

The Role of Affective Involvement and Knowledge in Processing Mixed Evidence for Social Issues

Megan D Bardolph (mbardolph@ucsd.edu)

Department of Cognitive Science (0515), 9500 Gilman Drive
La Jolla, CA 92093 USA

Seana Coulson (scoulson@ucsd.edu)

Department of Cognitive Science (0515), 9500 Gilman Drive
La Jolla, CA 92093 USA

Abstract

Exposure to mixed evidence can lead to polarization, or adopting a more extreme version of one's initial attitude. One potential reason for this is attitude congruency bias, rating evidence that supports one's attitude as stronger than evidence that undermines it. Here we explore factors associated with this bias and their relationship to attitude change following exposure to mixed evidence. We conducted several tests, including an attitude survey on two controversial social issues, a poll regarding participants' affective involvement in each issue, an argument rating task, and assessments of knowledge about social issues and political sophistication. We replicated the attitude congruency bias. Ratings bias was associated with affective involvement, but not with measures of topic knowledge or political sophistication. Attitude change was predicted by a linear combination of objective argument strength and rating bias. Participants' sensitivity to objective argument strength suggests the attitude congruency bias does not inevitably lead to polarization.

Keywords: decision making; reasoning; motivated reasoning; rationality; language and thought; attitude congruency bias

Introduction

Over time, people form attitudes about objects, people, and issues. These attitudes may change through any number of affective, cognitive, and/or behavioral mechanisms (see Petty & Wegener, 1998 for an account of the Elaboration Likelihood Model of persuasion). New information can lead to attitude change, but often that information is judged in light of extant attitudes (prior opinions).

Judging information differentially in consideration of its agreement with prior opinions may be natural and beneficial in environments where one's beliefs are true most of the time (Alloy & Tabachnik, 1984). One study showed that scientists judged research as being of higher quality when its conclusions were in agreement with their own prior opinions (Koehler, 1993). Researchers rejected the quality of studies based only on their outcomes, not the merits of the design. This can be viewed as rational when findings are inconsistent with a body of scientific knowledge, but outside of established fields could clearly present difficulties in reaching consensus.

Fortunately, polarization, where exposure to the same evidence leads people to opposite attitude adjustments in the direction of their prior attitudes, is relatively rare (Kahan, Peters, Dawson, & Slovic, 2017), although repeated

discussion within a group with a strong sense of shared identity can lead individuals to hold more extreme attitudes (Sunstein, 2002). The domain of social and political issues may be one such special case, since in many areas there can arise "two sides" of a seemingly factual issue, with supporters on each side failing to change their beliefs toward an objective consensus (Kahan, 2016). For example, in a survey of American voters, 75% of self-identified liberals believed that climate change was due primarily to human activity, whereas only 45% of conservatives shared this belief (McCright & Dunlap, 2011).

One explanation for polarization on social issues is motivated reasoning (Ditto & Lopez, 1992; Kunda, 1990). On such accounts, attitude polarization occurs because people with opposing views draw opposite conclusions from the very same evidence. In a classic study, Lord, Ross, and Lepper (1979) queried participants about their views on capital punishment, and then presented them with the results of two studies, one that suggested the death penalty deters crime, and one that suggested the opposite conclusion. Participants were asked to rate the quality of each study, and then to re-characterize their views on the death penalty. Interestingly, participants tended to rate the study that supported their own beliefs as being objectively better than the one that undermined them, and each group adjusted their beliefs to more strongly favor their original position (Lord, Ross, & Lepper, 1979).

Subsequent studies have shown that exposure to mixed evidence can lead people to polarize, changing their beliefs to be more in line with their initial attitudes (Edwards & Smith, 1996; Taber & Lodge, 2006). One explanation for this change in belief is that people accept information congruent with their extant opinions without critical examination, while incongruent information is critically examined and judged more negatively in the presence of negative affect (Taber, Cann, & Cucusova, 2009).

If the process by which the attitude congruency bias leads to polarization does involve emotional processing of evidence, then people with emotional commitment to their attitudes may be more likely to display the bias and also to polarize. If, however, knowledge about a topic influences the way evidence is processed, we may see topic knowledge influencing belief change.

We first examine which variables are related to attitude congruency bias, specifically whether affective involvement leads to increased bias or whether cognitive factors of knowledge about the topic or political sophistication can predict argument ratings. Next, we examine the role of attitude congruency bias in attitude change and explore the contribution of other factors. Findings are interpreted in light of their consistency with accounts of motivated reasoning.

Methods

Participants

Participants were undergraduate students enrolled in Psychology, Linguistics, or Cognitive science courses at the University of California, San Diego (UCSD) ($n=141$, 99 female) participating as part of a course requirement. Participants ranged from 18 to 29 years of age (mean = 20). All participants provided informed consent, and procedures were approved by the Institutional Review Board (IRB) at UCSD.

Procedure

After consenting to participate in the study, participants first completed Initial Attitude and Affective involvement measurements for each issue (described in the Materials section). One of the issues was randomly assigned to the Mix condition, meaning that the participant would read arguments for and against the position articulated in the issue statement (henceforth: the issue). The other issue was assigned to the Control condition. Participants were not exposed to any arguments regarding the issue in the Control condition.

During the treatment phase of the study, each participant read 3 Pro and 3 Con arguments regarding the issue in the Mix condition. Arguments were presented in a random order. Following the presentation of each argument, participants used a 100-point slider to indicate the argument's strength on a scale from Weak to Strong (numbers not visible). After half of the arguments, participants were asked to describe their thoughts about the argument via a typed response in a text box.

After the treatment phase, participants again completed the attitude measurement survey for both issues to determine their Post-treatment Attitude scores. Next, participants completed a Topic Knowledge test for each issue and a brief political knowledge quiz to assess their political sophistication. Finally, they viewed a debriefing page that explained the goal of the study and provided links to the websites used for the argument texts.

Materials

The survey used for the present study contains a subset of materials used in a previous study (Bardolph & Coulson, 2017). Two socio-political issues were included: animal testing and the death penalty, selected from the most popular topics on two debate websites, www.procon.org and idebate.org. Text from both sides of debate arguments was

used to create one-paragraph arguments that either supported or opposed the related issue.

Attitude measurement For both issues, participant attitude was measured using 5 survey questions: A single policy statement (“Animal testing should be banned”; “The death penalty should be illegal”) with a rating slider from Disagree to Agree (0 to 100, numbers not visible), followed by four position statements for each issue. These position statements were selected from “Points for” and “Points against” on the idebate.org archive (e.g., “Animals involved in animal research are mostly well treated.”). Each position statement was rated using a 9-point scale of agreement/disagreement. Ratings from the policy statement and the four position statements were scaled and combined to form an average initial Attitude, ranging from -5 (most opposed to the issue) to 5 (most in favor of the issue).

After the experimental treatment, participants responded again to the same five statements for each issue. Responses were combined as before to form an average post-treatment attitude score.

Affective involvement For each issue, affective involvement was measured using 4 survey questions with a 9-point rating scale indicating: how much participants care about the issue, how strong their feelings are, how certain they are of their feelings, and how much they have thought about the issue. These four measurements were combined to form a measure of affective involvement.

Arguments Six supporting (Pro) and six opposing (Con) arguments were selected using text from the debate sites for each issue. Arguments were generally matched for content (i.e., if a Pro and a Con argument addressed the same point, both arguments were usually selected), and for length (mean argument length = 120 words, $sd = 11$). To create arguments of similar length, portions of longer arguments were edited. A study of ratings for these arguments that drew from the same participant pool indicated slightly higher ratings of arguments regarding animal testing than the death penalty, but revealed similar ratings for participants supporting and opposed to the position statement for each issue (Supporting: mean = 61.5 and Opposed: mean = 62.6 for animal testing arguments, and Supporting: mean = 54.4 and Opposed: mean = 55.8 for death penalty arguments).

Topic knowledge For each issue, topic knowledge was measured using eight multiple choice factual questions (e.g., “Which animal is used most frequently for research?”; “The death penalty was ruled to be constitutional in the US under which amendment?”). These questions were piloted in an earlier norming study. Items that were too easy or too difficult were not included in the present study. Topic knowledge for each issue is represented by the percentage of questions the participant answered correctly.

Analysis

Initial attitude and Post treatment attitude were scaled from -5 to 5, representing, respectively, the opinion most opposed to each issue, and most in favor of the issue. In this coding, a positive score represents an attitude in favor of the legality of animal testing and the legality of the death penalty. A measure of attitude change was created by subtracting each participant's Prior attitude from their Post treatment attitude. Consequently, Attitude change could range (in principle) from -10 to 10. Affective involvement ratings for each participant for each issue ranged from 1 to 9 (least to most strong). Items for which participants spent less than 3 standard deviations below the median log reading time or more than 3 standard deviations above the median log reading time were removed (9 items out of 846).

Argument ratings A linear mixed effects regression (LMER) model was used to analyze argument rating data. Models were constructed with the lme4 package in R (Bates, Maechler, Bolker, Walker, et al., 2014 R Core Team, 2015). Analysis involved construction of an LMER model to predict argument ratings and the use of backward model comparison using ANOVA. Models were fit with random intercepts for participants and for arguments (items). The use of random intercepts helps control for individual variability in participants' use of the rating scale, as well as for differences in the quality of particular arguments (some arguments are intrinsically better and consequently tend to be rated as stronger by all participants). Backward model comparison yielded the most parsimonious model that included all significant predictors ($p < .01$ used as cutoff).

Linear models Predictive relationships among experimental variables, including Attitude change, were analyzed with a linear model in R. Analysis involved backward model comparison using ANOVA to establish the optimal model. This is roughly equivalent to selecting all predictors below a threshold p value in the model ANOVA.

Objective argument rating To obtain an approximately objective rating of argument strength, we used ratings from two prior experiments where participants viewed the same arguments in a mixed condition (Pro and Con arguments presented together). $N=39$ for animal testing, $N=48$ for death penalty, with participants' opinion approximately evenly distributed for each issue (prior opinion ranged from -4.1 to 3.9, mean = -0.2 for animal testing; -4.7 to 5.0, mean = 0.1 for death penalty).

Ratings from the two prior experiments and the current study did not differ significantly by study. Argument label was used to predict an average rating of argument strength for each argument. These predicted values from the two prior studies were used as a measure of argument quality. In the model, this factor is referred to as Argument.

Rating Bias To create a single measure of bias for each participant, each argument's objective rating was subtracted from the participant's rating. For example, consider a participant who gave a Pro argument a rating of 90 (very strong). If this argument's objective rating were 65, the residual rating would be 25, indicating that this participant is biased toward the Pro position.

An average rating bias was created by subtracting each participant's bias in favor of Con arguments from bias in favor of Pro arguments. Using this scale, bias reflects how much stronger participants rated Pro arguments than Con arguments for a given issue. For example, if a participant were on average biased by 25 points in favor of Pro arguments and 5 points against Con arguments (average bias = -5 for Con arguments), their Bias score would be $(25 - -5) = 30$, indicating that they rate Pro arguments more highly than Con arguments. A participant who rated Con arguments more highly than Pro arguments would receive a negative Bias score. Bias scores were transformed to a z -score variable for modeling. A measure of Folded bias, the absolute value of the bias term, was also used when correlating the magnitude of bias with other variables.

Correcting Post-survey attitude regression to the mean

An overall regression toward the mean was present in both the Experimental and Control conditions. On average, participants' change from their Prior to Post-survey attitude was toward the center of the attitude scale. This means that participants in favor of an issue changed their opinion to be more opposed to that issue and vice versa, even when they do not read any arguments. For example, one participant in favor of animal testing reported a prior attitude of 3.59 and a post-survey attitude of 2.97 for animal testing. Although this participant was not exposed to any arguments about animal testing, their attitude moved toward the center of the attitude scale. We refer to such "changes" in attitude as regression to the mean.

To correct for regression to the mean, we used attitude scores in the Control condition to determine a correction factor that could be applied to the Treatment condition. Beginning with data from the Control condition in which participants read no arguments, we used linear regression to predict post-survey attitude change from their initial attitude reports. This average slope (-0.17) was subtracted from attitude change measures in both the Control and the Treatment conditions, effectively correcting for regression to the mean. For example, the attitude change score for the participant described above was corrected from -0.62 to 0.06.

Results

Histograms of participant variables used as predictors are shown in Figure 1. The frequency of values of the predictor is shown for Affective involvement, Topic knowledge, Political sophistication, and Rating Bias (z -scored). Table 1 shows a correlation matrix for these four variables. The only correlation trending toward significance is the correlation of Affective involvement and Topic knowledge ($p = .07$).




Figure 1: Histograms showing the frequency of (a) Affective involvement: two measures per participant, one for each issue; (b) Topic knowledge: two measures per participant, one for each issue; (c) Political sophistication: one measure per participant; and (d) Bias: one measure per participant, calculated for the Treatment condition and transformed to a z-score variable.

Argument rating

This analysis tests for the attitude congruency effect in participants' ratings of arguments regarding animal testing and the death penalty by exploring how each participant's rating for a given argument varies as a function of their prior opinion on the issue (opposed to supportive), and argument polarity, that is, whether the argument itself was Pro or Con. On such an analysis, attitude congruency bias is revealed by a cross-over interaction of these factors, as participants rank arguments congruent with their prior opinions as stronger than arguments that are incongruent with those opinions. Further, to see if attitude congruency was related to participants' affective involvement with the issue, their political sophistication, or their degree of extant knowledge regarding the topic, we included these factors as additional predictors in the model.

Participants' argument ratings were analyzed with an LMER model (as described in the Analysis section). The initial model included predictors of Prior opinion, Affective involvement, Argument polarity (Pro/Con), Issue (Animal testing, Death penalty), Political sophistication, and Topic knowledge, with random intercepts for participants and for arguments. More complex models were compared to models with fewer predictors using model ANOVA, yielding the most parsimonious model that still contained all significant predictors of Argument rating (using $p < .01$ cutoff).

$$\text{Argument rating} \sim \text{Prior opinion} * \text{Argument polarity} (1)$$

Table 1: Correlations of predictor variables.

	Affective Involvement	Topic Knowledge	Political Sophistication
Topic knowledge	0.11		
Political sophistication	-0.06	0.03	
Bias	-0.09	-0.12	0.06

The model that best predicts argument rating is shown in Equation 1.¹ Results of the model are listed in Table 2. The relationship between Prior opinion and Argument polarity is shown in Figure 2. The predicted cross-over interaction reflects the fact that participants who were extreme supporters (5 on the Prior opinion axis) rated attitude-congruent Pro arguments as stronger than incongruent Con arguments. Similarly, extreme opponents (-5 on the Prior opinion axis) rated the attitude congruent Con arguments as stronger than incongruent Pro arguments. However, while the data suggest argument ratings were indeed subject to attitude congruency bias, we failed to detect a relationship between argument ratings and any of our other measures, including affective involvement, topic knowledge, or political sophistication. Affective involvement was correlated with Prior opinion ($R = 0.53$, $p < .001$), but did not have an additional effect on Argument rating.




Figure 2: Interaction of argument polarity and prior opinion (-5 most opposed, 5 most in favor of the issue). Circles represent individual argument ratings. Green and red lines represent average rating of Pro and Con arguments respectively.

¹ This same analysis can be performed using Position and Folded prior opinion (magnitude). The results are the same as Equation 1, and the two models of Argument rating are not statistically different.

Table 2: Model results for Equation 1.

Factor	df	F value
Prior opinion	1	4.7
Argument polarity	1	1.3
Prior opinion x Argument polarity	1	86.9

Attitude change

Apart from the cause of attitude congruency bias, another question of interest concerns its effects on attitude change. Were participants who displayed attitude congruency bias in their argument ratings more or less likely to change their opinions? On some accounts, biased assimilation of the evidence can have an undue effect on belief change and lead to attitude polarization. To explore the impact of biased assimilation on belief updating, we used linear models to test whether there was a relationship between attitude change and rating bias.

If biased ratings lead to polarizing, we expect a positive relationship bias and attitude change: that is, positive Bias, rating Pro arguments more highly, will lead to positive Attitude change (more supportive of the issue). A negative Bias, rating Con arguments more highly, will lead to negative Attitude change (more opposed to the issue).

Participants' argument ratings, however, do not reflect bias alone. Each argument may have a degree of strength relative to other arguments, or an objective quality. An individual's rating, therefore, may reflect the objective argument strength and individual bias. For this reason, argument ratings were split into a measure of objective argument quality and individual bias.

Individual bias (labeled Bias) was calculated as described in the Methods section and used to predict corrected Attitude change in a linear model. This model tests whether biased assimilation leads to polarizing. Objective argument rating was included as an additional, separate predictor (Argument). Further, to assess the impact of affective involvement, topic knowledge, and political sophistication on Attitude change, we included these factors as additional predictors.

Analysis involved construction of a linear model to predict Attitude change and the use of backward model comparison to establish the optimal model. The initial linear model included factors of Bias, Argument, Affective involvement, Issue (Animal testing, Death penalty), Political sophistication, and Topic knowledge. Nested linear models were compared using model ANOVA in R as described in the Methods section.

The model that best predicts Attitude change is shown in Equation 3. There was a significant main effect of Argument ($p < .005$), and a main effect of Bias ($p = .013$). The coefficient for both predictors was positive: Argument ratings in favor of Pro arguments predict opinion change in support of the issue, Bias in favor of Pro arguments predicts opinion change in support of the issue, and vice versa for Argument/Bias in favor of Con arguments. No significant

effects were found for other measures, including affective involvement, topic knowledge, or political sophistication.

$$\text{Attitude change} \sim \text{Argument} + \text{Bias} \quad (3)$$

Table 3: Model results for Equation 3.

Factor	Estimate	df	F value	P value
Argument	0.27	1	10.9	< .005
Bias	0.20	1	6.3	.013

Discussion

Here we explored the importance of topic knowledge and political sophistication on the one hand and affective involvement on the other to different phenomena related to reasoning about controversial social issues.

Our initial analyses explored the role of these factors in how participants evaluate arguments that are congruent vs. incongruent with their prior attitudes on the issue. We found attitude congruency bias, but no evidence for contribution of knowledge, political sophistication, or affective involvement as moderators of this phenomenon. Replication of the attitude congruency bias is consistent with previous findings (Lord, Ross, & Lepper, 1979; Edwards & Smith, 1996; Taber, Cann, & Cusova, 2009; Bardolph & Coulson, 2017) and consistent with accounts of motivated reasoning. The present study does not indicate that this bias is related to how knowledgeable individuals are about the topic under discussion. Because prior attitudes are highly correlated with affective involvement, the degree to which individuals care about an issue may contribute to their bias, although a precise relationship cannot be established by these data.

Further, we explored the role of topic knowledge, political sophistication, and affective involvement in ratings bias. Although we found no relationship between either topic knowledge or political sophistication in participants' degree of ratings bias, we did find a positive association between bias and affective involvement. The more affectively involved participants were with a given issue, the more biased their argument ratings were. These data are in keeping with motivated reasoning accounts.

Finally, we explored the relative importance of ratings bias, objective argument quality, affective involvement, topic knowledge, and political sophistication for attitude change. Of these factors, only objective argument quality and ratings bias were significant predictors of attitude change. While the relationship between ratings bias and attitude change is in line with motivated reasoning, our models suggest objective argument quality is a slightly better predictor of attitude change. The latter finding indicates participants were sensitive to the quality of the evidence, changing their opinions more when they were exposed to strong arguments than when they were exposed to weak ones.

These data argue against prior studies that suggest the attitude congruency bias demonstrated in the present study is likely to lead to the polarization of opinions (Taber, Cann, & Cucusova, 2009). Attitude change in the present study was in fact more influenced by the objective quality of the arguments than the participants' ratings bias. The present study suggests that while people are more skeptical of evidence that contradicts their existing attitudes, they are also sensitive to the quality of that evidence.

One limitation of the present study was the use of a highly-educated sample from a leading public university in the United States. The behavior of these student participants may not generalize to a larger sample. It is also possible that measurements of participants' attitudes do not reflect a single, stable opinion, but a combination of response instability and multiple response effects (see Zaller, 1992 for a model of survey response). Our method of correcting for regression to the mean addresses some of this variability, but there may be aspects of participants' opinions that are not fully captured by the survey methods.

Overall, these data are consistent with accounts of motivated reasoning, replicating the phenomenon of attitude congruency bias and revealing a relationship between bias and affective involvement in a controversial social issue. However, they also reveal participants' rational sensitivity to the quality of the evidence with which they are presented. This sensitivity to argument quality could potentially mitigate attitude polarization. It also highlights the possible impact of exposure to media of varying quality: although consumers of media may indeed be biased by their own attitudes, persuasive arguments from quality sources may have an impact on attitude change.

Acknowledgments

This research was supported in part by a grant from the Frontiers of Innovation Scholars Program (FISP) at UC San Diego.

References

Alloy, L. B., & Tabachnik, N. (1984). Assessment of covariation by humans and animals: The joint influence of prior expectations and current situational information. *Psychological review*, 91.

Bardolph, M. D. & Coulson, S. (2017). Belief updating and argument evaluation. In G. Gunzelmann, A. Howes, T. Tenbrink, & E. J. Davelaar (Eds.), *Proceedings of the 39th Annual Conference of the Cognitive Science Society* (pp. 1586-1591). London, UK: Cognitive Science Society.

Bates, D., Maechler, M., Bolker, B., Walker, S., et al. (2014). lme4: Linear mixed-effects models using eigen and s4. *R package version*, 1(7).

Ditto, P.H., & Lopez, D.F. (1992). Motivated skepticism: Use of differential decision criteria for preferred and non-preferred conclusions. *Journal of Personality and Social Psychology*, 63(4).

Edwards, K., & Smith, E. E. (1996). A disconfirmation bias in the evaluation of arguments. *Journal of Personality and Social Psychology*, 71(1).

Kahan, D. M. (2016). The politically motivated reasoning paradigm, part 1: What politically motivated reasoning is and how to measure it. *Emerging Trends in the Social and Behavioral Sciences: An Interdisciplinary, Searchable, and Linkable Resource*.

Kahan, D. M., Peters, E., Dawson, E. C., & Slovic, P. (2017). Motivated numeracy and enlightened self-government. *Behavioural Public Policy*, 1(1).

Koehler, J. J. (1993). The influence of prior beliefs on scientific judgments of evidence quality. *Organizational behavior and human decision processes*, 56(1).

Kunda, Z. (1990). The case for motivated reasoning. *Psychological bulletin*, 108(3), 480.

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).

McCright, A. M., & Dunlap, R. E. (2011). Cool dudes: The denial of climate change among conservative white males in the United States. *Global environmental change*, 21(4).

Petty, R. E., & Wegener, D. T. (1998). Attitude change: Multiple roles for persuasion variables (In D. Gilbert, S. Fiske, & G. Lindzey (Eds.). *The handbook of social psychology* (Vol. 1).

R Core Team. (2015). R: A language and environment for statistical computing [Computer software manual]. Vienna, Austria.

Sunstein, C. R. (2002). The law of group polarization. *Journal of political philosophy*, 10(2).

Taber, C. S., Cann, D., & Kucusova, S. (2009). The motivated processing of political arguments. *Political Behavior*, 31(2).

Taber, C. S., & Lodge, M. (2006). Motivated skepticism in the evaluation of political beliefs. *American Journal of Political Science*, 50(3).

Zaller, J., & Feldman, S. (1992). A simple theory of the survey response: Answering questions versus revealing preferences. *American journal of political science*.