Fact or fiction? How repeated claims influence belief updating

Greta Arancia Sanna (greta.sanna.23@ucl.ac.uk)[1], David Lagnado (d.lagnado@ucl.ac.uk)[1] & Hugo Mercier (hugo.mercier@gmail.com)[2]

[1]Department of Experimental Psychology, University College London

[2]Institut Jean Nicod, École Normale Supérieure

Abstract

Individuals often encounter the same claims from multiple sources in complex informational environments. Whether such repetition is treated as independent evidence or as redundant testimony depends on assumptions about source dependence. The present study investigates how prior beliefs about the claims shape perceptions of source dependence and how these perceptions influence belief updating. Using a preregistered experimental design, participants evaluated repeated claims presented in vignettes about an online forum. Participants reported both belief change and perceived dependence among sources. We find that lower prior belief in a claim leads to stronger perceptions of source dependence, indicating that prior beliefs shape not only initial judgments but also inferences about the structure of the informational environment. Participants also tended to treat dependence as a global property applying to several sources in the environment. Together, these findings suggest that belief updating reflects rational simplification of informational structure, with important implications for the persistence of belief polarisation.

Keywords: Belief updating, Source dependence, Bayesian modelling

Introduction

In everyday life, people routinely rely on information provided by others, from mundane judgements about the weather to evaluations of political events. Often, multiple sources offer convergent testimony on a single issue. Such convergence is typically taken to increase our confidence in the claim, as agreement among independent sources is widely regarded as a hallmark of evidential reliability (e.g. Pfänder et al., 2025). However, this inference depends on the assumption of source independence. When multiple sources draw on a shared informational source, the apparent consensus may be misleading. Failing to recognise such dependence can therefore lead to systematically overestimating evidential strength. For example, if several sources report an event, this may reasonably be treated as stronger evidence than a single report. Yet if it later emerges that all reports trace back to a common source, the additional confidence afforded by the consensus should be substantially discounted, as the corroboration could be illusory rather than informative.

In this domain, previous research indicates that people sometimes discount consensus information when they perceive sources to be dependent (Couch, 2022; Harkins & Petty, 1987; Mercier & Morin, 2019; Whalen et al., 2018; Xie & Hayes, 2022), although this sensitivity is not consistently observed across contexts (Yousif et al., 2019). Prior work has also examined which cues people use to infer source dependence (Desai et al., 2022; Young et al., 2025). Despite this emerging evidence that perceived dependence can influence belief revision (J. K. Madsen et al., 2020), the cognitive processes by which individuals form dependence judgements in the first place remain poorly understood. In particular, it is unclear whether prior beliefs about the world shape the dependence structures individuals infer within informational environments, and how perceived dependence and belief updating interact over time. Understanding these dynamics is critical to explaining how people evaluate evidence in complex informational settings and how features of those environments ultimately shape belief formation.

If people's understanding of the world is shaped by the dependence structures they perceive in their informational environments, they might discount evidence they judge to originate from dependent sources, even when others regard the same evidence as strengthening their beliefs. For instance, individuals with high confidence in the legitimacy of democratic processes may interpret multiple claims that an election was rigged as deriving from a single misleading source, and therefore discount these claims as uninformative. Other individuals, holding different prior beliefs, may instead interpret the same convergence as compelling evidence of electoral fraud. Consistent with prior research (e.g. Young et al., 2025), this pattern would support the notion of a self-reinforcing cycle in which prior beliefs shape perceived source dependence, and these dependence judgements in turn guide how new evidence is incorporated in ways that favour existing priors. Such a circular process may help explain why individuals struggle to disengage from polarised informational environments or to revise their beliefs even when revision would appear warranted.

To examine these processes, we focus on contexts involving repeated testimony. Such informational environments provide a particularly suitable paradigm for investigating belief and dependence dynamics, as they naturally give rise to a fundamental ambiguity: repeated claims may indicate increased veracity, or instead signal dependence on a common source. Resolving this ambiguity requires an account of how individuals mentally represent source dependence and how these representations interact with belief updating.

To address this gap, we develop two a priori Bayesian network (BN) models that formalise alternative representations

In G. Boleda, I. Dautriche, M. Maldonado, C. Saldana and T. Torrent (Eds.), *Proceedings of the 48th Annual Conference of the Cognitive Science Society*

of source dependence. The two models differ in the degree to which individuals are capable and willing to update rationally. Model 1 (local model) assumes that individuals integrate new pieces of information rationally, sequentially weighing the probability of each source being independent based on their current beliefs and their priors about the likelihood of dependence in the present context. Here, regardless of prior beliefs, as information in support of the claim is gathered, belief in the claim should increase and perceived dependence of the sources on a common informant should decrease. Conversely, Model 2 (global model) posits that individuals do not update rationally as this requires a significant amount of cognitive resources; instead individuals simplify their understanding of the informational environment by assuming that each source has the same likelihood of being dependent on the common informant throughout the evidence evaluation process. As a result, individuals quickly stop updating their beliefs as they reach an equilibrium between the likelihood of dependence of the sources and their belief in the claim. Now, although Bayesian networks are often employed as normative benchmarks, our aim is not to assess whether participants reason optimally. Rather, the models are used to characterise the space of inferences a reasoner could plausibly draw (e.g. Lieder & Griffiths, 2020) and to derive behavioural predictions about perceived dependence and belief updating. These predictions, in turn, motivate our hypotheses and guide our empirical tests.

Bayesian Network model

We adapted a Bayesian Network introduced by Young et al. (2025) to model how individuals can update their beliefs as repeated claims emerge. Bayesian networks (BNs) comprise two components: (i) a directed acyclic graph that represents probabilistic dependencies among variables, and (ii) conditional probability tables for each variable that quantify these dependencies and enable belief updating via Bayes' rule. Bayesian Networks (Pearl, 1988) have been widely applied to model how source characteristics influence belief updating and have received empirical support across multiple domains (Collins et al., 2018; Harris et al., 2012, 2016; J. Madsen, 2016; J. K. Madsen, 2019; Young et al., 2023). As such, they provide a natural framework for modelling how perceptions of source dependence shape belief updating in the present context. Building on this framework, we specified two competing models that differ in how dependence is assumed to operate across sources. The *local dependence model*, illustrated in Figure 1, represents a scenario in which three informants provide the same testimony, which may either reflect a direct observation of the world (H) or second-hand information originating from a latent source (LT). When all three informants report evidence generated directly from H, the resulting agreement constitutes strong evidence in support of the claim. In contrast, if the testimonies of all three informants are derived from LT, their agreement is substantially less diagnostic of the claim's truth. Crucially, whether each informant's testimony is

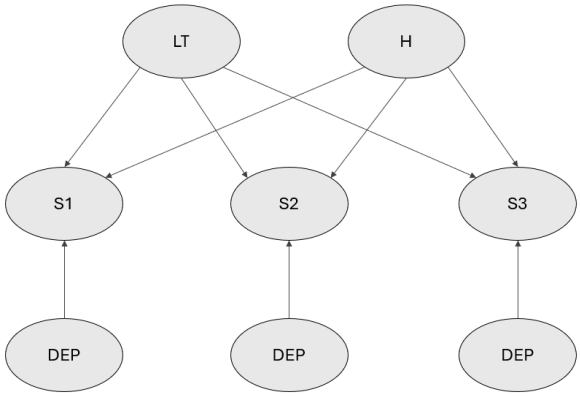

Figure 1: Local dependence model: each informant's testimony may be independently generated or dependent on a latent source.

generated from the world or from the latent source is governed by individual dependence nodes (DEP), which determine the evidential origin of each informant's report. This structure allows dependence to vary across informants. Thus, an accurate representation of dependence requires individual nodes of dependence for each informant. However, this conceptualisation may be cognitively effortful beyond individuals' capabilities as it requires them to track, after each claim, both their belief in the world (H) and their perceived dependence of each claimant. In contrast, constraints on cognitive resources (e.g. Zhu & Griffiths, 2024), may encourage individuals to simplify aspects of the information environment. Treating dependence as a global property rather than applying to each source individually can reduce cognitive strain and simplify the belief updating process. Building on this assumption, and consistent with prior modelling work (Young et al., 2025), we specified a *global dependence model*, illustrated in Figure 2, which was developed in light of results from a pilot study prior to the present experiment (see Methods section). This is the same as the local dependence model, except dependency is captured by using a single unifying node across sources instead of individual nodes for each source. Under this global representation, dependence is simplified to promote cognitive effectiveness; however, now additional repeated claims may be dismissed as mere products of dependence rather than as flexible and potentially informative evidence. This, in turn, could further entrench existing beliefs and exacerbate processes of polarisation.

We use the local and global dependence models to provide the theoretical framework to develop hypotheses about how individuals represent dependence in contexts involving repeated claims. Both models predict that prior belief in the claim influences inferences about source dependence (see Figure 5 for predicted trajectories). When a claim is initially plausible, the local dependence model predicts that belief in the claim should increase with each additional corroborating report, ultimately approaching a high degree of certainty. The

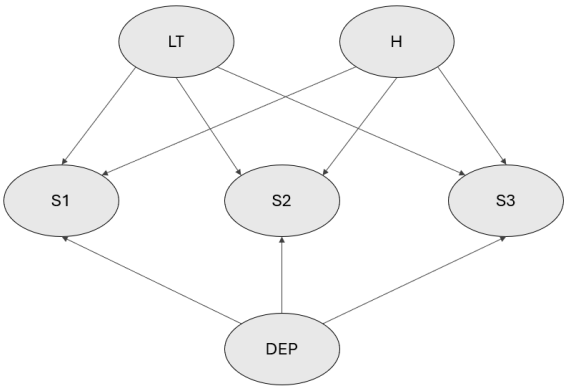

Figure 2: Global dependence model: all informants' testimonies are either independently generated or jointly dependent on a single latent source.

global dependence model similarly predicts an initial increase in belief; however, due to the global representation of source dependence, belief updating is expected to plateau, maintaining a residual level of uncertainty that reflects the inferred degree of dependence among sources.

By contrast, when a claim is initially implausible, both models predict that repeated testimony will be attributed to a high degree of source dependence. Under the local dependence model, this inferred dependence is expected to be overridden in the long run by the cumulative weight of repeated claims, leading belief in the claim to eventually increase despite initial scepticism. Under the global dependence model, however, belief is not expected to increase with repetition; instead, it should stabilise at an equilibrium level, as additional testimony is discounted once it is judged to be dependent and therefore non-informative.

The aim of the present study is not to assess whether participants match either model in a strict normative sense, but rather to examine whether their inferences qualitatively align with the high-level structural assumptions embedded in the simplified model.

To test these assumptions, we developed a study in which individuals read a vignette describing several sources making the same high (or low) plausibility claim sequentially. After each statement, participants rate their belief in the claim and the extent to which they judge the sources to base their testimony on independent first hand experiences versus a common latent source. This approach enables us to examine how prior knowledge shapes both the inferred structure of relationships between sources and belief updating regarding the claim itself.

The Study

Participants

200 US participants were recruited from Prolific (105 identified as female, 95 as male: mean age 47, range 23–85), and paid £0.75 for a 5 minute survey. In line with pre-registration criteria no participants were excluded from the sample. The study was pre-registered with supplementary material provided on OSF.

Design

Participants were presented with a vignette about an online neighbourhood forum where users post information about potential sightings of a lost cat. Three users (Beth, Carol, and Daisy) sequentially post information stating that the cat is near the local school. Participants are informed that posts either reflect genuine independent sightings of the cat or repeat previously shared information to increase visibility. Prior to seeing any forum posts, participants were given contextual information manipulating the prior likelihood of the cat being near the school. Due to the cat's physical health, participants were told that it was either very likely or very unlikely that the cat would be in that area. Baseline belief in the claim was measured immediately after this manipulation. After each post, participants reported (a) their belief that the cat is near the local school on a scale from 0 to 10, and (b) their judgement of whether the user was reporting their own sighting versus repeating earlier information on a scale from 0 to 10, anchored such that higher values indicate a personal sighting. These judgements were collected separately for each informant immediately after that informant's post. The study therefore used a 2 (Claim probability: High vs. Low, between participants) × 4 (Timepoint: baseline, after Informant 1, Informant 2, Informant 3, within participants) mixed factorial design.

Procedure

The study was delivered on Prolific through survey platform Qualtrics and took approximately 5 minutes to complete. Participants first completed a consent form, followed by an attention check. They were then presented with background information describing an online neighbourhood forum in which users are attempting to support an elderly lady, Amber, to find her cat Whiskers. Participants are told that some members of the community are out looking for the cat while others are keeping track of the forum, repeating earlier information to increase visibility. Therefore any given post could either be a legitimate sighting or could just be repeating previously delivered information. Participants are then told that, due to its physical health, the cat is either likely (high probability condition) or unlikely (low probability condition) to be near the school. Following this manipulation, participants are asked to rate the likelihood of Whiskers being located near the school to check for the effectiveness of the manipulation as well as to collect baseline likelihood ratings. Participants are then presented with three users making posts indicating that Whiskers is near the school. These posts are presented sequentially but are described as occurring at the same time, to avoid the possibility of dependency between them. After each post participants are asked to rate the likelihood that the cat is near the school as well as the likelihood that the post is a genuine sighting or a repetition of previous information

posted in the forum. At the end of the task participants are asked to complete some demographics questions (gender, age, education) and are debriefed about the condition they were in.

Hypotheses

In line with predictions from the Bayesian Network model, our pre-registered hypotheses were:

H1: Participants in the low prior probability condition will judge the set of claims as more likely to be based on a common latent source than participants in the high prior probability condition.

H2: Conditional on H1 being supported, inferred dependence on the latent source will moderate belief updating such that belief will increase from baseline following the first claim when dependence on the latent source is perceived to be low, but will show little or no increase when dependence on the latent source is perceived to be high.

H3a: Participants will treat dependence on the latent source as a global property of the informational environment, such that increased perceived dependence on the latent source will apply uniformly to all claimants rather than selectively to individual claimants.

H3b: Given that we anticipate people might rely on a global model, in both conditions, belief in the claim will not increase as additional claims are made, such that belief after the third claim will be indistinguishable from belief after the first claim.

H4: The global dependence model will provide a better fit to the data than the local dependence model.

Pilot

In advance of the main study, we developed a similar vignette as in the present study but exploring a different online environment. Participants were introduced to a horse racing scenario in which users posted their predicted winners for an upcoming race. Users could either make their predictions based on their own research or copy information that had previously been posted in the forum. Crucially, the plausibility of the repeated claims about the winning horse was manipulated such that the predicted outcome was either likely or unlikely, with participants informed that they possessed relevant insider information. Participants judged users to be more likely to be copying previously posted information when the claim was unlikely (M = 54.84) compared to when it was likely (M = 35.29), $t(186) = 6.29$, $p < .001$, $d = 0.89$, 95% CI [0.60, 1.18]. In contrast, belief in the predicted winner did not increase from the first to the third repeated claim in either probability condition. In the low-probability condition, a paired-samples t-test revealed no change in belief, $t(100) = -0.29$, $p = .77$, $M_{\text{diff}} = -0.36$, 95% CI [−2.76, 2.05], with the Bayesian analysis providing evidence in favour of the null hypothesis, $\text{BF}_{10} = 0.11$. Similarly, in the high-probability condition, belief ratings did not differ between the first and third claims, $t(98) = -0.54$, $p = .59$, $M_{\text{diff}} = -0.56$, 95% CI [−2.61, 1.50], with the corresponding Bayesian analysis again supporting the null hypothesis, $\text{BF}_{10} = 0.13$. This pattern was inconsistent with the predictions of the local dependence model and instead suggested that participants may have been treating source dependence as a shared property of the informational environment. To capture this possibility, we specified a global dependence model, which we evaluated in the present study.

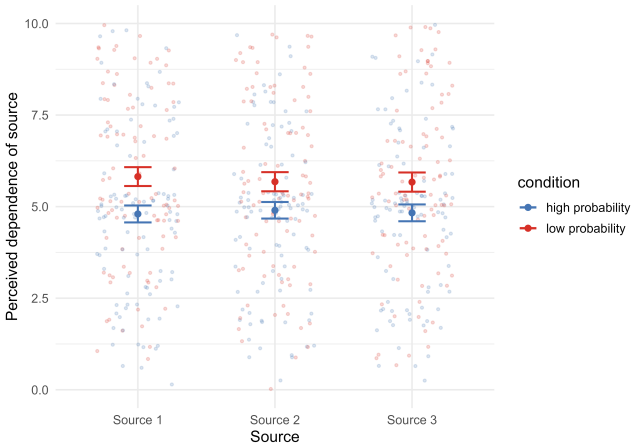

Figure 3: Perceived dependence of each source collected after their repeated claim Note. Ratings ranged from 0 (definitely dependent on another user's message) to 10 (definitely saw Whiskers) with error bars representing ±1 standard error of the mean. Individual data points have been jittered along the x-axis to improve clarity and visibility.

Results

Ahead of the main analyses, an independent-samples t-test confirmed that the probability manipulation was successful: participants in the low-probability condition reported significantly lower baseline belief in the claim (M = 1.91) than participants in the high-probability condition (M = 5.86), $t(197.99) = 14.66$, $p < .001$, 95% CI [3.41, 4.48]. Importantly, beliefs in the high-probability condition did not reach ceiling levels, leaving scope for belief updating across the experiment (Figure 3).

Regarding H1, an additional independent-samples t-test confirmed that prior beliefs affected perceived dependence on the latent source with participants in the low probability condition (M=5.20) rating the likelihood of dependence as higher than participants in the high probability condition (M=4.18), t(195.77) = 2.95, p=.0036, 95% CI [0.34, 1.70] (see Figure 3).

Regarding H2, a belief-change score was computed for each participant by subtracting baseline belief (TP0) from belief after the first claim (TP1) (see Figure 4). An independent-samples t-test revealed no statistically significant difference in the magnitude of belief change between the low- and high-probability conditions, $t(194.39) = -0.98$, $p = .33$, 95% CI [−0.66, 0.22]. A Bayesian t-test provided support for the null

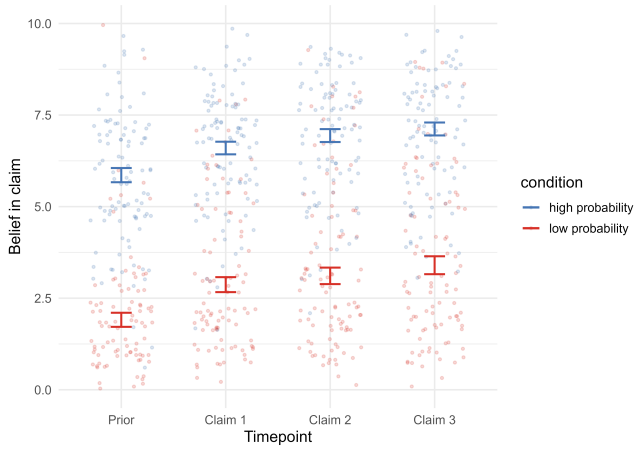

Figure 4: Belief updating across repeated claims. Note. Ratings ranged from 0 (very unlikely) to 10 (very likely) with error bars representing ±1 standard error of the mean. Individual data points have been jittered along the x-axis to improve clarity and visibility.

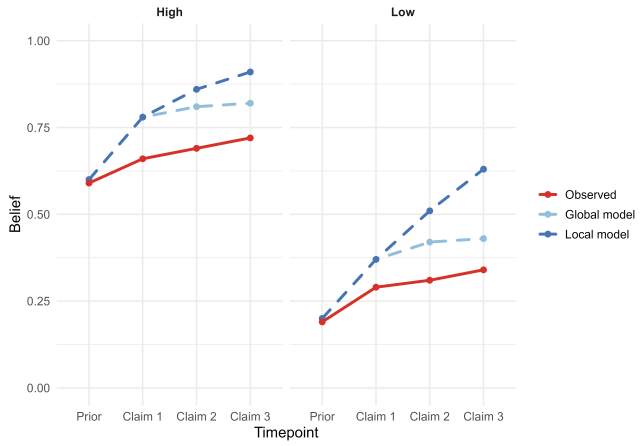

Figure 5: Observed beliefs compared with predictions from global and local dependence models for high and low probability conditions.

hypothesis, $BF_{10} = 0.24$. Together, these results indicate that the extent of belief updating following the first claim did not differ as a function of claim probability.

To assess whether perceived dependence on the latent source changed as additional claims were presented (H3a), dependence ratings after the first and third claims were compared within participants, separately for the low- and high-probability conditions. In the low-probability condition, a paired-samples t-test revealed no significant difference in dependence ratings between the first and third claims, $t(99) = -0.97$, $p = .33$, $M_{\text{diff}} = -0.15$, 95% CI $[-0.46, 0.16]$. A Bayesian paired-samples t-test provided moderate evidence in favour of the null hypothesis, $BF_{10} = 0.17$. In the high-probability condition, dependence ratings likewise did not differ between the first and third claims, $t(99) = 0.17$, $p = .87$, $M_{\text{diff}} = 0.03$, 95% CI $[-0.32, 0.38]$. The corresponding Bayesian analysis again supported the null hypothesis, $BF_{10} = 0.11$. Across both conditions, dependence judgements remained stable as additional claims were presented. To further substantiate this pattern, individual dependence trajectories were examined. Of the 200 participants, 104 assigned identical dependence ratings to the first and third informants, increasing to 154 participants whose ratings differed by no more than one scale point. This distribution indicates, in line with H3a predictions, that the modal response strategy was to maintain largely stable dependence judgements across the sequence of claims. To test whether belief in the claim increased as additional claims were presented (H3b), paired-samples t-tests compared belief after the first claim with belief after the third claim, conducted separately for the low- and high-probability conditions. In the low-probability condition, belief ratings were significantly higher after the third claim than after the first claim, $t(99) = -3.61$, $p < .001$, $M_{\text{diff}} = -0.53$, 95% CI $[-0.82, -0.24]$. A Bayesian paired-samples t-test provided strong evidence in favour of the alternative hypothesis, $BF_{10} = 43.73$. A similar pattern was observed in the high-probability condition. Belief ratings again increased from the first to the third claim, $t(99) = -3.89$, $p < .001$, $M_{\text{diff}} = -0.52$, 95% CI $[-0.79, -0.25]$. The corresponding Bayesian analysis yielded very strong evidence for the alternative hypothesis, $BF_{10} = 105.71$. Thus, contrary to H3b, belief in the claim increased as additional claims were presented in both probability conditions.

Finally, to test H4, empirical belief trajectories were compared with predictions generated by the local and global dependence models to assess which model provided the better fit to the data (see Figure 5). Model fit was evaluated at the participant level by computing the mean absolute deviation between observed belief ratings and model predictions across claims. Deviations were significantly smaller for the global dependence model than for the local dependence model, $t(199) = -9.31$, $p < .001$, mean difference $= -0.028$, 95% CI $[-0.034, -0.022]$. A Bayesian paired-samples analysis provided strong evidence in favour of the global dependence model, $BF_{10} = 2.1 \times 10^{14}$.

Discussion

The present study examined how prior beliefs shape perceived source dependence and how such perceptions, in turn, influence belief updating. Consistent with our hypotheses, we found that claims associated with lower prior belief were more likely to elicit perceptions of dependence among sources than claims associated with higher prior belief. When considered alongside prior work showing that perceived dependence constrains belief updating (e.g. Alister et al., 2025; Young et al., 2025), these findings suggest that prior beliefs might do more

than anchor initial judgments. They also shape how individuals construe the structure of the informational environment itself. By influencing whether multiple sources are treated as independent or redundant, priors may indirectly regulate the evidential weight assigned to repeated testimony. This dynamic creates the potential for recursive effects, whereby initial beliefs guide inferences about source dependence, which in turn determine how strongly subsequent information is incorporated.

Turning to our second hypothesis, we did not find evidence that perceived dependence modulated belief updating following the first claim. One interpretation of this null effect is that participants did not reliably integrate their dependence judgments into subsequent belief revision, consistent with prior work suggesting that individuals may struggle to fully incorporate structural properties of evidence into belief updating (Yousif et al., 2019). However, an alternative explanation is methodological rather than cognitive. The manipulation may not have induced a sufficiently strong contrast between low and high perceived dependence to yield detectable differences in belief change. In this case, perceived dependence may still be integrated into belief updating, but only when differences in dependence are more pronounced (Desai et al., 2022). Future work could adjudicate between these possibilities by employing stronger or more explicit manipulations designed to elicit clearer separation between high and low dependence conditions.

With regards to the perceptions of source dependence over time, this study points to the adoption of a global rather than local dependence model. In line with the global model there was no discernable difference between the dependence attributed to the first rather than the third source and model trajectory comparisons for belief updating were strongly in favour of the simplified model over the more complex one. In line with the literature (Lieder & Griffiths, 2020) this suggest participants might be following a rational model of belief updating but simplifying certain aspects to account for cognitive resources and capacities. At this stage, however, we do not exclude the possibility that at least some individuals might be engaging in a more heuristic approach or that the alignment of their updating patterns with the global model might be the result of conservative belief updating (Kovach, 2021). Contrary to our expectations, participants kept marginally updating their beliefs as new evidence came to light. While this is in contradiction to our original model's predictions it could also be the case that the indicative conditional probability values we used ahead of the study might not have matched the conditional probability values elicited by participants. This not only would affect belief updating patterns but would also shape the predictions that should be made on the sample. Adding a higher degree of uncertainty to the diagnosticity of the evidence (e.g. 0.8, 0.2, rather than 0.9, 0.1) might make individuals more aligned with the model regardless of our original predictions. To test this assumption future studies should aim to elicit conditional probability questions directly from participants to inform more accurate hypotheses and to use them for model fitting at the individual rather than aggregate level.

Taken together, these findings suggest that prior beliefs shape how individuals infer dependencies within complex informational environments and that people may simplify their mental representations of such environments to reduce cognitive demands. Crucially, this form of simplification has important consequences for belief updating. Although both the use of priors and the adoption of simplified dependence structures are compatible with a broadly rational framework, their interaction can nonetheless support self-reinforcing dynamics of belief polarisation. Under a local dependence model, individuals holding different priors should, in principle, converge on a shared belief space as independent evidence accumulates. By contrast, simplifying the informational structure in a way that approximates a global dependence model effectively limits the evidential value of repeated testimony, causing belief updating to attenuate early and reinforcing belief stagnation. Understanding how such cognitive processes operate is therefore essential for explaining the persistence of polarised beliefs and for informing interventions aimed at mitigating them, for example by emphasising cues to source independence. More broadly, these results highlight the value of a priori model simulations as conceptual tools for clarifying expected behavioural patterns and for interpreting empirical belief trajectories.

Several limitations of the present study warrant consideration. First, the specific setting in which we set our vignette limits the generalisability of the findings. While our pilot study situated in a different domain provides some grounds for generalisability, future work should examine whether similar patterns emerge across different domains and communication contexts, and identify instances in which qualitatively different updating patterns would emerge. Second, participants were explicitly asked to judge the degree of dependence between sources. Although this approach allows direct measurement of perceived dependence, it remains an open question whether individuals would spontaneously infer such dependence in more ecologically valid settings. Finally, the present study did not directly elicit participants' conditional probability judgments, which would provide more accurate insight into their assumptions about the diagnosticity of repeated testimony. Incorporating such measures in future research would allow closer alignment between empirical data and formal model predictions.

Notwithstanding these limitations, the present study provides evidence that individuals rely on their prior beliefs when forming judgments about source dependence and that they tend to simplify belief updating by treating dependence as a global rather than a local property of the informational environment.

References

Alister, M., Ransom, K., Connor Desai, S., Soh, E. V., Hayes, B., & Perfors, A. (2025). How convincing is a crowd? quantifying the persuasiveness of a consensus for different individuals and types of claims. *Psychological Science*, 09567976251344549.

Collins, P. J., Hahn, U., Von Gerber, Y., & Olsson, E. J. (2018). The bi-directional relationship between source characteristics and message content. *Frontiers in psychology*, *9*, 18.

Couch, N. (2022). The diversity principle and the evaluation of evidence. *Psychonomic Bulletin & Review*, *29*(4), 1270–1294.

Desai, S. C., Xie, B., & Hayes, B. K. (2022). Getting to the source of the illusion of consensus. *Cognition*, *223*, 105023. https://doi.org/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.

Harris, A. J. L., Hsu, A. S., & Madsen, J. K. (2012). Because Hitler did it! quantitative tests of Bayesian argumentation using ad hominem. *Thinking & Reasoning*, *18*(3), 311–343. https://doi.org/10.1080/13546783.2012.670753

Harris, A. J., Hahn, U., Madsen, J. K., & Hsu, A. S. (2016). The appeal to expert opinion: Quantitative support for a bayesian network approach. *Cognitive Science*, *40*(6), 1496–1533.

Kovach, M. (2021, January). Conservative Updating [arXiv:2102.00152 [econ]]. https://doi.org/10.48550/arXiv.2102.00152

Lieder, F., & Griffiths, T. L. (2020). Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources. *Behavioral and Brain Sciences*, *43*, e1. https://doi.org/10.1017/S0140525X1900061X

Madsen, J. (2016). Trump supported it?! A Bayesian source credibility model applied to appeals to specific American presidential candidates' opinions. *Cognitive Science*. Retrieved June 3, 2024, from https://www.semanticscholar.org/paper/Trump-supported-it! - A - Bayesian - source - credibility - Madsen / ccdbba50f93e81fdcb88b94828d68a760568d367

Madsen, J. K. (2019). Voter reasoning bias when evaluating statements from female and male political candidates. *Politics & Gender*, *15*(2), 310–335.

Madsen, J. K., Hahn, U., & Pilditch, T. D. (2020). The impact of partial source dependence on belief and reliability revision. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *46*(9), 1795–1805. https://doi.org/10.1037/xlm0000846

Mercier, H., & Morin, O. (2019). Majority rules: How good are we at aggregating convergent opinions? *Evolutionary Human Sciences*, *1*, e6. https://doi.org/10.1017/ehs.2019.6

Pearl, J. (1988). Embracing causality in default reasoning. *Artificial Intelligence*, *35*(2), 259–271.

Pfänder, J., De Courson, B., & Mercier, H. (2025). How wise is the crowd: Can we infer people are accurate and competent merely because they agree with each other? *Cognition*, *255*, 106005. https://doi.org/10.1016/j.cognition.2024.106005

Whalen, A., Griffiths, T. L., & Buchsbaum, D. (2018). Sensitivity to shared information in social learning. *Cognitive science*, *42*(1), 168–187.

Xie, B., & Hayes, B. (2022). Sensitivity to evidential dependencies in judgments under uncertainty. *Cognitive Science*, *46*(5), e13144.

Young, D. J., Madsen, J. K., & de-Wit, L. H. (2025). Belief polarization can be caused by disagreements over source independence: Computational modelling, experimental evidence, and applicability to real-world politics. *Cognition*, *259*, 106126. https://doi.org/10.1016/j.cognition.2025.106126

Young, D. J., Madsen, J. K., & Yousefi, S. (2023). The epistemic weight of silence. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *45*(45).

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.

Zhu, J.-Q., & Griffiths, T. L. (2024, November). Computation-Limited Bayesian Updating: A Resource-Rational Analysis of Approximate Bayesian Inference. https://doi.org/10.31234/osf.io/6aw3f