

Bayesian Belief Polarization due to Differential Perceptions of Source Independence

David J. Young (dy286@cam.ac.uk)

Department of Psychology,
University of Cambridge, Downing Place, CB2 3EB, Cambridge, UK

Jens Koed Madsen (j.madsen2@lse.ac.uk)

Department of Psychological and Behavioural Science,
London School of Economics and Political Science, Houghton Street, WC2A 2AE, London, UK

Lee H. de-Wit (lhd26@cam.ac.uk)

Department of Psychology,
University of Cambridge, Downing Place, CB2 3EB, Cambridge, UK

Abstract

Belief polarization represents a puzzling and important dynamic in belief updating. There is growing awareness that belief polarization can be Bayesian. We provide pre-registered experimental evidence that beliefs can polarize when people receive conflicting testimony from two groups of sources if they have different beliefs about which group's members have greater independence in terms of the factors which affect their testimony. We show this is predicted by a Bayesian Network model of belief updating.

Keywords: belief polarization; source independence; dependence; Bayesian updating.

Introduction

One of those most perplexing dynamics in belief updating is polarization, where people update their beliefs in opposing directions in response to the same evidence (Bullock, 2009; Jern et al., 2014). Belief polarization is sometimes interpreted as evidence for irrational reasoning (Gerber & Green, 1999; Mandelbaum, 2019), but in numerous scenarios can be a consequence of Bayesian updating (Benoît & Dubra, 2019; Bullock, 2009; Gerber & Green, 1999; Jern et al., 2014; Levy, 2021; Mandelbaum, 2019; Olsson, 2013; Pallavicini et al., 2021).

One scenario in which Bayesian updating can lead to belief polarization is when people are exposed to conflicting testimony from two sources about the same issue, but have different prior beliefs about which source is more *credible*. For example, if two people witness a Democrat politician argue that the prosecution of Donald Trump for election interference is valid and a Republican politician argue it is not, each would be expected to update their belief towards the position taken by the politician they deem more trustworthy, which will cause polarization if they disagree over which politician that is. Given that perceptions of source reliability do moderate belief updating in response to political claims (Madsen, 2016), and evidence suggests real-world US political belief polarization has increased in recent decades (DellaPosta, 2020; Dimock et al., 2014; Webster & Abramowitz, 2017), understanding whether source perception theories of belief polarization within politics are

valid is important for determining how effective depolarization strategies might be designed.

One factor which might contribute to real-world belief polarization, but has not yet been studied in the context of polarization, is source independence. In real-world politics, citizens typically receive convergent testimony from *groups* of sources – a party's politicians, supporters, and media allies may all make similar claims about a given topic. Understanding the extent to which members of these groups are *independent* is crucial for determining how much their testimony should be weighted. Condorcet's Jury Theorem (see, e.g., Boland, 1989) demonstrates that as the size of a group of unbiased sources providing testimony about an issue increases, the chance of their majority opinion being correct increases too, but does so at a slower rate when the sources are non-independent.

Dependence between sources can arise for two reasons: *shared factors*, and *intra-group transmission*. With *shared factors*, the same factors influence the sources' testimony, which could be factors affecting their *information* about the issue (e.g., getting their information from the same newspaper), their *reasoning* (e.g., sharing the same background assumptions), or their *motivation* about what to say on the issue (e.g., wanting to spread the same 'party line'). With *intra-group transmission*, some sources base their testimony on what they believe other members in the group to believe or have said, perhaps because they have discussed the issue beforehand or have heard each other's testimony. In both cases, source dependence creates *correlated error*, meaning that one source being inaccurate is more likely to mean other sources are also inaccurate; this diminishes the marginal effect of an increase in the number of sources saying the same thing upon persuasion.

Concretely, suppose people can choose from three newspapers to read, each with an independent probability of 0.2 of making an error. If three sources make a claim, and we know they all read the same newspaper we need to down-weight their testimony by the probability of that newspaper being incorrect, which is 0.2. Equally, if we know that one of the sources has read a newspaper, and the other two are just copying that source's testimony, we need to down-weight by the probability of that newspaper being wrong – also 0.2 But

if the three sources all read *different* newspapers, we would need to down-weight by the probability of *all three* being incorrect, which is much less, at 0.008. That is, dependency assumptions change the outcome by an order of magnitude.

Therefore, testimony from more-dependent groups should be down-weighted relative to testimony from less-dependent groups. One pathway by which perceptions of source dependence may contribute to real-world political belief polarization is that members of political groups may perceive members of opposing groups to have greater dependence than themselves. Therefore if exposed to conflicting testimony from both groups about a particular issue, partisans will place a greater weight on their own group's testimony than their opponents', and so tend to always shift their beliefs closer to that of their own group, even if entirely Bayesian.

Indeed, that partisan may perceive their opponents to possess greater dependence than themselves is predicted by Bayesian updating. To see why, firstly recall that Bayes' Rule tells us to update our belief in any hypothesis H after observing evidence E in proportion to the likelihood ratio $p(E|H)/p(E|\neg H)$ – how much more likely it is, that we would observe this evidence if the hypothesis were true than if it were false. Suppose we observe a group of sources make a claim we think is *false*, and use this as evidence to update our perception of that group's dependence – in this case E is the false claim, and H is the group's dependence. Since the *more dependent* a source group is, the *more likely* they are to make a false claim, $p(E|H)/p(E|\neg H) > 1$, and so we should *increase* our belief in the group's dependence – it is more likely they would make a false claim if they were dependent, so the false claim implies they are dependent.

Of course, opposing political groups often make claims that conflict with each other's priors, and will therefore appear to be incorrect to each other – Republicans argue Trump is innocent of wrongdoing, which Democrats disbelieve, and Democrats argue gun control laws should be tightened, which Republicans dispute. So, political groups will often be exposed to evidence which can be legitimately interpreted in Bayesian terms as providing evidence for their opponents' dependence. At the same time, the claims made by their own group will tend to cohere with their priors, providing weaker evidence for their own group's dependence. Scaled up, this will lead people to attribute greater dependence to opposing political groups than their own group.

We have developed a Bayesian Network (Pearl, 1988) model to provide a specific computational rationale for this general intuition, by adapting a standard source reliability Bayesian Network introduced by Bovens and Hartmann (2003). The network represents a model a person could use to make inferences using Bayes' Rule when exposed to testimony (TES) about a hypothesis (HYP) from a group of sources, accounting for their dependence. In this case, there are three sources, who all claim that the hypothesis is true.

Each source is modelled as having an 'effective reliability' (ER), where high reliability (ER = 1) means their testimony is the same as the ground truth of the hypothesis, and low

reliability (ER = 0) means they testify at random. Their effective reliability is determined by an interaction between their *latent individual reliability* (LIR), which is different for each source, a *latent shared reliability* factor (LSR), which affects all sources, and a dependence node (DEP) which controls whether the effective reliability is more influenced by the LIR or the LSR. With high dependence (DEP = 1), each source's effective reliability is equal to the latent shared reliability factor, with low dependence (DEP = 0), each source's reliability is equal to their latent individual reliability factor – see Table 2.

Therefore the more there are shared factors affecting the sources' testimony, or intra-group transmission, the higher the prior for Dependence should be set, as in these cases the effective reliability of the sources is primarily determined by factors which they all share; conversely when the sources are independent, their effective reliability is determined by factors unique to each source. Notably, if the hypothesis at hand is a highly subjective matter, relating to normativity or values, the model can still be used, but we may consider that no source can reliably speak to the ground truth as there may be no ground truth, in which case all reliability nodes can be set near to 0. The overall network of factors is shown in Figure 1 – the .1, .2, and .3 refer to whether the factor corresponds to Source 1, Source 2 or Source 3 respectively, where required. Tables 1 and 2 show the Conditional Probability Tables for the network.




Figure 1. The Bayesian Network used for simulations.

Table 1. Conditional Probability Table for TES.x nodes

HYP	ER.x	p(TES.x = "TRUE")	p(TES.x = "FALSE")
TRUE	1	1	0
TRUE	0	0.5	0.5
FALSE	1	0	1
FALSE	0	0.5	0.5

Table 2. Conditional Probability Table for ER.x nodes

DEP	LSR	LIR.x	p(ER.x = 1)	p(ER.x = 0)
1	1	1 or 0	1	0
1	0	1 or 0	0	1
0	1 or 0	1	1	0
0	1 or 0	0	0	1

We performed simulations to probe the behavior of the model in *R* (R Core Team, 2022) using the packages *gRain* (Højsgaard, 2012) and *purrr* (Wickham & Henry, 2023). All code is available via our OSF project (link below). We verified that agents update less when sources are more dependent. Then, we varied the prior for the hypothesis across simulations, keeping all other priors at 0.5, and exposing the agents to testimony from all three sources. As Figure 2 shows, the lower the prior for the hypothesis, the higher the perception of the source's dependence. This demonstrates, as expected, that greater dependence is attributed to source groups who make incorrect claims.




Figure 2. Simulation results: greater dependence is attributed to groups who make less-plausible claims.

Overall, there is a clear rationale for supposing that real-world political belief polarization could emerge from Bayesian reasoning about source dependence. To our knowledge, whether political groups do perceive their opponents to have greater testimonial dependence has not been tested. But more importantly, there is mixed empirical evidence that people are sensitive to source independence information when exposed to testimony from groups. Several studies suggest people do down-weight testimony from high-dependence groups. Madsen et al. (2020) finds that when expert biologists (Study 1) and economists (Study 2) are presented as having studied at the same school and subscribing to the same school of thought, people update less than when they are presented as independent. Similarly, across two studies which were both internally replicated, Mercier & Miton, (2019) find that source groups who provide testimony about the quality of a restaurant are believed less when they are presented as having shared information and intentions. These are instances of *shared factors* dependence, but Pilditch et al. (2020) provide evidence *intra-group transmission* cases – a plane crash investigator's report was less persuasive when the author was known to have read another investigator's report beforehand compared to when both were written independently.

But participants sometimes deviate from what researchers regard as normative. Mercier and Miton (2019) twice found

no evidence that shared *cognitive* factors affected belief updating, and regarding intra-group transmission, both Whalen et al. (2018) and Xie and Hayes (2022) find mostly null results. Therefore it is not safe to presume that people's reasoning about source dependence will follow our Bayesian model. Before proceeding any further with the study of whether real-world political belief polarization then, it seems necessary to experimentally establish whether differential perceptions of dependence can cause belief polarization when people are exposed to conflicting testimony from two source groups.

We present evidence from such an experimental study, which was pre-registered on the OSF: https://osf.io/ysqa8/?view_only=709c465855f04e509c6903852b8a6006. Our data analysis code, for both the Study and two pilot studies, as well as the code for our simulations above, is also available here. All aspects of the study were as pre-registered except that the study took 4 minutes rather than 5, and we recruited $N = 351$ rather than $N = 350$.

The Study

Participants

We recruited British participants from Prolific, and paid them £9/hour for 4 minutes, with 357 completing the study and six excluded (for failing attention checks – see Exclusions below), leaving a sample of $N = 351$. $N = 350$ was targeted, and pre-registered, on the basis it provides 80% power to detect effects as small as $d = 0.15$ for both one-sample and paired-samples *t*-tests (see Planned Analyses below), and 95% power for effects as small as $d = 0.20$ for the same. It therefore exceeds the power needed to detect the smallest effect observed in the pilot studies ($d = 0.23$ – see Pilot Studies below).

We inspected the demographic information participants had already provided to Prolific. Of the 351 retained participants, all provided gender data with 175 female and 176 male, 350 provided age data with a median of 39 and a range of 19-77, and 344 provided ethnicity data with 295 White, 24 Asian, 11 'Mixed', 10 Black, and 4 'Other'.

Design

Participants were presented with conflicting testimony from two groups of sources about whether a politician is guilty of a scandal they have been accused of. Alleged scandals are common and consequential topics of debate in politics, and typically involve disputes between groups of sources who may not always be reliable – it is therefore a useful domain for an experimental study of source dependence's contribution to polarization for the sake of real-world applicability. One group of sources, the 'Innocent' group, say the politician is not guilty, whereas the other group, the 'Guilty' group, says they are. Their claims are presented one at a time, with beliefs about the politician's guilt measured before and after each claim. One group is presented as having higher Dependence than the other, in terms of them sharing more factors which influence their

testimony (their ideology, background, and their own information sources).

We therefore have a 2x2 within-subjects factorial design across our four trials: Presentation Order (Guilty first vs Innocent first) x Greater Dependence (Guilty vs Innocent Group). We measure people's belief that the politician is guilty before any testimony and then after each group's testimony. This allows us to analyze the effect of the group's Dependence (High vs Low) on the persuasiveness of each claim, and the effect of which group has greater dependence on the overall direction of people's updating for each trial. We expected that people should be less persuaded by more-dependent source groups, and this should cause belief polarization between participants who have different beliefs about which group is more dependent: when the Guilty group is more dependent, participants should shift their opinion towards believing the politician is innocent, but when the Innocent group is more dependent, they should shift towards believing the politician is guilty.

Procedure

Before the experimental trials, participants chose to participate on Prolific and provided informed consent. They were given brief instructions that the scenarios concerned politics in a made-up Western democracy, but that they should respond as if they were real. They were warned about the attention checks. Then the four experimental trials were presented in a random order. After completing, participants were debriefed, thanked, and redirected back to Prolific (demographic data are taken from Prolific).

Trials

Figure 3 shows what happens in each trial.




Figure 3. Schematic of the events in each trial.

Each trial is presented as a short report about the debate concerning the allegation of a politician's involvement in a scandal. Each trial begins with the preamble "A prominent politician has been accused of {engaging in tax evasion/making fraudulent expenses claims/lying about who was driving their car in order to avoid a speeding fine/bullying members of staff}. They proclaim their

innocence, but not everyone is convinced. We spoke to six political commentators about the allegations." The exact scandal is chosen randomly for each trial (without replacement).

After this preamble, prior beliefs are measured on a 0-100 slider scale, using the question "How likely do you think it is that the politician is guilty of {scandal}?" As participants have little information to use to answer this question to begin with, we precede this measure with the following qualification: "We will tell you what the commentators said on the next two pages. But we want to know your view of the allegation before receiving this information. Therefore, though you have been given very little information about the allegation, please answer the following question".

After the initial belief measurement, participants are shown two paragraphs. In each paragraph, the testimony of one of the groups is given followed by a description of their Dependence. The testimony is of the form: "Three of the commentators said they thought the politician was {innocent/guilty}". Then the Dependence is described: "These commentators get their information about political issues from very {similar/different} sources to one another, have very {similar/different} political ideologies to one another, and all come from very {similar/different} backgrounds to one another", with "similar" used throughout for the High Dependence group and "different" for the Low Dependence group. This manipulation should therefore engender differential perceptions of the extent to which the sources possess shared factors which influence their testimony.

At the end of each paragraph, we ask participants to give their belief in the effectiveness of the policy using the same question as for the prior (but without the qualification). This question is presented immediately below the corresponding paragraph on the same page. Participants have to click forward after giving their belief rating to view the second paragraph and cannot navigate backwards after having done so. Beneath the post-testimony belief measurement for the second group, an attention check is given. Attention checks just ask for a particular number to be given on the sliding scale (72, 24, 18, or 86).

Measures

Testimony-wise Persuasion

We calculate 'Testimony-wise Persuasion' for each piece of testimony by finding how much the participant's post-testimony belief has shifted in the direction of the provided testimony compared to their pre-testimony belief. Since beliefs are measured on a scale where high scores indicate a perception of guilt, this is Post - Pre when the testimony is "Guilty" and Pre - Post when the testimony is "Innocent".

Trial-wise Updating

We calculate 'Trial-wise Updating' for each trial by subtracting each person's prior belief in the politician's guilt from their final posterior for that trial.

Hypotheses

Our pre-registered hypotheses were:

H1: Testimony-wise Persuasion will be greater when Dependence is low vs high.

H2: Trial-wise Updating will be positive in the ‘Innocent’ Greater Dependence condition and negative in the ‘Guilty’ Greater Dependence condition (i.e., whichever group is less dependent, people’s beliefs will shift towards their position).

Planned Analyses

Our pre-registered planned analyses were:

H1: We will aggregate Testimony-wise Persuasion scores by finding the average score in each Dependency condition across trials for each participant. We will then perform a paired-samples *t*-test comparing these aggregated scores (two-tailed, $\alpha = 0.05$). We expect Testimony-wise Persuasion to be higher in the Low Dependence condition than the High Dependence condition.

H2: We will aggregate Trial-wise Updating scores across Presentation Order conditions by finding the average Trial-wise Updating score in each Greater Dependence condition for each participant. We will perform two one-sample *t*-tests where $\mu = 0$ (two-tailed, $\alpha = 0.05$), one for each Greater Dependency condition. We expect trial-wise updating to be positive when the Innocent group have greater dependency, and negative when the Guilty group have greater dependency. It should be noted that since *both* significant tests must be passed to affirm this hypothesis, there is no need to correct the alpha level – see Rubin (2021).

Pilot Studies

To inform our design and sampling we conducted two pilot experiments. One was identical to the proposed design with $N = 40$. The other was identical except the stimuli were policy proposals (devoid of any ideological content) and the testimony and judgments concerned whether they would be beneficial or damaging for the country, with some adjustments to the question wordings and task instructions to align with this. This pilot had $N = 39$ after one person was excluded, for failing attention checks in accordance with the criteria stated here.

We applied our planned analyses to the data collected in these pilot studies. The results were consistent with H1 in both studies. Testimony-wise Persuasion was greater in response to testimony from Low Dependence sources than High Dependence sources in both the scandals version (Low: $M = 8.66$, $SD = 11.42$ vs High: $M = 2.25$, $SD = 10.29$), with a significant difference, $t(39) = 4.201$, $p < .001$, and a moderate effect size, Cohen’s $d = 0.66$ [0.31, 1.00], and in the policies version (Low $M = 9.53$, $SD = 13.45$ vs High $M = 5.74$, $SD = 10.31$), with a significant difference $t(38) = 2.587$, $p = .014$, and a small effect, Cohen’s $d = 0.41$ [0.08, 0.74].

The results were consistent with H2 in three out of four analyses. Trial-wise Updating was positive in the condition where this was expected in the scandals version (Greater Dependence = “Innocent” Group: $M = 4.01$, $SD = 11.92$), with a significant difference from 0, $t(39) = 2.129$, $p = .040$,

and a small effect size, Cohen’s $d = 0.34$ [0.02, 0.65], as well as the policies version (Greater Dependence = “Damaging” Group: $M = 4.63$, $SD = 12.43$), with a significant difference from 0, $t(38) = 2.326$, $p = .025$, and a small effect size, Cohen’s $d = 0.37$ [0.05, 0.69]. Trial-wise Updating was negative in the condition where this was expected in the scandals version (Greater Dependence = “Guilty” Group: $M = -8.80$, $SD = 13.11$), with a significant difference from 0, $t(39) = 4.245$, $p < .001$, and a moderate effect size, Cohen’s $d = 0.67$ [0.32, 1.01], and was also negative where expected in the policies version (Greater Dependence = “Beneficial” Group: $M = -2.94$, $SD = 12.63$), but this difference was *not* significantly different from 0, $t(38) = 1.452$, $p = .155$, and had a small effect size, Cohen’s $d = 0.23$ [-0.09, 0.55].

These results provide preliminary evidence for H1 and H2. They also show larger effect sizes in the scandals version than the policies version (0.66, 0.34, and 0.67 vs 0.41, 0.37, 0.23), which is why we used scandals as stimuli in the main study.

Results

Regarding H1, Testimony-wise Persuasion was, as expected, greater in response to testimony from Low Dependence sources than High Dependence sources (Low: $M = 9.34$, $SD = 10.89$ vs High: $M = 3.92$, $SD = 8.38$), with a significant difference, $t(350) = 11.006$, $p < .001$, and a moderate effect size, Cohen’s $d = 0.59$ [0.47, 0.70]. The distribution of responses with means and confidence intervals is shown in Figure 4.




Figure 4. The distribution of Testimony-wise Persuasion scores by Dependency condition, with means and 95% confidence intervals.

Regarding H2, Trial-wise Updating was, as expected, positive when the Innocent group had greater dependence ($M = 2.85$, $SD = 10.94$), with a significant difference from 0, $t(350) = 4.876$, $p < .001$, and a small effect size, Cohen’s $d = 0.26$ [0.15, 0.37]. Trial-wise Updating was, as expected, negative when the Guilty group had greater dependence (M

= -8.00, SD = 13.77), with a significant difference from 0, $t(350) = 10.893$, $p < .001$, and a moderate effect size, Cohen's $d = 0.58$ [0.47, 0.69]. The distribution of Trial-wise Updating scores is shown in Figure 5.




Figure 5. The distribution of Trial-wise Updating scores by Dependency condition, with means and 95% confidence intervals.

Discussion

Both of our pre-registered hypotheses were affirmed – people were less persuaded by political sources who had higher dependence due to shared backgrounds, ideologies, and information sources (H1), and because participants were presented with conflicting testimony from both a low-dependence and high-dependence group for each issue, the direction in which they updated their belief overall was towards the group they thought was less dependent (H2). This created belief polarization between participants who believed different groups had higher dependence for a given issue, who updated their beliefs in opposing directions in response to what was otherwise identical testimonial evidence.

This work contributes to the literatures on belief updating, belief polarization, and source dependence. Firstly, we provide further evidence that perceptions of source dependence moderate belief updating, but extend this evidence to the important real-world domain of politics for the first time. Secondly, we present a novel mechanism by which belief polarization can occur due to differential perceptions of source group dependence and provide empirical evidence of its validity.

Our results strengthen the empirical evidence that belief polarization can be Bayesian (Benoît & Dubra, 2019; Bullock, 2009; Gerber & Green, 1999; Jern et al., 2014; Levy, 2021; Mandelbaum, 2019; Olsson, 2013; Pallavicini et al., 2021). Future research could explore whether perceptions of source dependence provides a viable explanation for well-known cases of belief polarization following exposure to mixed evidence like Lord et al.'s (1979) 'biased assimilation' study – this study has been cited over 6500 times but people's attitudes only rarely polarize in this paradigm (Anglin, 2019;

Greitemeyer, 2014; Kuhn & Lao, 1996; Munro & Ditto, 1997), suggesting the conventional explanation that people polarize when exposed to mixed evidence because they look for ways to discredit the evidence which conflicts with their prior is unlikely to be true. If the effect actually depends upon an additional criterion being satisfied, such as differential priors for source dependence (or other relevant source perceptions), the patchy replication record would make sense, as without such differential background priors, which will not always be in place and have never been measured in these studies, the effect would not be expected to occur.

As for limitations, we cannot be sure differential perceptions of source dependence will create belief polarization 'in the wild' – participants were given timely cues to consider source dependence in this study, so whether source dependence is spontaneously integrated into belief updating in response to real-world evidence will require further study and clever designs to avoid confounding. What's more, the dependency information was 'served on a plate' to participants – we told them whether the sources were 'very similar' or 'very different' regarding factors affecting their testimony, whereas in the real world, people would have to keep track of such factors and judge their similarity for a given group for themselves. Nevertheless, the fact that participants regarded shared factors as a reason to discredit testimony shows this critical cognitive building block for differential source dependence perceptions to cause belief polarization is in place. Given that real-world political sources do form groups who share dependencies and provide converging testimony – e.g., a party's politicians share motivations to defend their party's policies, a party's supporters may all use similar media sources – but also that diverse political groups sometimes argue for the same end goals for independent reasons, accurately keeping track of the relevant dependencies between political sources would be useful so that people can properly gauge how persuasive a group's testimony should be. Clearly however, doing this in the real world presents additional cognitive demands which have been elided in this study and which are deserving of further investigation.

It should also be noted that we have not empirically tested the accuracy of our Bayesian Network in this study, merely whether its qualitative claims are borne out. People may well *procedurally* reason in a different, simpler way, but our results suggest the Network may at least broadly capture the logic underlying how people reason about source dependence, though apparent deviations from normativity regarding dependence elsewhere (Mercier & Miton, 2019; Whalen et al., 2018; Xie & Hayes, 2022) still require explanation. More specific quantitative testing of the model would be useful in future studies.

To conclude, this work establishes a computational Bayesian rationale for hypothesizing that when exposed to conflicting testimony from two groups, as is common in politics, differential perceptions of source dependence might cause belief polarization. Furthermore, we provide pre-registered experimental evidence that they do.

Acknowledgments

This research was funded by a Templeton World Charity Foundation grant, TWCF Number 31453, <https://doi.org/10.54224/31453>. We thank Adam Harris, Hugo Mercier, Gordon Brown, and three anonymous reviewers for helpful feedback.

References

- Anglin, S. M. (2019). Do beliefs yield to evidence? Examining belief perseverance vs. change in response to congruent empirical findings. *Journal of Experimental Social Psychology*, 82, 176–199. <https://doi.org/10.1016/j.jesp.2019.02.004>
- Benoît, J., & Dubra, J. (2019). Apparent Bias: What Does Attitude Polarization Show? *International Economic Review*, 60(4), 1675–1703. <https://doi.org/10.1111/iere.12400>
- Boland, P. J. (1989). Majority Systems and the Condorcet Jury Theorem. *The Statistician*, 38(3), 181. <https://doi.org/10.2307/2348873>
- Bovens, L., & Hartmann, S. (2003). *Bayesian Epistemology*. Oxford University Press.
- Bullock, J. G. (2009). Partisan Bias and the Bayesian Ideal in the Study of Public Opinion. *The Journal of Politics*, 71(3), 1109–1124. <https://doi.org/10.1017/S0022381609090914>
- DellaPosta, D. (2020). Pluralistic Collapse: The “Oil Spill” Model of Mass Opinion Polarization. *American Sociological Review*, 85(3), 507–536. <https://doi.org/10.1177/0003122420922989>
- Dimock, M., Doherty, C., Kiley, J., & Oates, R. (2014). Political Polarization in the American Public: How Increasing Ideological Uniformity and Partisan Antipathy Affect Politics, Compromise, and Everyday Life. Pew Research Center.
- Gerber, A., & Green, D. (1999). Misperceptions about Perceptual Bias. *Annual Review of Political Science*, 2(1), 189–210. <https://doi.org/10.1146/annurev.polisci.2.1.189>
- Greitemeyer, T. (2014). I Am Right, You Are Wrong: How Biased Assimilation Increases the Perceived Gap between Believers and Skeptics of Violent Video Game Effects. *PLoS ONE*, 9(4), e93440. <https://doi.org/10.1371/journal.pone.0093440>
- Hahn, U. (2023). Individuals, Collectives, and Individuals in Collectives: The Ineliminable Role of Dependence. *Perspectives on Psychological Science*, 17456916231198479. <https://doi.org/10.1177/17456916231198479>
- Højsgaard, S. (2012). Graphical Independence Networks with the gRain Package for R. *Journal of Statistical Software*, 46(10), 1–26. <https://www.jstatsoft.org/v46/i10/>
- Jern, A., Chang, K. K., & Kemp, C. (2014). Belief Polarization is not Always Irrational. *Psychological Review*, 121(2), 206–224. <https://doi.org/10.1037/a0035941>
- Kuhn, D., & Lao, J. (1996). Effects of Evidence on Attitudes: Is Polarization the Norm? *Psychological Science*, 7(2), 115–120. <https://doi.org/10.1111/j.1467-9280.1996.tb00340.x>
- Levy, N. (2021). *Bad Beliefs: Why They Happen to Good People* (1st ed.). Oxford University Press. <https://doi.org/10.1093/oso/9780192895325.001.0001>
- Lord, C. G., Ross, L., & Lepper, M. R. (1979). Biased Assimilation and Attitude Polarization: The Effects of Prior Theories on Subsequently Considered Evidence. *Journal of Personality and Social Psychology*, 37(11), 2098–2109.
- 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>
- Mandelbaum, E. (2019). Troubles with Bayesianism: An introduction to the psychological immune system. *Mind & Language*, 34(2), 141–157. <https://doi.org/10.1111/mila.12205>
- Mercier, H., & Miton, H. (2019). Utilizing simple cues to informational dependency. *Evolution and Human Behavior*, 40(3), 301–314. <https://doi.org/10.1016/j.evolhumbehav.2019.01.001>
- Munro, G. D., & Ditto, P. H. (1997). Biased Assimilation, Attitude Polarization, and Affect in Reactions to Stereotype-Relevant Scientific Information. *Personality and Social Psychology Bulletin*, 23(6), 636–653. <https://doi.org/10.1177/0146167297236007>
- Olsson, E. J. (2013). A Bayesian Simulation Model of Group Deliberation and Polarization. In F. Zenker (Ed.), *Bayesian Argumentation* (Vol. 362, pp. 113–133). Springer Netherlands. https://doi.org/10.1007/978-94-007-5357-0_6
- Pallavicini, J., Hallsson, B., & Kappel, K. (2021). Polarization in groups of Bayesian agents. *Synthese*, 198(1), 1–55. <https://doi.org/10.1007/s11229-018-01978-w>
- Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann Publishers: San Francisco, CA.
- Pilditch, T. D., Hahn, U., Fenton, N., & Lagnado, D. (2020). Dependencies in evidential reports: The case for informational advantages. *Cognition*, 204, 104343. <https://doi.org/10.1016/j.cognition.2020.104343>
- R Core Team (2022). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL www.R-project.org/.
- Rubin, M. (2021). When to adjust alpha during multiple testing: A consideration of disjunction, conjunction, and individual testing. *Synthese*, 199(3–4), 10969–11000. <https://doi.org/10.1007/s11229-021-03276-4>
- Van Bavel, J. J., & Pereira, A. (2018). The Partisan Brain: An Identity-Based Model of Political Belief. *Trends in Cognitive Sciences*, 22(3), 213–224. <https://doi.org/10.1016/j.tics.2018.01.004>
- Webster, S. W., & Abramowitz, A. I. (2017). The Ideological Foundations of Affective Polarization in the U.S. Electorate. *American Politics Research*, 45(4), 621–647. <https://doi.org/10.1177/1532673X17703132>

Wickham H, Henry L (2023). purrr: Functional Programming
Tools. R package version 1.0.1,
CRAN.R-project.org/package=purrr