Differential Memory for Belief-Congruent versus Belief-Incongruent Arguments Cannot Explain Belief-Driven Argument Evaluation

Calvin Deans-Browne
(calvin.deans-browne.20@ucl.ac.uk)

Pia Roth
(pia.roth.23@alumni.ucl.ac.uk)

Carolin Echterbeck
(caro.echterbeck.22@ucl.ac.uk)

Henrik Singmann
(h.singmann@ucl.ac.uk)

University College London, Department of Experimental Psychology, 26 Bedford Way, London, WC1H 0AP, UK

Abstract

A well-established phenomenon is that people often rely more on their prior beliefs than on the presented evidence when evaluating arguments. In the present study, we are interested in the cognitive mechanisms underlying this phenomenon. We hypothesise that when individuals encounter an argument that is congruent with their beliefs, it activates related information in memory. As a consequence, for arguments in line with their beliefs, people should be more likely to correctly recognise previously encountered information but also more likely to incorrectly recognise new information as previously seen. To investigate this hypothesis, our two-part experiment first investigated the effect of participants' beliefs about political claims on their evaluation of corresponding arguments that varied in quality. In the second part, a surprise memory test was employed to assess their recognition memory for these arguments. While we replicated the finding that prior beliefs drive argument evaluations, prior beliefs did not affect memory performance for all arguments in the same way. Our results indicate that individuals may rely on prior beliefs to aid memory only when faced with a difficult memory task.

Keywords: belief bias; recognition memory; reasoning

Introduction

The internet and social media in particular have vastly increased the availability and variety of information we can access daily. With this abundance of information, individuals carry greater responsibility for deciding what to consume and what to trust. One interesting phenomenon is that the same information can be perceived quite differently from different people. For example, some people do not trust the information provided by governmental institutions regarding important topics even if they reflect the best available evidence (e.g., vaccine information provided by institutions such as the WHO). In this manuscript we are concerned with the cognitive mechanisms underlying this phenomenon: Why are some people persuaded by certain arguments and other people are not?

Numerous studies have shown that when asked to evaluate the quality of presented information or arguments, people not only consider the presented evidence, but also their prior beliefs about the content being discussed (e.g., Deans-Browne et al., 2024; McCrudden et al., 2017; Stanovich & West, 1997; Taber & Lodge, 2006). For example, in our recent research (Deans-Browne & Singmann, 2025; Deans-Browne et al., 2024) we found strong correlations between people's prior beliefs about a topic and their evaluation of related arguments. In our studies, participants first had to rate their beliefs about politically charged claims such as "Abortion should be illegal" or "Cancel Culture is good for society" on a scale from *extremely false* to *extremely true*. Participants were then shown arguments about these claims and had to rate them on a scale from *extremely bad* to *extremely good*. We also manipulated the arguments people saw so that they were either of good quality, internally inconsistent, or based on appeals to authority. Results showed that whilst participants perceived good arguments as best and authority-based arguments as worst, the quality rating of all arguments was significantly correlated with participants' prior beliefs. Importantly, prior beliefs were a better predictor of the perceived quality of the argument than was the type of argument each participant saw. With regards to evaluating arguments representative of those that people see in their everyday life, beliefs largely determined the subjective quality of the argument to an individual.

From a normative perspective, an effect of prior beliefs on perceived argument quality is not necessarily a sign of irrationality. For example, Hahn and Oaksford (2007) presented a Bayesian model of argument evaluation showing that argument quality, prior beliefs, and the type of argument should all have sizable effects on perceived argument quality.

In the present work, we investigate the cognitive mechanisms underlying the effects of prior beliefs on perceived argument quality. One potential mechanism could be the well-evidenced interaction between belief and memory (e.g., Alba & Hasher, 1983; Brod et al., 2013). According to Integration Theory (Heit, 1993), when a person processes new information, related existing information that is held in memory is reactivated. Beliefs are one such form of existing information that can get reactivated. For example, a meta-analysis by Stangor and McMillan (1992) analysed studies investigating stereotype effects on memory. The authors found that memory items congruent with the presented stereotype, and as such congruent with participants' prior beliefs, are more likely to be both correctly recognised when the items had previously been seen (i.e., more hits) and incorrectly recognised when the items had not previously been seen (i.e., more false alarms) compared to incongruent memory items. In line with these results, in a study investigating the recall of complex emotive arguments, Wiley

In D. Barner, N.R. Bramley, A. Ruggeri and C.M. Walker (Eds.), *Proceedings of the 47th Annual Conference of the Cognitive Science Society*

(2005) found that participants better recalled arguments which were congruent with their beliefs.

Applying Integration Theory to argument evaluation we make the following hypothesis: When reading an argument about a topic for which we hold congruent prior beliefs, associated prior knowledge is activated. This activation can embellish the argument in our heads; the activated knowledge increases the argument's perceived quality by serving as additional supporting evidence. Thus, people do not only use the information provided to judge the quality of the argument but also the associated prior knowledge that becomes activated during argument processing. We can test this hypothesis by testing participants' memory of arguments they have read. If our hypothesis is correct there should be more hits and false alarms for belief-congruent compared to belief-incongruent argument memory items.

The Current Study

The aim of the current study is to investigate whether the effect of belief on argument evaluations can be explained by differential memory for belief-congruent versus belief-incongruent information. Our hypothesis drawing from Integration Theory is that if reading a newly presented argument activates related belief-congruent information, then participants should show both more hits and more false alarms for belief-congruent information.

To test this hypothesis, the study is split into two parts. In the first part, the Everyday Argument Assessment Task, we first measure participants' beliefs about political claims that vary in political ideology and then ask them to rate the quality of arguments related to these claims. In the second part of the study, participants are presented with a surprise recognition memory test for the arguments they were just asked to rate.

Further details of stimuli and analysis mentioned in this manuscript can be found on OSF: https://osf.io/c48er/?view_only=dfd7f7aa4138402490f20a2a2ef2a3d9

Method

Participants

A total of 125 participants were recruited through Prolific. Our sample was restricted to native English speakers in the US as our stimuli were related to political topics that are polarizing within the US (e.g., gun control laws). 13 participants failed attention checks (similar to those from Deans-Browne et al., 2025) and were excluded. This left us with N = 112 participants (53 male and 59 female; average age = 37 years). The experiment was approved by the Experimental Psychology Department's ethics committee of University College London.

Design

Everyday Argument Assessment Task The task consisted of initial ratings of belief statements followed by subsequent argument quality ratings. Participants first rated eight political belief statements (e.g. "Abortions should be legal in the US") on a 7-point Likert scale ranging from 1 (*extremely false)* to 7 (*extremely true)*. Each participant received four left-leaning and four right-leaning belief statements, so that each participant saw a roughly equal number of statements that were congruent and incongruent with their beliefs. *Belief ratings* were our first (continuous) independent variable.

After all belief ratings had been made, participants were presented with eight arguments. Each argument's conclusion was identical to a previously shown belief statement. Thus, participants received four left-leaning and four right-leaning arguments corresponding to the previously presented belief statements. Participants were instructed to rate the quality of arguments on a 6-point Likert scale (from 1 = *extremely bad* to 6 = *extremely good*) based on how well the evidence supported the argument's conclusion.

Crucially, *argument quality* was manipulated so that arguments were either good or inconsistent. For both left-leaning and right-leaning arguments, participants saw two good arguments and two inconsistent arguments. Argument quality served as the second independent variable.

Together, the Everyday Argument Assessment Task employed a continuous (belief; extremely false to extremely true) × 2 (argument quality; good vs inconsistent) design. Each participant provided one belief rating for each belief-statement for the eight topics and saw an equal number of good and inconsistent arguments (i.e., four each). The dependent variable was the argument quality rating each participant provided for each argument.

The materials in the Everyday Argument Assessment Task revolved around ten political topics (abortion, climate change, kneeling during the national anthem, private prisons, fracking, three strikes laws, gun control, cancel culture, secularization, affirmative action). For each topic, there were two belief statements, one left-leaning and one right-leaning. For each belief statement there were two arguments: one good argument and one inconsistent argument (i.e., a total of $10 \times 2 \times 2 = 40$ arguments). For each participant, eight of the ten topics were randomly selected with the condition assigned to topics randomly given the design outlined above (i.e., 2 left-leaning & good, 2 left-leaning & inconsistent, 2 right-leaning & good, and 2 right-leaning & inconsistent).

Good arguments presented evidence that consistently supported their conclusion and consisted of summarized versions of established arguments from educational bipartisan sources (procon.org). Conversely, inconsistent arguments contained evidence for and against the arguments' overall conclusion and were presented in the following manner: The argument starts by providing evidence in support of the given conclusion, followed by contradictory evidence against the conclusion, and ends by once again providing evidence in support of the conclusion. Examples of both good and inconsistent arguments can be seen in Table 1. All arguments were 75 words in length.

Table 1: Example of a good and inconsistent argument including the separation into old memory items

Good Argument	In-line	In-line	In-line
Abortions under Roe v. Wade balanced two fundamental rights; the right of the pregnant woman to bodily autonomy and the right of the unborn child to life. The unborn child only has the potential for life as we know it when they can survive outside the womb, and abortions had to occur before this stage under this ruling. Consequently, abortions can be consistent with both fundamental rights. Abortion should therefore be legal in the US.	(Item 1) Abortions under Roe v. Wade balanced two fundamental rights; the right of the pregnant woman to bodily autonomy and the right of the unborn child to life.	(Item 2) The unborn child only has the potential for life as we know it when they can survive outside the womb, and abortions had to occur before this stage under this ruling [Roe v. Wade].	(Item 3) Consequently, abortions can be consistent with both fundamental rights [to bodily autonomy and to life].
Inconsistent Argument	**In-line**	**Not in-line**	**In-line**
Abortions are safe procedures that protect lives. Women who are denied abortions are also more likely to later have poorer mental and physical health, alongside financial problems. Instead of promoting abortions, increased access to birth control, health insurance, and sexual education would make abortions unnecessary. Abortions promote the idea that human lives are disposable when inconvenient. Abortions protect the bodily autonomy of women – a fundamental human right. Therefore, abortions should be legal in the US.	(Item 1) Abortions are safe procedures that protect lives. (Item 2) Women who are denied abortions are also more likely to later have poorer mental and physical health, alongside financial problems.	(Item 3) Instead of promoting abortions, increased access to birth control, health insurance, and sexual education would make abortions unnecessary. (Item 4) Abortions promote the idea that human lives are disposable when inconvenient.	(Item 5) Abortions protect the bodily autonomy of women – a fundamental human right.

Recognition Memory Test In the surprise recognition memory test, participants were presented with ten memory items per argument. A memory item consisted of one statement that could have been part of the argument they had previously read (see Table 1 for examples). Crucially, memory items varied in two ways. Firstly, they were either in-line or not-in-line with the conclusion of a previously shown argument. Secondly, the memory item was either old (i.e., part of the argument participants read) or new (i.e., not part of the argument they read), which we refer to as the *memory status* of an item. For each memory item, participants had to indicate whether it was part of a previously shown argument, using a 6-point memory confidence rating scale ranging from 1 (*definitely new)* to 6 (*definitely old)*.

Participants were shown 10 memory items for each argument. Of those 10 memory items, six were in-line with the argument's conclusion and four not-in-line with the conclusion. For good arguments, participants saw six in-line memory items, of which three were old, three were new, and four not-in-line memory items, of which all were new. Consequently, for good arguments, 3 out of 10 memory items were old. For inconsistent arguments, participants were presented with six in-line memory items, of which two were old and four were new, and four not-in-line memory items, of which two were old and two were new. Consequently, for inconsistent arguments, 4 out of 10 memory items were old. Because only inconsistent arguments contained in-line and not-in-line information, we could not simultaneously balance the number of in-line and not-in-line memory items and the number of old and new memory items across both argument types. We chose to balance the number of in-line and not-in-line memory items across good and inconsistent arguments. This way, participants could not use the number of in-line and not-in-line memory items to infer whether they had seen a good or inconsistent argument before (i.e., they had to rely on their memory for that inference).

To simplify the analysis of the recognition memory test, we combined participants' belief ratings from the Everyday Argument Assessment Task as well as the alignment of the memory item with the argument's conclusion into a new variable, *belief congruency* that captures if the memory item was congruent or incongruent with participants' beliefs. To do so, we reverse coded participants' belief ratings for all not-in-line memory items. Belief congruency ranged from 1 (memory item is incongruent with participants' beliefs) to 7 (memory item is congruent with participants' beliefs).

In sum, we had a continuous (belief congruency; incongruent to congruent) × 2 (memory status; old vs new) × 2 (argument quality; good vs inconsistent) design for the recognition memory test. Participants' memory confidence ratings was the dependent variable.

To construct the old memory items, each argument was split into statements consisting of 1-2 sentences each. Good arguments were separated into 3-4 statements, all in-line with the argument's conclusion. Inconsistent arguments were split into 4-5 statements: 2-3 in-line statements that support the argument's conclusion and 2-3 not-in-line statements that contradict the argument's conclusion (see Table 1). To ensure each memory item made sense without the context of the full

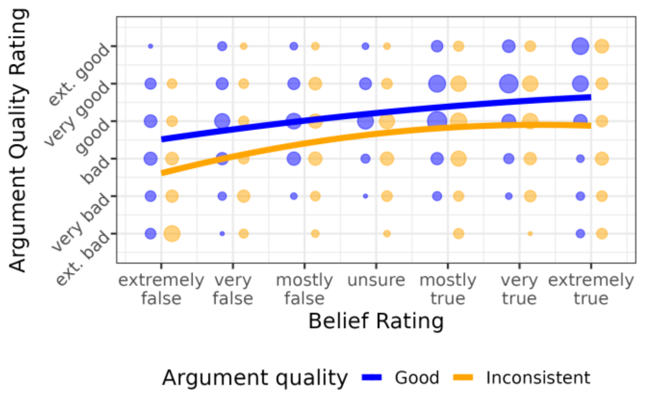

Figure 1: Participants' argument quality ratings for good or inconsistent arguments across belief ratings. Results are collapsed across left-leaning and right-leaning belief statements. The dots represent individual responses, while the lines depict the LMM predictions. The size of each dot represents proportionally the number of argument quality rating responses.

argument, some minor changes were made and presented in square brackets (which participants were told about), as shown in Table 1.

Some of the new memory items were taken from the old memory items of other arguments on the same topic that were not seen by the participant (e.g., if the participant saw a left-leaning inconsistent argument about abortion, a new memory item might be from a good left-leaning argument about abortion). As this did not always provide enough new items (e.g., because some arguments used overlapping statements), we created additional new memory items for some arguments.

Procedure

The experiment started with the Everyday Argument Assessment Task. Participants were first asked to rate their belief for each of the eight topics. For each topic, participants were first presented with some background information followed by a statement related to the topic (e.g., "Abortions should be legal in the US"). After rating each belief statement, participants saw one argument per topic with a matching conclusion and had to rate its quality. The order of arguments was randomly determined for each participant (i.e., it did not need to match the order in which participants rated their beliefs).

After the Everyday Argument Assessment Task, participants were given the surprise recognition memory test. Each of the 10 memory items per argument was presented individually and participants had to rate their memory confidence for each item. All memory items for one argument were presented in a blocked form in a random order. Furthermore, the order of the item blocks for each argument was also randomly determined. After participants had completed the memory recognition test, they answered standard demographic questions and were debriefed.

The same two attention-check items as in Deans-Browne and Singmann (2025) were included: two pairs of belief statements and aligning arguments. Participants who answered both attention checks incorrectly were excluded from the study.

Results

The results section is split into two parts following the design of the experiment. In the first part, we analyse the responses to the Everyday Argument Assessment Task. The aim is to check if we replicate the effect of belief on argument evaluations. In the second part, we analyse the responses to the recognition memory task to see whether we find evidence for the hypothesis that differential memory can explain the effect of belief on argument evaluations.

Everyday Argument Assessment Task

To investigate the effect of belief on argument evaluations, we employed a linear mixed model predicting argument quality ratings from fixed-effects of belief (with ratings from - 3 [= extremely false] to 3 [= extremely true]), belief squared (the squared values of belief), argument quality (good vs inconsistent), and all interactions and crossed random effects for participants and items (i.e., the specific argument). We started the analysis with the maximal model justified by the design. As this showed convergence issues, we reduced the random-effects structure until no convergence issues appeared. Here, we report the results of the reduced model, which showed a similar pattern of significant and non-significant effects as the maximal model, except that the reduced model also exhibited a significant effect for belief squared.[1]

The main results of the experiment are shown in Figure 1 and replicate those from Deans-Browne et al. (2024). We can see a clear main effect of argument quality, $F(1, 110.41) = 11.17$, $p = .001$. Participants rated good arguments on average 0.68 (95% CI [0.40, 0.96]) points higher than inconsistent arguments.

Figure 1 also shows a clear effect of prior beliefs on argument quality ratings. Participants who rated the quality of arguments as better tended to be in greater agreement with the corresponding belief statement. In the model this is captured by a significant main effect of belief rating, $F(1, 121.14) = 55.95$, $p < .001$. For every additional point of belief rating, argument quality ratings increased by 0.20 points on average (95% CI [0.15, 0.25]). The belief rating scale had seven points, encompassing six increments between the lowest and highest belief ratings. Thus, the average difference in argument quality ratings between the lowest and

[1] In the maximal model, the by-participant and by-item random-effects structure mirrored the fixed-effect structure, with the exception that the by-item random effects structure excluded the term for argument quality. The final model employed by-participant random intercepts with by-participant random slopes for all three main effects (i.e., belief, belief squared, and argument quality), as well as a by-item random intercept.

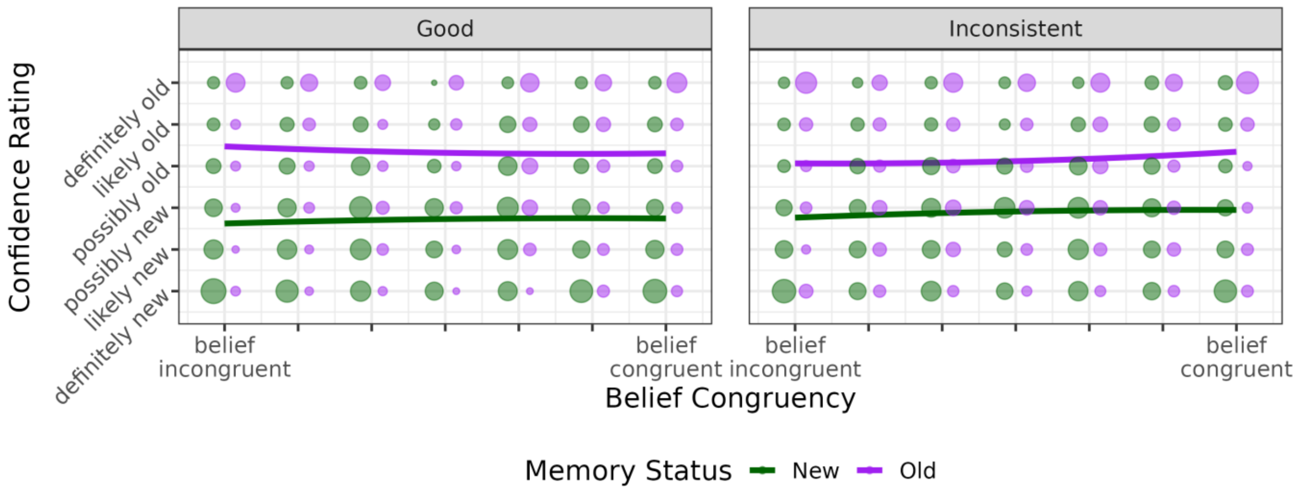

Figure 2. Memory confidence ratings for old and new memory items across belief congruency, illustrated separately for good and inconsistent arguments. Results are collapsed across left-leaning and right-leaning belief statements as well as in-line and not in-line memory items. The dots represent aggregated responses across participants, while the lines depict the regression lines as fitted by the model. Violet dots show memory confidence ratings for old memory items, while green dots show memory confidence ratings for new memory items. The size of each dot represents proportionally the number of memory confidence rating responses across the corresponding memory items.

highest belief scores was 1.2 points (0.29 x 6). Conversely, the main effect of argument quality showed a difference of 0.68 points between ratings for good and inconsistent arguments. Our results, therefore, replicate the finding that prior beliefs have a larger effect on argument quality than does the argument quality itself (Deans-Browne & Singmann, 2025; Deans-Browne et al., 2024).

Further in line with Deans-Browne et al. (2024), we found no evidence for an interaction between argument quality and prior beliefs, $F(1, 738.47) = 0.35$, $p = .553$.

Recognition Memory Test

Now that we have established the effect of prior beliefs on participants' argument quality ratings, we check whether we can explain it with differential memory for belief congruent versus incongruent information. As a reminder, our hypothesis derived from Integration Theory is that belief congruency (i.e., memory items aligning with participants' beliefs) should lead to both more hits and more false alarms. In our design, this would entail a positive relationship between belief congruency and recognition confidence for both old and new memory items (i.e., higher recognition confidence for old and new items for items congruent with beliefs compared to those incongruent with beliefs).

To test this, we estimated a linear mixed model predicting participants' memory confidence ratings from fixed effects belief congruency (from -3 [= *incongruent*] to 3 [= *congruent*]), belief congruency squared (consisting of the squared values of belief congruency), argument quality (good vs inconsistent), memory status (old vs new), and all interactions. We started with the maximal model justified by the design. As this showed convergence issues, we reduced the random-effects structure until no convergence issues appeared. Here, we report the results of the reduced model, which showed the same pattern of significant and non-significant effects as the maximal model[2].

The main results are shown in Figure 2. As expected, we saw a main effect of memory status, $F(1, 143.61) = 99.17$, $p < .001$, indicating that participants reported higher confidence for old compared to new memory items. Confidence ratings for old memory items were on average 1.45 (95% CI [1.19, 1.72]) points higher than confidence ratings for new items. This shows that despite this being a surprise recognition memory test participants could distinguish old from new memory items.

In contrast to our main hypothesis, Figure 2 does not suggest a consistent effect of belief congruency on memory confidence ratings. In line with this visual impression, we did not find a significant main effect of belief congruency $F(1, 123.98) = 2.10$, $p = .150$. This finding speaks against the hypothesis derived from Integration Theory that differential memory for belief congruent versus incongruent items can explain the effect of belief on argument evaluations.

However, Figure 2 suggests a potentially more complicated relationship between belief congruency and memory. For inconsistent arguments, confidence ratings seem to slightly increase with belief for old and new memory items. However, for good arguments, we saw no consistent pattern. In line

[2] The maximal model justified by the design employed crossed random effects with by-participant and by-memory item grouping factors, each containing random intercepts and slopes mirroring the fixed-effect structure with the exception that the argument quality term was excluded within the memory item factor. The reduced model employed by-participant random intercepts and random slopes for all main effects except for belief congruency squared and the same random effect structure for the by-item term as the maximal model.

with this impression, we found a significant interaction between belief congruency and argument type $F(1, 122.45) = 4.64$, $p = .033$. The slope for confidence ratings for inconsistent arguments was positive, $\beta = 0.04$ (95% CI [0.01, 0.07]) and different from zero. As the belief congruency scale had seven points, encompassing six increments between the lowest and highest belief ratings, the average difference in confidence rating scores for inconsistent memory items was 0.24 points (0.04 x 6). The average slope for good arguments was not significant, $\beta = 0.00$ (95% CI [-0.04, 0.28]).

Further inspection of Figure 2, as well as the model results, did not provide evidence for a main effect of argument quality, $F(1, 334.67) = 0.03$, $p = .869$. The regression lines for old and new memory items for good arguments are not noticeably higher or lower than those for inconsistent arguments.

Interestingly, we found a significant interaction between argument quality and memory status, $F(1, 357.51) = 9.96$, $p = .002$. This can be seen in Figure 2 as the regression lines for old and new responses are closer together for inconsistent memory items than for good memory items. For both argument qualities, we can see a clear positive effect of memory status ($p < .001$). However, the effect of memory status was 0.37 (95% CI [0.18, 0.55]) confidence rating points smaller for inconsistent arguments compared to good arguments. Thus, it was more difficult for participants to distinguish old from new memory items for inconsistent arguments compared to good arguments (i.e., the memory task was more difficult for inconsistent arguments). Importantly, the differential difficulty of the memory task was not confounded with an obvious shift in response bias: new items received on average higher and old items received on average lower confidence ratings for inconsistent compared to good arguments.

Finally, we did not see any noticeable effects of belief congruency squared, neither as a main effect, $F(1, 343.83) = 0.07$, $p = .794$, nor in interaction with memory status, $F(1, 117.48) = 1.35$, $p = .247$.

Discussion

The aim of this study was to test whether we can explain the effect of prior beliefs on perceived argument quality through differential memory for belief congruent versus incongruent arguments. To this end, we first replicated the effect of prior beliefs on perceived argument quality following the design of Deans-Browne et al. (2024; see also Deans-Browne & Singmann, 2025). We found that participants, when tasked with rating the quality of socio-political arguments from an objective standpoint, tended to rate arguments that were in line with their beliefs as being of higher quality than those that were not. Participants did this for both arguments of good quality, and those that were internally inconsistent, even though they could distinguish both types of arguments (i.e., good arguments were perceived as being of higher quality than inconsistent arguments).

However, we did not find compelling evidence for our main hypothesis. Participants did not show more hits and more false alarms for *all* belief-congruent compared to belief-incongruent information. Instead, we found more hits and more false alarms for belief-congruent information only for memory items based on inconsistent arguments and not for memory items based on good arguments. In addition, the effect of beliefs on memory for inconsistent argument was of small magnitude (i.e., the difference in confidence ratings for memory items that were most belief-congruent vs least belief congruent was 0.24 points, a quarter of a point on the scale). This suggests that belief-driven memory activation is *not* the general mechanism behind belief-driven argument evaluation.

One possible explanation for the finding that beliefs drive increased hits and false alarms for inconsistent memory items only could be related to task difficulty. Distinguishing between old and new memory items is harder for inconsistent compared to good arguments. One possible reason for this difference in difficulty is that only for the former can both in-line and not-in-line items be plausibly old or new.

The finding that beliefs only play a role for the more difficult memory items is in line with the broader reasoning literature showing that people use their beliefs to a greater degree for more complex tasks. For instance, in the belief bias task by Evans and colleagues (1983), participants were more likely to conflate argument validity with believability for more complex syllogisms compared to simple ones (Newstead et al., 1992; Oakhill et al., 1989; Klauer et al., 2000). We suspect the overarching explanation for this set of phenomena is *attribute substitution* (Kahneman and Fredrick, 2002) – where people unconsciously substitute difficult attributes (e.g., whether the item was seen in a previous argument) with simpler ones (e.g., whether the item is believable). Thus, the observed effect of prior beliefs on perceived argument quality is likely a combination of different cognitive processes: potentially a rational effect of evidence updating in line with Bayesian accounts of argument evaluation (e.g., Hahn & Oaksford, 2007) combined with occasional attribute substitution if arguments are too complex for reasoners to effectively synthesise the evidence contained within them.

We acknowledge that through our experimental design, good and inconsistent arguments may inherently possess varying levels of memorability. For example, some inconsistent arguments might be more memorable due to their unconventional nature. However, this does not interact with our hypothesis which is about the different memorability of belief-congruent vs belief-incongruent arguments regardless of argument quality. Thus, even if inconsistent arguments are more memorable than good arguments, we can still compare if belief-congruent arguments led to more hits and false alarms compared to belief-incongruent arguments for good and inconsistent arguments respectively. As our results showed, this pattern was only observed, to a relatively small degree, for the inconsistent arguments.

References

Alba, J. W., & Hasher, L. (1983). Is memory schematic? *Psychological Bulletin*, *93*(2), 203–231. https://doi.org/10.1037/0033-2909.93.2.203

Brod, G., Werkle-Bergner, M., & Shing, Y. L. (2013). The Influence of Prior Knowledge on Memory: A Developmental Cognitive Neuroscience Perspective. *Frontiers in Behavioral Neuroscience*, *7*. https://doi.org/10.3389/fnbeh.2013.00139

Deans-Browne, C. C. J. L., Băitanu, A., Dubinska, Y., & Singmann, H. (2024). Inconsistent Arguments are Perceived as Better Than Appeals to Authority: An Extension of the Everyday Belief Bias. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *46*. https://escholarship.org/uc/item/1w3885q9

Deans-Browne, C., & Singmann, H. (2025, April 25). For Everyday Arguments Prior Beliefs Play a Larger Role on Perceived Quality Than Argument Quality Itself. https://doi.org/10.31234/osf.io/x6ehu_v2

Evans, J. St. B. T., Barston, J. L., & Pollard, P. (1983). On the conflict between logic and belief in syllogistic reasoning. *Memory & Cognition*, *11*(3), 295–306. https://doi.org/10.3758/BF03196976

Hahn, U., & Oaksford, M. (2007). The rationality of informal argumentation: A Bayesian approach to reasoning fallacies. *Psychological Review*, *114*(3), 704–732. https://doi.org/10.1037/0033-295X.114.3.704

Heit, E. (1993). Modeling the Effects of Expectations on Recognition Memory. *Psychological Science*, *4*(4), 244–252. https://doi.org/10.1111/j.1467-9280.1993.tb00268.x

Kahneman, D., & Frederick, S. (2002). Representativeness revisited: Attribute substitution in intuitive judgment. *Heuristics and biases: The psychology of intuitive judgment*, *49*(49–81), 74.

Klauer, K. C., Musch, J., & Naumer, B. (2000). On belief bias in syllogistic reasoning. *Psychological Review*, *107*(4), 852–884. https://doi.org/10.1037/0033-295X.107.4.852

McCrudden, M. T., Barnes, A., McTigue, E. M., Welch, C., & MacDonald, E. (2017). The effect of perspective-taking on reasoning about strong and weak belief-relevant arguments. *Thinking & Reasoning*, *23*(2), 115–133. https://doi.org/10.1080/13546783.2016.1234411

Newstead, S. E., Pollard, P., Evans, J. St. B. T., & Allen, J. L. (1992). The source of belief bias effects in syllogistic reasoning. *Cognition*, *45*(3), 257–284. https://doi.org/10.1016/0010-0277(92)90019-E

Oakhill, J., Johnson-Laird, P. N., & Garnham, A. (1989). Believability and syllogistic reasoning. *Cognition*, *31*(2), 117–140. https://doi.org/10.1016/0010-0277(89)90020-6

Stangor, C., & McMillan, D. (1992). Memory for expectancy-congruent and expectancy-incongruent information: A review of the social and social developmental literatures. *Psychological Bulletin*, *111*(1), 42–61. https://doi.org/10.1037/0033-2909.111.1.42

Stanovich, K. E., & West, R. F. (1997). Reasoning independently of prior belief and individual differences in actively open-minded thinking. *Journal of Educational Psychology*, *89*(2), 342–357. https://doi.org/10.1037/0022-0663.89.2.342

Taber, C. S., & Lodge, M. (2006). Motivated Skepticism in the Evaluation of Political Beliefs. *American Journal of Political Science*, *50*(3), 755–769. https://doi.org/10.1111/j.1540-5907.2006.00214.x

Wiley, J. (2005). A fair and balanced look at the news: What affects memory for controversial arguments? *Journal of Memory and Language*, *53*(1), 95–109. https://doi.org/10.1016/j.jml.2005.02.001