high number of participants sharing no tweets and thus highly skewed outcome distribution ($M = 1.54$, $SD = 3.43$, $\sigma^2 = 11.74$), we estimated the following negative binomial regression model to test *Hypothesis 1A*:

$$\begin{aligned} \ln(\text{shared-tweets}) = & \alpha + \beta_1(\text{social-reference-frame}) + \\ & \beta_2(\text{misinformation-flag}) + \beta_3(\text{combined}) + \\ & \beta_4(\text{age}) + \beta_5(\text{education}) + \\ & \beta_6(\text{digital-literacy}) + \beta_7(\text{political-orientation}) + \ln_t \end{aligned}$$

The model was estimated using Maximum Likelihood estimation, and it had low explanatory power (Nagelkerke's $R^2 = 0.05$). The model intercept, corresponding to the count of shared misinformation tweets at the hypothetical point of all predictors with a value of zero, was 1.60 (95% CI [0.49, 2.75], $p = 0.004$). We included a dummy coded predictor for each intervention group in the model ($\beta_1, \beta_2, \beta_3$). The control group served as reference. Additionally, all models included participant age and education, digital health literacy and political orientation as predictors. We report regression coefficients, standard errors, and incidence rate ratios (*IRR*) as change in Y per one-unit increase in a predictor as effect size (see Table 2). The analyses were pre-registered and can be retrieved from AsPredicted (<https://aspredicted.org/ju2c7.pdf>). In the preregistration, we had also planned to include intention to not share misleading information, their momentary self-control as well as distrust of experts and intellectuals (anti-intellectualism). However, none of these variables contributed significantly as a predictor and all other estimates also

remained unchanged across both models (Table SI6). All others divergences from the pre-registration are described above and we here now focus on inferences based on the more parsimonious model described above. Descriptives and a correlation matrix of all predictors included in this final model can be found in the Supplementary Information (SI1). The *second outcome* was the count of shared control tweets (sum of liked and retweeted control tweets per participant). Again, we had to account for a highly skewed outcome distribution ($M = 0.60$, $SD = 1.45$, $\sigma^2 = 2.11$) and estimated a negative binomial regression model to test *Hypothesis 1B*. This model included the same predictors as the model described above and also had low explanatory power (Nagelkerke's $R^2 = 0.11$). The model intercept, corresponding to the count of shared control tweets at the hypothetical point of all predictors with a value of zero, was 1.20 (95% CI $[-0.11, 2.54]$, $p = 0.053$). The *third outcome* was the number of shared control tweets subtracted from the the number of shared misinformation tweets. Here, we estimated a linear regression model to test *Hypothesis 1C*. This model included the same predictors as the model described above and explained only a small part of the total variance ($R^2 = 0.01$, $F(7, 809) = 1.47$, $p = 0.175$). The model intercept was 1.91 (95% CI $[0.50, 3.31]$, $p = 0.008$).

Next, we describe two sets of post-hoc analyses we conducted to explore the mechanisms potentially underlying the observed intervention effects.

Studies 2 and 3: Social norms and identity

Participants. The crowdsourcing platform Prolific Academic was also used to recruit all participants for study 2 and 3. We invited 650 English-speaking participants in total. We collected all the data for studies 2 and 3 on November 25, 2021. As several participants had to be dropped from the final sample because they e.g., did not provide consent after being briefed about the true study aims, we report a final sample size for study 2 of $N = 322$ and for study 3 of $N = 278$. We did not drop any participants that provided full consent. Again, the minimum age was 18 years and we only invited users with

more than 100 posts submitted to Twitter in the last 12 months. Participants were between 18 and 76 years ($M_{age} = 28.02$, $SD_{age} = 9.77$) old and either unemployed (and job seeking; $N = 151$, 25.17%), in full-time employment ($N = 207$, 34.50%) or in part-time employment ($N = 108$, 18.00%).

Material. *For both studies*, we again obtained demographic data (e.g., age, first language, employment status) from Prolific.

In study 2, we assessed subjective descriptive and injunctive norms with one item each after each trial. Regarding injunctive norms, we asked participants to rate the following statement: “The people I care about in my personal Twitter network approve of me sharing Covid-19-related information such as the ones I just saw”. Regarding descriptive norms, the statement read: “The people I care about in my personal Twitter network ... share Covid-19-related information such as the ones I just saw”.

In study 3, we assessed inter-group bias on positive and negative traits after each trial as well as their overall in-group-identification post-assessment. Inter-group bias-related traits were, for example, “intelligent”, “trustworthy” or “gullible” and participants rated how well those described their personal Twitter network as well as others with similar or different opinions on COVID-19 (range from 1- 7). Additionally, participants were asked to answer whether they agreed with three statements regarding their in-group-identification post-assessment on a seven-point scale (e.g., “I am similar to others who have opinions on COVID-19 similar to mine”, Leach et al. (2008)). The three-item scale exhibited good reliability with $\omega_u = 0.83$ representing the proportion of total-score variance due to a common single factor (Cronbach’s $\alpha = 0.83$).

Procedure. Procedure and cover-up story for studies 2 and 3 were similar to those of the intervention study. Again, participants were briefed that they were testing a Tweet recommendation system, and after providing consent, they were randomly assigned to one of four experimental conditions. Each of the three intervention groups was shown Twitter’s misinformation flag and one of three different social cues that either referenced only the

user’s personal network, all Twitter users or combined both references. Thus, the four

conditions were: (1) Control (no flag or frame),

(2) Social reference cues (reference: all Twitter users) and misinformation flag,

(3) Social reference cues (reference: personal network) and misinformation flag,

(4) Combined (reference with both social reference points and misinformation flag).

After being redirected to Chrome, participants were instructed to browse and interact with their personal Twitter feed as usual for three 5 minute trials. Again, participants in all groups had approximately 50% of their Twitter feed replaced with augmented tweets (misinformation and controls as described regarding Study 1). All other content remained unchanged. If participants had seen all pieces of augmented tweets, they would repeat in randomized order. After each trial, participants were redirected to the survey platform and asked to answer questions on social norms (study 2) or inter-group bias-related traits (study 3). After three trials, participants were debriefed about the true study aims and provided consent again. Additionally, participants in study 3 answered post assessment measures on in-group identification. As described above, we ensured that misperceptions did not spread among participants after being exposed to misinformation by providing them with screenshots of all misinformation tweets at the end of the study.

Analysis. To test hypotheses 2A, 2B and 2C we pooled the experimental data of study 2 and study 3. As we not only had to account for the high number of participants sharing no tweets, and a thus highly skewed outcome distributions (2A and 2B), but also for the non-independency of repeated assessments across experimental trials within participants. We thus estimated negative binomial generalized linear mixed-effects models to examine hypotheses 2A and 2B, and a linear mixed-effects model to examine hypothesis 2C. The *main outcome measures* were again (1) the amount of misinformation tweets users shared (H2A),

(2) the amount of control tweets users shared (H2B),

(3) the ratio of misinformation and control tweets users shared (calculated as

misinformation tweets minus control tweets shared; $H2C$).

The models were estimated using Maximum Likelihood estimation and, in the case of the two negative binomial mixed-effects models with the BOBYQA optimizer. The models total explanatory power was moderate to substantial (conditional $R^2 = 0.54$, $R^2 = 0.40$ and $R^2 = 0.25$), and again included a dummy coded predictor for each intervention group in the model (γ_1^α , γ_2^α , γ_3^α) with the control group serving as reference.

To test hypotheses $3A$, $3B$ and $4A$, respectively, we estimated three separate first-order linear latent growth curve models (for a graphical representation of the first model for Study 2 see supplementary Figure SI1). Here, we tested whether the intervention effects was associated with the average (i.e., intercept) of participants' injunctive norms ($H3A$), descriptive norms ($H3B$), intergroup-bias ($H4A$) or change (i.e., slope) thereof. Dummy coded predictors for each intervention group were included in the model to test associations with the latent intercept or slope.

To test hypothesis $4B$, we estimated a linear mixed-effects regression model with ingroup identification as the outcome.

Data availability

Anonymized data are available on the Open Science Framework:
(<https://osf.io/49raq/>)

Code availability

The analytic code and this reproducible manuscript are available on the Open Science Framework: (<https://osf.io/49raq/>)

Acknowledgements

The authors thank Sarah Janetzki for her assistance in conducting the study and data collection. Work on this study was supported by a grant from *VolkswagenStiftung* (call: “Corona crisis and beyond”) to Johannes Schöning.

Additionally, we thank all authors of the following software packages we used to wrangle, analyze or visualize data and findings: R (Version 4.2.1; R Core Team, 2020b) and the R-packages *broom.mixed* (Version 0.2.9.4; Bolker & Robinson, 2022), *car* (Version 3.1.0; Fox & Weisberg, 2019; Fox, Weisberg, & Price, 2020), *carData* (Version 3.0.5; Fox et al., 2020), *citr* (Version 0.3.2; Aust, 2019), *corx* (Version 1.0.7.1; Conigrave, 2020), *cowplot* (Version 1.1.1; Wilke, 2020), *dplyr* (Version 1.0.10; Wickham, François, Henry, & Müller, 2021), *equatiomatic* (Version 0.3.1; Anderson, Heiss, & Sumners, 2021), *forcats* (Version 0.5.2; Wickham, 2021a), *foreign* (Version 0.8.82; R Core Team, 2020a), *ggplot2* (Version 3.3.6; Wickham, 2016), *gridExtra* (Version 2.3; Auguie, 2017), *hrbrthemes* (Version 0.8.0; Rudis, 2020), *lavaan* (Version 0.6.12; Rosseel, 2012), *lme4* (Version 1.1.30; Bates, Mächler, Bolker, & Walker, 2015), *lmerTest* (Version 3.1.3; Kuznetsova, Brockhoff, & Christensen, 2017), *lmtest* (Version 0.9.40; Zeileis & Hothorn, 2002), *MASS* (Version 7.3.57; Venables & Ripley, 2002), *Matrix* (Version 1.5.1; Bates & Maechler, 2019), *papaja* (Version 0.1.1; Aust & Barth, 2020), *patchwork* (Version 1.1.2; Pedersen, 2020), *performance* (Version 0.10.0; Lüdtke, Ben-Shachar, Patil, Waggoner, & Makowski, 2021), *png* (Version 0.1.7; Urbanek, 2013), *pscl* (Version 1.5.5; Zeileis, Kleiber, & Jackman, 2008), *purrr* (Version 0.3.5; Henry & Wickham, 2020), *readr* (Version 2.1.3; Wickham & Hester, 2021), *report* (Version 0.5.5; Makowski, Ben-Shachar, Patil, & Lüdtke, 2021), *sandwich* (Zeileis, 2004, 2006; Version 3.0.2; Zeileis, Köll, & Graham, 2020), *see* (Version 0.7.3; Lüdtke, Patil, et al., 2021), *semPlot* (Version 1.1.6; Epskamp, 2019), *semTable* (Version 1.8; Johnson & Kite, 2020), *semTools* (Version 0.5.6; Jorgensen, Pornprasertmanit, Schoemann, & Rosseel, 2021), *sjmisc* (Version 2.8.9; Lüdtke, 2018), *sjPlot* (Version 2.8.11; Lüdtke, 2021), *stringr* (Version 1.4.1; Wickham, 2019), *table1* (Rich, 2021), *tibble* (Version 3.1.8; Müller &

Wickham, 2021), *tidyr* (Version 1.2.1; Wickham, 2021b), *tidyverse* (Version 1.3.2; Wickham et al., 2019), *tinylabals* (Version 0.2.3; Barth, 2022), *trackdown* (Version 1.1.1; Kothe, Callegher, Gambarota, Linkersdörfer, & Ling, 2021), *viridis* (Version 0.6.2; Garnier et al., 2021a, 2021b), *viridisLite* (Version 0.4.1; Garnier et al., 2021b), and *zoo* (Version 1.8.11; Zeileis & Grothendieck, 2005).

Author Contributions

C.J. conceived of the study, designed the study, collected data, analyzed data and contributed to drafting of the manuscript. *D.D.* conceived of the study, designed and programmed the study paradigm and contributed to the drafting of the manuscript. *J.S.* conceived of the study, contributed to the design of the study and study paradigm and contributed to the drafting of the manuscript. *R.S.* contributed to study design and the drafting of the manuscript. *T.J.* contributed to study design and the drafting of the manuscript. *B.S.* conceived of the study, data analysis plan and contributed to drafting of the manuscript.

References

- Anderson, D., Heiss, A., & Sumners, J. (2021). *Equationomatic: Transform models into 'LaTeX' equations*. Retrieved from <https://CRAN.R-project.org/package=equationomatic>
- Andi, S., & Akesson, J. (2021). Nudging Away False News: Evidence from a Social Norms Experiment. *Digital Journalism*, 9(1), 106–125. <https://doi.org/10.1080/21670811.2020.1847674>
- Auguie, B. (2017). *gridExtra: Miscellaneous functions for "grid" graphics*. Retrieved from <https://CRAN.R-project.org/package=gridExtra>
- Aust, F. (2019). *Citr: 'RStudio' add-in to insert markdown citations*. Retrieved from <https://github.com/crsh/citr>
- Aust, F., & Barth, M. (2020). *papaja: Create APA manuscripts with R Markdown*. Retrieved from <https://github.com/crsh/papaja>
- Avram, M., Micallef, N., Patil, S., & Menczer, F. (2020). Exposure to Social Engagement Metrics Increases Vulnerability to Misinformation. *Harvard Kennedy School Misinformation Review*. <https://doi.org/10.37016/mr-2020-033>
- Barth, M. (2022). *tinylabels: Lightweight variable labels*. Retrieved from <https://cran.r-project.org/package=tinylabels>
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1), 1–48. <https://doi.org/10.18637/jss.v067.i01>
- Bates, D., & Maechler, M. (2019). *Matrix: Sparse and dense matrix classes and methods*. Retrieved from <https://CRAN.R-project.org/package=Matrix>
- Bolker, B., & Robinson, D. (2022). *Broom.mixed: Tidying methods for mixed models*. Retrieved from <https://CRAN.R-project.org/package=broom.mixed>
- Bond, R. M., Fariss, C. J., Jones, J. J., Kramer, A. D. I., Marlow, C., Settle, J. E., & Fowler, J. H. (2012). A 61-million-person experiment in social influence and

political mobilization. *Nature*, 489(7415), 10.1038/nature11421.

<https://doi.org/10.1038/nature11421>

Charquero-Ballester, M., Walter, J. G., Nissen, I. A., & Bechmann, A. (2021).

Different types of COVID-19 misinformation have different emotional valence on Twitter. *Big Data & Society*, 8(2), 20539517211041279.

<https://doi.org/10.1177/20539517211041279>

Conigrave, J. (2020). *Corx: Create and format correlation matrices*. Retrieved from

<https://CRAN.R-project.org/package=corx>

Conner, M., & Norman, N. (2015). *Predicting and Changing Health Behaviour:*

Research and Practice with Social Cognition Models, 3rd Edition. Open

University Press. <https://doi.org/10.1036/9780335263790>

Dryhurst, S., Schneider, C. R., Kerr, J., Freeman, A. L. J., Recchia, G., van der

Bles, A. M., . . . van der Linden, S. (2020). Risk perceptions of COVID-19

around the world. *Journal of Risk Research*, 23(7-8), 994–1006.

<https://doi.org/10.1080/13669877.2020.1758193>

Epskamp, S. (2019). *semPlot: Path diagrams and visual analysis of various SEM*

packages' output. Retrieved from

<https://CRAN.R-project.org/package=semPlot>

Fan, X. (2003). Power of Latent Growth Modeling for Detecting Group Differences

in Linear Growth Trajectory Parameters. *Structural Equation Modeling: A*

Multidisciplinary Journal, 10(3), 380–400.

https://doi.org/10.1207/S15328007SEM1003_3

Fox, J., & Weisberg, S. (2019). *An R companion to applied regression* (Third).

Thousand Oaks CA: Sage. Retrieved from

<https://socialsciences.mcmaster.ca/jfox/Books/Companion/>

Fox, J., Weisberg, S., & Price, B. (2020). *carData: Companion to applied regression*

data sets. Retrieved from <https://CRAN.R-project.org/package=carData>

- Freeman, D., Waite, F., Rosebrock, L., Petit, A., Causier, C., East, A., . . . Lambe, S. (2020). Coronavirus conspiracy beliefs, mistrust, and compliance with government guidelines in England. *Psychological Medicine*, 1–13.
<https://doi.org/10.1017/S0033291720001890>
- Garnier, Simon, Ross, Noam, Rudis, Robert, . . . Cédric. (2021a). *viridis - colorblind-friendly color maps for r*. <https://doi.org/10.5281/zenodo.4679424>
- Garnier, Simon, Ross, Noam, Rudis, Robert, . . . Cédric. (2021b). *viridis - colorblind-friendly color maps for r*. <https://doi.org/10.5281/zenodo.4679424>
- Geeng, C., Yee, S., & Roesner, F. (2020). Fake News on Facebook and Twitter: Investigating How People (Don't) Investigate. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–14. Honolulu HI USA: ACM. <https://doi.org/10.1145/3313831.3376784>
- Gimpel, H., Heger, S., Olenberger, C., & Utz, L. (2021). The Effectiveness of Social Norms in Fighting Fake News on Social Media. *Journal of Management Information Systems*, 38(1), 196–221.
<https://doi.org/10.1080/07421222.2021.1870389>
- Grinberg, N., Joseph, K., Friedland, L., Swire-Thompson, B., & Lazer, D. (2019). Fake news on Twitter during the 2016 U.S. Presidential election. *Science (New York, N.Y.)*, 363(6425), 374–378. <https://doi.org/10.1126/science.aau2706>
- Guess, A., Nagler, J., & Tucker, J. (2019). Less than you think: Prevalence and predictors of fake news dissemination on Facebook. *Science Advances*, 5(1), eaau4586. <https://doi.org/10.1126/sciadv.aau4586>
- Henry, L., & Wickham, H. (2020). *Purrr: Functional programming tools*. Retrieved from <https://CRAN.R-project.org/package=purrr>
- Imhoff, R., & Lamberty, P. (2020). A Bioweapon or a Hoax? The Link Between Distinct Conspiracy Beliefs About the Coronavirus Disease (COVID-19) Outbreak and Pandemic Behavior. *Social Psychological and Personality Science*,

- 11(8), 1110–1118. <https://doi.org/10.1177/1948550620934692>
- Janssen, J., Stoyanov, S., Ferrari, A., Punie, Y., Pannekeet, K., & Sloep, P. (2013). Experts' views on digital competence: Commonalities and differences. *Computers & Education*, 68, 473–481. <https://doi.org/10.1016/j.compedu.2013.06.008>
- Johnson, P., & Kite, B. (2020). *semTable: Structural equation modeling tables*. Retrieved from <https://CRAN.R-project.org/package=semTable>
- Jorgensen, T. D., Pornprasertmanit, S., Schoemann, A. M., & Rosseel, Y. (2021). *semTools: Useful tools for structural equation modeling*. Retrieved from <https://CRAN.R-project.org/package=semTools>
- Kothe, E., Callegher, C. Z., Gambarota, F., Linkersdörfer, J., & Ling, M. (2021). *Trackdown: Collaborative writing and editing of r markdown (or sweave) documents in google drive*. <https://doi.org/10.5281/zenodo.5167320>
- Kozyreva, A., Lewandowsky, S., & Hertwig, R. (2020). Citizens Versus the Internet: Confronting Digital Challenges With Cognitive Tools. *Psychological Science in the Public Interest*, 21(3), 103–156. <https://doi.org/10.1177/1529100620946707>
- Krause, N. M., Freiling, I., Beets, B., & Brossard, D. (2020). Fact-checking as risk communication: The multi-layered risk of misinformation in times of COVID-19. *Journal of Risk Research*, 23(7-8), 1052–1059. <https://doi.org/10.1080/13669877.2020.1756385>
- Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). lmerTest package: Tests in linear mixed effects models. *Journal of Statistical Software*, 82(13), 1–26. <https://doi.org/10.18637/jss.v082.i13>
- Leach, C. W., van Zomeren, M., Zebel, S., Vliek, M. L. W., Pennekamp, S. F., Doosje, B., . . . Spears, R. (2008). Group-level self-definition and self-investment: A hierarchical (multicomponent) model of in-group identification. *Journal of Personality and Social Psychology*, 95(1), 144–165.

<https://doi.org/10.1037/0022-3514.95.1.144>

Lerman, K., Yan, X., & Wu, X.-Z. (2016). The "Majority Illusion" in Social Networks. *PLOS ONE*, 11(2), e0147617.

<https://doi.org/10.1371/journal.pone.0147617>

Lewandowsky, S., Ecker, U. K. H., Seifert, C. M., Schwarz, N., & Cook, J. (2012). Misinformation and Its Correction: Continued Influence and Successful Debiasing. *Psychological Science in the Public Interest*, 13(3), 106–131.

<https://doi.org/10.1177/1529100612451018>

Lobato, E. J. C., Powell, M., Padilla, L. M. K., & Holbrook, C. (2020). Factors Predicting Willingness to Share COVID-19 Misinformation. *Frontiers in Psychology*, 11, 2413. <https://doi.org/10.3389/fpsyg.2020.566108>

Lüdecke, D. (2018). Sjmisc: Data and variable transformation functions. *Journal of Open Source Software*, 3(26), 754. <https://doi.org/10.21105/joss.00754>

Lüdecke, D. (2021). *sjPlot: Data visualization for statistics in social science*.

Retrieved from <https://CRAN.R-project.org/package=sjPlot>

Lüdecke, D., Ben-Shachar, M. S., Patil, I., Waggoner, P., & Makowski, D. (2021). performance: An R package for assessment, comparison and testing of statistical models. *Journal of Open Source Software*, 6(60), 3139.

<https://doi.org/10.21105/joss.03139>

Lüdecke, D., Patil, I., Ben-Shachar, M. S., Wiernik, B. M., Waggoner, P., & Makowski, D. (2021). see: An R package for visualizing statistical models. *Journal of Open Source Software*, 6(64), 3393.

<https://doi.org/10.21105/joss.03393>

Makowski, D., Ben-Shachar, M. S., Patil, I., & Lüdecke, D. (2021). Automated results reporting as a practical tool to improve reproducibility and methodological best practices adoption. *CRAN*. Retrieved from <https://github.com/easystats/report>

- 769 Meshi, D., Tamir, D. I., & Heekeren, H. R. (2015). The Emerging Neuroscience of
770 Social Media. *Trends in Cognitive Sciences*, 19(12), 771–782.
771 <https://doi.org/10.1016/j.tics.2015.09.004>
- 772 Montgomery, J. M., Nyhan, B., & Torres, M. (2018). How Conditioning on
773 Posttreatment Variables Can Ruin Your Experiment and What to Do about It.
774 *American Journal of Political Science*, 62(3), 760–775.
775 <https://doi.org/10.1111/ajps.12357>
- 776 Mosleh, M., Pennycook, G., & Rand, D. (2021). *Field experiments on social media*.
777 <https://doi.org/10.31234/osf.io/dgmc2>
- 778 Müller, K., & Wickham, H. (2021). *Tibble: Simple data frames*. Retrieved from
779 <https://CRAN.R-project.org/package=tibble>
- 780 O'Connor, C., & Weatherall, J. O. (2019). The Misinformation Age: How False
781 Beliefs Spread. In *The Misinformation Age*. Yale University Press.
782 <https://doi.org/10.12987/9780300241006>
- 783 Pedersen, T. L. (2020). *Patchwork: The composer of plots*. Retrieved from
784 <https://CRAN.R-project.org/package=patchwork>
- 785 Pennycook, G., Epstein, Z., Mosleh, M., Arechar, A. A., Eckles, D., & Rand, D. G.
786 (2021). Shifting attention to accuracy can reduce misinformation online. *Nature*,
787 592(7855, 7855), 590–595. <https://doi.org/10.1038/s41586-021-03344-2>
- 788 Pennycook, G., McPhetres, J., Zhang, Y., Lu, J. G., & Rand, D. G. (2020).
789 Fighting COVID-19 Misinformation on Social Media: Experimental Evidence for
790 a Scalable Accuracy-Nudge Intervention. *Psychological Science*, 31(7), 770–780.
791 <https://doi.org/10.1177/0956797620939054>
- 792 Pennycook, G., & Rand, D. G. (2021). The Psychology of Fake News. *Trends in*
793 *Cognitive Sciences*, 25(5), 388–402. <https://doi.org/10.1016/j.tics.2021.02.007>
- 794 Poynter. (2021). Covid-19 Misinformation. Retrieved December 18, 2021, from
795 <https://www.poynter.org/ifcn-covid-19-misinformation/>

R Core Team. (2020a). *Foreign: Read data stored by 'minitab', 's', 'SAS', 'SPSS', 'stata', 'systat', 'weka', 'dBase', ...* Retrieved from <https://CRAN.R-project.org/package=foreign>

R Core Team. (2020b). *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from <https://www.R-project.org/>

Rathje, S., Bavel, J. J. V., & van der Linden, S. (2021). Out-group animosity drives engagement on social media. *Proceedings of the National Academy of Sciences*, 118(26). <https://doi.org/10.1073/pnas.2024292118>

Rich, B. (2021). *table1: Tables of descriptive statistics in HTML*. Retrieved from <https://CRAN.R-project.org/package=table1>

Rosseel, Y. (2012). lavaan: An R package for structural equation modeling. *Journal of Statistical Software*, 48(2), 1–36. Retrieved from <https://www.jstatsoft.org/v48/i02/>

Rudis, B. (2020). *Hrbrthemes: Additional themes, theme components and utilities for 'ggplot2'*. Retrieved from <https://CRAN.R-project.org/package=hrbrthemes>

Schultz, P. W., Nolan, J. M., Cialdini, R. B., Goldstein, N. J., & Griskevicius, V. (2007). The Constructive, Destructive, and Reconstructive Power of Social Norms. *Psychological Science*, 18(5), 429–434. <https://doi.org/10.1111/j.1467-9280.2007.01917.x>

Shahi, G. K., Dirkson, A., & Majchrzak, T. A. (2021). An exploratory study of COVID-19 misinformation on Twitter. *Online Social Networks and Media*, 22, 100104. <https://doi.org/10.1016/j.osnem.2020.100104>

Shin, J., Yang, A., Liu, W., Min Kim, H., Zhou, A., & Sun, J. (2022). Mask-Wearing as a Partisan Issue: Social Identity and Communication of Party Norms on Social Media Among Political Elites. *Social Media + Society*, 8(1), 205630512210862. <https://doi.org/10.1177/20563051221086233>

- Smaldino, P. E. (2022). Models of Identity Signaling. *Current Directions in Psychological Science*, 096372142210756.
<https://doi.org/10.1177/09637214221075609>
- Sørensen, K., Van den Broucke, S., Pelikan, J. M., Fullam, J., Doyle, G., Slonska, Z., . . . Brand, H. (2013). Measuring health literacy in populations: Illuminating the design and development process of the European Health Literacy Survey Questionnaire (HLS-EU-Q). *BMC Public Health*, 13(1), 948.
<https://doi.org/10.1186/1471-2458-13-948>
- Tajfel, H., Turner, J. C., Austin, W. G., & Worchel, S. (1979). An integrative theory of intergroup conflict. *Organizational Identity: A Reader*, (65), 56–65.
- Urbanek, S. (2013). *Png: Read and write PNG images*. Retrieved from <https://CRAN.R-project.org/package=png>
- van Bavel, J. J., Baicker, K., Boggio, P. S., Capraro, V., Cichocka, A., Cikara, M., . . . Willer, R. (2020). Using social and behavioural science to support COVID-19 pandemic response. *Nature Human Behaviour*, 4(5, 5), 460–471.
<https://doi.org/10.1038/s41562-020-0884-z>
- Venables, W. N., & Ripley, B. D. (2002). *Modern applied statistics with s* (Fourth). New York: Springer. Retrieved from <http://www.stats.ox.ac.uk/pub/MASS4/>
- Vosoughi, S., Roy, D., & Aral, S. (2018). The spread of true and false news online. *Science*, 359(6380), 1146–1151. <https://doi.org/10.1126/science.aap9559>
- Wickham, H. (2016). *ggplot2: Elegant graphics for data analysis*. Springer-Verlag New York. Retrieved from <https://ggplot2.tidyverse.org>
- Wickham, H. (2019). *Stringr: Simple, consistent wrappers for common string operations*. Retrieved from <https://CRAN.R-project.org/package=stringr>
- Wickham, H. (2021a). *Forcats: Tools for working with categorical variables (factors)*. Retrieved from <https://CRAN.R-project.org/package=forcats>
- Wickham, H. (2021b). *Tidyr: Tidy messy data*. Retrieved from

850 <https://CRAN.R-project.org/package=tidyr>

851 Wickham, H., Averick, M., Bryan, J., Chang, W., McGowan, L. D., François, R., . . .

852 Yutani, H. (2019). Welcome to the tidyverse. *Journal of Open Source Software*,
853 4(43), 1686. <https://doi.org/10.21105/joss.01686>

854 Wickham, H., François, R., Henry, L., & Müller, K. (2021). *Dplyr: A grammar of*
855 *data manipulation*. Retrieved from <https://CRAN.R-project.org/package=dplyr>

856 Wickham, H., & Hester, J. (2021). *Readr: Read rectangular text data*. Retrieved
857 from <https://CRAN.R-project.org/package=readr>

858 Wilke, C. O. (2020). *Cowplot: Streamlined plot theme and plot annotations for*
859 *'ggplot2'*. Retrieved from <https://CRAN.R-project.org/package=cowplot>

860 Xitao, F., & Xiaotao, F. (2005). Power of Latent Growth Modeling for Detecting
861 Linear Growth: Number of Measurements and Comparison With Other Analytic
862 Approaches. *The Journal of Experimental Education*, 73(2), 121–139.

863 <https://doi.org/10.3200/JEXE.73.2.121-139>

864 Zarocostas, J. (2020). How to fight an infodemic. *The Lancet*, 395(10225), 676.
865 [https://doi.org/10.1016/S0140-6736\(20\)30461-X](https://doi.org/10.1016/S0140-6736(20)30461-X)

866 Zeileis, A. (2004). Econometric computing with HC and HAC covariance matrix
867 estimators. *Journal of Statistical Software*, 11(10), 1–17.

868 <https://doi.org/10.18637/jss.v011.i10>

869 Zeileis, A. (2006). Object-oriented computation of sandwich estimators. *Journal of*
870 *Statistical Software*, 16(9), 1–16. <https://doi.org/10.18637/jss.v016.i09>

871 Zeileis, A., & Grothendieck, G. (2005). Zoo: S3 infrastructure for regular and
872 irregular time series. *Journal of Statistical Software*, 14(6), 1–27.

873 <https://doi.org/10.18637/jss.v014.i06>

874 Zeileis, A., & Hothorn, T. (2002). Diagnostic checking in regression relationships. *R*
875 *News*, 2(3), 7–10. Retrieved from <https://CRAN.R-project.org/doc/Rnews/>

876 Zeileis, A., Kleiber, C., & Jackman, S. (2008). Regression models for count data in

877 R. *Journal of Statistical Software*, 27(8). Retrieved from
878 <http://www.jstatsoft.org/v27/i08/>
879 Zeileis, A., Köll, S., & Graham, N. (2020). Various versatile variances: An
880 object-oriented implementation of clustered covariances in R. *Journal of*
881 *Statistical Software*, 95(1), 1–36. <https://doi.org/10.18637/jss.v095.i01>

Table 1

*Descriptive summary of different forms of misinformation and control information sharing behaviors in Study 1:
Means, standard deviations (SD) and highest values within each experimental group.*

	Misinformation						Control			Ratio	
	Shared			Liked			Shared			Liked	
	Shared	Liked	Retweeted	Shared	Liked	Retweeted	Shared	Liked	Retweeted	Shared	Retweeted
Control	Mean	2.16	2.16	0.27	0.92	0.95	0.05	1.23	1.21	0.22	
	SD	3.91	3.83	0.77	1.77	1.85	0.25	3.41	3.36	0.72	
	Max	26.00	25.00	4.00	12.00	13.00	2.00	19.00	19.00	4.00	
Misinfo flag	Mean	1.37	1.27	0.17	0.40	0.40	0.03	0.97	0.87	0.14	
	SD	3.34	3.12	0.73	1.21	1.21	0.18	2.74	2.56	0.72	
	Max	26.00	26.00	6.00	10.00	10.00	2.00	19.00	19.00	6.00	
Social cue	Mean	1.46	1.28	0.28	0.49	0.45	0.08	0.97	0.83	0.21	
	SD	3.63	3.06	1.26	1.39	1.22	0.39	2.67	2.38	0.94	
	Max	23.00	22.00	14.00	9.00	8.00	4.00	18.00	15.00	10.00	
Combined	Mean	1.12	1.12	0.09	0.63	0.60	0.04	0.49	0.52	0.05	
	SD	2.29	2.17	0.40	1.34	1.29	0.26	1.84	1.79	0.37	
	Max	15.00	12.00	3.00	7.00	7.00	2.00	10.00	10.00	2.00	

Table 2

Summary of 1) negative binomial regression model predicting count of shared misinformation tweets, 2) negative binomial regression model predicting count of shared control tweets and, 3) linear regression model predicting ratio of shared misinformation and control tweets in Study 1.

	Misinfo				Control				Ratio			
	Estimate	SE	IRR	p	Estimate	SE	IRR	p	Estimate	SE	t	p
Intercept	1.60	0.56	4.97	.004	1.20	0.62	3.32	.053	1.91	0.72	2.66	.008
Group: Social cue	-0.40	0.21	0.67	.060	-0.54	0.23	0.58	.019	-0.26	0.27	-0.97	.332
Group: Misinfo flag	-0.43	0.21	0.65	.039	-0.78	0.23	0.46	.001	-0.25	0.27	-0.93	.351
Group: Combined	-0.66	0.23	0.52	.004	-0.32	0.24	0.72	.182	-0.75	0.29	-2.58	.010
Age	-0.02	0.01	0.98	.014	-0.06	0.01	0.94	< .001	-0.01	0.01	-0.48	.630
Education	-0.02	0.05	0.98	.702	0.00	0.06	1.00	.997	-0.03	0.07	-0.38	.701
Digital health literacy	-0.08	0.08	0.93	.365	-0.05	0.09	0.95	.552	-0.10	0.11	-0.96	.340
Political orientation	0.09	0.03	1.09	.009	0.11	0.04	1.11	.003	0.06	0.04	1.41	.158

Note. p values are bold if significant at $p < .05$; SE represents standard errors; IRR represents the incidence rate ratio (percentage change in the dependent variable per one-unit change in the predictor, either above or below 1); estimated Model 1 $\theta = 0.25$ and Model 2 $\theta = 0.26$.

Table 3

Summary of hierarchical 1) negative binomial regression model predicting count of shared misinformation tweets, 2) negative binomial regression model predicting count of shared control tweets and, 3) linear regression model predicting ratio of shared misinformation and control tweets with pooled data of Study 2 and Study 3.

	Misinfo			Control			Ratio		
	<i>Estimate</i>	<i>SE</i>	<i>p</i>	<i>Estimate</i>	<i>SE</i>	<i>p</i>	<i>Estimate</i>	<i>SE</i>	<i>p</i>
Intercept	-1.70	0.26	< .001	-5.74	0.62	< .001	0.55	0.11	< .001
Trial	0.00	0.08	> .999	0.68	0.17	< .001	0.00	0.04	.963
Reference: Personal network	-0.76	0.25	.003	-0.95	0.52	.067	-0.24	0.11	.024
Reference: Combined	-1.10	0.26	< .001	-0.11	0.41	.797	-0.31	0.11	.005
Reference: All users	-0.29	0.22	.199	0.02	0.36	.958	0.03	0.10	.787

Note. Predictors are bold if significant at $p < .05$.

Table 4

Summary of univariat latent growth curve model with dummy-coded experimental conditions predicting change in participants' injunctive norms across experimental trials: Estimates and standard errors.

	Model	
	Estimate (SE)	p
<u>Regression Slopes</u>		
<u>Intercept</u>		
Group: Combined	0.16(0.33)	0.636
Group: Personal network	0.15(0.34)	0.658
Group: All users	0.31(0.34)	0.354
<u>Slope</u>		
Group: Combined	−0.16(0.12)	0.177
Group: Personal network	−0.29(0.12)*	0.013
Group: All users	−0.18(0.12)	0.135
<u>Intercepts</u>		
Injunctive norms (t1)	0.00 ⁺	
Injunctive norms (t2)	0.00 ⁺	
Injunctive norms (t3)	0.00 ⁺	
Group: Combined	0.26 ⁺	
Group: Personal network	0.25 ⁺	
Group: All users	0.24 ⁺	
<u>Residual Variances</u>		
Injunctive norms (t1)	1.07(0.26) ^{***}	0.000
Injunctive norms (t2)	0.92(0.14) ^{***}	0.000
Injunctive norms (t3)	−0.36(0.26)	0.157
Group: Combined	0.19 ⁺	

Group: Personal network	0.19 ⁺	
Group: All users	0.18 ⁺	
<u>Latent Intercepts</u>		
Intercept	3.33(0.23) ^{***}	0.000
Slope	−0.16(0.08) [*]	0.050
<u>Latent Variances</u>		
Intercept	3.26(0.37) ^{***}	0.000
Slope	0.45(0.13) ^{***}	0.001
<u>Fit Indices</u>		
$\chi^2(\text{df})$	11.45(4) [*]	0.022
CFI	0.99	
TLI	0.97	
RMSEA	0.08	

⁺Fixed parameter

^{*} p<0.05, ^{**} p<0.01, ^{***} p<0.001

A

B

Figure 1. A) Example of Twitter's misinformation flag and B) of misinformation content with social reference cue (personal network) and standard misinformation flag.

Figure 2. Density plots, boxplots and stacked data points for model-based predictions of the A) count of shared misinformation tweets, B) count of shared control tweets and, C) ratio of shared misinformation and control tweets per experimental condition.

Figure 3. Studies 1 to 3: Effect plot depicting model estimates across models for shared misinformation tweets, shared control tweets and ratio thereof. A) depicts main effects of experimental conditions on all outcomes in Study 1, B) depicts main effects of different social references combined with misinformation flag in Studies 2 and 3.