

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Source independence affects argument persuasiveness when the relevance is clear

Permalink

<https://escholarship.org/uc/item/5hg4p8cm>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 44(44)

Authors

Alister, Manikya
Ransom, Keith James
Perfors, Andrew

Publication Date

2022

Peer reviewed

Source independence affects argument persuasiveness when the relevance is clear

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

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

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

School of Psychological Sciences, University of Melbourne

Abstract

Making inferences about claims we do not have direct experience with is a common feature of everyday life. In these situations, it makes sense to consult others: an apparent consensus may be a useful cue to the truth of a claim. This strategy is not without its challenges. The utility of a consensus should depend in part on the sources of evidence that underlie it. If each person based their conclusion on independent data then the fact that they agree is informative. If, instead, everyone relied on the same primary source, the consensus is less meaningful. However, the extent to which people are actually sensitive to this kind of source independence is still unclear. Here, we present the results of three experiments that examine this issue in a social media setting, by varying the sources of primary data cited via retweets. In each experiment, participants rated their agreement with 12 different claims before and after reading four tweets that were retweeted on the basis of either the same or different primary data. We found that people were sensitive to source independence only when it was clear that the tweeters had relied on the primary data to reach their conclusion. Implications for existing research are discussed.

Keywords: consensus; persuasion; source independence; social reasoning; induction

Introduction

In today's information-rich world, we often must evaluate the legitimacy of a claim without having enough direct experience ourselves to be certain about it. In such a situation, it can be useful to consider what other people think. For example, suppose you heard the claim that narcissists are more politically engaged. Although you might have an opinion about this, most of us do not know enough narcissists or politicians to be very sure about that opinion; in that case, it would make sense to see what other people think as well. If you found many people who had each independently reached the same conclusion, that might constitute good evidence that the claim was true (and better evidence than if you had found just one). Indeed, a wide literature suggests that people are more likely to agree with claims when they are endorsed by many other people (e.g., Asch, 1956; Lewandowsky, Gignac, & Vaughan, 2013; Ransom, Perfors, & Stephens, 2021).

Taking into account the consensus of others is a sensible reasoning strategy, especially if everyone's conclusions were reached independently from each other. Indeed, rational Bayesian models have demonstrated that a dependent consensus should carry less weight than an independent consensus (e.g., Whalen, Griffiths, & Buchsbaum, 2018). People have been shown to be sensitive to such reasoning dependencies: when people are explicitly told that consensus opinions were not formed independently (for instance, everyone was on the same committee) people do not give more weight to multiple sources than a single source (Harkins & Petty, 1987).

But in everyday life, assessing the independence of a consensus from each individual's statements alone can be difficult. For example, what if all of the people involved agreed with a claim and mentioned the same news article or study? What you conclude may be influenced by your assumptions about how people provide evidence or explanations in support of their views. If, for example, you assumed that people were drawing on personal experience or extensive knowledge of the subject at hand and mentioning the article to summarise their view, then the common reference may simply reflect the quality of the source or the availability of the content. In contrast, if you assume that people had formed their view largely on the basis of the common information, then this sort of *dependent consensus* should be less convincing than an *independent consensus* with the same number of people, since the claim is based on fewer distinct sources of knowledge.

In the real world, of course, even when we know what sources people relied on, it may not be clear how independent *those* sources were. This can lead to different kinds of errors. On one hand, people might wrongly believe that an independent consensus is dependent and thus discount the consensus when they shouldn't: for instance, if scientists independently reach the same conclusions about climate change but the public believes that they did not. On the other hand, people might erroneously believe that a dependent consensus is independent, and thus weight it more highly than they should. For instance, the majority of science-denying blogs rely on the same few primary sources (Harvey et al., 2018) and most COVID-19 anti-vaccination views originated from the same few people (Center for Countering Digital Hate, 2021).

How sensitive are people to source dependence, and thus how good are we at avoiding these reasoning errors? Several studies have investigated this question experimentally, but the findings are mixed. Yousif, Aboody, and Keil (2019) presented participants with several texts (either student essays or news articles) making different claims. They found that people did not think an independent consensus was more convincing when the sources were economists and the claims were about a new tax policy or a country's economic prospects. Yet when the sources were eyewitnesses and the claim was that a bear had been seen at a local school, people did weight information from independent sources more highly. The authors concluded that people may not be sensitive to source dependence unless the claim is easily knowable to the sources. However, an alternative explanation is that the participants may not have realised that the consensus of economists was actually independent, since it would be plausible for multiple economists to have relied on the same data.

Desai, Xie, and Hayes (2022) addressed these concerns by testing U.S.-based participants in a twitter-like paradigm where the independence between sources was clearer. Participants read tweets by independent users (news companies in one experiment, people in another) who posted a polling company’s prediction for an upcoming fictitious election. The tweets reached a near unanimous consensus (four tweets agreeing, one opposing) that “Candidate X” would win the election, but the four agreeing tweets either all mentioned the same polling company (dependent consensus) or four distinct polling companies (independent). People were more convinced by the independent consensus than the dependent one, suggesting that people are sensitive to source dependence as long as it is made sufficiently clear and salient.

This is promising, but it is still worth considering how far this sensitivity might extend. Arguably, the polling company scenario is unusually clear: not only is the underlying question (“who will win the election?”) precisely framed, it is also readily apparent what data would bear directly on this question (voting intentions) and how it should be assessed (tally them). This is far more straightforward than the problem faced by the reasoner evaluating the claim about future economic performance: the nature of the relevant evidence, the means of assessment, and the appropriate yardstick of measurement are not completely clear, even to an expert. This kind of uncertainty surrounding the evaluation of many everyday claims can mean that a potentially complex web of data inter-dependencies remains hidden to the reasoner even when the primary data sources are clearly identified. In those situations, people may find it difficult (or simply not attempt) to make strong inferences about the underlying independence of the data based on the apparent independence of the sources.

An alternative possibility is that for sufficiently complex issues, what people infer from a consensus is that the overall reasoning process was sound. For example, in the polling scenario employed by Desai et al. (2022), the tweets made it clear that the tweeter had actively considered the primary source themselves, since they directly referenced the pollster. It is possible that people were sensitive to the independence of the pollsters because it was obvious that the tweeters’ conclusions were based on the polls. Had that been less salient or more implicit, would people still have used that information?

Our work addresses two main questions. First, how much does sensitivity to source dependence depend on the topic or nature of the sources? Second, are people sensitive to source dependence if the link between the conclusions and the source is not obvious? We address these questions across 12 real-world topics using a realistic twitter paradigm. Over three experiments, we consider two different kinds of source independence (whether the tweeters are distinct individuals or one person posting multiple times, as well as whether all of the tweets reference the same primary information source or use different ones) as well as the strength of the link between source and conclusion (whether the primary source is retweeted without engagement or not).

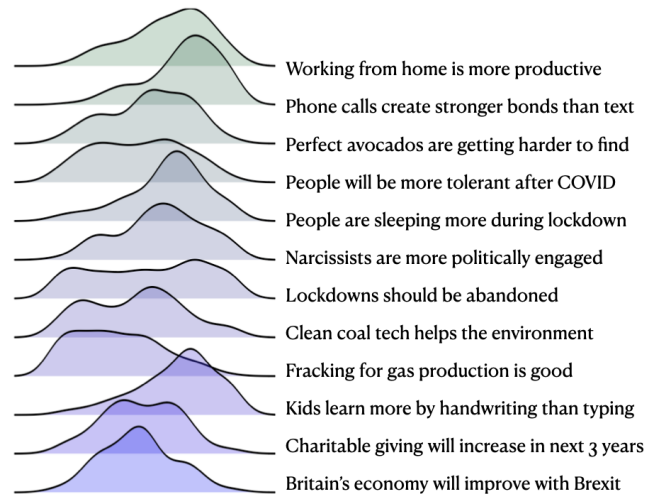


Figure 1: **Topics.** In all experiments people rated what they thought about a claim before seeing four tweets arguing for or against it. People varied considerably in the strength and range of their priors on the 12 topics (distributions range from 0 to 100 indicating how much they agree with the claim).

Experiment 1

In this experiment we tested whether participants were more persuaded by a consensus of four people who relied on the same source, or a consensus of four people who relied on different sources. Given previous work suggesting that people are sensitive to source dependence when the independence of sources is clear (Desai et al., 2022), we predicted that they would be more convinced by the independent consensus.

Method

Participants 111 participants were recruited from Amazon Mechanical Turk and paid \$3 for the 15-20 minute task.¹² Ages ranged from 20 to 78 years old ($M = 38$) and 45% were female. 88% reported being native English speakers, and all passed a qualification assessing English proficiency.

Procedure After providing consent and passing a short quiz regarding the instructions of the task, each participant saw 12 trials, all of which followed the same basic structure. Each 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. As is evident from Figure 1, the range and strength of prior beliefs endorsed by participants varied considerably across topics.

After providing their prior beliefs, people then viewed four tweets by four distinct twitter users who always formed a consensus, either all for (PRO) or all against (CON) the claim. As shown in Figure 2, each tweet consisted of the twitter users themselves, the news source they retweeted, and the primary

¹All predictions, methods, and analysis were preregistered here: <https://aspredicted.org/me3fg.pdf>.

²This sample size was sufficient to find similar consensus effects in Ransom et al. (2021), from which our experimental structure, stimuli, and topics were adapted.



Figure 2: **Experiment 1 stimuli.** Sample tweets from each condition. The two on the left are from the INDEPENDENT condition, where each person retweeted a different source (here, arguing against the claim). Those on the right are from the DEPENDENT condition, where each person retweeted the same source (here, arguing for it). Each trial in the actual experiment showed four tweets rather than two, and the sources, names, photos, text, dependence, and pro/con direction was randomised for each person and topic.

data referred to by the news source. The four tweeters were always distinct people with unique profile photos and names, and each tweeter always accompanied the re-tweet with their own opinion about the claim. Regardless of condition, all four tweeters and all sources in a given trial gave essentially the same reason in different words (e.g., those arguing PRO on the narcissist question all pointed out that narcissists want attention, and those arguing CON all pointed out that narcissists don't want to spend the time). Everything was completely randomised across participants and topics.

After clicking on each tweet to confirm they had read it, participants then indicated (using the 0-100 slider) how much they agreed with the claim. The difference between their rating before and after represents their degree of belief revision due to the tweets, and is the dependent variable of interest.

Design Our primary manipulation was whether each of the four tweeters in any given trial cited the same primary news source (DEPENDENT consensus) or all different sources (INDEPENDENT consensus). This was randomised within-participant, so each person saw six trials in each condition in random order. In the INDEPENDENT condition, both the news source of the post that was retweeted and the data that the source referred to were distinct for each of the four tweeters: Person A cited Source X, Person B cited Source Y, and so forth. Thus, each was retweeting an independent source and referring to independent primary data. In the DEPENDENT condition, each of the four tweeters re-tweeted the same article by the same news source (hence the same primary data).

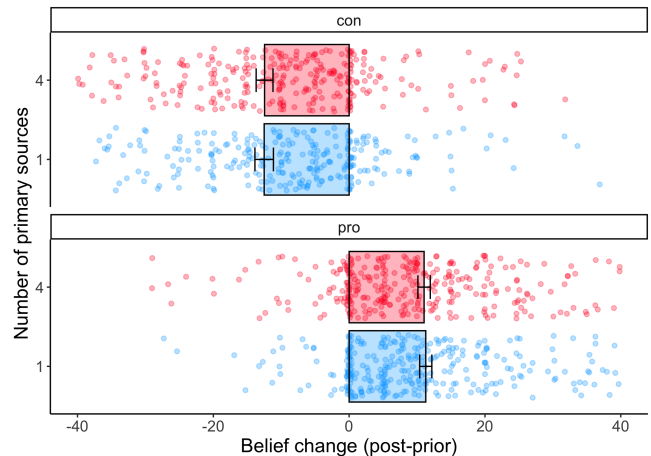


Figure 3: **Experiment 1 results.** Difference in agreement with the claim as a function of number of primary sources (4 is INDEPENDENT consensus, 1 is DEPENDENT consensus) and whether the consensus was supporting (PRO) or opposing (CON) the claim. Although people adjusted their beliefs in the direction of the consensus, consistent with Ransom et al. (2021), there was no difference in the magnitude of adjustment whether the consensus was independent or not. Error bars denote standard error.

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 companies were attached to the different primary data at random for each participant and were always unique in each trial. Like the tweeters, the news companies always had a profile photo, full name, and username, which matched the company’s actual twitter account where possible. The news companies also had a “verified” tick to signal authenticity.

The primary data referenced by the news company consisted of either a university study or an expert’s testimony. In the INDEPENDENT consensus condition, there were always three university studies and one expert’s testimony. The expert’s testimony was always dated before the university studies so that it would have been impossible for the expert testimony to have drawn from the university studies. The universities were a sample of real universities ranked between 100 and 200 by the QS World University Rankings 2021.

Results and Discussion

The purpose of this experiment was to see whether the independence of a consensus influenced how persuaded people were by that consensus. To assess this, we compared how much people’s ratings of 12 claims changed after viewing four tweets pertaining to each respective claim, having varied whether or not the four tweets cited four distinct sources. As Figure 3 shows, people did tend to change their score in the direction of the consensus: when the consensus supported the claim, they become more favourable to it, and vice-versa. However, contrary to our prediction, there are no evident differences in the magnitude of adjustment if the consensus

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

Table 1: Comparison of four models that predict people’s support for a claim in Experiment 1. The model with Direction but not Independence is marginally preferred (lower LOOIC value), suggesting that people did not reason differently when the consensus was INDEPENDENT vs DEPENDENT, but that they did respond differently to PRO vs CON arguments.

Model	LOOIC	SE
M1. Prior	12082	67
M2. Prior + Direction	11648	84
M3. Prior + Direction + Independence	11649	83
M4. Prior + Direction × Independence	11651	83

was INDEPENDENT instead of DEPENDENT: people were not more persuaded by four independent sources of evidence than four instances of the same source.

To quantitatively assess how the independence of the consensus influenced people’s ratings, we compared four nested generalised linear models in which the outcome variable was the rating after reading the four tweets. All included separate error terms using a random intercept term for each participant and topic. M1 was a baseline whose only fixed term was the prior; it thus represents the null hypothesis that ratings were only influenced by their prior beliefs. M2 also incorporated a term indicating the direction of the consensus (PRO or CON); this model is favoured if participants were not sensitive to independence but did reason differently for PRO and CON tweets. M3 added consensus independence as a predictor, and thus reflects different reasoning in the INDEPENDENT and DEPENDENT consensus conditions, but no interaction between this and tweet direction (which M4 added). All models were run using the *brms* package (v2.16.3) in R (v4.1.2).

To assess the relative performance of each model, we compared leave-one-out cross-validation criterion (LOOIC), which has several advantages over simpler information criteria such as AIC and DIC (Vehtari, Gelman, & Gabry, 2017). As shown in Table 1, M2 was preferred (i.e., had the lowest LOOIC), suggesting participants were sensitive to whether the tweets were PRO or CON, but were not sensitive to the independence of the consensus. This conflicts with the findings of Desai et al. (2022), which suggested that people will give more weight to an independent consensus if it is clear that the sources are independent.

An obvious alternative explanation for why we obtained a null result is that participants may not have been actually reading the tweets. However, this is unlikely. First, participants would not have responded in the direction of the consensus (increasing support when they saw PRO tweets and decreasing it when they saw CON) had they not been reading the tweets enough to understand their content. Second, we found no relationship between reading speed and sensitivity to consensus independence ($r = -.06, p = .510$). Finally, excluding participants who failed to adjust their beliefs in the expected direction made no difference: there was still no significant difference between the INDEPENDENT and DEPENDENT conditions.

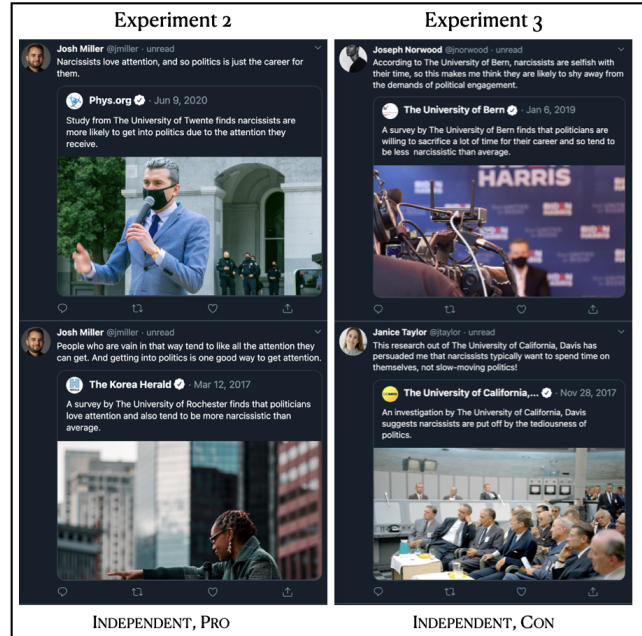


Figure 4: **Experiment 2 and 3 stimuli.** Sample tweets from Experiment 2 and 3. The two on the left are from Experiment 2, where the same person made all the tweets on each trial. The two on the right are from Experiment 3, where different people made the tweets but (unlike Experiment 1) explicitly discussed the source and made it clear that it had influenced their reasoning; moreover, the source tweet was always the primary source rather than a news agent acting as intermediary. In both of these examples the sources were INDEPENDENT, but in the experiment this varied across trials.

Another possibility for why our results conflicted with Desai et al. (2022) may stem from how the consensus was presented in our experiment. A key difference between our paradigm and theirs was that we had an extra source level (the tweeters themselves) who were independent. It could be that this surface layer of independence distracted from the independence of the primary data and news companies.

Experiment 2

We suspected that the separate tweeters in each trial of Experiment 1 were more salient than the primary sources, causing them to be overlooked. To test this hypothesis, in Experiment 2 we held the tweeters constant, such that the same person always tweeted the four news articles in any given trial.

Method

The methods for Experiment 2 were identical to the previous experiment except that on each trial the same twitter user tweeted four times (Figure 4). As before, each trial had different users and we used a within-subject design in which the tweeter referenced four distinct sources on INDEPENDENT trials and the same source four times on DEPENDENT ones.

109 participants were recruited from Amazon Mechanical Turk and paid \$3 for the 15-20 minute task. They were 20 to 71 years old ($M = 38$) and 44% were female. 88% were native English speakers, and all passed a qualification assessing English proficiency. None participated in Experiment 1.

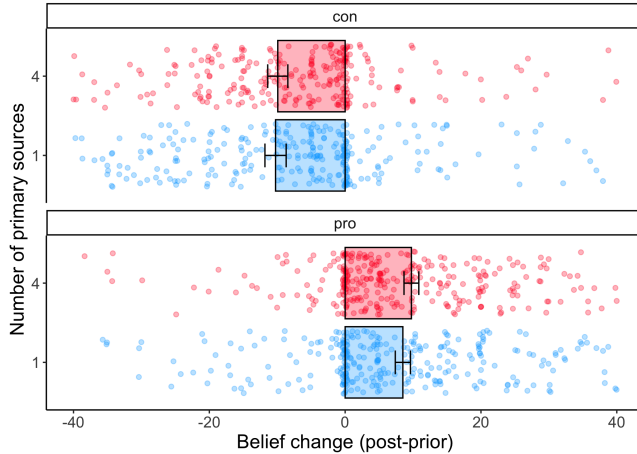


Figure 5: **Experiment 2 results.** Difference in agreement with the claim as a function of number of primary sources (4 is INDEPENDENT, 1 is DEPENDENT) and whether the consensus was supporting (PRO) or opposing (CON) the claim. Although people adjusted their beliefs in the direction of the consensus, there was no difference in the magnitude of adjustment by whether the consensus was independent or not. Error bars denote standard error.

Results and Discussion

As is shown in Figure 5, the results in Experiment 2 were almost identical to Experiment 1.⁴ People adjusted their beliefs in different direction depending on whether the consensus was PRO or CON, but there was no discernible difference between consensus independence conditions. We quantitatively evaluated this result using the same procedure and models in Experiment 1. As before, the model containing only the prior and direction (M2) had the lowest LOOIC (Table 2). This indicates that there was no difference in performance as a function of the independence of the sources. It also suggests that the null results from Experiment 1 cannot be due to the independence of the tweeters distracting people from noticing the independence of the primary source.

How can we explain the lack of sensitivity to consensus independence in both experiments? One possibility is that it may not have been clear that the tweeters were actually reasoning from the primary data. In Desai et al. (2022) the twitter user directly referenced the primary data in the same tweet, making it obvious that their inference occurred *because of* the source. In contrast, because our twitter users only made a general statement on the topic, that may have given the impression that they did not actually read the primary data and were simply passively sharing something they saw online. Further, because we used a news company as the intermediary between the tweeter and the primary data, it might not have been clear to the participant that the tweeter’s claims were based on the primary data itself (rather than independent news articles which might have been using the same primary data). We address this possibility in Experiment 3.

Table 2: Comparison of four models that predict people’s support for a claim in Experiment 2. The preferred model (M2) contains a term for both direction of support (PRO/CON) and prior beliefs, suggesting that people did reason differently depending on the direction of support, but not when the consensus in the four tweets was INDEPENDENT vs DEPENDENT.

Model	LOOIC	SE
M1. Prior	12024	67
M2. Prior + Direction	11796	75
M3. Prior + Direction + Independence	11799	75
M4. Prior + Direction \times Independence	11799	76

Experiment 3

This experiment was identical to Experiment 1 except that the tweeters actively referenced the primary data, which was sourced directly from the university or expert’s own twitter account rather than through an intermediary news agent.

Method

The methods were identical to Experiment 1 apart from the following changes to the stimuli (right panel of Figure 4). Firstly, the primary data was reported directly from the twitter account of the corresponding university or expert; this eliminated any ambiguity about whether the tweeters were reasoning directly from the independent primary source rather than a news article that may not be truly independent. Secondly, we rephrased the tweeters’ accompanying message so that they were directly referring to the tweet, making it clear that their reasoning had been based on the source and they were not just passively sharing it. These messages were phrased to indicate that the tweeters had formed their opinion *after* seeing the data (e.g., ‘Having read about this study from The University of Ohio, I now think...’).

105 participants were recruited from Amazon Mechanical Turk and paid \$3 for the 15-20 minute task.⁵ Ages ranged from 20 to 70 years old ($M = 37$) and 56% were female. 78% reported being native English speakers, and all passed a qualification assessing English proficiency. None of these people participated in Experiments 1 or 2.

Results and Discussion

As Figure 6 shows, people in this experiment *did* appear to be sensitive to the independence of the consensus. As before, people changed their beliefs to support the claim more when they saw PRO tweets and to oppose it more when they saw CON tweets. Unexpectedly, however, the effect of independence was in a different direction for the PRO and CON tweets. When the consensus supported the claim (PRO), people changed their beliefs more when the tweets were independent; this makes sense because independent evidence is normatively more persuasive. However, when the consensus opposed the claim (CON), people changed their beliefs more when the tweets were *not* independent. This is a more surprising result which calls for further investigation.

⁴Pre-registration: <https://aspredicted.org/re7cs.pdf>

⁵Pre-registration: <https://aspredicted.org/ns5ks.pdf>




Figure 6: **Experiment 3 results.** Difference in agreement with the claim as a function of number of primary sources (4 is INDEPENDENT consensus, 1 is DEPENDENT consensus) and whether the consensus was supporting (PRO) or opposing (CON) the claim. Unlike the other experiments, people adjusted their beliefs differently depending on whether the consensus was independent or not. Unexpectedly, the adjustment was in different directions depending on whether the consensus supported the claim; people adjusted more for an independent consensus for PRO arguments, but more for a *dependent* consensus for CON arguments. That said, this interaction was (marginally) not statistically supported.

The obvious question is whether this apparent interaction is supported quantitatively by comparing the same four models as before. As Table 3 shows, the LOOIC values indicate that the preferred model was M3 (with a term for direction and independence, but without the interaction). That said, the value is quite similar to M4 (which is identical but includes the interaction), suggesting that we should be wary of drawing strong conclusions about the presence of the interaction.

Regardless, these results show that in Experiment 3 people *did* reason differently when their information was based on independent sources than when it was not. Taken together, these three experiments support the conclusion that people *can* take source independence into account when reasoning; however, they only do so when it is obvious that the opinions they hear were actually formed on the basis of the sources.

General Discussion

Our result here is consistent with previous work investigating people’s sensitivity to source independence, but suggests that the explanation may be different to what was initially thought. For example, Desai et al. (2022) also found the strongest evidence for an effect of source independence when they created a paradigm where it was obvious that members of the consensus were reasoning based on the primary source. However, while they attributed their results to the necessity that the relationship between source independence and the consensus be clear, our findings suggest that their results may have had more to do with whether it was obvious that tweeters were actually reasoning based on the sources. The same could be said for the results of Yousif et al. (2019). It is plausible that a student essay or article about the economy might selectively

Table 3: Comparison of four models that predict people’s support for a claim in Experiment 3. The preferred model (M3) contains a term for both direction of support (PRO/CON) and independence condition but no interaction, suggesting that people did reason differently when the consensus in the four tweets was INDEPENDENT vs DEPENDENT, as well as depending on the direction of support.

Model	LOOIC	SE
M1. Prior	11666	55
M2. Prior + Direction	11255	66
M3. Prior + Direction + Independence	11253	67
M4. Prior + Direction \times Independence	11254	67

choose sources that support a view that the author had decided on for other reasons. Conversely, a journalist reporting a bear sighting is more likely to actually have reasoned about what happened based on the eye-witness accounts of the bear.

We posit that it is actually quite intelligent to only be persuaded by source independence when it is clear that the reasoning was truly based on that source. After all, source independence does not actually imply stronger evidence if the sources didn’t truly influence the conclusions. Indeed, it is often the case – particularly in social media contexts – that people will share or endorse a source without having properly read it, simply because it supports or confirms an opinion that they already had. If people are aware of this reality, it makes sense for them to default to disregarding the source unless its relevance to the conclusion is made extremely clear.

It is important to note that even though we found an effect of source independence in one experiment, the effect was small and went in different directions for PRO and CON arguments (though the quantitative evidence for this difference was uncertain at best). Future work is necessary to substantiate whether the opposite direction effect is robust.

Overall, our results provide a unified explanation for the growing body of literature investigating how people reason about source independence. Taken together with our results, it appears that people tend to discount the independence of a consensus under most conditions, but *do* show sensitivity to it if the relationship between the source and the conclusions is clear (Yousif et al., 2019; Desai et al., 2022). An important next step is to integrate these findings into a formal theoretical account for how people reason based off a consensus. Modelling attempts so far have shown that source independence is an important cue that people *should* consider (e.g., Whalen et al., 2018), but existing models generally do not incorporate the kinds of nuances that our work explored (like how to determine whether a cited source really shaped someone’s reasoning). Nor do they incorporate the kind of cognitive and temporal limitations that people must deal with, where it might make sense to take certain shortcuts even if they cause errors in some situations. More work is therefore required in order to develop a precise theoretic account of how people use source independence as a reasoning cue in the real world.

Acknowledgments

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

References

- Asch, S. E. (1956). Studies of independence and conformity: I. A minority of one against a unanimous majority. *Psychological Monographs*, 70(9), 1–70.
- Center for Countering Digital Hate. (2021). *The disinformation dozen* (Tech. Rep.). Retrieved from www.counterhate.com/disinformationdozen
- Desai, S. C., 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
- Harkins, S. G., & Petty, R. E. (1987). Information utility and the multiple source effect. *Journal of Personality and Social Psychology*, 52(2), 260–268.
- 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.
- 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.
- Ransom, K., Perfors, A., & Stephens, R. (2021). Social meta-inference and the evidentiary value of consensus. In H. L. T Fitch C Lamm & K. Teßmar-Raible (Eds.), *Proceedings of the 43rd Annual Conference of the Cognitive Science Society* (p. 833-839).
- 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.
- 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
- 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.