

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Sensitivity to Online Consensus Effects Within Individuals and Claim Types

Permalink

<https://escholarship.org/uc/item/2rm2155p>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 46(0)

Authors

Alister, Manikya

Ransom, Keith James

Connor Desai, Saoirse

et al.

Publication Date

2024

Peer reviewed

Sensitivity to Online Consensus Effects Within Individuals and Claim Types

Manikya Alister¹ (alisterm@student.unimelb.edu.au)

Keith Ransom¹ (keith.ransom@unimelb.edu.au)

Saoirse Connor Desai² (saoirse.c.d@gmail.com)

Ee Von Soh¹ (eevons@student.unimelb.edu.au)

Brett Hayes³ (b.hayes@unsw.edu.au)

Andrew Perfors¹ (andrew.perfors@unimelb.edu.au)

¹School of Psychological Sciences, University of Melbourne

²School of Psychology, University of Sydney

³School of Psychology, University of New South Wales

Abstract

When reasoning about a claim, it makes sense to be more persuaded if lots of other people agree. But, there are many factors that make weighing the evidence behind a consensus complicated. For example, a consensus might be more or less informative depending on the type of claim, or whether each consensus member formed their opinions independently. These factors might also influence people differently depending on their own assumptions or preferences. In this study we used a mock social media paradigm to assess how persuaded people were by two factors: the presence of consensus (no consensus vs. consensus), and source independence (a consensus based on independent information sources vs. a consensus formed off shared, dependent sources). We varied these factors at both the group and individual level. At the group level, we assessed a third factor: whether people were influenced by the type of claim being reasoned about (we assessed 60 different claims divided into 4 categories). Almost everyone was more persuaded by consensus trials compared to no consensus trials. However, the strength of this effect was credibly stronger if the claim was likely to have a ground truth. We found that around one third of participants were sensitive to source independence. Of these, three quarters were more persuaded by a consensus based on independent sources, but the quarter who were more persuaded by dependent sources were persuaded just as strongly.

Keywords: consensus; persuasion; source independence; social reasoning; individual differences

Introduction

When encountering a new opinion or claim, there is often no obvious way to quickly and accurately determine its veracity. One of the ways we approach this situation is to rely on cues like how many people agree with it or not. For example, suppose you read a social media post claiming that “perfect avocados are getting harder to find”. While you may have your own thoughts on this, you know that your personal experience with avocados may not be sufficient: not only do you not know other peoples’ experiences, you also lack expertise in agriculture or information about avocado quality trends. Given this, it makes sense to give this claim more weight if there is a consensus of opinions supporting it.

These kinds of “consensus” effects, where people tend to be more convinced by something that multiple people agree with, have been demonstrated extensively (e.g., Asch, 1956; Lewandowsky, Gignac, & Vaughan, 2013; Ransom, Perfors, & Stephens, 2021; Franzen & Mader, 2023). However, reasoning about consensus in everyday life can be complex. One such complexity is that the persuasiveness of the consensus might be influenced by the topic or type of claim being reasoned about. For instance, Yousif, Aboody, and Keil (2019) found that people weighted consensus evidence differently

for different claims. When the claim was about a new tax policy and the sources were economists, people were less persuaded than when it was about an event at a local school and the sources were eyewitnesses. The authors concluded that people may reason differently about claims that are less knowable (like an economic prediction) compared to those that have a clear ground truth (like an eyewitness case). This is consistent with work demonstrating that the expertise of the sources matters (Maddux & Rogers, 1980; Simmonds, Stephens, Searston, Asad, & Ransom, 2023), probably in part because experts have more insight into the ground truth of a situation. While it is reasonable to think that people reason differently about consensus based on the “knowability” of the claim, knowability has not previously been systematically manipulated over a wide variety of claims, so it remains unclear how much it matters or how robust the difference is.

Another issue is that *many* factors go into evaluating complex claims, and individuals differ in the extent to which they weight those factors. These factors include the prestige (Atkisson, O’Brien, & Mesoudi, 2012) or confidence (Sah, Moore, & MacCoun, 2013) of the source of the claim, or the complexity of the arguments given in support of it (Zemla, Sloman, Bechlivanidis, & Lagnado, 2017). Moreover, reasoners realise that people making arguments in support of a claim (e.g., on social media) not only have different levels of competence (Lin, Spence, & Lachlan, 2016), they also vary in their goals, which include persuasion, identity signalling, trolling, and others (Pucci, Kashima, & Perfors, 2023). If individuals have different assumptions about the nature or importance of these factors, this will mean that they differ in how sensitive they are to consensus effects in the first place. However, there is very little work evaluating individual differences in susceptibility to consensus effects rather than group-level aggregate behaviour.

Yet another issue is that a consensus should be more convincing when everyone within the consensus reached their opinions *independently* from each other (Whalen, Griffiths, & Buchsbaum, 2018; Harkins & Petty, 1987; Yousif et al., 2019; Connor Desai, Xie, & Hayes, 2022; Xie & Hayes, 2022). If multiple social media posts from different people agree that perfect avocados are getting harder to find, and they all reference different, independent sources supporting this claim (e.g., different scientific studies or surveys), one would think that this sort of “independent consensus” *should* be more convincing than if each of these people referenced the same source (“dependent consensus”).

This idea that people should give more weight to claims corroborated by multiple independent sources than dependent (repeated) sources is supported by normative models of decision-making (Whalen et al., 2018; Xie & Hayes, 2022). However, in the real world, it is often unclear which primary sources have influenced someone’s opinions. Even if the primary sources that influenced people in a consensus are known, one might be unsure whether those sources are truly independent, since those sources could have collaborated or used the same underlying data. These difficulties mean that people might believe that a consensus is not independent when it actually is, or that what looks like an independent consensus actually is not. For instance, most COVID-19 anti-vaccination views originated from the same few people (Center for Countering Digital Hate, 2021) and the majority of climate denial blogs rely on the same few primary sources (Harvey et al., 2018).

Consistent with this, Connor Desai et al. (2022) argued that the results from Yousif et al. (2019) could be explained by participants doubting that the economists were truly independent, since different experts could have been relying on the same primary data or sources. Indeed, Connor Desai et al. (2022) found that when the independence of sources was emphasised, people took the consensus into account more. Alister, Perfors, and Ransom (2022) sought to investigate whether this explanation would generalise across a wider range of real-world claims using a paradigm that emphasised source independence via a realistic social media re-post design. In it, the nature of the consensus could be naturally portrayed as retweets; when the consensus was dependent, multiple users retweeted the same primary study, and when it was independent they retweeted different primary studies. People were insensitive to source independence unless it was clear that the primary source had influenced the re-poster’s opinion, in which case the effect very small.

In summary, the emerging consensus¹ in this literature is that the independence of a consensus bears little weight as a reasoning cue except in specific contexts and scenarios, and even then, the effect is usually small. However, there are several large limitations to be aware of. Firstly, it is usually assumed that an independent consensus *should* be the more persuasive, particularly when compared to a consensus where multiple people relied on the same source. However, given that in the real world the reliability of the source is unknown, a dependent consensus may indicate higher source reliability: a single source that has influenced several people might be considered to be more reliable, since multiple independent people have evaluated it and decided that it was worth considering (Pilditch, Hahn, Fenton, & Lagnado, 2020). Indeed, when asked specifically about their preferences, people often do not say that they prefer an independent consensus (Connor Desai et al., 2022). At the very least, people might differ in the extent to which they use source independence to judge reliability or competence.

¹Ha! See what we did there?



Figure 1: **Experiment stimuli.** Sample posts from each condition. The top panel is from the INDEPENDENT condition, where each person re-posted a different source (here, arguing in favour of the claim). Those in the middle are from the DEPENDENT condition, where each person re-posted the same source (here, against it). The bottom row shows the CONTESTED condition, where an equal number of people support/oppose the claim. Each trial in the actual experiment showed four posts rather than two, and the sources, names, photos, text, dependence, and pro/con direction were randomised for each person and claim. In this example, the sources are news organisations, but for some claims they were universities.

This leads us to a second limitation, which is that most research has focused on group level effects rather individual differences. As discussed, many factors go into the evaluation of how much weight to give to a consensus; if people evaluate those factors differently, they may vary widely in when or how much they pay attention to different kinds of consensus. This individual variation may look like a null or weak effect on the group level but reflect interesting and sophisticated reasoning at the individual level (Xie & Hayes, 2022).

A third limitation is that most consensus studies have only looked at a limited number of claims or claim types at a time. Although this idea that the people might reason about consensus differently for different kinds of claims is not new (Yousif et al., 2019; Alister et al., 2022), no study has systematically investigated this question across a large set of claims with predefined claim types. Only using a small set of claims also makes examining individual differences difficult, as there is not enough power to get reliable estimates of a single person’s behaviour (Smith & Little, 2018).

The current study aimed to address the limitations discussed above by being the first study, to our knowledge, to focus specifically on claim type and individual level differences in consensus effects. This was achieved by extending the realistic social media stimulus set used in Alister et al. (2022) to include more claims (60 instead of 12). Because each participant evaluated all of these claims, we obtained a sufficient quantity of data to enable reliable analyses for each person. Including more claims also meant we were able to see whether different kinds of claims elicited different effects.

Method

Participants 115 participants were recruited from Prolific Academic and paid £5.25 per session for up to two 35 minute sessions.² 40 were removed based on our pre-registered exclusion criteria; 27 for having lower than 90% accuracy on our comprehension checks (see Procedure) and 13 due to only completing the first session. Ages ranged from 18 to 72 years old ($M = 36$) and 54% were female. All were pre-screened as being both native English speakers and English being their primary language.

Procedure After providing consent and passing a short quiz regarding the instructions of the task, each participant saw 60 trials over the course of two sessions over two separate days (30 trials per day). Each trial began with participants viewing a claim (e.g., “Narcissists are more politically engaged”) after which they were asked to rate the extent to which they agreed with that claim using a slider from 0 to 100. They then viewed four social media posts by four distinct users. To ensure that people were engaging with the task properly, they had to indicate whether each post was arguing for or against the claim. The 27 participants who were less than 90% accurate at this task were removed from the analysis.

As shown in Figure 1, each post took the form of a re-post and included the primary source being re-posted, the primary data referred to by the source, and the user’s own words explaining how the primary source persuaded them. Because the four users were always distinct people with unique profile photos and names, this made it clear that the user had read the source and that it had influenced their opinion about the claim; this was shown to be important in Alister et al. (2022).

After confirming each post had been read (by rating its stance towards the claim), participants were once again asked to indicate how much they agreed with the claim (using the same 0-100 slider as before, but initialised to indicate their initial rating). The difference between their rating before and after represented their degree of belief revision due to the posts. After completing all 60 trials of the experiment, they were asked which strategy they used to evaluate the claims.

Consensus conditions Our primary manipulation (within participant) was whether each of the four social media posts cited the same primary source (DEPENDENT consensus, 20 trials) or all different (INDEPENDENT consensus, 20 trials).

In these *full consensus* trials, all four of the posts agreed with each other, arguing either for (PRO) or against (CON) the claim. The stance of the consensus was randomised, although we ensured that there would be an equal number of PRO and CON trials in each claim type and consensus condition.

We also had a CONTESTED condition (20 trials), where instead of the four posts all agreeing with each other, two of them agreed with the claim and two opposed it. Assignment of claim to condition was randomised within-participant. This condition was included as a baseline so that we could identify participants who did not change their beliefs more in the presence of consensus. It also served to reduce demand effects and ensure participants read all of the posts (otherwise they could pass the manipulation check after reading only one post and assuming all the others agreed with it).

Regardless of condition, all four posters and all sources in a given consensus gave essentially the same reason in different words (e.g., those arguing PRO on the Avocado claim all pointed out that climate change was making avocados worse, and those arguing CON all pointed out that genetic modification has improved avocado quality). Order of claims and posts as well as assignment of avatars and names was completely randomised across participants and claims.

In all conditions, the primary source was always either a news organisation or a university, and the primary data was always some kind of study or investigation carried out specifically by that organisation (made clear through the wording of the post). For example, the primary data could be a study by the University of Springfield, and the primary source would be the official account of the University of Springfield.

In the DEPENDENT condition, each of the four posts (by four different users) re-posted the same article by the same source (hence the same primary data). In the INDEPENDENT and CONTESTED conditions, the source of the post that was re-posted and the data that the source referred to were both distinct for each of the four posts: Person A cited Source X, Person B cited Source Y, and so forth. Thus, each was re-posting an independent source and referring to independent primary data. There were always three primary sources and one expert testimony, which was included to add some variety and reduce demand effects. We tried to maintain a balance between how many trials used each source type, but in some cases it only made sense to have a particular source type (e.g., a university would not conduct a study about whether a mayor ran into a burning building).

The news companies were real media companies chosen via the website AllSides,³ which allows people to rate the bias of different news companies. We chose news companies that were mid-range in popularity and deemed “centrist” by the raters. The universities were a sample of real universities ranked between 100 and 200 by the QS World University Rankings. Although whether the sources were news companies or institutions on each trial was deliberately chosen based on the appropriateness to the claim, the actual

²Preregistration: https://aspredicted.org/979_7LC

³<https://www.allsides.com/unbiased-balanced-news>

company/institution included in each trial was randomised for each participant and were always unique. The companies/institutions always had a profile photo, full name, and a “verified” tick to signal authenticity. All content was fictitious. The content that was not part of Alister et al. (2022) was first generated using chatGPT-3.5-turbo-1106 and then refined manually to ensure realism and variety.

Claim type conditions In addition to varying the nature of the consensus across conditions within-participant, we also varied the nature of the claims themselves. Given that previous research suggests more “knowable” claims are more likely to induce consensus independence effects (Yousif et al., 2019), we selected 30 *knowable* claims and 30 *unknowable* ones. Each category was further subdivided based on the way in which it was knowable or not. For instance, KNOWABLE EYEWITNESS claims include something that somebody could have seen (e.g., a mayor saving a child from a burning building). A KNOWABLE FACT is something that while verifiable in principle, you would need to be an expert in order to do so (e.g., a new species of jellyfish being discovered). An UNKNOWNABLE EXPERT claim is one which does not have a known ground truth at the moment, but nevertheless expertise would be helpful in evaluating (e.g., economic forecasting). Lastly, UNKNOWNABLE PREFERENCE claims do not have a ground truth and expertise is less likely to be important (e.g., whether flying is a better super power than invisibility). Each participant saw 15 of each of the four claim types and claim types were randomly assigned to a consensus condition for each person. The range and strength of prior beliefs endorsed by participants varied considerably across claims, with people tending to be less certain about the KNOWABLE claims.⁴

Results

Aggregate behaviour

We first explore the extent to which belief revision is affected by the knowability of a claim and the type of consensus. To account for the fact that CON trials (where the consensus argued against the claim) would shift beliefs in the opposite direction, the pre and post belief scores for CON trials were reversed so that the update would be in the same direction as the PRO trials. As Figure 2 shows, the largest belief updates occurred in the full consensus trials (INDEPENDENT and DEPENDENT, where all four posts took the same stance towards the claim (all in favour, or all against)). People were much more persuaded by these trials than CONTESTED ones in which half of the posts argued in one direction and half in the other. In the CONTESTED trials, people actually tended to believe the claim *less* after seeing the posts, although the degree of belief revision was very small.

The figure suggests the presence of a small but consistent difference between INDEPENDENT and DEPENDENT consensus trials, with participants more convinced when sources are

⁴The full set of claims and their prior distributions can be found here: tinyurl.com/33xws5x2.

independent. It also suggests that although people are convinced by a consensus on all claim types, they changed their beliefs more when the claims were KNOWABLE.

To quantitatively assess the persuasiveness of different kinds of consensus as a function of the type of claim, we compared four nested Bayesian generalised linear models using the brms (2.20.4; Bürkner, 2018) package in R (4.2.2) in which the outcome variable was the rating after reading the four posts. To assess the relative performance of each model, we compared them using the leave-one-out cross-validation criterion (LOOIC; Vehtari, Gelman, & Gabry, 2017). Our focus was on two kinds of consensus effects, so we ran each of the models on the subset of data relevant to them, as described in the two subsections below.

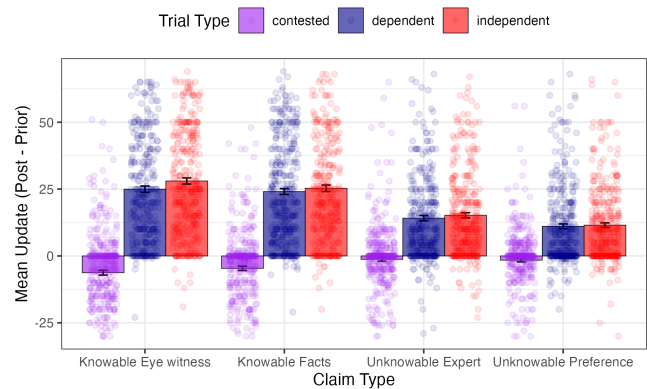


Figure 2: Belief update as a function of trial type and claim type. Positive belief update indicates belief change in the direction of the consensus for the full consensus trials. INDEPENDENT trials (red) showed four users who shared the same opinion and cited *different* primary sources. DEPENDENT trials (blue) showed four users who shared the same opinion but cited the *same* primary source. CONTESTED trials (purple) showed four users, half agreeing with the claim and half disagreeing, with each user citing different sources. The axes have been slightly constrained to better show the mean differences, so some individual data points are not visible.

Contested vs full consensus We first asked whether people changed their belief more when the four posts agreed (a full consensus) than when half argued in one direction and half in the other (a CONTESTED consensus). In order to quantitatively test this, we compared the INDEPENDENT condition trials to the CONTESTED condition trials which were directly comparable on the other factor manipulated: in both conditions, the four posts had distinct, independent sources. DEPENDENT trials differed from CONTESTED trials on both factors, and were thus excluded from this comparison.

Results of the full model comparison for this data set are in Table 1. The best model (M4) reported credible main effects of both consensus type (INDEPENDENT vs CONTESTED) as well as an interaction with claim type⁵. The main effect of consensus suggests that people were more convinced by INDEPENDENT trials compared to CONTESTED trials (see purple and red bars in Figure 2.). The main effect of claim type

⁵Space constraints meant we were unable to report the coefficients, but the size and direction of the effects is shown qualitatively in Figure 2

Table 1: **Model comparison of belief in the claim (0-100) after INDEPENDENT versus CONTESTED trials.** M1 only considered participants’ prior beliefs about the claim. M2 also considered whether there was a consensus or not, and M3 added claim type. While M1-M3 only considered main effects, M4 also considered the interaction between the presence of a consensus and claim type and was favoured by LOOIC.

Model	LOOIC	SE	Rank
M1. Prior	27693	100	4
M2. Prior + Consensus	26633	116	3
M3. Prior + Consensus + Claim Type	26565	117	2
M4. Prior + Consensus \times Claim Type	26335	121	1

suggests people were most convinced by KNOWABLE eye-witness claims, and follow up comparisons revealed credible differences between each claim type. The model also suggested an interaction between the nature of the consensus and the type of claim. The interaction suggests that the difference in belief between a full consensus and no consensus was larger when the claims were more knowable.

Dependent vs. Independent Consensus Our second question was whether people reasoned differently for INDEPENDENT and DEPENDENT trials, so we excluded the CONTESTED trials for this analysis. The full model comparison results favouring model M3, shown in Table 2, reveal that the independence of the consensus and the type of claim were both important to belief revision (with no interaction⁶). Although the main effect of independence was credibly greater than zero in the winning model, it was quite small relative to the other consensus comparison (see the difference between the blue and red bars in Figure 2).

Individual Differences

A key aim of this experiment was to quantify how different consensus effects emerged within individuals. Specifically, we were interested in both *how many* people displayed the two different consensus effects explored above, and *how much* of an effect there was. We therefore looked at individual behaviour on the same two questions: To what extent do people change their beliefs more when there is a full vs CONTESTED consensus? Or an INDEPENDENT vs DEPENDENT one? We can identify four types of participants. First, some people might be insensitive to consensus at all, reasoning similarly when there is a full consensus (four people agreeing) as when it is CONTESTED (two people on each side). Second, some people might be sensitive to the presence of a consensus but insensitive to source independence. Third, some people might be more convinced by a consensus when the sources are INDEPENDENT. And finally, some might be more convinced when the sources are DEPENDENT.

We relied on Bayesian linear models to quantitatively classify the different kinds of participants. Each participant was

⁶As per the pre-registration, we also ran all of the group level models with just two claim types (UNKNOWNABLE vs. KNOWABLE) but there was still no interaction.

Table 2: **Model comparison of belief in the claim (0-100) after INDEPENDENT versus DEPENDENT trials.** The models were the same as Table 1, but considered independence rather than the standard consensus. M3 was favoured by LOOIC.

Model	LOOIC	SE	Rank
M1. Prior	26952	103	4
M2. Prior + Independence	26946	103	3
M3. Prior + Independence + Claim Type	26493	111	1
M4. Prior + Independence \times Claim Type	26496	111	2

fit to a model in which the outcome variable was that person’s belief in a claim after seeing the posts. We then compared two models, as shown in Figure 3. As a baseline, we considered a Null model in which the only predictor was that person’s prior beliefs. This was compared to an Alternative model that also included a predictor corresponding to the effect in question.

The top panel of Figure 3 shows, for each person, whether they were best fit by the Null model (blue) or by the Alternative (red) that included a predictor corresponding to Consensus (CONTESTED vs INDEPENDENT). All but three participants (96%) were best described by the Alternative model, suggesting that the vast majority of people were more convinced by a four people agreeing than by two people on each side. The median size of the effect was 23, meaning that belief change on the full consensus (INDEPENDENT) trials tended to be around 23 points higher than than on the CONTESTED trials. There was a lot of individual variation, however, ranging from close to zero change to over 40 points of shift.

The bottom panel of Figure 3 explores sensitivity to independence. It shows, for each person, whether they were best fit by the Null or an Alternative model that included a predictor corresponding to Independence (DEPENDENCE vs INDEPENDENT). The majority of people were best fit by the Null model, but 22% were more persuaded on INDEPENDENT trials (positive on the y axis) and 8% were more persuaded on DEPENDENT trials (negative). Although the majority of participants showed no sensitivity to independence, these results support the idea that there are substantial individual differences and that *amount* of sensitivity, for at least some people, is considerably higher than what group level estimates would suggest. For example, at the group level, belief change was only 2 points higher on INDEPENDENT trials, but the median estimate for the 22% best fit by the Alternative model was 9 points. In addition, the median estimate for belief change for the 8% of participants who were more convinced by a DEPENDENT consensus was just as strong (also 9).⁷

⁷We also ran the model on each session separately, to see if this was a reliable indicator of individual differences. There was a strong positive correlation ($r = .57$) between estimates in the two sessions for participants who were best fit by the Alternate model overall. In other words, people who changed their beliefs more on INDEPENDENT trials in the first session also did so in the second session, and vice-versa for people who changed more on DEPENDENT trials.

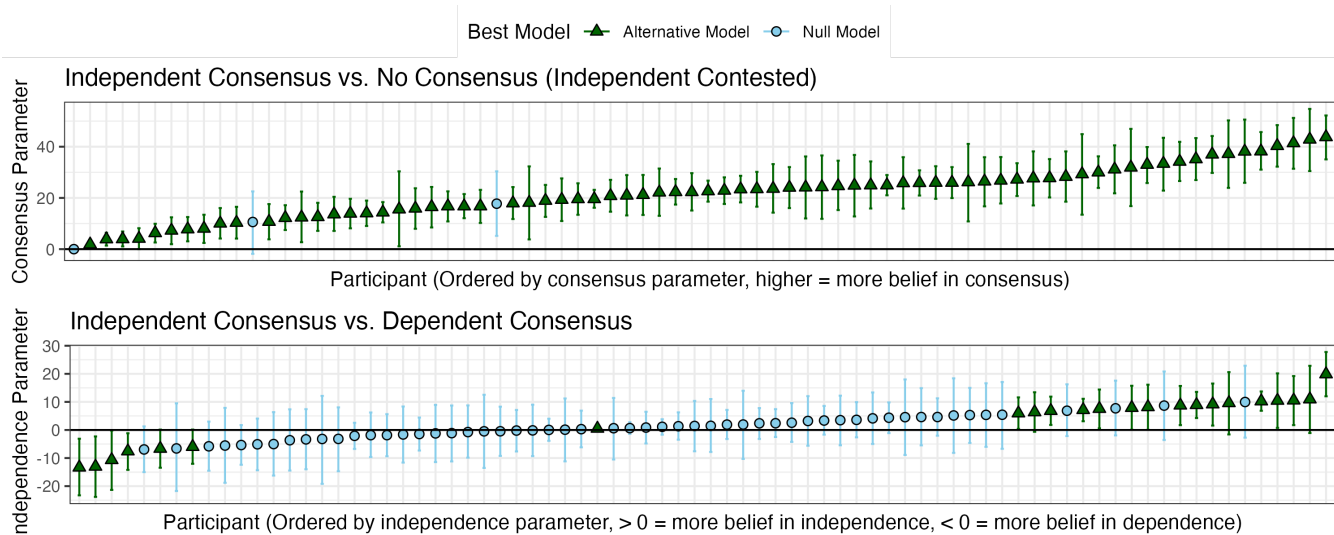


Figure 3: **Individual-level results.** (Top panel) For each person (x axis), the estimated persuasiveness of an INDEPENDENT consensus relative to CONTESTED consensus. Green indicates those participants whose behaviour was best captured by an Alternative model containing a predictor for Consensus (i.e., INDEPENDENT vs CONTESTED). The y axis indicates the extent to which a participant was more persuaded by a full consensus (in the INDEPENDENT trials). The parameter is shown by the point and its 89% credible interval (Kruschke, 2014) by the lines. (Bottom panel) For each person, the estimated persuasiveness of an INDEPENDENT consensus relative to DEPENDENT consensus, with red participants being sensitive to independence in either direction.

General Discussion

In line with previous literature (e.g., Ransom et al., 2021), there was a large consensus effect such that a consensus of posts quoting diverse sources was more convincing than no consensus where the posts disagreed with each other. Given the substantial number of claims that participants assessed, we were in a novel position to quantify this effect at the individual level, showing that almost everyone updated their beliefs in line with the consensus.

We were also able to systematically demonstrate that this effect was stronger for claims that were more “knowable”—that is, claims that were more likely to have a ground truth. While this finding has been suggested in reference to consensus *independence* effects, we are the first to show that this applies to standard consensus effects. This finding makes sense in that if it is impossible for any one person to know the truth of a claim, an aggregate of opinions might not be that convincing. However it could also be considered counter intuitive, since if no single person can know the answer definitively, it should make more sense to consult a range of people. This finding has important real world implications. It means that if a bad actor wanted to reduce the public’s belief in a claim that has a consensus among experts, they do not necessarily need to reduce peoples’ perception that a consensus exists so long as they can create enough doubt such that the public begins to believe that that the claim is unknowable.

Consistent with a number of recent studies, we found a small effect of consensus independence, such that at the group level, people tended to be more convinced by a consensus where members all cited different sources compared to those that all cited the same sources (Yousif et al., 2019; Simmonds et al., 2023; Connor Desai et al., 2022). However, our individual-level analyses provided important insight into what underlies these small group-level effects. Indeed, the majority of participants were completely insensitive to source

independence, but those that were sensitive changed their beliefs much more than the group-level effects would suggest.

Further, although most literature in this area has argued that people *should* be more convinced by an independence consensus, a small subset of participants were more convinced by a dependent consensus (see Xie & Hayes, 2022). Not only did they prefer this kind of consensus, but they were convinced to an equivalent degree as those who preferred an independent consensus. These individual differences provide insight into why there might have been such small and inconsistent group-level effects in previous studies as not accounting for these participants who prefer a dependent consensus weakens any group-level independence effects will weaken any group level effects. What (if anything) is driving these individual differences in sensitivity to consensus independence? One explanation is that some people care more about the reliability of a source compared to the number of source, however our analyses do not allow us to infer this explanation directly.

Although we found a small group level effect of independence, this effect was not influenced by the claim type (unlike for the standard consensus). This finding therefore fails to support the hypothesis posited by Yousif et al. (2019) who suggested that people are more sensitive to source independence when the claim is more knowable. Importantly however, due to insufficient power we were not able to assess these claim type effects at the individual level. It is possible that the extent to which different claim types influence participants’ sensitivity to source independence differs at the individual level, which should be examined in future research. Further, we chose which claims fit into which claim type category somewhat arbitrarily, based on the opinions of some of the authors and chatGPT, which might mean that some of the claims were not very representative of the true category. That being said, we were still able to differentiate between them for the standard consensus effect comparisons.

Acknowledgments

This work was supported by the Defence Science & Technology Group.

References

- Alister, M., Perfors, A., & Ransom, K. (2022). Source independence affects argument persuasiveness when the relevance is clear. Retrieved 2023-11-02, from <https://osf.io/5yx36> doi: 10.31234/osf.io/5yx36
- Asch, S. E. (1956). Studies of independence and conformity: I. A minority of one against a unanimous majority. *Psychological Monographs: General and Applied*, 70(9), 1–70. doi: 10.1037/h0093718
- Atkisson, C., O'Brien, M. J., & Mesoudi, A. (2012). Adult Learners in a Novel Environment Use Prestige-Biased Social Learning. *Evolutionary Psychology*, 10(3), 147470491201000309. doi: 10.1177/147470491201000309
- Bürkner, P.-C. (2018). Advanced Bayesian Multilevel Modeling with the R Package brms. *The R Journal*, 10(1), 395. doi: 10.32614/RJ-2018-017
- Center for Countering Digital Hate. (2021). *The Disinformation Dozen* (Tech. Rep.). Center for Countering Digital Hate.
- Connor Desai, S., Xie, B., & Hayes, B. K. (2022). Getting to the source of the illusion of consensus. *Cognition*, 223, 105023. doi: 10.1016/j.cognition.2022.105023
- Franzen, A., & Mader, S. (2023). The power of social influence: A replication and extension of the Asch experiment. *PLOS ONE*, 18(11), e0294325. doi: 10.1371/journal.pone.0294325
- Harkins, S. G., & Petty, R. E. (1987). Information utility and the multiple source effect. *Journal of Personality and Social Psychology*, 52(2), 260–268. doi: 10.1037/0022-3514.52.2.260
- Harvey, J. A., van den Berg, D., Ellers, J., Kampen, R., Crowther, T. W., Roessingh, P., ... Mann, M. E. (2018). Internet Blogs, Polar Bears, and Climate-Change Denial by Proxy. *BioScience*, 68(4), 281–287. doi: 10.1093/biosci/bix133
- Kruschke, J. (2014). *Doing Bayesian Data Analysis: A Tutorial with R, JAGS, and Stan*. Academic Press.
- Lewandowsky, S., Gignac, G. E., & Vaughan, S. (2013). The pivotal role of perceived scientific consensus in acceptance of science. *Nature Climate Change*, 3(4), 399–404. doi: 10.1038/nclimate1720
- Lin, X., Spence, P. R., & Lachlan, K. A. (2016). Social media and credibility indicators: The effect of influence cues. *Computers in Human Behavior*, 63, 264–271. doi: 10.1016/j.chb.2016.05.002
- Maddux, J. E., & Rogers, R. W. (1980). Effects of source expertise, physical attractiveness, and supporting arguments on persuasion: A case of brains over beauty. *Journal of Personality and Social Psychology*, 39(2), 235–244. doi: 10.1037/0022-3514.39.2.235
- Pilditch, T. D., Hahn, U., Fenton, N., & Lagnado, D. (2020). Dependencies in evidential reports: The case for informational advantages. *Cognition*, 204, 104343. doi: 10.1016/j.cognition.2020.104343
- Pucci, V., Kashima, Y., & Perfors, A. (2023). Online communication to the ingroup and the outgroup: the role of identity in the “what” and “why” of information sharing. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 45(45).
- Ransom, K. J., Perfors, A., & Stephens, R. (2021). Social meta-inference and the evidentiary value of consensus. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 43(43).
- Sah, S., Moore, D. A., & MacCoun, R. J. (2013). Cheap talk and credibility: The consequences of confidence and accuracy on advisor credibility and persuasiveness. *Organizational Behavior and Human Decision Processes*, 121(2), 246–255. doi: 10.1016/j.obhdp.2013.02.001
- Simmonds, B. P., Stephens, R., Searston, R. A., Asad, N., & Ransom, K. J. (2023). The Influence of Cues to Consensus Quantity and Quality on Belief in Health Claims. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 45(45).
- Smith, P. L., & Little, D. R. (2018). Small is beautiful: In defense of the small-N design. *Psychonomic Bulletin & Review*, 25(6), 2083–2101. doi: 10.3758/s13423-018-1451-8
- Vehtari, A., Gelman, A., & Gabry, J. (2017). Practical bayesian model evaluation using leave-one-out cross-validation and waic. *Statistics and Computing*, 27(5), 1413–1432. doi: 10.1007/s11222-016-9696-4
- Whalen, A., Griffiths, T. L., & Buchsbaum, D. (2018). Sensitivity to Shared Information in Social Learning. *Cognitive Science*, 42(1), 168–187. doi: 10.1111/cogs.12485
- Xie, B., & Hayes, B. (2022). Sensitivity to Evidential Dependencies in Judgments Under Uncertainty. *Cognitive Science*, 46(5), e13144. doi: 10.1111/cogs.13144
- Yousif, S. R., Aboody, R., & Keil, F. C. (2019). The Illusion of Consensus: A Failure to Distinguish Between True and False Consensus. *Psychological Science*, 30(8), 1195–1204. doi: 10.1177/0956797619856844
- Zemla, J. C., Sloman, S., Bechlivanidis, C., & Lagnado, D. A. (2017). Evaluating everyday explanations. *Psychonomic Bulletin & Review*, 24(5), 1488–1500. doi: 10.3758/s13423-017-1258-z