IN PRESS AT APPLIED COGNITIVE PSYCHOLOGY

Others (Dis-) Endorse This So It Must (Not) Be True: High Relative Endorsement Increases Perceived Misinformation Veracity But Not Correction Effectiveness

Lucy H. Butler, Nicolas Fay, & Ullrich K. H. Ecker

School of Psychological Science

University of Western Australia

Word Count: 8495 (including references, excluding figures and figure captions)

Abstract

People increasingly rely on social-media platforms to access information; thus, understanding how platform characteristics influence belief in misinformation is important. Recent findings indicate perceived social endorsement of information (e.g., number of likes) can influence misinformation belief and correction acceptance. However, how the influence of endorsement may be modulated by concurrent *disendorsement* information (e.g., dislikes) is unclear. Across two experiments, we assessed the influence of relative endorsement on misinformation belief and correction acceptance. Experiment 1 exposed participants to claims and fact-checks with a high or low likes-to-dislikes ratio. Experiment 2 simplified the relative-endorsement of misinformation significantly increases misinformation belief, particularly when the endorsement information is presented as a single value. Conversely, relative endorsement had a negligible impact on correction effectiveness. This suggests perceived relative endorsement may influence belief primarily when other cues of information veracity are unavailable.

Key words: Misinformation; belief updating; social endorsement; social-media

Others (Dis-) Endorse This So It Must (Not) Be True: High Relative Endorsement Increases Perceived Misinformation Veracity But Not Correction Effectiveness

Social-media platforms have fundamentally changed the information environment. Information production has become decentralised, with the general population now having the capacity to create and share information with an undefined number of others (Flanagin, 2017). This is in direct contrast to the high barriers of information production and dissemination characteristic of traditional media, which, at least in principle, provide a mechanism for maintaining a set standard of information quality (Metzger & Flanagin, 2013). The lack of regulation on social-media is often cited as a key driver for the increased prevalence of misinformation—that is, false or otherwise misleading information presented as factual—in the information environment (Ecker et al., 2022). In fact, although blatant misinformation makes up a relatively small portion of the total information people encounter online (Allen et al., 2020), recent analyses of content on popular social-media platforms have highlighted that for some topics misinformation is relatively pervasive. For example, in an analysis of user-generated content on TikTok, 52 of the 100 most-viewed videos containing information about attention-deficit/hyperactivity disorder were classed as misleading (Yeung et al., 2022).

Not only has social-media increased the quantity of misinformation in the information environment, information propagated on social-media often lacks traditional cues of information credibility (Mena et al., 2020; Metzger & Flanagin, 2013). That is, conventional judgements of information credibility take source characteristics into account—specifically the source's trustworthiness and expertise—and information credibility tends to be evaluated in light of the credibility of the source (Metzger & Flanagin, 2013). In fact, perceived source credibility has been found to have a significant influence on belief in, and reliance on, misinformation, with people being more susceptible to misinformation from sources they deem to be credible (Swire et al., 2017). Additionally, high source trustworthiness, though not expertise, has also been linked to increased correction effectiveness (Ecker & Antonio, 2021). However, source information is often missing or ambiguous on social-media platforms, making it difficult, and in many cases impossible, to gauge source credibility (Mena et al., 2020). Given people's relatively high reliance on social-media to access news (Pew Research Center, 2022), it is important to better understand how people make judgements of information credibility on social-media, specifically when source credibility information is unavailable.

One key factor proposed to influence judgements of information credibility on socialmedia is perceived social endorsement of information. Based on social-validation theory, researchers have proposed that people may use social-endorsement information, such as quantity of post engagement (i.e., likes, shares), to gauge information credibility and, in turn, veracity (Jucks & Thon, 2017; Mena et al., 2020). Indeed, a study by Jucks and Thon (2017) assessed credibility judgements of health-related information when it was (vs. was not) associated with an expert source (quality cue), and when it was (vs. was not) socially validated by 'the masses' (quantity cue). Perceptions of credibility were found to be greater in both cue conditions compared to a no-validation control condition, with no statistical difference between quality-cue and quantity-cue conditions. This suggests people evaluate endorsement by a number of unknown others relatively equivalently to endorsement by a source with high credibility, despite no clear indication of how quantity of endorsement relates to information quality.

This has potentially negative implications when it comes to people's susceptibility to misinformation on social-media. Specifically, anyone can engage with social-media posts irrespective of subject expertise, or for reasons other than signalling information quality (e.g., for entertainment purposes; Madrid-Morales et al., 2021). People will thus not only be more

likely to see misinformation on social-media when it is highly engaged with (Weeks et al., 2017), but may also be more likely to believe said misinformation due to the perception that engagement signals credibility. Recent empirical research provides some support for the influence of perceived social endorsement on misinformation susceptibility, with a number of studies finding susceptibility to believe or share misinformation to be higher when false claims are associated with a high level of positive engagement (particularly "likes" and/or "shares") than with a low level of positive engagement (Avram et al., 2020; Butler et al., 2022; Shin et al., 2022). However, there is also evidence that belief in false information following a correction may be lower, at least initially, when the correction has a high (vs. low) level of positive engagement (Butler et al., 2022; Vlasceanu & Coman, 2021). This is potentially encouraging, as it suggests endorsement of fact-checks may at least somewhat counteract endorsement of misinformation.

Although these studies provide some evidence that endorsement information may influence belief in false information, other studies have found no effect of perceived endorsement on (mis)information belief (Koch et al., 2023; Mena et al., 2020). Further, almost all studies have exclusively focused on *positive* engagement (i.e., likes and shares), neglecting how negative engagement (which may be perceived as disendorsement; e.g., dislikes) might impact belief. From an applied perspective, the focus on positive engagement is unsurprising, given most social-media platforms only allow positive engagement (e.g., Instagram, TikTok) or do not make information about disendorsement easily available (e.g., YouTube). Though restricting reaction options that signal disendorsement has some benefits (e.g., arguably reducing large-scale online harassment; Wojcicki, 2022), the bias towards only presenting endorsement information may artificially inflate people's perceptions of actual information endorsement. Specifically, in isolation the number of likes provides minimal tangible information about the true level of social endorsement a piece of information has. That is, although a number of people may "like" a piece of information, this may only represent a small portion of the population that is exposed to said information. In fact, especially for false information, it is plausible that there is a greater number of people who disagree with the information presented in a post than agree with it (Altay, Hacquin, et al., 2022). If this is the case, the inclusion of disendorsement options on social-media platforms may help improve people's judgements of information veracity. By contrast, however, if misinformation receives a large degree of positive (relative to negative) endorsement (or corrections receive a large degree of negative endorsement relative to positive), relative endorsement may make people more susceptible to misinformation than they otherwise would be.

To the best of our knowledge, no studies have investigated the concurrent influence of positive and negative (mis)information engagements. However, some studies have assessed the influence of positive and negative *comments* on how people evaluate news content on social-media type platforms (Boot et al., 2021; Lewandowsky et al., 2019). These studies typically find negative comments to be more persuasive than positive comments, suggesting a negativity bias. If such a negativity bias is also seen with post engagements, disendorsement may have a stronger impact (i.e., enhance scepticism towards low-quality information) than endorsement, which may have net benefits to the extent that mainly low-quality information is disendorsed. Given some social-media platforms do allow for, and readily present, both positive and negative endorsement information (e.g., downvotes on Reddit or 9Gag; to a degree angry emoticons on Facebook), and other platforms have considered the inclusion of dislikes (e.g., X [formerly Twitter], TikTok), it is important to understand if and how disendorsement (vs. endorsement) information may impact people's ability to discern truth in the face of misinformation and corrections.

Understanding how people appraise both endorsements and disendorsements may also be useful given their potential application as a social-norm proxy. Specifically, explicit normative information can reduce the negative impacts of misinformation at least temporarily (Jones et al., 2023). However, adding normative information to posts is highly resource intensive and practically unviable at scale. Thus, although somewhat crude, making others' endorsement *and* disendorsement of information salient could be a cheap and scalable alternative to norm-based misinformation interventions.

As such, the broad aim of the current research was to assess the influence of relative endorsement (i.e., quantity of endorsement vs. quantity of disendorsement) on (1) belief in false claims and (2) fact-check effectiveness. In Experiment 1, both positive (likes) and negative (dislikes) endorsement information was presented to assess how people appraise the believability of misinformation and corrections with a high vs. low likes-to-dislikes ratio. In Experiment 2, the study design was simplified by presenting endorsement information as a single, salient percentage figure, thereby reducing participants' cognitive load.

Experiment 1

Experiment 1 aimed to assess the influence of relative endorsement (i.e., likes-todislikes ratio) on belief in misinformation. To this end, we implemented likes-to-dislikes ratios that reflected either high or low endorsement. In the high-endorsement condition, participants were exposed to false claims and corrective fact-checks with a high likes ratio (more likes than dislikes); in the low-endorsement condition, participants were exposed to false claims and fact-checks with a low likes ratio (fewer likes than dislikes). Apart from the addition of dislikes, the design of Experiment 1 was identical to that used in Butler et al. (2022). Specifically, participants received both false claims and true (filler) claims, which were subsequently fact-checked (i.e., a correction for false claims and an affirmation for true claims). Claim belief was measured both before and after the fact-check, as well as after a delay; participants also responded to inferential-reasoning questions designed to measure claim beliefs indirectly. It was hypothesized that there would be a main effect of relative claim endorsement on false-claim beliefs and misinformation-consistent inferential reasoning, with belief being greater when claims are associated with high (vs. low) relative endorsement. It was further hypothesized that there would be a main effect of relative fact-check endorsement on post-correction false-claim beliefs and misinformation-consistent inferential reasoning, with belief being lower (i.e., greater belief updating) when fact-checks were associated with high (vs. low) relative endorsement. Experiment 1 was pre-registered at https://osf.io/8ah4g/?view_only=4579e1f28f1446fda9ce8e94daf62857.

Method

Experiment 1 adopted a 2 × 2 within-subjects design, with factors claim endorsement (high vs. low relative endorsement) and fact-check endorsement (high vs. low relative endorsement). Belief in false claims was measured at three time-points: (1) during initial claim exposure (pre-fact-check), (2) immediately following exposure to the associated fact-check (post-fact-check), and (3) at the end of the experiment, following a short distraction (post-delay). Claim-consistent inferential reasoning was additionally measured at time-point 3. Claim veracity (false, true) was technically an additional within-subjects factor. However, as the aim of the current research was to assess belief in false claims, true claims were included primarily as fillers to provide some balance in participant experience. Analyses regarding true filler claims were conducted separately and for exploratory purposes only. Results for true filler claims are briefly described in the Results section; however, full analyses are available in the Supplement B available at

https://osf.io/ds6up/?view_only=c1b129c4846e4a6ab04c1811316f970f.

Participants

U.S.-based adult participants were recruited through Prolific. Minimum sample size was set to 352 (in line with Butler et al., 2022). Prolific workers who had participated in any of the studies in Butler et al. (2022) were excluded from the study invitation. To ensure adequate sample size post-exclusions, a total of 376 participants were recruited. Participants were excluded in accordance with the following a-priori criteria: Self-reported English proficiency rated as only "fair" or "poor" (n = 0); self-reported lack of effort (n = 0); completion time < 10 min. (n = 0); uniform responding (identical response to ≥ 90 out of 120 items; n = 0); inconsistent responding (where the difference between mean responses to standard vs. reverse-coded inference questions was identified as an outlier, using the outlier labelling rule with a 1.5 multiplier; n = 16). This resulted in a final sample size of N = 360; the sample comprised 92 males, 259 females, and 7 non-binary and 2 genderqueer individuals ($M_{age} = 37.34$, $SD_{age} = 14.21$, age range = 18-82).

Materials

The claims used (16 false; 8 true) were identical to those used in Butler et al. (2022). The claims were non-political and primarily health-based (e.g., *"Lemon and ginger tea is an effective immune detox"*). This was done to avoid any potential interaction between level of endorsement and political congruence, and because health-based misinformation can (1) have a substantial negative impact on people's behavior (e.g., reliance on alternative over traditional medicine in the treatment of life-threatening diseases), and is (2) relatively prolific on social-media (Borges do Nascimento et al., 2022). Claims used were sampled from an initial pool of claims rated on believability, familiarity, and engagingness in a previous material-validation study (see Supplement A for full details). Claims were presented in a mock social-media format, and associated with one of four fictitious sources. Fact-checks were presented in the same format, and provided a simple correction (for false claims) or

affirmation (for true claims) of the initial claim. Fact-checks were also associated with one of four (different) fictitious sources. Pairings of posts (claims and fact-checks) and sources was fixed such that each of the 16 possible combinations of false-claim and fact-check sources occurred exactly once for each participant; average claim believability was relatively equivalent across sources.

Each claim and each fact-check was associated with a number of likes and dislikes. Dependent on condition, posts were associated with more likes than dislikes (high relative endorsement) using a ratio of approximately 5:1, or fewer likes than dislikes (low relative endorsement) using a ratio of approximately 1:5. Some minor variability in specific values was included to increase realism. Quantity of both likes and dislikes ranged between 10 and 1000 across posts, and high and low endorsement conditions were either (1) matched on the number of likes, or (2) approximately matched on the total number of engagements (i.e., total number of likes and dislikes) to reduce any potential confound between the absolute and relative level of engagements across conditions. Please see Supplement A for specific details regarding how the quantity of likes and dislikes were generated.

The relative-endorsement levels of claims and fact-checks were fully crossed, resulting in four condition combinations (high/high; high/low; low/high; low/low). Assignment of claims and fact-checks to endorsement conditions was fully counterbalanced across participants such that each claim and fact-check combination was shown in each of the four endorsement conditions. See Figure 1 for example false claims and corrections across all high and low endorsement condition combinations.

Figure 1

Example False Claims and Associated Fact-Checks used in Experiment 1.

Claim	Fact-Check
InfoHQ @InfoHQ Most people only use between 10 and 50% of their brain capacity. 111 Q 556	B The Beacon @Beacon_Online X NO @InfoHQ: Most people do NOT only use between 10 and 50% of their brain capacity. 138 □ 28
Cranberry juice is an effective treatment for urinary-tract infections.	The Post The Post This NOT TRUE that cranberry juice is an effective treatment for urinary-tract infections. @Cactus_Mag 137 Q 686
Vibe Magazine @VIBE_Magazine Hypnosis is an effective way to retrieve repressed memories. 497 99	S Straight Up @Straight_Up FALSE: "Hypnosis is an effective way to retrieve repressed memories." @VIBE_Magazine 288 \$ 58
B Buzz News Buzz News	E The Enquire @TheEnquire_ @BuzzNews_: "Vitamin E enriched oil will make scars fade." FALSE ★ 111 □ 557

Note. Examples show all claim and fact-check sources, and all four condition combinations, that is, from top to bottom, the low/high (low claim, high fact-check endorsement), low/low, high/high, and high/low combinations. Note that true claims (paired with affirmative fact-checks) were also included (see Supplement).

Measures. Outcome variables were identical to those used in Butler et al. (2022).

Specifically, claim belief was measured through direct ratings on an 11-point Likert scale

ranging from 0 ("certainly false") to 10 ("certainly true"). Claim-specific reasoning was measured through 32 inferential-reasoning questions, with two questions (one reverse-coded) per claim. Each question presented participants with a statement related to the claim (an example reverse-coded item for the false claim shown in Figure 1 is "*A person's functioning would be reduced if they only used half of their brain capacity*") and participants rated their agreement on a Likert scale ranging from 0 ("strongly disagree") to 10 ("strongly agree")—in the following, these ratings are referred to as inference scores. See Supplement A for all claims and associated inference questions.

Procedure

The experiment was run using Qualtrics software (Qualtrics, Provo, UT). Participants first received an ethics-approved information sheet, and completed a short demographics questionnaire. All claims (including the eight true filler claims) were presented individually and in a randomised order. For each claim, participants rated their initial belief (pre-fact-check; time 1) before being presented with the associated fact-check and then providing a second belief rating (post-fact-check; time 2), on a separate page, before moving to the next claim. After presentation of all claims, participants were presented with a 1-minute distraction task (a word puzzle). This was followed by the inference questions (with question pairs presented in a randomised order) and a final belief rating for each claim (post-delay; time 3). See Figure 2 for a visual depiction of the experimental procedure.

Upon completion of the main task, participants were asked to self-report whether their data should be discarded due to a lack of effort (participants were aware that a "yes" response would have no negative consequences), before being debriefed. The debrief informed participants that the corrections and affirmations were accurate, to the best of our knowledge, whereas the endorsement information was simulated. Median completion time was 20 minutes; participants were compensated £2.50 (US\$3.40 at time of data collection).

Figure 2

Experimental Procedure for Experiments 1 and 2



Results

All data and additional analyses are available at

https://osf.io/ds6up/?view_only=c1b129c4846e4a6ab04c1811316f970f. Given our outcome variables were measured on ordinal scales, analyses were conducted using cumulative-link mixed-effects modelling (CLMM; McElreath, 2020). All analyses were conducted using the *clmm* function of the *ordinal* R package (Christensen, 2018), and all data visualizations were created using *ggplot2* (Wickham, 2016).

Prior to statistical analyses, the fixed effects of claim endorsement and fact-check endorsement were centered, and time (pre-fact-check, post-fact-check, and post-delay) was factor-coded. Dependent variables—belief and inference scores—were coded as ordinal factors with 11 levels (0-10). The maximum random-effects structure justified by the experimental design was included for each analysis where possible (Barr et al., 2013). Specifications of random-effects structures are provided in the Supplement B.

Belief in False Claims

Mean belief ratings across time-points and endorsement levels are presented in Figure 3. We first assessed the results of claim endorsement and fact-check endorsement at time 1 (pre-fact-check). This was done to assess the effect of claim endorsement immediately, and to ensure fact-check-endorsement conditions did not significantly differ prior to fact-check exposure. The best-fitting model specified claim endorsement as a fixed effect, $\beta = .28$, SE = .05, z = 5.46, p < .001, with belief in false claims with high relative endorsement greater than belief in false claims with low relative endorsement. As expected, there was no significant effect of fact-check endorsement, nor an interaction between claim and fact-check endorsement (ps > .898).

Figure 3



Mean False-Claim Belief Ratings Across Timepoints and Endorsement Level



To assess the effect of fact-check endorsement, false-claim belief was assessed across time-points 1 (pre-fact-check) and 2 (post-fact-check). The best-fitting model included fixed effects for claim endorsement, $\beta = .29$, SE = .06, z = 4.72, p < .001, and time, $\beta = -1.05$, SE = .03, z = -30.68, p < .001. There was no significant claim endorsement by time interaction, $\beta = -.09$, SE = .07, z = -1.40, p = .162. The effect of claim endorsement mirrored

that at time 1, and the effect of time indicated that belief in false claims reduced following the fact-checks. By contrast, there was no significant effect of fact-check endorsement, nor a fact-check endorsement by time interaction (ps > .724).

To assess whether the effect of claim endorsement was maintained over time, postfact-check and post-delay belief was assessed across time-points 2 and 3. The best-fitting model included a fixed effect of claim endorsement, $\beta = .20$, SE = .06, z = 3.20, p = .001, as well as a claim endorsement × time interaction, $\beta = -.14$, SE = .07, z = -2.01, p = .044. The interaction effect suggests the influence of claim endorsement significantly reduced over time. To quantify this reduction, we assessed the effect of claim-endorsement level only at time 3 (post-delay) and found it to be nonsignificant, $\beta = .07$, SE = .06, z = 1.32, p = .187. There was additionally no significant effect of fact-check endorsement, nor a fact-check by time interaction across time-points 2 and 3 (ps > .085).

Inference Scores for False Claims

There was no significant effect of claim endorsement, fact-check endorsement, nor an interaction between claim endorsement and fact-check endorsement on participants' reliance on the false claims in their inferential reasoning (i.e., inference scores; all ps > .165)

True claims

Results of the exploratory analyses of true claims are briefly described below; see Supplement B for details. Relative claim endorsement had a significant influence on belief in true claims at time-point 1; belief in true claims was significantly higher in the high (vs. low) relative endorsement condition. However, this effect reduced to nonsignificant after the factcheck, and remained nonsignificant after the delay. There was a significant fact-check endorsement × time interaction between time-points 1 and 2, indicating that the influence of fact-check endorsement was significantly greater following the fact-check. The magnitude of this effect significantly reduced after the delay.¹

True-claim inference scores were also analysed. There was a fixed effect of factcheck endorsement, indicating reliance on true information was greater when affirmative fact-checks had a high (vs. low) level of endorsement.

Discussion

Experiment 1 assessed the influence of relative endorsement on belief in misinformation and correction effectiveness. The results indicate that relative endorsement can influence belief in misinformation: Consistent with predictions, we observed an effect of relative claim endorsement at time-points 1 (pre-fact-check) and 2 (post-fact-check), with belief in false claims being higher when claims were associated with a high compared to low level of relative endorsement. However, all other effects (i.e., claim-endorsement effect after a delay; all fact-check-endorsement effects) were nonsignificant, suggesting the effect of relative endorsement on misinformation appraisal was limited to initial exposure. As such, on the whole the current findings suggest the influence of relative endorsement on misinformation belief and correction acceptance is small and relatively fleeting.

A potential explanation for the modest effect of relative endorsement, that could explain why only initial claim-endorsement level influenced belief, is the cognitive load placed on participants. Specifically, taking endorsement levels into account required participants to attend to (1) the claim presented in the post, (2) the number of likes, and (3) the number of dislikes. Participants were also required to contrast the number of likes to the number of dislikes to determine the relative level of endorsement (i.e., was there a greater

¹ Although the influence of fact-check endorsement remained significant at timepoint 3, we note the size of this effect is relatively equivalent to pre-fact-check levels (where any effect would by definition be spurious). Thus, the post-delay effect may reflect differences in true claim belief that are independent of actual endorsement effects.

number of likes than dislikes, or vice versa) before using this information to evaluate the claim. On top of this, for the fact-check, participants were required to encode whether the initial claim was deemed to be true or false, as well as the relative number of likes and dislikes associated with the fact-check (and how this level of relative endorsement compared to the initial claim endorsement). This is a non-trivial amount of information processing, especially given the number of claims (16 false and 8 true) participants were exposed to. Thus, it is plausible the majority of participants' cognitive resources were allocated to information that was central to the task (i.e., the claim itself, and whether the claim was corrected or affirmed) rather than information that is peripheral (here, the relative level of endorsement). To assess whether this was the case, we ran Experiment 2 where the relative-endorsement information was simplified into a single value (percentage endorsement) to reduce cognitive load.

Another possible explanation for the pattern of results in Experiment 1 is that the level of endorsement presented in the high-endorsement condition (approximately 80%) may have in fact signalled relatively high disendorsement. In fact, an analysis of videos on YouTube— one of the few social-media platforms to allow dislikes—found that only around 17% of videos contained less than 80% likes (Emplifi, 2014). Given this, it is plausible that people might consider approximately 80% likes as relatively low endorsement. In support of this, a study by Aklin and Urpelainen (2014) found that presenting expert consensus on non-polarised scientific issues as 80% significantly *decreased* participants' support for related policies compared to when this endorsement information was not provided. As such, an auxiliary aim of Experiment 2 was to assess whether increasing the level of endorsement in the high-endorsement conditions would meaningfully increase belief and belief updating.

Experiment 2

The design of Experiment 2 was identical to that of Experiment 1 bar three changes. Most notably, to reduce cognitive load, the level of endorsement was presented as a percentage of likes (as a proportion of total endorsements), which expressed either a high or low relative level of endorsement. Secondly, to allow for an exploratory test of whether 80% endorsement is insufficiently high to produce an enduring effect on belief and belief updating, participants were randomly split into two high-endorsement groups. In the first group, high endorsement was set in line with Experiment 1 (percentage endorsement was sampled from a normal distribution with M = 80, SD = 7.5, truncated at 65 and 100). In the second group, high endorsement was set to a mean of approximately 90% (percentage endorsement was sampled from a normal distribution with M = 90, SD = 7.5, truncated at 75 and 100). Finally, six filler claims (four false, two true) and fact-checks with a moderate endorsement level were added to the experimental design (endorsement for these claims was sampled from a normal distribution with M = 50, SD = 10, truncated at 35 and 65). This was done to increase believability and reduce potential demand characteristics; these filler claims were not analysed. The primary aims and hypotheses for Experiment 2 were identical to Experiment 1. Experiment 2 was pre-registered at

https://osf.io/3rs57/?view_only=99ae177aca9c43419f2aaa32137e2f1f.

Method

As with Experiment 1, Experiment 2 adopted a 2×2 within-subjects design, with factors claim endorsement (high vs. low endorsement) and fact-check endorsement (high vs. low endorsement). Belief in false claims was measured at three time-points (pre-fact-check; post-fact-check; post-delay). Claim-congruent inferential reasoning was measured at timepoint 3. Claim veracity (false, true) was again technically an additional within-subjects factor. As in experiment 1 true-claim results are briefly described in the Results section; full analyses are provided in the Supplement B.

Participants

Adult participants based in the U.S. were recruited through Prolific. Minimum sample size was again set to 352. We recruited a total of 380 participants; in line with Experiment 1, participants were excluded in accordance with a-priori criteria: self-reported English proficiency rated as only "fair" or "poor" (n = 1); self-reported lack of effort (n = 0); completion time < 10 min. (n = 0); uniform responding (identical response to $\ge 75\%$ of 138 items²; n = 0); and inconsistent responding (n = 19). This resulted in a final sample size of N = 360 (n = 185 in the 80% high-endorsement condition; n = 175 in the 90% high-endorsement condition); the sample comprised of 183 males, 170 females, and 7 non-binary individuals ($M_{age} = 39.65$, $SD_{age} = 13.91$, age range = 18-81).

Materials

Target claims (16 false, 8 true) were identical to Experiment 1. The additional four myths and two facts used as moderate-endorsement filler items were similar in style to the target claims (an example false filler claim is "*Hair and nails continue to grow after death*"). The full list of target and filler claims is available in Supplement A.

Presentation of claims and fact-checks was identical to Experiment 1 with the exception of the presentation of the endorsement information. Specifically, claims and fact-checks were associated with a level of endorsement expressed as a single percentage of likes (compared to total endorsement, i.e., theoretical likes and dislikes) (see Figure 4). Dependent on condition, claims and fact-checks were presented with a low level of endorsement (sampled from a normal distribution with M = 20, SD = 7.5, truncated at 0 and 35) or a high

² Due to an oversight of the addition of filler items, the pre-registration specified exclusion based on uniform responding to ≥ 90 of 120 items, as in Experiment 1. To account for this oversight, this exclusion criterion slightly deviates from the pre-registration.

level of endorsement (sampled from a normal distribution with M = 80, SD = 7.5, truncated at 65 and 100 for approximately half of the participants, and from a normal distribution with M = 90, SD = 7.5, truncated at 75 and 100 for the other half of participants). The moderateendorsement filler claims and fact-checks had values sampled from a normal distribution with M = 50, SD = 7.5, truncated at 35 and 65. Endorsement levels of claims and fact-checks were fully crossed, resulting in four condition combinations (high/high; high/low; low/high; low/low). Assignment of claims and fact-checks to endorsement conditions was fully counterbalanced across participants. Endorsement level for filler claims and fact-checks remained constant across participants.

Figure 4

Example Claim and Associated Fact-Check used in Experiment 2



Note. The above example illustrates the high/low (high claim, low fact-check endorsement) condition.

Measures. Outcome variables were identical to Experiment 1. Additional belief ratings were included for filler claims; however, inference questions for the filler claims were not included.

Procedure

Bar the addition of filler items, the procedure was identical to Experiment 1. Median completion time was approximately 22.75 minutes, and participants were compensated £2.70 for their time.

Results

For the primary analyses, data were prepared and analysed in line with Experiment 1. Data and supplementary analyses are available at

https://osf.io/ds6up/?view_only=c1b129c4846e4a6ab04c1811316f970f.

Belief in False Claims

Mean belief ratings across time-points and endorsement levels are presented in Figure 5. We first assessed results of claim endorsement and fact-check endorsement at time 1 (pre-fact-check), to assess the effect of claim endorsement immediately and ensure factcheck-endorsement conditions did not differ prior to the fact-checks. The best-fitting model specified claim endorsement as a fixed effect, $\beta = .52$, SE = .06, z = 8.44, p < .001, with belief in false claims with high endorsement greater than belief in false claims with low endorsement. As expected, there was no significant effect of fact-check endorsement, nor an interaction between claim and fact-check endorsement (ps > .346).

To assess the initial effect of fact-check endorsement, and to assess whether the influence of claim endorsement persisted, false-claim belief was assessed across time-points 1 (pre-fact-check) and 2 (post-fact-check). The best-fitting model included fixed effects of claim endorsement, $\beta = .55$, SE = .07, z = 8.21, p < .001, and time, $\beta = -1.08$, SE = .03, z = -31.42, p < .001, as well as a claim endorsement × time interaction, $\beta = -.17$, SE = .07, z = -2.55, p = .011. The effect of claim endorsement mirrored that at time 1, and the effect of time indicated that belief in false claims reduced following fact-checks. To scrutinize the claim endorsement × time interaction, we assessed the influence of claim-endorsement level

isolated at time 2. Claim-endorsement level remained significant at time 2, $\beta = .39$, SE = .06, z = 6.78, p < .001, suggesting the interaction was driven by the size claim-endorsement effect decreasing post-fact-check. By contrast, there was no significant effect of fact-check endorsement, nor a fact-check endorsement by time interaction across time-points 1 and 2 (*ps* > .183).

Figure 5

Mean False Claim Belief Ratings Across Timepoints and Endorsement Levels



Percentage Fact-Check Endorsement 🗾 High 📃 Low



To assess whether the effect of claim endorsement remained significant after a delay, post-fact-check and post-delay belief was assessed across time-points 2 and 3. The bestfitting model included a fixed effect of claim endorsement, $\beta = .38$, SE = .06, z = 6.24,

p < .001, as well as a claim endorsement × time interaction, $\beta = -.21$, SE = .07, z = -3.17, p < .001. To scrutinize this interaction, we assessed the influence of claim-endorsement level isolated at time 3. Claim-endorsement level remained significant at time 3, $\beta = .15$, SE = .05, z = 2.87, p = .004, suggesting again that the interaction was driven by the claim-endorsement effect decreasing further after the delay.

Finally, there was also a significant fixed effect of fact-check endorsement across time-points 2 and 3, $\beta = -.16$, SE = .06, z = -2.67, p = .007, with belief in corrected false claims higher when fact-checks were associated with low compared to high endorsement. There was no significant effect of time, $\beta = .02$, SE = .03, z = 0.66, p = .508, nor a fact-check by time interaction, $\beta = .07$, SE = .07, z = 1.09, p = .275. Given the non-significant effect of fact-check endorsement across times 1 and 2, but significant effect across times 2 and 3, we assessed the influence of fact-check endorsement at time 2, $\beta = -.15$, SE = .06, z = -2.56, p = .010, with the effectiveness of fact-checks being greater (i.e., lower belief in false claims) when they were associated with high endorsement. However, this effect of fact-check endorsement on belief was nonsignificant at time 3 (post-delay), $\beta = -.08$, SE = .05, z = -1.61, p = .108.

Inference Scores for False Claims

Mean inference scores across conditions are presented in Figure 6. There was a significant effect of claim endorsement on inference scores, $\beta = .20$, SE = .05, z = 4.05, p < .001. In line with claim belief, reliance on false claims was significantly higher when claims were associated with a high compared to a low level of relative endorsement. There was no significant effect of fact-check endorsement, nor an interaction between claim endorsement and fact-check endorsement on false-claim inference scores (ps > .293).

Figure 6





Percentage Fact-Check Endorsement 🗾 High 📃 Low

Note. Error bars represent 95% confidence intervals.

Influence of 80% vs. 90% Endorsement

To provide insight into whether the increased influence of endorsement, specifically claim endorsement, in Experiment 2 relative to Experiment 1 was driven by the reduced cognitive load, the increased level of high endorsement in the 90%-endorsement condition, or some combination of the two, exploratory cumulative-linked mixed-effects models were run that included high-endorsement level (80% vs. 90%) as an additional between-subjects factor; see supplement B for details. There was no influence of high-endorsement level, nor an interaction between claim endorsement and high-endorsement level, at any belief-rating time-

point or on the inference scores. This suggests that false-claim belief was not significantly increased by the higher 90% level of endorsement. That being said, numerically the difference in mean claim belief between the high and low claim-endorsement conditions at time-points 1 and 3, as well as on the inference scores, was greater in the 90% compared to the 80% endorsement condition. As such, the inclusion of the 90% condition likely increased the size of the effect of claim endorsement in Experiment 2 slightly.

Influence of Endorsement on Belief in True Claims

For completeness, results of the exploratory analyses of true claims are briefly described below; see Supplement B for details. Percentage claim endorsement had a significant influence on belief in true claims at all time-points; belief in true claims was significantly higher in the high (vs. low) endorsement condition. There was also a significant claim endorsement \times time interaction between time-points 1 and 2, indicating a reduction in the claim-endorsement effect from time 1 (pre-fact-check) to time 2 (post-fact-check).

There were additionally significant fact-check endorsement × time interactions between time-points 1 and 2—indicating high fact-check endorsement being associated with greater belief updating—and time-points 2 and 3—indicating the influence of fact-check endorsement reduced after the delay; however, the effect remained significant at time 3.

For true-claim inference scores, there was a significant influence of claimendorsement level, with reliance on true claims significantly greater when associated with high compared to low relative endorsement. There was no significant influence of fact-check endorsement, nor a claim \times fact-check endorsement interaction effect on true-claim inference scores.

Discussion

Experiment 2 was a conceptual replication of Experiment 1, with endorsement information reduced to a single (percentage) value to reduce participant cognitive load. As an

auxiliary factor, a stronger manipulation was used for the high endorsement condition for half of the participants (i.e., an average endorsement level of 90% rather than 80%). This was done to provide insight into whether the modest effect of claim endorsement level observed in Experiment 1 was due to the high endorsement condition signalling substantial disagreement.

Overall, the influence of claim endorsement level on false claim belief was consistent with predictions: Belief in false claims was significantly greater when claims were associated with a high compared to low level of relative endorsement. Although the strength of the effect reduced over time, the effect remained significant after the fact-check and after the delay. Claim endorsement level also had a significant effect on the inference scores, with indirect reliance on false claims significantly higher when the original claim was associated with a high compared to low level of endorsement. The size and persistence of the effect of claim endorsement in the current study was greater than in Experiment 1: Although this increase seemed to be partially driven by the reduced cognitive load placed on participants in the study, exploratory analyses suggest the increase may also have been enhanced by the slightly stronger manipulation provided to half of the participants. However, regardless of the specific mechanism driving the effect, the current study provides clear evidence that belief in false claims can be modulated by perceived level of social endorsement.

By contrast, the effect of fact-check endorsement on direct and on indirect false claim belief was minimal: Although there was an effect of fact-check endorsement across times 2 (post-fact-check) and 3 (post-delay), seemingly driven by differences immediately post-factcheck, the effect of fact-check endorsement did not significantly increase after fact-check exposure, nor was the effect significant after the delay. Additionally, there was no significant effect of fact-check endorsement on indirect belief in the corrected false claims (i.e., inference scores). These findings are relatively consistent with Experiment 1, and suggests endorsement of fact-checks may have a negligible impact on their effectiveness.

General Discussion

The overarching aim of the current research was to assess whether perceived level of relative endorsement of misinformation and corrections has a meaningful influence on belief in (corrected) misinformation. It was hypothesized that belief in and reliance on false claims would be higher when the claims were associated with a high compared to a low level of relative endorsement. It was also hypothesized that belief in and reliance on corrected misinformation would be lower (i.e., greater belief updating) when fact-checks were associated with a high compared to a low level of relative endorsement.

Our results provide evidence that endorsement of false claims can modulate belief in misinformation, with high (vs. low) relative claim endorsement associated with significantly greater false-claim belief. This finding is consistent with hypotheses and past research (e.g, Avram et al., 2020; Butler et al., 2022; Shin et al., 2022) and suggests people's belief in false information can be meaningfully influenced by the perception that there is (or is not) a consensus level of agreement with the information. We note that enduring effects of claimendorsement level on false-claim belief, and an influence on inferential reasoning, were only observed in Experiment 2, suggesting the influence of endorsement information is particularly strong when the level of endorsement is both (1) expressed as a single value (and thus likely more salient) and (2) high (i.e., > 80%). Although the effect of claim-endorsement is somewhat modest in size even in experiment 2 (i.e., approx. half the size of the effect of the fact-check), it is noteworthy given that participants had to process a non-trivial amount of information and the endorsement information was peripheral to the task. We do note, however, that in the current study the endorsement information provided was fictional, and thus the differences in endorsement levels may have provided a clearer signal to that

observed in the real world. While this may have enhanced the size of the effect on belief in misinformation, it is equally plausible that endorsement level may have been disregarded in part due to it being perceived as implausible. Future research that uses actual levels of social endorsement (as has been done in studies investigating the impact of expert-consensus levels, for example in the climate-change domain; Cook, 2016) can provide a robustness test of the findings reported here.

The current findings suggest the presence of simple endorsement and disendorsement information has the capacity to meaningfully impact online belief formation. If true, allowing easy disendorsement of information on social-media platforms (i.e., "disliking") could be a beneficial mechanism for reducing belief in low-quality information, especially for information where some level of endorsement is to be expected (e.g., endorsement of vaccine misinformation by anti-vaccinationists), or where the barriers to providing explicit corrective information are sufficiently high (e.g., due to fear of negative social implications). In fact, the current findings suggest the way both endorsement and disendorsement information is presented on certain social-media platforms, such as the relative level of endorsement present on Reddit (displayed as a single figure representing level of endorsement minus level of disendorsement), may help people be less susceptible to misinformation than they would be on platforms that only provide positive engagement information (e.g., TikTok, X, Instagram). Given it is relatively cheap and does not require direct content moderation by platforms, adapting platforms that currently only allow users to positively endorse information to include a disendorsement option (e.g., a dislike button) may thus be a simple, non-obtrusive way to reduce the negative impact of misinformation spread online.

However, we add the caveat that the efficacy of such a mechanism will depend on users engaging with content (i.e., liking or disliking) based at least to some extent on the veracity of the information in question. It is of course well-established that people engage with information for reasons unrelated to veracity (Altay, de Araujo, et al., 2022), and in many cases veracity may be unknown to the user or even undeterminable in principle (e.g., due to lack of evidence or because a claim is not verifiable or falsifiable; also see Moberger, 2020). We also note that exploratory analyses of the true claims suggest that (low) relative endorsement may have equally strong (negative) effects on true-claim beliefs. Thus, in situations where information veracity is not prioritised (e.g., political echo-chambers; Törnberg, 2018), or where the prevalence of low-quality information is low, providing users with both endorsement and disendorsement options may have a net negative impact. Given such situations are relatively common, it may in fact be beneficial to remove information about quantity of engagements from social-media platforms altogether.

In contrast to the influence of claim endorsement on belief in misinformation, the influence of fact-check endorsement across both studies was negligible: In Experiment 1, belief in corrected false claims did not significantly differ between the high and low fact-check endorsement conditions. In Experiment 2, although there was an effect of fact-check endorsement immediately post-fact-check in the hypothesized direction, the endorsement effect did not significantly differ between pre- and post-fact-check, and there was no significant effect after the delay. This suggests that the significant effect of fact-check endorsement seen immediately post-fact-check may have been partially driven by slight, spurious condition differences in belief levels *pre*-fact-check. Thus, we found no strong evidence that the effectiveness of a fact-check is meaningfully influenced by others' endorsement of it. This result is inconsistent with past research (e.g., Butler et al., 2022; Vlasceanu & Coman, 2021), although previous studies only looked at the influence of positive endorsement (i.e., likes and shares) on fact-check effectiveness. The inclusion of negative endorsement information may have made participants more sceptical of the validity

of both types of endorsement information, though further research is required to draw strong conclusions regarding this.

The influence of endorsement on claim belief but not fact-check effectiveness could also be explained by differences in the level of ambiguity associated with claims and factchecks. That is, claims in the current study were relatively ambiguous and not associated with any cues of information quality other than the social-endorsement information. Thus, it seems rational for participants to rely to some extent on the endorsement information provided to gauge claim veracity. That is, if a participant has no knowledge about a claim, the claim will provide negligible impetus towards any particular belief rating; thus, the endorsement information may have relatively large utility during the decision-making process. By contrast, fact-checks contain an explicit signal, namely, to increase or reduce one's belief relative to one's initial rating (depending on whether the fact-check is affirmative or corrective). This clear signal likely diminishes the utility of the social-endorsement information. One implication of this is that the influence of social-endorsement information on belief formation and updating may vary based on the availability and strength of other quality cues, either inherent in the specific post itself (e.g., message persuasiveness) or its context (e.g., sourcecredibility cues). It follows then that when other quality cues are available, the influence of generic social-endorsement information on belief formation may be weaker or nonmeaningful. Therefore, future research that includes multiple-and potentially conflictingcues of information quality may be beneficial to understand whether, and if so how, any effect of claim endorsement persists across contexts.

Given the large effect of fact-checks per se on belief in misinformation, the absence of an effect of fact-check endorsement may also be desirable in many situations, as high disendorsement of fact-checks may not meaningfully decrease their effectiveness. This would have positive implications for countering misinformation that a vocal minority agrees with (e.g., climate-change misinformation; Dinan et al., 2022). However, the claims used in the current study were non-controversial and beliefs were reported privately; thus, participants may have placed low importance on the belief decisions (Chaiken & Maheswaran, 1994), and results may not generalize to situations where endorsement or rejection of information could have personal or social implications.

In summary, this study provides evidence that the level of perceived endorsement and disendorsement can influence people's belief in misinformation, particularly when information quality is ambiguous. This may have implications for the design of social-media platforms, which at present often deny users' the opportunity to disendorse information. In fact, under certain conditions allowing and displaying low-friction endorsement *and* disendorsement of information on social-media, as is possible on the likes of Reddit, could have a beneficial impact on belief formation, especially when fact-checks are not present or easy to provide.

References

- Aklin, M., & Urpelainen, J. (2014). Perceptions of scientific dissent undermine public support for environmental policy. *Environmental Science and Policy*, 38, 173–177. https://doi.org/10.1016/j.envsci.2013.10.006
- Allen, J., Howland, B., Mobius, M., Rothschild, D., & Watts, D. J. (2020). Evaluating the fake news problem at the scale of the information ecosystem. *Science Advances*, 6(14). https://doi.org/10.1126/sciadv.aay3539F
- Altay, S., de Araujo, E., & Mercier, H. (2022). "If this account is true, it is most enormously wonderful": Interestingness-if-true and the sharing of true and false news. *Digital Journalism*, 10(3), 373–394. https://doi.org/10.1080/21670811.2021.1941163
- Altay, S., Hacquin, A. S., & Mercier, H. (2022). Why do so few people share fake news? It hurts their reputation. *New Media and Society*, 24(6), 1303–1324. https://doi.org/10.1177/1461444820969893
- Avram, M., Micallef, N., Patil, S., & Menczer, F. (2020). Exposure to social engagement metrics increases vulnerability to misinformation. *The Harvard Kennedy School Misinfromation Review*, 1(5). https://doi.org/10.37016/mr-2020-033
- Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, 68(3), 255–278. https://doi.org/10.1016/j.jml.2012.11.001
- Boot, A. B., Dijkstra, K., & Zwaan, R. A. (2021). The processing and evaluation of news content on social media is influenced by peer-user commentary. *Humanities and Social Sciences Communications 2021 8:1*, 8(1), 1–11. https://doi.org/10.1057/s41599-021-00889-5
- Borges do Nascimento, I. J., Beatriz Pizarro, A., Almeida, J., Azzopardi-Muscat, N., André Gonçalves, M., Björklund, M., & Novillo-Ortiz, D. (2022). Infodemics and health

misinformation: a systematic review of reviews. *Bulletin of the World Health Organization*, *100*(9), 544–561. https://doi.org/10.2471/BLT.21.287654

- Butler, L. H., Fay, N., & Ecker, U. K. H. (2022). Social endorsement influences the continued belief in corrected misinformation. *Journal of Applied Research in Memory* and Cognition. https://doi.org/10.1037/MAC0000080
- Chaiken, S., & Maheswaran, D. (1994). Heuristic processing can bias systematic processing: Effects of source credibility, argument ambiguity, and task importance on attitude judgment. *Journal of Personality and Social Psychology*, 66(3), 460–473. https://doi.org/10.1037//0022-3514.66.3.460
- Christensen, R. H. B. (2018). Cumulative link models for ordinal regression with the R Package ordinal. *Unsubmitted Manuscript*, 1–40. www.jstatsoft.org/
- Cook, J. (2016). Countering Climate Science Denial and Communicating Scientific Consensus. In Oxford Research Encyclopedia of Climate Science. Oxford University Press. https://doi.org/10.1093/acrefore/9780190228620.013.314
- Dinan, W., Bernardi, C. L., Esteves, V., & Harkins, S. (2022). Online climate denialism. In *The Routledge Handbook of Environment and Communication* (pp. 292–305). Routledge. https://doi.org/10.4324/9781003119234-23
- Ecker, U. K. H., & Antonio, L. M. (2021). Can you believe it? An investigation into the impact of retraction source credibility on the continued influence effect. *Memory & Cognition*, 49(4), 631–644. https://doi.org/10.3758/s13421-020-01129-y
- Ecker, U. K. H., Lewandowsky, S., Cook, J., Schmid, P., Fazio, L. K., Brashier, N., Kendeou,
 P., Vraga, E. K., & Amazeen, M. A. (2022). The psychological drivers of
 misinformation belief and its resistance to correction. *Nature Reviews Psychology*, 1(1),
 13–29. https://doi.org/10.1038/s44159-021-00006-y
- Emplifi. (2014). YouTube Videos: What's not to like?

https://emplifi.io/resources/blog/youtube-videos-what-s-not-to-

like?utm_source=socialbakers.com

- Flanagin, A. J. (2017). Online social influence and the convergence of mass and interpersonal communication. *Human Communication Research*, 43(4), 450–463. https://doi.org/10.1111/hcre.12116
- Jones, C. M., Diethei, D., Schöning, J., Shrestha, R., Jahnel, T., & Schüz, B. (2023). Impact of Social Reference Cues on Misinformation Sharing on Social Media: Series of Experimental Studies. *Journal of Medical Internet Research*, 25, e45583. https://doi.org/10.2196/45583
- Jucks, R., & Thon, F. M. (2017). Better to have many opinions than one from an expert? Social validation by one trustworthy source versus the masses in online health forums. *Computers in Human Behavior*, 70, 375–381. https://doi.org/10.1016/J.CHB.2017.01.019
- Koch, T. K., Frischlich, L., & Lermer, E. (2023). Effects of fact-checking warning labels and social endorsement cues on climate change fake news credibility and engagement on social media. *Journal of Applied Social Psychology*, *53*(6), 495–507. https://doi.org/10.1111/jasp.12959
- Lewandowsky, S., Cook, J., Fay, N., & Gignac, G. E. (2019). Science by social media: Attitudes towards climate change are mediated by perceived social consensus. *Memory* and Cognition, 47(8), 1445–1456. https://doi.org/10.3758/s13421-019-00948-y
- Madrid-Morales, D., Wasserman, H., Gondwe, G., Ndlovu, K., Sikanku, E., Tully, M.,
 Umejei, E., & Uzuegbunam, C. (2021). Motivations for Sharing Misinformation: A
 Comparative Study in Six Sub-Saharan African Countries. *International Journal of Communication*, 15, 1200–1219. http://ir.lsu.ac.zw/handle/123456789/532

McElreath, R. (2020). Statistical rethinking: A bayesian course with examples in R and stan.

In *Statistical Rethinking: A Bayesian Course with Examples in R and Stan*. CRC Press. https://doi.org/10.1201/9781315372495

- Mena, P., Barbe, D., & Chan-Olmsted, S. (2020). Misinformation on Instagram: The impact of trusted endorsements on message credibility. *Social Media + Society*, 6(2), 205630512093510. https://doi.org/10.1177/2056305120935102
- Metzger, M. J., & Flanagin, A. J. (2013). Credibility and trust of information in online environments: The use of cognitive heuristics. *Journal of Pragmatics*, 59, 210–220. https://doi.org/10.1016/j.pragma.2013.07.012
- Moberger, V. (2020). Bullshit, Pseudoscience and Pseudophilosophy. *Theoria (Sweden)*, 86(5), 595–611. https://doi.org/10.1111/theo.12271
- Pew Research Center. (2022). *More Americans are getting news on TikTok, bucking social media trend*. https://www.pewresearch.org/fact-tank/2022/10/21/more-americans-are-getting-news-on-tiktok-bucking-the-trend-on-other-social-media-sites/
- Shin, I., Wang, L., & Lu, Y. (2022). Twitter and endorsed (fake) news: The influence of endorsement by strong ties, celebrities, and a user majority on credibility of fake news during the covid-19 pandemic. *International Journal of Communication*, 16, 2573–2595.
- Swire, B., Berinsky, A. J., Lewandowsky, S., & Ecker, U. K. H. (2017). Processing political misinformation: comprehending the Trump phenomenon. *Royal Society Open Science*, 4(3), 160802. https://doi.org/10.1098/rsos.160802
- Törnberg, P. (2018). Echo chambers and viral misinformation: Modeling fake news as complex contagion. *PLoS ONE*, *13*(9). https://doi.org/10.1371/journal.pone.0203958
- Vlasceanu, M., & Coman, A. (2021). The impact of social norms on health-related belief update. *Applied Psychology: Health and Well-Being*, 85–91. https://doi.org/10.1111/aphw.12313

Weeks, B. E., Lane, D. S., Kim, D. H., Lee, S. S., & Kwak, N. (2017). Incidental Exposure,

Selective Exposure, and Political Information Sharing: Integrating Online Exposure Patterns and Expression on Social Media. *Journal of Computer-Mediated Communication*, 22(6), 363–379. https://doi.org/10.1111/JCC4.12199

Wickham, H. (2016). Data Analysis. 189-201. https://doi.org/10.1007/978-3-319-24277-4_9

- Wojcicki, S. (2022). *Letter from Susan: Our 2022 Priorities*. https://blog.youtube/insideyoutube/letter-susan-our-2022-priorities/
- Yeung, A., Ng, E., & Abi-Jaoude, E. (2022). TikTok and Attention-Deficit/Hyperactivity Disorder: A Cross-Sectional Study of Social Media Content Quality. *Canadian Journal* of Psychiatry. https://doi.org/10.1177/07067437221082854