



Belief polarization can be caused by disagreements over source independence: Computational modelling, experimental evidence, and applicability to real-world politics

David J. Young^{a,*}, Jens Koed Madsen^b, Lee H. de-Wit^a

^a Department of Psychology, University of Cambridge, Downing Site, Downing Street, Cambridge, Cambridgeshire CB2 3EB, United Kingdom

^b Department of Psychological and Behavioural Science, London School of Economics and Political Science, Connaught House, 63-65 Aldwych, London WC2B 4EJ, United Kingdom

ARTICLE INFO

Keywords:

Polarization
Source
Dependence
Independence
Bayesian
Updating

ABSTRACT

A large literature debates whether belief polarization, in both experiments and real-world political opinion data, is the result of biased forms of reasoning like biased assimilation and motivated reasoning, or if it can be caused by rational reasoning. We present evidence for the plausibility of a novel Bayesian mechanism of experimental and real-world political belief polarization involving perceptions of source independence. We show, using a novel Bayesian network, that when presented with conflicting testimony from two source groups, Bayesians should update towards the position of the group they deem to be more independent, meaning those who disagree about which group that is should polarize. We find in a pre-registered experiment ($N = 351$) that human participants polarize under these conditions. We then find in a UK study ($N = 507$) and a pre-registered US replication ($N = 300$) that, using a novel scale instrument, real-world partisans (Labour, Conservative, Republican, and Democrat) perceive their party's supporters to be more independent than the opposing party's supporters, with large average effect sizes ($d = 0.87$ UK, $d = 0.82$ US), suggesting the conditions are in place for such polarization to occur in the real world. Accordingly, we find that those who view their party's supporters as more independent than their opponents to the greatest extent have the most polarized beliefs, even after controlling for partisanship and affective polarization. Overall, our results highlight perceptions of testimonial independence as a plausible mediator of experimental and real-world belief polarization.

1. Introduction

Polarization, where people update their beliefs in opposing directions after observing the same evidence (Bullock, 2009; Jern et al., 2014), is one of those most perplexing dynamics in belief updating. Belief polarization is often ascribed to irrational reasoning (Baron, 2008; Gerber & Green, 1999; Kahan et al., 2012; Lord et al., 1979; Mandelbaum, 2019), but can theoretically be a consequence of Bayesian updating too (Benoît & Dubra, 2019; Bullock, 2009; Gerber & Green, 1999; Jern et al., 2014; Levy, 2021; Mandelbaum, 2019; Olsson, 2013; Pallavicini et al., 2021). But rarely have attempts been made to empirically test the predictions these Bayesian theories make about the circumstances under which polarization will occur. Furthermore, it is unknown whether Bayesian belief polarization processes contribute to the real-world increases in political issue polarization that have been

well-documented in the US (DellaPosta, 2020; Dimock et al., 2014; Webster & Abramowitz, 2017). This paper presents and tests a Bayesian explanation for belief polarization which involves perceptions of testimonial independence, and applies it to real-world political polarization.

According to Condorcet's Jury Theorem (Condorcet, 2014 [1785]; see, e.g., Boland, 1989), when a group of sources provide testimony about an issue, the chance their majority opinion is correct increases as the size of the group increases, so long as they are unbiased (i.e., their individual chance of correctly stating whether the hypothesis at issue is true or false is greater than 50%), and they are independent. This is because, generally speaking, the greater the number of unbiased individuals in agreement, the more unlikely it is they could all be wrong simultaneously, however, there is often a greater risk of error when a group's members are less independent.

Non-independence between sources can arise for two reasons: shared

* Corresponding author.

E-mail addresses: dy286@cam.ac.uk (D.J. Young), j.madsen2@lse.ac.uk (J.K. Madsen), lhd26@cam.ac.uk (L.H. de-Wit).

<https://doi.org/10.1016/j.cognition.2025.106126>

Received 1 August 2024; Received in revised form 10 March 2025; Accepted 17 March 2025

Available online 22 March 2025

0010-0277/© 2025 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

factors, and intra-group transmission. Shared factors could refer to factors affecting their *information* about the issue (e.g., getting their information from the same newspaper or data), 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’). Intra-group transmission refers to cases where the group’s members directly influence each other’s testimony, perhaps because some copy others’ testimony or beliefs, or because they pressure each other into saying particular things.

To see why non-independence carries greater epistemic risk, consider a case where a group of n people have the choice of reading n independent and equally-reliable newspapers to learn about an issue. If they all read a different newspaper and all end up having the same belief on that issue, all n newspapers would have to be wrong for the group’s consensus position to be wrong. In this case independence is high, and the group’s risk of error is low. But, if all of the group read the *same* newspaper, only that one newspaper would need to be incorrect for the whole group to be wrong, which is much more likely. In this case, independence is low and the group’s risk of error is high.

This is an example of non-independence occurring due to shared factors, but intra-group transmission has the same effect. Consider a case where a group of otherwise-independent and equally-reliable sources all make an independent determination as to the truth regarding a particular matter, and arrive at the same conclusion. All of their individual processes for determining the truth would have to simultaneously fail for their position to be wrong (high independence, low risk of error). But, if the group simply copied the conclusion arrived at by one of their members, only that member’s process for determining the truth would need to fail to make everybody’s position wrong, which is much more likely (low independence, high risk of error). Perceptions of source independence may contribute to real-world political polarization, in the sense of disagreement between political groups over policy issues (Dimock et al., 2014; Webster & Abramowitz, 2017). In politics, citizens often witness conflict between groups of sources who provide internally-consistent but mutually-conflicting testimony about political issues, where the testimonial independence of the groups is dubious. For example, a party’s politicians, supporters, and media allies may all push a narrative which is sympathetic to the party when its leader is embroiled in a scandal, whereas the equivalent figures from opposing parties may push a more critical narrative. Each group’s members may share internal dependencies by consuming the same media sources and sharing the same assumptions about political reality; furthermore they may increase their level of dependence by transmitting opinions to one another through discussion and from pressuring each other to conform to the group’s position (Jönsson et al., 2015). People’s beliefs about the relative testimonial independence of the conflicting groups should impact which group they trust more, and consequently whether they come to believe those who dispute or endorse the allegations.

Critically, members of political groups may perceive their opponents to have greater dependence than themselves. For example, partisans may believe their opponents to have less diverse media diets, to be more strongly biased in the way they interpret political information, and to be under greater pressure to conform to the party line. If partisans do perceive their own political groups to be more independent than opposing political groups, this should cause supporters of opposing groups to shift their beliefs closer to those of their own group over time, as they effectively place a greater weight on their own group’s testimony than their opponents’ when exposed to conflicting testimony about contested issues. This would create polarization while being purely Bayesian.

Further, Bayesian reasoning about source independence predicts that partisans *should* infer that their opponents are more dependent than themselves. To see why, 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. With this in mind, 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, we know that $p(E|H)/p(E|\neg H) > 1$ – and so we should *increase* our belief in the group’s dependence. This is because it is more likely they would make a false claim if they were dependent, so the false claim implies, probabilistically, that they *are* dependent. In other words, when people encounter a group of sources making a claim they consider to be false, they can infer that the group has a high level of dependence to explain why so many people would simultaneously make the same error.

Importantly, opposing political groups routinely make claims that conflict with each other’s priors. This means that political groups’ testimonies often appear to be incorrect from their opponents’ perspective. 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 should lead people to attribute greater dependence to opposing political groups than their own group over time, thus setting in place the conditions for belief polarization to occur.

Thus, there are, normatively, reciprocal paths whereby implausible claims imply higher dependence, and higher dependence implies lower plausibility. In the real world, these paths may connect to create a mass belief polarization cycle. If citizens are exposed to conflicting claims from two source groups repeatedly over time, it follows that citizens will attribute lower dependence to the group whose claims lie closer to their initial priors. Then, because they view them as less-dependent, they will trust that group more when exposed to further conflicting testimony, and shift their beliefs further to that group’s position. This will cause those whose priors initially lie closer to different groups to shift further apart, polarizing, because they each attribute lower dependence to a different group. These processes will continue in a polarizing feedback loop. Therefore, we should investigate perceptions of testimonial independence as a possible mediator of political belief polarization in the mass public.

1.1. The model

We have developed a Bayesian Network (Pearl, 1988) model to provide a specific computational rationale for these general intuitions. Specifying theories in precise computational terms has been proposed as a way to improve the replicability of psychological science (Guest & Martin, 2021; Muthukrishna & Henrich, 2019; Oberauer & Lewandowsky, 2019; Rollwage et al., 2019), one reason for this being that it avoids ambiguity about what theories predict, improving falsifiability. Our model also makes it objective and verifiable that the process by which we predict polarization to occur is indeed Bayesian.

We adapt a Bayesian Network introduced by Bovens and Hartmann (2003) to model how people can integrate perceptions of a source’s reliability when learning from their testimony. This network has already been adapted by others to develop Bayesian Network models of how people learn from testimony provided by sources taking into consideration different source characteristics, which have been empirically validated (Collins et al., 2018; Harris et al., 2012, 2016; Madsen, 2016, 2019; Young et al., 2023); it therefore provides a natural foundation to begin modelling the impact of dependence perceptions.

The network represents a model a person could use to make inferences using Bayes’ Rule when exposed to several pieces of testimony (“TES”) about a hypothesis (some true-or-false claim about the state of the world, “HYP”) from a group of sources, accounting for their level of dependence on a factor which might influence all their testimonies simultaneously. This factor could be a kind of shared-factors dependency, i.e. a media source, a reasoning assumption, or an intention

about what to say, or it could be an intra-group transmission dependency, i.e. the reliability of one source from whom all the other sources take their testimony.

In our model we have three sources for tractability’s sake, but more could easily be added. 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 a three-way interaction between three factors: their Latent *Individual* Reliability factor (“LIR”), which refers to how reliable the source would be in a counterfactual world where they made their determination of what to say completely on their own, and is therefore different for every source, with one node per source; a Latent *Shared* Reliability factor (“LSR”), which refers to the reliability of the factor on which the other sources may be relying to determine what to say, and is therefore the same for every source, with one node shared between all three sources; and a dependence factor (“DEP”), which refers to how reliant on this shared factor the sources are, which again is the same for all sources, with one node shared between all three sources. 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.

A novel feature of this model is that it allows participants to *flexibly* update their perception of the sources’ dependence, and it can be used to model dependencies of any kind. Previous studies using Bayesian Networks to study reasoning about dependencies have ‘hard-coded’ the dependencies into the structure of the model (e.g., Madsen et al., 2020; Pilditch et al., 2020), meaning different networks have to be used to model independent and dependent cases (Madsen et al., 2020); this approach is therefore not easily translatable to cases where we want to make normative predictions about when people should *infer* dependencies, as inferring the structure of Bayes Nets lies beyond the usual scope of Bayesian updating. By encoding beliefs about dependencies in people’s priors rather than the model structure, our model can be used to model Bayesian reasoning about any kind of dependency and at any level of dependency. Admittedly, the model makes some simplifying assumptions by assuming that all sources are equally dependent on the shared factor and that there is only one shared factor, but more complex models could be built-out as variations on this one.

The overall network of factors is shown in Fig. 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.

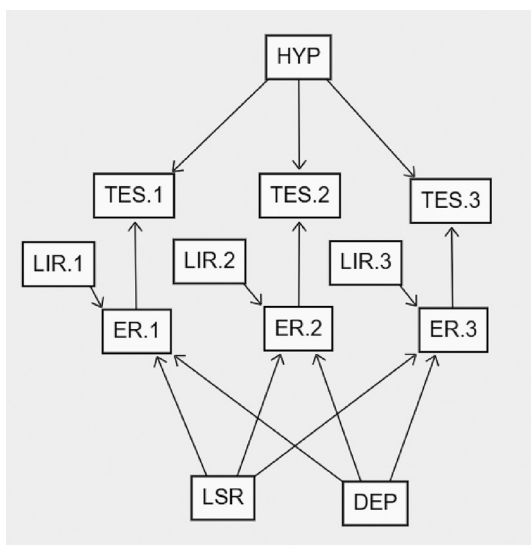


Fig. 1. The Bayesian Network used for simulations.

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

Note. x denotes the specific source to whom the node belongs, 1, 2, or 3.

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

Note. x denotes the specific source to whom the node belongs, 1, 2, or 3.

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). In the first simulation, we varied only the Dependence prior, holding all other priors at 0.5. We then exposed the model to the evidence of all three sources testifying ‘True’, and measured the posterior belief in the hypothesis. As Fig. 2a shows, the posterior is lower when the sources are more dependent. For the second simulation, we varied the prior for the Hypothesis, keeping all other priors at 0.5, then exposed the model to the evidence of all three sources testifying ‘True’, and measured the posterior belief in Dependence. As Fig. 2b shows, greater dependence is attributed when the prior for the hypothesis is lower.

Thus, the model predicts a reciprocal relationship between perceived credibility and perceived dependence: implausible claims imply higher dependence, and lower dependence implies plausibility.

Extrapolating from the trend in Fig. 2A, the model predicts belief polarization when people hear conflicting testimony from two source groups, and disagree about which group has higher dependence. We verify this by simulating updating for an agent who receives testimony about the same hypothesis from *two* groups, each comprising three sources, where one group’s sources all say the hypothesis is true and the other’s all say it is false. We expand the network in Fig. 1 to accommodate the second group, by forming two copies of the network, one for each group, but connecting them to the same hypothesis node. The agent begins with priors of 0.5 for all nodes except the two group’s dependence nodes, whose priors we vary independently of each other in intervals of 0.05 between 0.05 and 0.95. We measure the amount the agent updates their prior for the hypothesis after hearing all six pieces of testimony.

As Fig. 3A shows, updating is positive (pink) when agents think the group who says the hypothesis is false is more dependent than the group who says it is true (above the diagonal), whereas agents update negatively (turquoise) when they think the group who says the hypothesis is true are more dependent than the group who says it is false (below the diagonal). Therefore two agents on different sides of the diagonal confronted with this scenario would update in opposing directions, polarizing. We see a concrete example of this in Fig. 3B which contrasts the updating of two agents: one (pink) with a higher dependence prior for the group who say the hypothesis is false (0.8) than the group who says it is true (0.2), and so updates positively, and one (turquoise) with a higher dependence prior for the group who says the hypothesis is true (0.8) than the group who says it is false (0.2).

1.2. The present studies

Our model demonstrates that belief polarization could emerge from




Fig. 2. Results of the Bayesian Network simulations: more-dependent groups are less persuasive (Panel A) and groups who make less-plausible claims are seen as more dependent (Panel B).

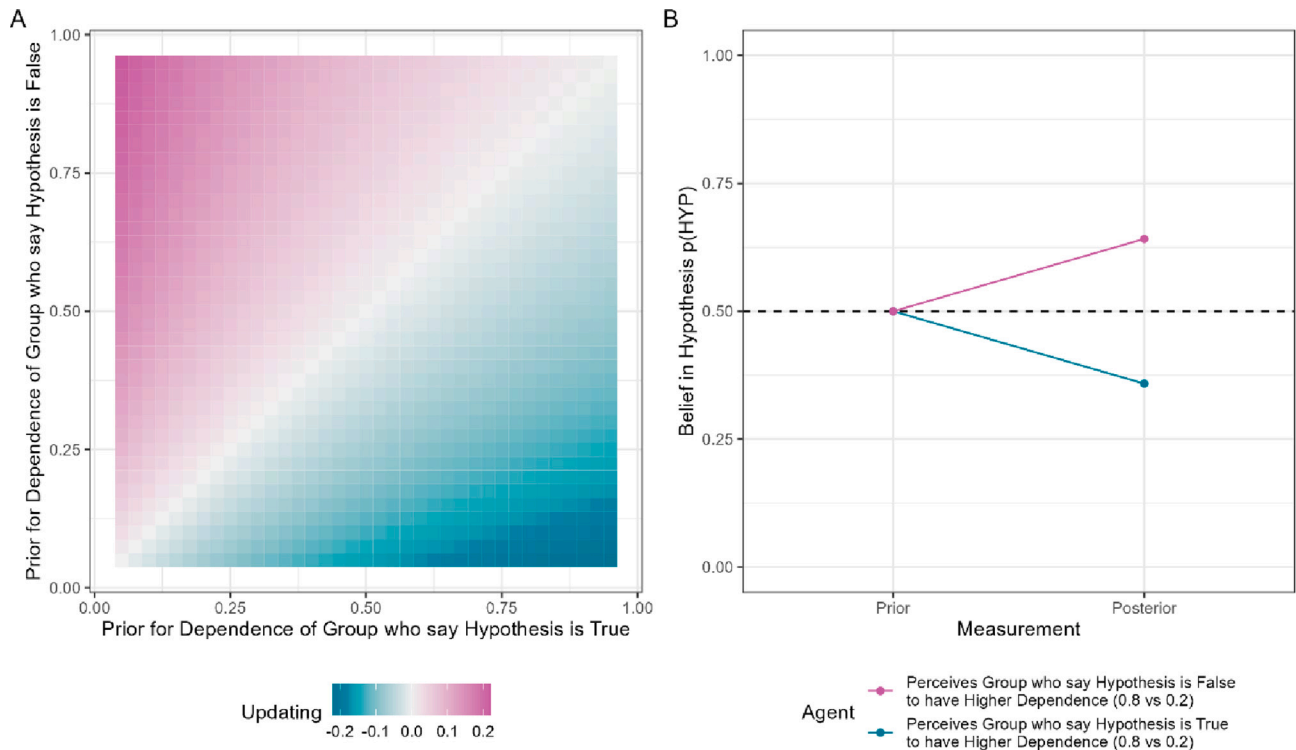


Fig. 3. Panel A: When exposed to conflicting testimony from two groups of sources, the direction of updating is towards the group deemed less dependent. Along the diagonal, both groups are seen as equally dependent, so there is no updating; above, the group who says the hypothesis is true are less dependent, so updating is positive; below, the group who says the hypothesis is false are less dependent, so updating is negative. Panel B: An example using just two agents, one who has a prior of 0.8 for the Dependence of the group who say the Hypothesis is true and 0.2 for the group who says it is false, and one with the opposite priors; these agents update in opposing directions.

Bayesian reasoning about source dependence when people are confronted with conflicting testimony and disagree about which group has a higher level of dependence. Importantly, the circumstances under which this polarization occurs are not circumstances under which the dominant explanation for belief polarization – biased assimilation (Baron, 2008; Lord & Taylor, 2009) – would predict it to occur. This is because biased assimilation requires that people have relatively strong prior beliefs about the hypothesis-at-stake, which motivate them to look for

reasons to reject incongruent evidence or testimony (Baron, 2008; Lord & Taylor, 2009). According to our Bayesian model, prior beliefs play no causal role – in fact, the simulation in Fig. 3a shows polarization occurring between agents who share the same neutral prior belief. This Bayesian model, then, predicts a new set of circumstances under which belief polarization may occur, which may be of relevance to real-world political polarization.

However, we lack empirical evidence about whether belief

polarization does occur under the circumstances predicted by the model, let alone whether this kind of polarization could contribute to real-world political polarization. In fact, there is mixed empirical evidence that people are even sensitive to source independence information in a normative way, meaning the predictions of our Bayesian model may well not be borne out in people's actual reasoning.

Several studies suggest people do down-weight testimony from higher-dependence groups. Madsen et al. (2020) finds that when expert biologists (Study 1) and economists (Study 2) were presented as having studied at the same school and subscribing to the same school of thought, people were less persuaded by their conclusions than when they were presented as independent. Similarly, across two studies which were both internally replicated, Mercier and Miton (2019) found that source groups who provided testimony about the quality of a restaurant were believed less when they were presented as having shared information or intentions. Pilditch et al. (2020) found evidence that intra-group transmission dependences were treated appropriately – 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 reports were written independently. Finally, across three experiments, Whalen et al. (2018) found evidence for both shared-factors and intra-group transmission dependences being treated appropriately.

However, participants sometimes deviate from what researchers consider normative regarding source dependence. Yousif et al. (2019) found participants were, in most scenarios, no less trusting of groups who formed a consensus after all consulting the same primary source than those who all consulted different primary sources. However, there is evidence this could occur just because participants doubted the independence of the primary sources (Connor Desai et al., 2022), or did not necessarily believe the group's members had based their opinions on the primary sources (Alister et al., 2022). Mercier and Miton (2019) twice found that the testimony of individuals who judged the quality of a restaurant was not discounted if they were presented as sharing the same culinary taste, and Xie and Hayes (2022) found mixed but mostly null results when a group's consensus testimony either was or was not influenced by one individual's testimony. Therefore, it is not safe to presume that people's reasoning about source dependence will lead to belief polarization under the predicted circumstances.

Before proceeding to test whether perceptions source dependence can underpin real-world political belief polarization then, we experimentally test whether differential perceptions of dependence can cause belief polarization when people are exposed to conflicting testimony from two source groups.

We do this in Study 1, finding evidence that differential independence perceptions *can* cause belief polarization. In Studies 2 and 3, we test, in Great Britain and the United States respectively, whether source dependence perceptions are associated with real-world political polarization. We test whether perceptions of the testimonial independence of real-world political groups show the expected pattern whereby partisans view their own group as more independent than their opponents, finding they do, and whether the extremity of this view is correlated with the extremity of people's political beliefs, finding it is.

All our data and code, including for two pilot studies for Study 1, and the code for the simulations above, is available via our OSF project: https://osf.io/ysqa8/?view_only=709c465855f04e509c6903852b8a6006, where pre-registrations for Study 1 and Study 3 are also available alongside print-outs of the Qualtrics surveys for all studies.

2. Study 1

Study 1 tested whether participants polarize when exposed to conflicting testimony from two source groups if they have conflicting priors for which group has a higher level of dependence. All aspects of the study were as pre-registered except that the study took 4 min rather than 5, and we recruited $N = 351$ rather than $N = 350$.

2.1. Methods

2.1.1. Participants

We recruited British participants from Prolific, and paid them £9/hour for 4 min, with 357 completing the study and six excluded (for failing attention checks – see Exclusions below), leaving a final sample of $N = 351$. $N = 350$ was targeted because it provides 80 % power to detect effects as small as $d = 0.15$ for both one-sample and paired-samples t -tests (see 2.3. Planned analyses), 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 our pilot studies, $d = 0.23$ (see Supplementary Materials Appendix A).

We inspected the demographic information participants had previously provided to Prolific. Of the 351 non-excluded 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'.

2.1.2. Design

Participants were presented with conflicting testimony from two groups of sources about whether a politician was guilty of a scandal. 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 – we therefore believe it is an ecologically valid domain for experimentally testing whether testimonial independence perceptions may influence polarization. One group of sources, the 'Innocent' group, said the politician was not guilty, whereas another group, the 'Guilty' group, said they were. Their claims were presented one at a time, with beliefs about the politician's guilt measured before and after each claim. One group was presented as having higher Dependence than the other, by virtue of them sharing more factors which influenced their testimony: their ideology, background, and their own information sources.

We have a 2×2 within-subjects factorial design, which we factorially manipulate across four trials: Greater Dependence ('Guilty Group' condition: the group who argue the politician is *guilty* is presented as being more Dependent than the group who argue the politician is *innocent*, vs 'Innocent Group' condition: the group who argue the politician is *innocent* is presented as being more Dependent than the group who argue the politician is *guilty*) \times Presentation Order ('Guilty First' condition: the group who argue the politician is guilty provide their testimony before the group who argue the politician is innocent, vs 'Innocent First' condition: the group who argue the politician is innocent provide their testimony before the group who argue the politician is guilty).

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 both the effect of the group's Dependence (High vs Low) on the persuasiveness of each claim ('Testimony-Wise Persuasion'), and the effect of which group has greater dependence on the overall direction of people's updating for each trial ('Trial-Wise Updating'). We expected that people should be less persuaded by more-dependent source groups, and that this would lead to participants shifting their opinion towards believing the politician is innocent when the Guilty group is more dependent, but shifting their opinion towards believing the politician is guilty when the Innocent group is more dependent.

2.1.3. 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).

2.1.3.1. Trials. Fig. 4 shows what happened in each trial.

Each trial was presented as a short report about the debate concerning the allegation of a politician's involvement in a scandal. Each trial began 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." For each participant, the scandal was chosen randomly without replacement for a given trial.

After the preamble, we measured prior beliefs 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 preceded 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 were shown two paragraphs. In each paragraph, the testimony of one of the groups was given, followed by a description of their Dependence. The testimony was of the form: "Three of the commentators said they thought the politician was {innocent/guilty}". Then the Dependence of the group was 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 was intended to engender differential perceptions of the extent to which the sources possess shared factors which influence their testimony.

At the end of each paragraph, we asked participants to estimate their belief in the guilt of the politician using the same question as for the prior (but without the qualification). This question was presented immediately below the corresponding paragraph on the same page. Participants had to click forward after giving their estimate to view the second paragraph, and could not navigate backwards after having done so. Beneath the post-testimony belief measurement for the second group, an attention check was given. Attention checks just asked for a particular number to be given on the sliding scale.

2.1.4. Measures

2.1.4.1. 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".

2.1.4.2. 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.

2.2. 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).

2.3. 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 Dependence condition. We expect trial-wise updating to be positive when the Innocent group have greater dependence, and negative when the Guilty group have greater dependence. It should be noted that since *both* significance tests must be passed to affirm this hypothesis, there is no need to correct the alpha level – see Rubin (2021).

Two pilot studies ($Ns = 40, 39$), informed our design and provided preliminary evidence mostly consistent with our hypotheses. For H1, there was sufficient evidence to reject the null in both, and supporting qualitative evidence. For H2, there was sufficient evidence to reject the null in one, and the other had a small non-significant effect in the expected direction. Our main study emulates the design of the pilot where the effects were larger. See Supplementary Materials Appendix A for further details.

2.4. Results

2.4.1. Analyses of H1

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 Fig. 5.

2.4.2. Analyses of H2

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 Fig. 6, with means and confidence intervals.

2.5. Discussion

We were able to reject the null for both hypotheses – 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. This provides evidence that belief polarization can occur under the circumstances predicted by the Bayesian model presented in this paper, as findings were in line with directional predictions.

It should be noted that the results obtained here are mostly consistent with results obtained from two pilot studies in both direction and magnitude, which strengthens the robustness of the results. Further

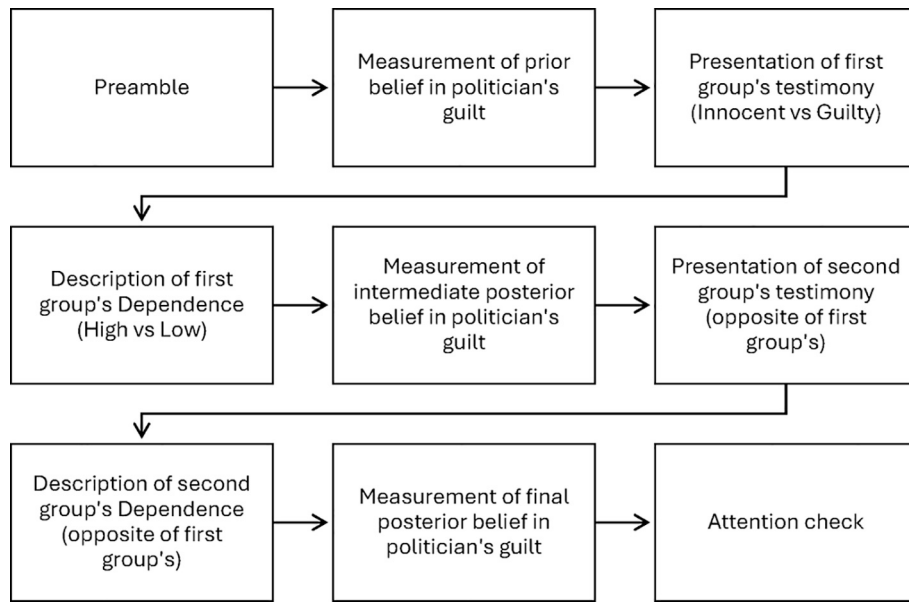


Fig. 4. Schematic of the events in each trial.




Fig. 5. The distribution of Testimony-wise Persuasion scores by Dependency condition, with means and 95 % confidence intervals.

details about the pilots are given in Supplementary Materials Appendix A, but they had very similar designs to the main study, with *N*s of 39 and 40. Both allowed us to reject the null for H1, with effect sizes of $d = 0.66$ [0.31, 1.00] and $d = 0.41$ [0.08, 0.74]. Of the four tests of H2, all obtained effects in the expected direction, with three of these effects significantly different from 0, and one not.

One possible critique of this evidence is that the results could be dependent on a form of contrastive learning arising from our repeated-measures design. That is, perhaps participants only updated more strongly to testimony from low-dependence groups when they encountered it in the context of having previously encountered a high-




Fig. 6. The distribution of Trial-wise Updating scores by Dependency condition, with means and 95 % confidence intervals.

dependence group, and vice versa, rather than recognizing the importance of the dependence information outright. Of course, a repeated-measures design is advantageous because of the greater statistical power it affords, but this is a limitation worth considering.

As it turns out, we can reject the null for both hypotheses when we restrict the analyzed data to rule out any effect of contrastive learning. To reassess H1 in this way, we analyzed participants' Testimony-wise Persuasion for the very first piece of testimony they encountered. Testimony-wise Persuasion was, as expected, lower when the testimony came from the High Dependence group ($M = 1.34$, $SD = 11.4$) than the Low Dependence group ($M = 13.2$, $SD = 15.8$), which was a significant

difference, independent-samples $t(349) = 8.142, p < .001$, with a large effect size, $d = 0.87$ [0.65, 1.09].

To reassess H2, we analyzed participants' Trial-wise Updating for the very first trial they encountered. As expected, updating was positive when the 'Innocent' group had greater dependence ($M = 4.19, SD = 16.6$), with a significant difference from 0, one-sample $t(165) = 3.25, p = .001$, and a small effect size, $d = 0.25$ [0.10, 0.41], whereas updating was negative when the 'Guilty' group had greater dependence ($M = -10.8, SD = 17.9$), with a significant difference from 0, one-sample $t(184) = 8.196, p < .001$, and a moderate effect size, $d = 0.60$ [0.45, 0.76].

Given that all these effects were in the expected direction, and their effect sizes were either greater than or very similar to the effect sizes found in the primary analyses, this suggests contrastive learning need not be a meaningful concern – our participants recognized the importance of the dependence information and responded appropriately at the first time of asking.

3. Study 2

Study 1 shows that differences of opinion about which of two conflicting source groups has higher testimonial independence can lay the groundwork for belief polarization. In Study 2, we investigate whether the pre-conditions are in place for such differences of opinion to cause real-world belief polarization in Great Britain, by examining whether supporters of the Labour Party and supporters of the Conservative Party each perceive their own group to have higher testimonial independence than the other.

We also test a further prediction of the theory, supported by our Bayesian Network modelling, that there is a reciprocal relationship between a person's perceptions of source groups' testimonial independence and their political beliefs. If true, we should see that the more a person's beliefs lean towards one end of a given political axis (e.g., to the left), the more they should perceive groups on the same side to have higher testimonial independence than groups on the opposing side. We therefore also examine the relationship between relative perceptions of in-group vs out-group testimonial independence and the extremity of political beliefs for supporters of the Labour Party and supporters of the Conservative Party.

We further test whether any such association is robust to controlling for two plausible confounders: affective polarization and party identity. Since it is somewhat disparaging to describe a group as having lower testimonial independence, believing a rival group to have lower testimonial independence than your own group could just be a reflection of affective polarization, i.e. disliking a political outgroup relative to your own group. Affective polarization is known to encompass disparaging trait ratings (Druckman & Levendusky, 2019), and is likely to correlate with belief positions as people typically dislike others more when they disagree on policy issues (Orr et al., 2023). Similarly, it is widely argued that disliking of opposing political groups emerges from strongly identifying with one's own party, which triggers automatic in-group biases that cause partisans to show in-group love and out-group hate (e.g., Iyengar et al., 2019; Mason, 2018) and, furthermore, that partisanship directly influences people's policy positions because people want to adopt positions which align them with their political in-group (e.g., van Bavel & Pereira, 2018). Thus both partisan identification and affective polarization could create confounding paths linking dependence perceptions to political beliefs in the absence of any direct relationship between the two. Accordingly, we control for both variables in our analyses. No element of Study 2 was pre-registered.

3.1. Methods

3.1.1. Participants

We analyze data from the $N = 507$ participants who completed Wave 7 of the MHP UK Polarization Tracker (a biannual survey which surveys

polarization-related beliefs and attitudes in the UK) who identified as a supporter of either the Labour Party ($n = 322$) or Conservative Party ($n = 185$). The survey took place between 7th March and 4th April 2024, and was quota-sampled to be near-representative of the Great British population in terms of nationality (English, Scottish, Welsh), gender, ethnicity, age, degree status, 2019 General Election vote and 2016 EU Referendum vote. This wave of the MHP UK Polarization Tracker had a total sample size of $N = 1018$, which excludes 47 participants who completed the survey but failed at least one of three attention checks. Attention checks were embedded within batteries of questions which used Likert scales throughout the survey, and asked participants to choose a particular response using the same scale as the questions immediately adjacent to it.

Demographically, our participants contained 277 females, 228 males, and 2 trans or non-binary people. To describe their ethnicity, 446 participants identify as "White", 25 as "Asian", 20 as "Mixed", 13 as "Black", and 3 as "Other". The median age was 47 with a range of 18–79. The number of participants who had completed the preceding wave of the MHP UK Polarization Tracker, which took place between 26th July and 6th August 2023, was 405. These returning participants were not asked to provide demographic data as it had been collected in the preceding wave, so their Gender, Age, and Ethnicity were presumed not to have changed – while this creates some imprecision, exact demographic information is not required for any of our analyses. Participants were informed the study had received ethical approval from the <Department of Psychology Research Ethics Committee at the University of Cambridge>.

3.1.2. Measures

3.1.2.1. Perceived Testimonial Independence. To measure perceptions of Testimonial Independence, we used a novel eight-item battery, the Testimonial Independence Perceptions Scale ("TIPS"). Each item consisted of a pair of opposing statements about the level of a target group's independence with respect to a particular factor which could affect their testimony, one statement corresponding to a perception of low independence for that factor and the other to a perception of high independence. Participants responded using a seven-point numeric scale, where 1 corresponded to the low independence perception and 7 to the high independence perception. The items were as shown in Table 3:

Participants completed the TIPS twice, once with "People who support the Labour Party" as the target group, and then again with "People who support the Conservative Party" as the target group. To obtain each participant's overall perception of each group's testimonial independence, we averaged the responses they gave for each group respectively.

The items were sufficiently reliable for aggregation into a single scale, demonstrating a 'good' Cronbach's alpha of 0.88, and an 'acceptable' McDonald's Omega-hierarchical of 0.78.¹ The mean inter-item correlation was 0.47, range 0.28–0.69.

3.1.2.2. Testimonial Independence Differential. We calculate a 'Testimonial Independence Differential' by subtracting people's rating of Labour supporters from their rating of Conservative supporters. This measure therefore indicates which group they perceive to have higher testimonial independence, and by what margin, scoring positively when Conservatives are perceived to have higher Testimonial Independence and negatively when Labour supporters are perceived to have higher

¹ McDonald's omega-hierarchical is a suitable measure of scale reliability when, as in this case, the scale may be hierarchically structured, with multiple dimensions which load onto a single underlying factor whose measurement is sought (Flora, 2020). It measures the proportion of variance in total scores attributable to this single underlying factor, in this case 78%. We use the function `omega()` from the package {psych} (Revelle, 2023) to calculate omega-hierarchical.

Table 3
Items in the Testimonial Independence Perceptions Scale (TIPS).

Item	Low Independence (1)		High Independence (7)
Information	They all get their information about politics from the same source	vs	They all get their information about politics from different sources
Experiences	They have all been through the same kind of life experiences	vs	They have all had different kinds of life experiences
Background	They all come from the same background	vs	They all come from different backgrounds
School of Thought	They all subscribe to the same school of thought on politics	vs	They all subscribe to different schools of thought on politics
Opinions	They all have the same opinions on political issues	vs	They all have different opinions on political issues
Goals	They all have the same goals when it comes to politics	vs	They all have different goals when it comes to politics
Opinion Transmission	They form their political opinions by picking them up from each other	vs	They form their opinions about politics independently of each other
Belief Adoption	If other people who support the same party believe something, they will just adopt that belief	vs	They are not influenced by what other people who support their party believe

Testimonial Independence.

3.1.2.3. *Party ID and strength.* Participants were asked the following to establish their party ID: “Which of the following political parties do you feel you identify with the most? Even if you don’t identify with any party very strongly, try to pick one that you identify with a little more than the others. If you really don’t feel that you identify with any of these parties, please pick ‘None’. If you identify with a party that isn’t on this list, please pick ‘Other’ and write in their name.” with the provided options: “The Labour Party/Scottish Labour Party”, “The Conservative Party/Scottish Conservative Party”, “The Liberal Democrats/Scottish Liberal Democrats”, “The Green Party”, “The Scottish National Party (SNP)”, “Plaid Cymru”, “Alba”, “Reform UK”, “None”, and “Other”.

For those who selected any option other than “None” for Party ID, they then completed the four-item version of Huddy and Bankert’s (2017) Partisan Identity Scale for that party. This scale is a measure of partisan identity strength that has been empirically validated in the UK (Huddy & Bankert, 2017). Participants responded to questions about their degree of affiliation with the party using a four-point Likert scale, which we convert to a linear numeric 1–4 scale; we then average across items to obtain the overall ‘Strength’ score. To reiterate, only those who identified with either the Labour or Conservative parties were analyzed in this study. The mean Strength was $M = 2.13$ ($SD = 0.57$).

3.1.2.4. *Political Belief Position.* To measure participants’ political beliefs, we analyze responses to a set of eight statements about important issues and debates in contemporary British politics concerning policy and government performance. The statements are shown in Table 4. Each item consists of a statement with which participants indicate their level of agreement or disagreement using a seven-point Likert scale. We convert responses to a linear numeric 1–7 scale, reverse-code item 6, then find the mean to give a score we call ‘Belief Position’. After reverse-scoring, these items have an ‘acceptable’ McDonald’s omega-hierarchical of 0.70,² a ‘good’ Cronbach’s alpha of 0.86, and a mean inter-item correlation of 0.51 (range 0.32–0.81). High Political Belief Position scores indicate right-wing, pro-Conservative Party beliefs whereas low Political Belief Position scores indicate left-wing, anti-Conservative Party beliefs (as this study was conducted several months before Labour’s General Election, while the Conservative Party were still in power). As Table 4 shows, Labour and Conservative supporters had different mean positions on every item, with each difference being in the expected direction with statistical significance (independent-samples *t*-tests, all $ps < .001$), with a large effect size (all Cohen’s $ds \geq 0.88$). The mean position for Conservative supporters was $M = 4.08$ ($SD = 0.78$) and for Labour supporters $M = 2.43$ (0.81), which is a significant difference, $t(505) = 22.457$, $p < .001$, with a large effect size, Cohen’s $d =$

² Originally we planned to use another item from the same section of the survey as our eight items, “The government should be doing more to address the issue of climate change”, however including this item reduced the McDonald’s omega-hierarchical to 0.50, so we excluded it.

Table 4
Political Belief Position items – Study 2.

Item	Statement	Mean position M (SD)		Cohen’s d
		Labour supporters	Conservative supporters	
<i>To what extent do you agree or disagree with this statement? (Strongly disagree = 1, Strongly agree = 7)</i>				
1	Britain is heading in the right direction as a country	2.08 (1.09)	3.76 (1.40)	1.42 [1.22, 1.62]
2	The government is handling the problems that Britain faces well	1.85 (0.97)	3.64 (1.29)	1.63 [1.42, 1.84]
3	Britain is a force for good in the world	3.60 (1.50)	4.84 (1.22)	0.88 [0.69, 1.07]
4	The British economy is doing well	1.96 (0.97)	3.60 (1.28)	1.49 [1.29, 1.70]
5	The UK made the right decision to leave the European Union	2.23 (1.60)	4.69 (1.92)	1.42 [1.22, 1.62]
6	More should be done to address racial inequalities in the UK (R)	5.24 (1.45)	3.83 (1.44)	0.98 [0.79, 1.17]
7	The government is doing a good job on the issue of healthcare in the UK	1.70 (1.03)	3.16 (1.28)	1.30 [1.10, 1.49]
8	Overall, immigration does more harm than good to this country	3.26 (1.78)	4.74 (1.50)	0.88 [0.69, 1.07]

2.07 [1.85, 2.29].

3.1.2.5. *Affective polarization.* Participants rated how much they like ‘Supporters of the Labour Party’ and ‘Supporters of the Conservative Party’ on a 0–10 scale. To obtain a directional measure of Affective Polarization, we subtract the rating they gave to Labour supporters from the rating they gave to Conservative supporters. This measure, therefore, like the Testimonial Independence Differential and Political Belief Position variables, scores higher when participants have a pro-Conservative/right-wing attitude, and lower when they have a pro-Labour/left-wing attitude. The mean affective polarization score for Conservatives was $M = 2.24$ ($SD = 2.31$) and $M = -3.68$ ($SD = 2.72$) for Labour supporters.

3.1.3. *Procedure*

Participants in Wave 7 of the MHP UK Polarization Tracker completed a 30-minute online survey. It was advertised on Prolific as taking 30 min and paying £3.90, and titled “Views on British politics and society” with the description “In this study we will ask you to share your views on British politics and society.” To participants from the previous wave, the study was additionally advertised as being the latest wave of a longitudinal study they had previously completed.

At the beginning of the survey, participants read task instructions

and provided consent. The survey asked many questions about beliefs and feelings about contemporary British politics, policies, parties, and politicians, as well as containing demographic questions, questions about British society more generally, and several psychometric scales. The Party ID and Strength questions were in the same block, which was presented early on after a few blocks containing questions about demographics and political identities. Shortly afterwards, separated by questions about Vote Intention and ‘Feeling Thermometer’ like ratings of parties, politicians, and political groups, a block containing the nine political belief items was presented. The TIPS blocks for each target group were situated in a section of the survey which formed its latter two-thirds, in which 24 blocks were presented in a random order, and which began immediately after the political belief items block. The TIPS blocks for each group were therefore presented in a random order, and not necessarily consecutively. After completing the survey, participants were debriefed, thanked, and paid.

3.2. Hypotheses

We tested the following hypotheses: **H1**: Participants will perceive supporters of their own party to have greater testimonial independence than supporters of the out-party. **H2**: Relative perceptions of partisans’ testimonial independence will explain variance in political beliefs beyond partisanship and affective polarization.

3.3. Results

Fig. 7 shows the mean ratings each participant group gave to both target groups, both for each item in the TIPS individually, and overall.

3.3.1. Planned analyses

3.3.1.1. Hypothesis 1. People rated their in-group as having higher testimonial independence, $M = 4.35$ [4.28, 4.43], than their out-group, $M = 3.30$ [3.22, 3.39]. A paired-samples t -test revealed this to be a significant difference with a large effect size, $t(506) = 19.614$, $p < .001$, $d = 0.87$ [0.77, 0.97].

3.3.1.2. Hypothesis 2. We performed an OLS regression analysis with Political Belief Position as the dependent variable, and the following as predictor variables: Testimonial Independence Differential, Party ID, Strength, the Party ID: Strength interaction, and Affective Polarization. The Testimonial Independence Differential was a significant, positive predictor $b = 0.134$ ($se = 0.029$), $t(501) = 4.634$, $p < .001$.

3.3.2. Exploratory analyses

We conducted exploratory analyses to test H1 and H2 for those in each party separately.

3.3.2.1. Hypothesis 1. Labour Party supporters rated other Labour Party supporters as more independent, $M = 4.36$ [4.26, 4.46], than Conservative Party supporters, $M = 3.11$ [3.01, 3.22]. A paired-samples t -test found this was a significant difference with a large effect size, $t(321) = 18.174$, $p < .001$, $d = 1.10$ [0.88, 1.15]. Similarly, Conservative Party supporters rated other Conservative Party supporters as more independent, $M = 4.33$ [95 % CI: 4.21, 4.45], than Labour Party supporters, $M = 3.63$ [95 % CI: 3.50, 3.76]. A paired-samples t -test found this was also a significant difference, with a moderate effect size, $t(184) = 8.915$, $p < .001$, $d = 0.66$ [0.50, 0.81].

3.3.2.2. Hypothesis 2. We performed the same regressions as in the planned analysis for Hypothesis 2 within each party-specific sub-sample separately, but with the omission of Party ID and the Party ID: Strength interaction terms as these become redundant. Among Labour supporters, the Testimonial Independence Differential was a significant positive predictor, $b = 0.118$ ($se = 0.35$), $t(318) = 3.383$, $p < .001$, as it was among Conservative supporters, $b = 0.186$ ($se = 0.052$), $t(181) = 3.567$, $p < .001$.

3.4. Discussion

We examined real-world perceptions of testimonial independence among British supporters of the Labour Party and the Conservative Party. Consistent with H1, we found people perceived their political in-group to possess greater testimonial independence than their out-group, with a large overall effect size of $d = 0.87$. As demonstrated in Study 1,

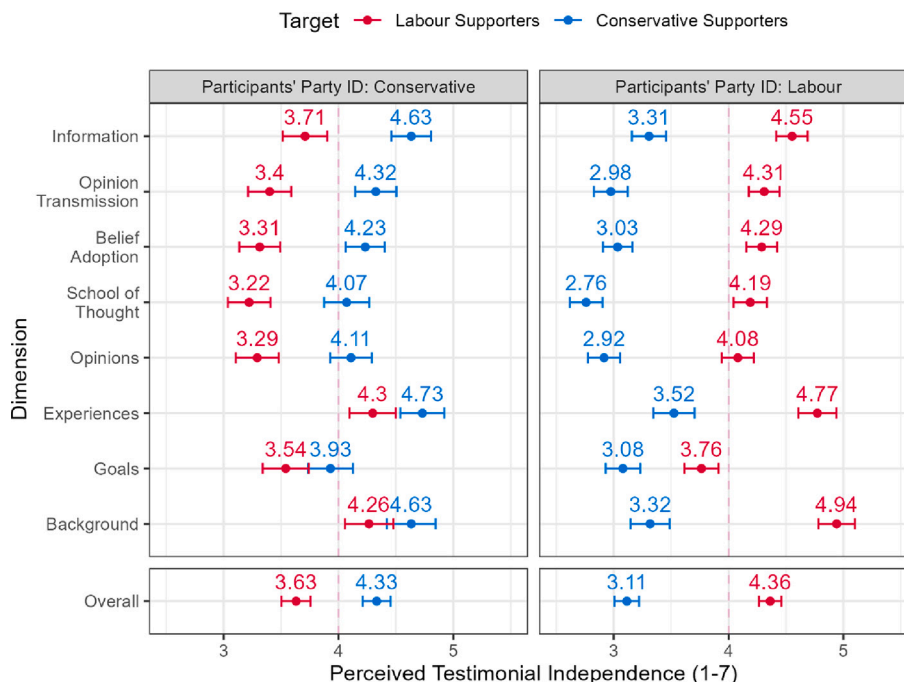


Fig. 7. Perceptions of testimonial independence by the group participants identify with, and the group they are making judgments about (the “target group”).

such differential beliefs may contribute to polarization, and accordingly, consistent with H2, we found that those who perceived there to be a greater difference in terms of testimonial independence in favor of their party's supporters also tended to have more extreme political beliefs in the direction of their party. This result was robust to controlling for their partisanship and affective polarization.

These results show that the pre-conditions are in place for perceptions of relative testimonial independence to drive political belief polarization in the UK. Moreover, they show that more-polarized people do indeed have more-polarized testimonial independence perceptions. One limitation of Study 2 is that while we had very straightforward model-derived hypotheses consistent with the results of Study 1, we did not preregister them. Consequently, in Study 3, we attempt a pre-registered replication of these results in the US.

4. Study 3

The study's design, analyses, and hypotheses were pre-registered in our OSF project. All details were as pre-registered except where noted.

4.1. Methods

4.1.1. Participants

We analyze data from $N = 300$ Americans who identified as a supporter of either the Democratic Party ("Democrats", $n = 150$) or Republican Party ("Republicans", $n = 150$). The survey took place from 22 to 24th May 2024 via the online survey participation platform Prolific. The survey was targeted to participants with particular characteristics based on Prolific's pre-screening data: they had to have previously answered "United States" when asked "What is your nationality?", and either "Democrat" or "Republican" when asked "In general, what is your political affiliation?". These questions were asked again at the beginning of the survey and participants had to give the same answer to their on-file information to participate further. A total of 310 people passed these criteria and completed the survey, but 9 were excluded for failing at least one attention check (see below), which was a pre-registered exclusion criterion, and 1 was excluded for completing the study twice, which had not been pre-registered as an exclusion criterion. When they first took the study, this participant was automatically rejected for saying they affiliated with a party different to that which they had previously told Prolific, but then took the study again giving a different affiliation; we therefore do not regard their data as trustworthy.

Demographically, the sample consisted of 165 women, 131 men, 3 "Non-binary/third gender" people, and 1 who preferred to self-describe as "Transgender Man". 241 identified as "White", 38 as "Black or African American", 11 as "Asian", 1 as "American Indian or Alaska Native", 4 as "Other", and 5 chose more than one category; to describe their 'Ethnicity', 281 chose "Not Hispanic or Latino" and 19 chose "Hispanic or Latino". The median age was 43 with a range of 18–79. Participants were informed the study had received ethical approval from <REDACTED FOR BLIND REVIEW>.

4.1.2. Measures

4.1.2.1. Perceived Testimonial Independence. Participants completed the TIPS questions from Study 2, once with the target group 'Supporters of the Democratic Party' and once with the target group 'Supporters of the Republican Party'. We found the scale had a McDonald's omega-hierarchical of 0.78 (the same as in Study 2), indicating it could be construed to measure a single latent construct, with a mean inter-item correlation of 0.50 (range 0.32–0.70). To obtain an overall measure of 'Perceived Testimonial Independence' for each target group for each participant, we averaged the scores they gave to that group.

4.1.2.2. Testimonial Independence Differential. To construct the 'Testimonial Independence Differential' measure for each participant, we subtract their 'Perceived Testimonial Independence' score for 'Supporters of the Democratic Party' from their score for 'Supporters of the Republican Party'.

4.1.2.3. Affective polarization. Participants rated how much they liked 'Supporters of the Democratic Party' and 'Supporters of the Republican Party' on a 0–10 scale. To obtain a directional measure of Affective Polarization, we subtracted the rating they gave to Democrats from the rating they gave to Republicans. This measure, therefore, like the Testimonial Independence Differential, scores positively when participants have a pro-Republican attitude, and negatively when they have a pro-Democrat attitude. The mean affective polarization score was $M = 3.74$ ($SD = 3.10$) for Republicans and $M = -4.47$ ($SD = 2.84$) for Democrats.

4.1.2.4. Party identity and strength. As in Study 2, participants completed the four-item version of Huddy and Bankert's (2017) Partisan Identity Scale for the party they support. Participants responded using a four-point Likert scale, which we converted to a linear 1–4 numeric scale; we then obtained an overall measure of partisan identification strength ("Strength") by averaging their responses across the four questions. The mean Strength was $M = 2.60$ ($SD = 0.64$).

4.1.2.5. Political Belief Position. Participants indicated their position towards ten political belief statements. All statements concerned a different political issue relating to policy matters, the performance of the current US government, or Donald Trump. Participants responded using a seven-point Likert scale (Strongly/Moderately/Slightly Agree/Disagree, Neither agree nor disagree), which was converted to a linear

Table 5
Political Belief Position items.

Item	Statement	Mean position M (SD)		Cohen's d
		Democrats	Republicans	
<i>To what extent do you agree or disagree with this statement? (Strongly disagree = 1, Strongly agree = 7)</i>				
1	The US economy is doing well.	3.99 (1.91)	2.42 (1.76)	0.85 [0.62, 1.09]
2	The Government is handling the problems that the US faces well.	3.56 (1.71)	2.23 (1.60)	0.80 [0.56, 1.04]
3	The US needs stricter gun control laws.	6.49 (1.12)	3.29 (2.16)	1.86 [1.59, 2.13]
4	The US needs stricter laws to limit abortion.	1.51 (1.25)	4.62 (2.06)	1.83 [1.56, 2.09]
5	The US owes it to African Americans to take special measures to improve their position in society.	5.49 (1.44)	2.85 (1.98)	1.53 [1.27, 1.78]
6	Overall, immigration does more harm than good to the US as a country.	2.03 (1.45)	4.79 (1.90)	1.64 [1.37, 1.90]
7	The US should do all it can to reduce its carbon emissions.	6.39 (0.88)	4.14 (1.92)	1.51 [1.25, 1.76]
8	The country works best when tax is low and so is government spending.	3.75 (1.80)	5.69 (1.34)	1.22 [0.98, 1.47]
9	Joe Biden was the legitimate winner of the 2020 US Presidential Election.	6.83 (0.67)	3.69 (2.20)	1.93 [1.65, 2.20]
10	The many legal cases against Donald Trump are legitimate and have nothing to do with political persecution.	6.29 (1.32)	2.57 (1.91)	2.26 [1.97, 2.55]

1–7 scale. The items were as shown in Table 5, which also shows the mean position taken by Democrats and Republicans, and the effect size of the difference between their positions.

Every item created division between Republicans and Democrats in the expected direction with a significant difference (all $ps < .001$) and with a large effect size (all $ds \geq 0.8$). The position items could be aggregated to measure a single underlying construct, with a McDonald's omega-hierarchical of 0.82, thereby satisfying the pre-registered criterion of omega hierarchical being at least 0.70. The mean inter-item correlation was 0.49 (range 0.21–0.73).

To calculate the overall 'Political Belief Position' measure for each participant, we reverse-coded items 1, 2, 3, 5, 7, 9, and 10, then averaged responses to all items. Thus, higher scores indicate a more right-wing/pro-Republican/anti-Democrat position. Accordingly, the mean Political Belief Position of Republicans ($M = 4.99$, $SD = 1.08$) was higher than that of Democrats ($M = 2.43$, $SD = 0.69$) with a large effect size, Cohen's $d = 2.83$ [2.51, 3.15].

4.1.2.6. Attention checks. Two attention check questions were included, one embedded in the partisan identification strength block, and one in the political beliefs block. The questions had the same response scale as the other questions in the block, but asked participants to give a particular response ("Agree" for one and "Slightly agree" for the other). Participants had to pass *both* to avoid exclusion; nine failed at least one and were excluded.

4.1.3. Procedure

The study was advertised on Prolific to take 8 min and pay the equivalent of £1.07 in US dollars. After reading task instructions and providing consent, participants completed a block of demographic questions, which included party affiliation. They then completed the two blocks of testimonial independence perceptions questions, one block per target group ("Supporters of the Democratic party" and "Supporters of the Republican party"), with the two blocks presented in a random order and the order of questions within each block randomized. They then completed the affective polarization block, the partisan identity strength block, and the political beliefs block in a random order, with the order of questions within each block also randomized. Within all blocks, all questions were simultaneously visible on the same page. After completing all blocks, participants were debriefed and thanked for taking part, and asked to click a button to return to Prolific.

4.2. Hypotheses

We pre-registered the hypotheses with associated analysis plans and expectations shown in Table 6.

4.3. Results

Fig. 8 shows the mean rating each group gave to each target group on every item in the scale and overall.

4.3.1. Pre-registered analyses

4.3.1.1. Hypothesis 1. Participants rated their own party's supporters as having higher testimonial independence, $M = 4.52$, $SD = 0.97$, than the out-party's supporters, $M = 3.20$, $SD = 1.17$. A paired-samples t -test found this was a significant difference, $t(299) = 14.141$, $p < .001$, with a large effect size $d = 0.82$ [0.69, 0.95].

4.3.1.2. Hypothesis 2. An OLS regression found that Political Belief Position was predicted by the Testimonial Independence Differential, $b = 0.180$ ($se = 0.032$), $t(294) = 5.713$, $p < .001$, after controlling for Party ID, Strength, the Party ID: Strength interaction, and Affective Polarization.

Table 6

Pre-registered hypotheses, planned analyses, and expectations for Study 3.

Hypothesis	Planned analysis	Expectation
H1: Participants will perceive supporters of their own party to have greater testimonial independence than supporters of the out-party.	A paired-samples t -test of Own Party vs Out-Party Perceived Testimonial Independence scores, using all 300 participants.	Own Party scores should be higher than Out-Party scores, with a significant difference.
H2: Relative perceptions of partisans' testimonial independence will explain variance in political beliefs beyond partisanship and affective polarization.	An OLS multiple regression with Political Belief Position as the dependent variable, and independent variables of Testimonial Independence Differential, Affective Polarization, Party ID, Strength and a Party ID: Strength interaction, using all 300 participants.	Testimonial Independence Differential should be a positive, significant predictor.

4.3.2. Exploratory analyses

We conducted exploratory analyses of H1 and H2 within samples of a) just the Democrats and b) just the Republicans.

4.3.2.1. Hypothesis 1. Democrats perceived fellow Democrats to have higher testimonial independence, $M = 4.76$, $SD = 0.87$, than Republicans, $M = 2.92$, $SD = 1.08$, which was a significant difference, paired-samples $t(149) = 15.838$, $p < .001$, with a large effect size, $d = 1.29$ [1.07, 1.51]. Republicans perceived fellow Republicans to have higher testimonial independence $M = 4.29$, $SD = 1.00$, than Democrats, $M = 3.47$, $SD = 1.19$, which was a significant difference, paired-samples $t(149) = 6.015$, $p < .001$, with a small effect size, $d = 0.49$ [0.32, 0.66].

4.3.2.2. Hypothesis 2. An OLS regression among Democrats found that Political Belief Position was predicted by the Testimonial Independence Differential, $b = 0.138$ ($se = 0.038$), $t(146) = 3.623$, $p < .001$, while controlling for Strength and Affective Polarization. Similarly, it was a significant predictor when the same analysis was applied to Republicans, $b = 0.211$ ($se = 0.048$), $t(146) = 4.442$, $p < .001$.

These results affirm all our hypotheses and replicate the effects of Study 2, but extend the documented phenomena to an important new context: the US.

5. General discussion

In this paper, we set out to study the role of perceptions of testimonial independence as a possible contributing factor to political belief polarization. We took a multi-pronged approach combining modelling, experiments, and surveys, and we make several contributions to the literature. We developed a novel Bayesian Network model of how people can infer the dependence of source groups, and integrate those perceptions when learning from groups' testimony. The model's simulated behavior suggested the possibility of a polarization cycle mediated by perceptions of testimonial independence, constituting a new theory of mass belief polarization, from which we derived our hypotheses. We then conducted a pre-registered experiment (Study 1), finding that people are less trusting of groups presented as having higher dependence, and polarize if exposed to conflicting testimony from groups when they possess different beliefs about which group has a higher level of dependence. This represents evidence for a novel polarization mechanism that is consistent with Bayesian reasoning, and which may have relevance for real-world politics. Then, using a newly-developed scale for measuring perceptions of testimonial independence whose basic properties we were able to validate (the TIPS), we found that, in

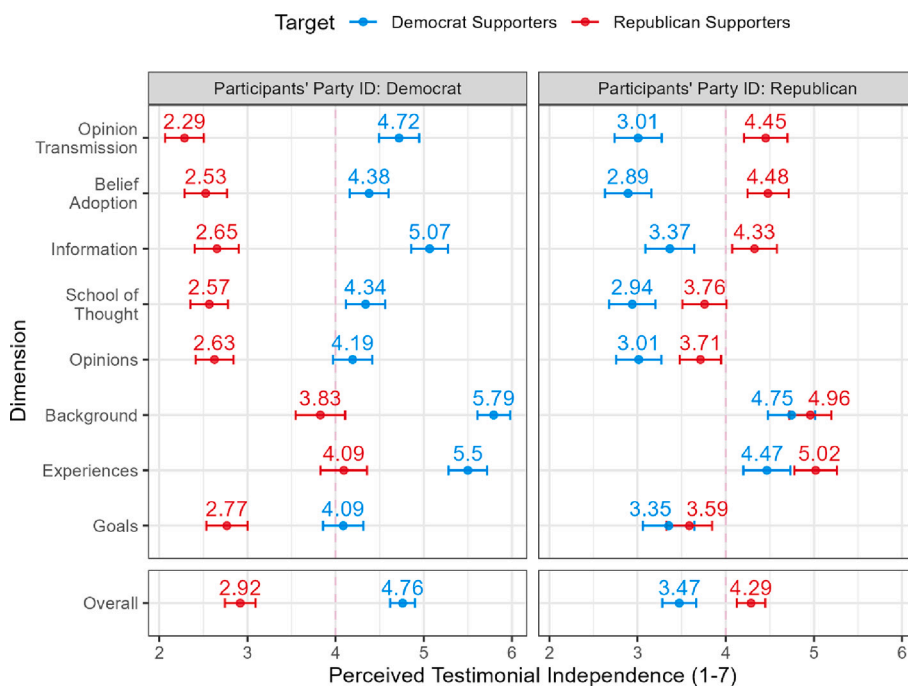


Fig. 8. Mean Testimonial Independence Perceptions by the group participants identify with, and the group they are making judgments about (the “target group”).

the UK (Study 2) and in a pre-registered replication study in the US (Study 3), real-world partisans perceive their own party’s supporters to have higher testimonial independence than their opponents, with large overall effect sizes of $d = 0.87$ and $d = 0.82$ respectively. This shows that the conditions are in place for our novel polarization mechanism to occur in the real world. Both Study 2 and Study 3 found evidence consistent with our hypothesized polarization cycle, as those with more polarized independence perceptions also had more polarized political beliefs.

Our results strengthen the case 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), both theoretically and empirically. While numerous explanations of how belief polarization might be caused by Bayesian reasoning have been proposed in prior literature, empirical tests of whether polarization actually occurs under the circumstances predicted by these accounts is very limited. Two exceptions are Jern et al. (2014), and Cook and Lewandowsky (2016). However, Jern et al. (2014) and Cook and Lewandowsky (2016) provide empirical evidence for models of Bayesian polarization that are relatively narrow in focus. Jern et al.’s (2014) model works for a complex fabricated scenario involving medical testing, whereas Cook and Lewandowsky’s (2016) model is of how belief in anthropogenic global warming is affected by knowledge of the 97 % consensus among climate scientists that it is real. The model of polarization presented here, however, is broadly generalizable across many political debates, as it is common for citizens to encounter conflicting testimony from two source groups about political issues.

However, one limitation of this work is that we do not prove that people reason in a Bayesian way at the procedural level in Study 1. We do not make any attempt to use our model to make individual-level predictions about how much people should update, and calculate fit statistics, or compare whether our Bayesian model outperforms alternatives, like strategies reliant on heuristics or boundedly rational models. We must therefore remain open to the possibility that people are not really reasoning as Bayesians in Study 1. For this article, our primary focus was on establishing whether the particular process by which polarization can occur due to differential perceptions of source dependence predicted by our model has any empirical validity, and whether there is

any evidence of this process affecting real-world political beliefs. Future research could do more to fill in the gaps in our understanding about the kind of reasoning that underlies this process at the procedural level. We would note, however, that it is worth acknowledging that given our participants appear to act as if they were Bayesians in this context, at least broadly, our model may therefore capture the computational level of people’s reasoning – i.e., the problem the brain is trying to solve – even if it turns out they use something other than Bayesian reasoning to solve the problem in practice.

We also stress that in Study 1, belief polarization occurred in conditions which it would not be predicted to occur in according to the biased assimilation account, which is the standard explanation for cases of belief polarization. While the term is not always consistently defined, we consider biased assimilation to refer to a theoretical phenomenon whereby people use their judgment of how much a piece of presented evidence is congruent with their prior belief about the issue to which the evidence pertains to determine how strongly to weight that evidence in belief updating, leading them to selectively weight belief-congruent evidence more heavily than belief-incongruent evidence. As Baron (2008) points out, this is incompatible with Bayesianism – while it is fine for us to doubt the reliability of evidence which conflicts with our priors, Bayesians should not use this conflict to justify ignoring that evidence. This is because under Bayes’ rule, we should update in proportion to the likelihood ratio $p(e|H)/p(e|\neg H)$, which does not contain any information about $p(H)$, our prior for the issue-at-stake; therefore how much we weight the evidence about any hypothesis H should never be directly influenced by our prior for H . Any systematic relationship between our prior and how much we update would need to be due to other factors, such as other beliefs we hold which affect our calculation of the likelihood ratio (see Benoît and Dubra, 2019, for a discussion of how such systematicity can occur). Biased assimilation therefore cannot explain the results of Study 1, because it requires that people have relatively strong prior beliefs in order for some of the evidence to be incongruent with their initial position. In our Study 1, all participants had neutral prior beliefs as the issue at stake was artificial, therefore none of the presented evidence (the groups’ testimonies) could have been incongruent. Our result therefore shows that biased assimilation is not a necessity for belief polarization to occur, which could prompt a re-

assessment of whether previous results attributed to biased assimilation could be explained by alternative factors. Reconsideration of biased assimilation is also warranted given the poor replicability of belief polarization under the conditions it predicts (Anglin, 2019; Greitemeyer, 2014; Kuhn & Lao, 1996; Munro & Ditto, 1997). Ultimately though, our results cannot offer any definitive adjudication on whether biased assimilation explains cases of belief polarization or not, because belief polarization can be caused by multiple different pathways, one of which could be biased assimilation in other contexts.

The fact that the polarization which occurred in Study 1 is consistent with a Bayesian model is also important because a prominent view within political psychology is that phenomena like belief polarization are caused by “groupishness”, “tribalism”, and identity-driven motivated reasoning (i.e., people ‘wanting’ to hold beliefs which align with their group, see, e.g., Ditto et al., 2019; Kahan et al., 2012; Van Bavel & Pereira, 2018; Williams, 2021). Our results suggest this view, while not necessarily incorrect, is not the only plausible account – Study 1 showed that belief polarization can occur in the absence of directional motivations, and in Studies 2 and 3, the relationship between polarized beliefs and independence perceptions was robust to controlling for partisanship. But while our results point to alternative explanations for some studies taken to provide evidence for motivated reasoning, like Lord et al. (1979), not all evidence for motivated reasoning can be explained away by appealing to the impact of source perceptions, as some evidence comes from designs where partisans respond differently to information provided by the same source (e.g., Kahan et al., 2012, 2017; Thaler, 2024). Therefore our results do not challenge the validity of many pieces of evidence for motivated reasoning, though other Bayesian counter-explanations, perhaps with some allowances made for bounded rationality, might be plausible (e.g., Melnikoff & Strohminger, 2024). Overall, our results do not disprove motivated reasoning theories, but they do justify continued efforts to critically scrutinize the evidence for motivated reasoning (Tappin et al., 2020; Tappin & Gadsby, 2019; Williams, 2023).

Our results provide further evidence that perceptions of source dependence moderate belief updating (Madsen et al., 2020; Mercier & Miton, 2019; Pilditch et al., 2020). However, apparent deviations from normativity regarding dependence elsewhere (e.g., Mercier & Miton, 2019; Xie & Hayes, 2022) still require explanation. Part of the difference could be the relative simplicity of our task, and how explicitly the dependence information was communicated. Our participants were given information about the sources’ testimonial independence immediately before making their judgments, which may have cued participants to consider the source’s independence when making the judgment. Similarly, an assessment of the sources’ independence was communicated directly to participants, as we told them whether the sources were ‘very similar’ or ‘very different’ for each factor. Other studies may have made it more complex for people to extract information about dependence, appreciate its significance, or access it when making relevant judgments. We think there is virtue in making the task as simple as possible so that we can assess our participants’ ability to make appropriate judgments based on dependence information. But in real world political judgments, these additional complexities are relevant, meaning the ecological validity of Study 1 is open to critique – in the real world, people would have to make judgments about how similar different sources of information are for themselves based on observations, rather than having summaries provided to them, and they would need to recall this information when listening to sources’ testimony. Therefore, it would be useful for future studies to reintegrate some of these complexities and test whether the effect persists. Nevertheless, the fact that participants regarded similarities in the sources’ backgrounds, ideologies, and information consumption as a reason to discredit their testimony shows this critical cognitive building block for differential source dependence perceptions to cause belief polarization is in place.

In Studies 2 and 3, we established that the cross-sectional relationship between political beliefs and testimonial independence perceptions

among is robust to controlling for two plausible confounders partisanship and affective polarization. This result suggests that independence perceptions actually do have a polarizing influence on real-world political beliefs, though alternative explanations remain possible. Longitudinal analysis will be necessary before causal inferences can be made, and we cannot rule out that the relationships we have seen are caused by uncontrolled confounding variables. One plausible uncontrolled confounder could be slanted exposure – perhaps people on the left (right) consume political information from sources who successfully convince them to adopt left-wing (right-wing) beliefs and convince them that Conservative/Republican (Labour/Democrat) supporters have lower testimonial independence. There is also a danger of residual confounding if we have not measured partisanship and affective polarization with sufficient accuracy, though we have used standard approaches to do so. So while the idea of a real-world direct association between testimonial independence perceptions and belief polarization has survived initial falsification, significant further research is required.

It is worth considering whether our finding from Studies 2 and 3 that partisans judge their political in-group to have higher testimonial independence than their political out-group is “just” an instance of the ‘outgroup homogeneity effect’, whereby outgroups are perceived to have more homogeneous characteristics than ingroups (Boldry et al., 2007; Quattrone & Jones, 1980). Konovalova and Le Mens (2020) have proposed a plausible cognitive explanation for the outgroup homogeneity effect: because people typically are exposed to larger- n samples of information about the characteristics of in-groups, it is just a statistical truism that they are likely to estimate the variability of the in-group to be greater. We of course hypothesized a different process for the emergence of differential testimonial independence perceptions – that people infer independence in proportion to the perceived plausibility of a group’s claims. The mechanism proposed by Konovalova and Le Mens’ (2020) could also contribute to differential testimonial independence perceptions, however, it seems that there is more to our result than just the outgroup homogeneity effect. For one, Boldry et al.’s (2007) meta-analysis of the outgroup homogeneity effect finds that when participants are asked to judge the level of *similarity* between members of a group, which is what the items in the TIPS mostly concern, there was no evidence for an effect whereby outgroups were judged to have more similar members than ingroups. Instead, the out-group homogeneity effect was mostly driven by the perception that outgroup members are more likely to *conform to stereotypes* than ingroup members, and by facial recognition errors (Boldry et al., 2007). Furthermore, the effect sizes we observed for the extent to which people perceive their out-group to have lower testimonial independence than their in-group, $d = 0.87$ in the UK and $d = 0.82$ in the US, are much larger than the average meta-analytic effect size for the outgroup homogeneity effect estimated for non-minimal groups by Boldry et al. (2007), $d = 0.23$. While a stronger effect size like this could arise because our measurements of perceived similarity are more precise than those used in Boldry et al.’s studies, these discrepancies imply our effect is not just an instance of the outgroup homogeneity effect. However, we stress that while it is a prediction of applying Bayes’ theorem to Condorcet’s Jury Theorem and a prediction of our Bayesian Network, we have not yet empirically tested whether people do indeed infer greater dependence when a group’s claims are seen as less plausible. Future studies should test this hypothesis.

Future studies could also consider the contribution of source dependence perceptions to other areas of politics. For instance, it is a trope of conspiracy theorists to regard those who disagree with them as a homogeneous mass of “sheeple”, which is effectively an attribution of very high dependence; this may be a result of, and contribute to, their holding such discrepant beliefs to the rest of society. Furthermore, a feature of populist rhetoric is to regard mainstream politicians as ‘all the same’ – for example, populist American right-wingers sometimes describe Republicans and Democrats as a “uniparty”; again, this is essentially a description of mainstream political agents as being high-

dependence, which may be in a feedback loop with populists' rejection of mainstream politics. Furthermore, dependence is a characteristic of a *group*, but characteristics of *individuals* are also known to affect whether people trust political testimony (e.g., [Sanna & Lagnado, 2025](#); [Traberg & van der Linden, 2022](#)), and since the credibility of a group will be affected by perceptions of the individuals within the group, individual-level source characteristics like reliability ([Bovens & Hartmann, 2003](#); [Collins et al., 2018](#); [Hahn, Merdes, & von Sydow, 2018](#); [Hahn, von Sydow, & Merdes, 2018](#); [Merdes et al., 2021](#)), bias ([Wallace et al., 2020b, 2020a, 2020c](#); [Young & de Wit, 2024](#)), and trustworthiness and expertise ([Harris et al., 2012, 2016](#); [Madsen, 2016, 2019](#)) could also create polarizing feedback loops in a similar way to how we imagine dependence to. We capture the importance of individual-level credibility perceptions in our Bayesian Network by including nodes for the latent reliability of each source individually, but we do not manipulate any characteristics of the individuals, so future research could explore in more detail whether individual-level perceptions of bias, trustworthiness, expertise, and reliability contribute to polarization, and how these interact with dependence perceptions.

Overall, our findings provide empirical evidence for a new mechanism by which belief polarization may occur. We are able to demonstrate that real-world partisans possess the priors necessary for this belief polarization to occur in practice, and those partisans with the priors most theoretically-predictive of belief polarization are indeed the most polarized.

CRediT authorship contribution statement

David J. Young: Writing – review & editing, Writing – original draft, Visualization, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization. **Jens Koed Madsen:** Writing – review & editing, Supervision, Resources, Project administration, Methodology, Funding acquisition, Conceptualization. **Lee H. de-Wit:** Writing – review & editing, Supervision, Resources, Project administration, Methodology, Funding acquisition, Conceptualization.

Declaration of competing interest

The authors declare no competing interests.

Acknowledgments

We thank Adam Harris, Hugo Mercier, and Gordon Brown for helpful comments on earlier versions of this work. This research was funded by a Templeton World Charity Foundation grant, TWCF Number 31453, <https://doi.org/10.54224/31453>, as well as MHP Group, who fund the MHP UK Polarization Tracker. An earlier version of this article, without Studies 2 and 3, was published in the *Proceedings of the Annual Meeting of the Cognitive Science Society*, Volume 46, 2024, under the title 'Bayesian Belief Polarization due to Differential Perceptions of Source Independence' (<https://escholarship.org/uc/item/19x670dt>); we thank the three anonymous reviewers who evaluated that submission for their feedback, which was integrated into this paper. We thank the Editor and two anonymous reviewers of this article.

Appendix A. Supplementary data

Supplementary data to this article can be found online at <https://doi.org/10.1016/j.cognition.2025.106126>.

Data availability

A link to our data repository is provided at the end of the Introduction.

References

- Alister, M., Perfors, A., & Ransom, K. (2022). Source independence affects argument persuasiveness when the relevance is clear. In J. Culbertson, A. Perfors, H. Rabagliati, & V. Ramenzoni (Eds.), *Proceedings of the 44th annual conference of the cognitive science society*. <https://escholarship.org/uc/item/5hg4p8cm>.
- 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>
- Baron, J. (2008). *Thinking and deciding (fourth edition)*. Cambridge University Press.
- 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>
- Boldry, J. G., Gaertner, L., & Quinn, J. (2007). Measuring the measures: A meta-analytic investigation of the measures of outgroup homogeneity. *Group Processes & Intergroup Relations, 10*(2), 157–178. <https://doi.org/10.1177/1368430207075153>
- 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>
- 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. <https://doi.org/10.3389/fpsyg.2018.00018>
- Condorcet, N.d. (2014). *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. Cambridge: Cambridge University Press.
- Connor Desai, S., Xie, B., & Hayes, B. K. (2022). Getting to the source of the illusion of consensus. *Cognition, 223*, Article 105023. <https://doi.org/10.1016/j.cognition.2022.105023>
- Cook, J., & Lewandowsky, S. (2016). Rational irrationality: Modeling climate change belief polarization using Bayesian networks. *Topics in Cognitive Science, 8*(1), 160–179. <https://doi.org/10.1111/tops.12186>
- 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. Retrieved from: <https://www.pewresearch.org/politics/2014/06/12/political-polarization-in-the-american-public/>.
- Ditto, P. H., Liu, B. S., Clark, C. J., Wojcik, S. P., Chen, E. E., Grady, R. H., ... Zinger, J. F. (2019). At least bias is bipartisan: A meta-analytic comparison of partisan bias in liberals and conservatives. *Perspectives on Psychological Science, 14*(2), 273–291. <https://doi.org/10.1177/1745691617746796>
- Druckman, J. N., & Levendusky, M. S. (2019). What do we measure when we measure affective polarization? *Public Opinion Quarterly, 83*(1), 114–122. <https://doi.org/10.1093/poq/nfz003>
- Flora, D. B. (2020). Your coefficient alpha is probably wrong, but which coefficient omega is right? A tutorial on using R to obtain better reliability estimates. *Advances in Methods and Practices in Psychological Science, 3*(4), 484–501. <https://doi.org/10.1177/2515245920951747>
- 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), Article e93440. <https://doi.org/10.1371/journal.pone.0093440>
- Guest, O., & Martin, A. E. (2021). How computational modeling can force theory building in psychological science. *Perspectives on Psychological Science, 16*(4), 789–802. <https://doi.org/10.1177/1745691620970585>
- Hahn, U., Merdes, C., & von Sydow, M. (2018). How good is your evidence and how would you know? *Topics in Cognitive Science, 10*(4), 660–678. <https://doi.org/10.1111/tops.12374>
- Hahn, U., von Sydow, M., & Merdes, C. (2018). How communication can make voters choose less well. *Topics in Cognitive Science, 11*(1), 194–206. <https://doi.org/10.1111/tops.12401>
- Harris, A. J. L., 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. <https://doi.org/10.1111/cogs.12276>
- 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>
- 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/>.
- Huddy, L., & Bankert, A. (2017). Political partisanship as a social identity. In *Oxford research encyclopedia of politics*. Oxford University Press. <https://doi.org/10.1093/acrefore/9780190228637.013.250>.
- Iyengar, S., Lelkes, Y., Levendusky, M., Malhotra, N., & Westwood, S. J. (2019). The origins and consequences of affective polarization in the United States. *Annual Review of Political Science, 22*(1), 129–146. <https://doi.org/10.1146/annurev-polisci-051117-073034>
- 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>
- Jönsson, M. L., Hahn, U., & Olsson, E. J. (2015). The kind of group you want to belong to: Effects of group structure on group accuracy. *Cognition, 142*, 191–204. <https://doi.org/10.1016/j.cognition.2015.04.013>

- Kahan, D. M., Hoffman, D. A., Braman, D., Evans, D., & Rachlinski, J. J. (2012). "They saw a protest": Cognitive illiberalism and the speech-conduct distinction. *Stanford Law Review*, 64, 57.
- Kahan, D. M., Peters, E., Dawson, E. C., & Slovic, P. (2017). Motivated numeracy and enlightened self-government. *Behavioural Public Policy*, 1(1), 54–86. <https://doi.org/10.1017/bpp.2016.2>
- Konovalova, E., & Le Mens, G. (2020). An information sampling explanation for the in-group heterogeneity effect. *Psychological Review*, 127(1), 47–73. <https://doi.org/10.1037/rev0000160>
- 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.
- Lord, Charles G., & Taylor, Cheryl A. (2009). Biased Assimilation: Effects of Assumptions and Expectations on the Interpretation of New Evidence: Biased Assimilation. *Social and Personality Psychology Compass*, 3(5), 827–841. <https://doi.org/10.1111/j.1751-9004.2009.00203.x>
- Madsen, J. K. (2016). Trump supported it?! A Bayesian source credibility model applied to appeals to specific American presidential candidates' opinions. In A. Papafragou, D. Grodner, D. Mirman, & J. C. Trueswell (Eds.), *Proceedings of the 38th annual conference of the cognitive science society* (pp. 165–170). Cognitive Science Society.
- Madsen, J. K. (2019). Voter reasoning bias when evaluating statements from female and male political candidates. *Politics & Gender*, 15(02), 310–335. <https://doi.org/10.1017/S1743923X18000302>
- 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>
- Mason, L. (2018). *Uncivil agreement: How politics became our identity*. University of Chicago Press.
- Melnikoff, D. E., & Strohminger, N. (2024). Bayesianism and wishful thinking are compatible. *Nature Human Behaviour*, 8(4), 692–701. <https://doi.org/10.1038/s41562-024-01819-6>
- 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>
- Merdes, C., von Sydow, M., & Hahn, U. (2021). Formal models of source reliability. *Synthese*, 198(S23), 5773–5801. <https://doi.org/10.1007/s11229-020-02595-2>
- 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>
- Muthukrishna, M., & Henrich, J. (2019). A problem in theory. *Nature Human Behaviour*, 3(3), 221–229. <https://doi.org/10.1038/s41562-018-0522-1>
- Oberauer, K., & Lewandowsky, S. (2019). Addressing the theory crisis in psychology. *Psychonomic Bulletin & Review*, 26(5), 1596–1618. <https://doi.org/10.3758/s13423-019-01645-2>
- Olsson, E. J. (2013). A Bayesian simulation model of group deliberation and polarization. In F. Zenker (Ed.), *Vol. 362. Bayesian argumentation* (pp. 113–133). Netherlands: Springer. https://doi.org/10.1007/978-94-007-5357-0_6
- Orr, L. V., Fowler, A., & Huber, G. A. (2023). Is affective polarization driven by identity, loyalty, or substance? *American Journal of Political Science*, 67(4), 948–962. <https://doi.org/10.1111/ajps.12796>
- 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*. San Francisco, CA: Morgan Kaufmann Publishers.
- Pilditch, T. D., Hahn, U., Fenton, N., & Lagnado, D. (2020). Dependencies in evidential reports: The case for informational advantages. *Cognition*, 204, Article 104343. <https://doi.org/10.1016/j.cognition.2020.104343>
- Quattrone, G. A., & Jones, E. E. (1980). The perception of variability within in-groups and out-groups: Implications for the law of small numbers. *Journal of Personality and Social Psychology*, 38(1), 141–152. <https://doi.org/10.1037/0022-3514.38.1.141>
- R Core Team. (2022). *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing. URL www.R-project.org/.
- Revelle, W. (2023). *psych: Procedures for psychological, psychometric, and personality research (2.3.3) [Computer software]*. Northwestern University. <https://CRAN.R-project.org/package=psych>.
- Rollwage, Max, Zmigrod, Leor, de-Wit, Lee, Dolan, Raymond J., & Fleming, Stephen M. (2019). What Underlies Political Polarization? A Manifesto for Computational Political Psychology. *Trends in Cognitive Sciences*, 23(10), 820–822. <https://doi.org/10.1016/j.tics.2019.07.006>
- 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>
- Sanna, G. A., & Lagnado, D. (2025). Belief updating in the face of misinformation: The role of source reliability. *Cognition*, 258, Article 106090. <https://doi.org/10.1016/j.cognition.2025.106090>
- Tappin, B. M., & Gadsby, S. (2019). Biased belief in the Bayesian brain: A deeper look at the evidence. *Consciousness and Cognition*, 68, 107–114. <https://doi.org/10.1016/j.concog.2019.01.006>
- Tappin, B. M., Pennycook, G., & Rand, D. G. (2020). Thinking clearly about causal inferences of politically motivated reasoning: Why paradigmatic study designs often undermine causal inference. *Current Opinion in Behavioral Sciences*, 34, 81–87. <https://doi.org/10.1016/j.cobeha.2020.01.003>
- Thaler, M. (2024). The fake news effect: experimentally identifying motivated reasoning using trust in New. *American Economic Journal: Microeconomics*, 16(2), 1–38. <https://doi.org/10.1257/mic.20220146>
- Traberg, C. S., & van der Linden, S. (2022). Birds of a feather are persuaded together: Perceived source credibility mediates the effect of political bias on misinformation susceptibility. *Personality and Individual Differences*, 185, Article 111269. <https://doi.org/10.1016/j.paid.2021.111269>
- 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>
- Wallace, L. E., Wegener, D. T., & Petty, R. E. (2020a). Consuming information from sources perceived as biased versus untrustworthy: Parallel and distinct influences. *Journal of the Association for Consumer Research*, 5(2), 137–148. <https://doi.org/10.1086/707732>
- Wallace, L. E., Wegener, D. T., & Petty, R. E. (2020b). Influences of source bias that differ from source untrustworthiness: When flip-flopping is more and less surprising. *Journal of Personality and Social Psychology: Attitudes and Social Cognition*, 118(4).
- Wallace, L. E., Wegener, D. T., & Petty, R. E. (2020c). When sources honestly provide their biased opinion: Bias as a distinct source perception with independent effects on credibility and persuasion. *Personality and Social Psychology Bulletin*, 46(3), 439–453. <https://doi.org/10.1177/0146167219858654>
- 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>
- Whalen, Andrew, Griffiths, Thomas L., & Buchsbaum, Daphna (2018). Sensitivity to Shared Information in Social Learning. *Cognitive Science*, 42(1). <https://doi.org/10.1111/cogs.12485>
- Wickham, H., & Henry, L. (2023). *purrr: Functional programming tools*. R package version 1.0.1. CRAN.R-project.org/package=purrr.
- Williams, D. (2021). Socially adaptive belief. *Mind & Language*, 36(3), 333–354. <https://doi.org/10.1111/mila.12294>
- Williams, D. (2023). The case for partisan motivated reasoning. *Synthese*, 202(3), 89. <https://doi.org/10.1007/s11229-023-04223-1>
- Xie, B., & Hayes, B. (2022). Sensitivity to evidential dependencies in judgments under uncertainty. *Cognitive Science*, 46, Article e13144. <https://doi.org/10.1111/cogs.13144>
- Young, D. J., & de Wit, L. (2024). *The bias-and-expertise model: A Bayesian network model of political source characteristics*. <https://doi.org/10.31234/osf.io/bf36a>
- 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. <https://escholarship.org/uc/item/57h4j672>.
- 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. <https://doi.org/10.1177/0956797619856844>