Research Article

Psychological Science
2021, Vol. 32(6) 916–933

DOI: 10.1177/0956797621995208
www.psychologicalscience.org/PS
SAGE

The Effect of Prediction Error on Belief Update Across the Political Spectrum

Madalina Vlasceanu[1,2], Michael J. Morais[2,3], and Alin Coman[1,4]
[1]Department of Psychology, Princeton University; [2]Princeton Neuroscience Institute, Princeton University; [3]Center for Statistics and Machine Learning, Princeton University; and [4]Princeton School of Public and International Affairs, Princeton University

Abstract

Making predictions is an adaptive feature of the cognitive system, as prediction errors are used to adjust the knowledge they stemmed from. Here, we investigated the effect of prediction errors on belief update in an ideological context. In Study 1, 704 Cloud Research participants first evaluated a set of beliefs and then either made predictions about evidence associated with the beliefs and received feedback or were just presented with the evidence. Finally, they reevaluated the initial beliefs. Study 2, which involved a U.S. Census–matched sample of 1,073 Cloud Research participants, was a replication of Study 1. We found that the size of prediction errors linearly predicts belief update and that making large errors leads to more belief update than does not engaging in prediction. Importantly, the effects held for both Democrats and Republicans across all belief types (Democratic, Republican, neutral). We discuss these findings in the context of the misinformation epidemic.

Keywords

belief update, belief change, prediction error, ideological beliefs, misinformation, open data, open materials, preregistered

Received 2/26/20; Revision accepted 11/19/20

> A wise man proportions his belief to the evidence.
> —David Hume (p. 73)

Political scientist Larry Bartels observed that "the political ignorance of the American voter is one of the best-documented features of contemporary politics" (Bartels, 1996, p. 194). A burgeoning literature across the social sciences has been dedicated to developing strategies to address this notorious limitation. Kuklinski and colleagues (2000) identified two conditions that are needed to assuage this problem: increased access to objective facts and their incorporation in individuals' mental models. The first condition is difficult to satisfy given that nearly half of Americans get their news from Facebook (Shearer & Gottfried, 2017), a social media platform known for providing access to a vast volume of misinformation (Shu et al., 2017). However, even if such organizations successfully implement strategies to diminish misinformation, a more daunting challenge arises: persuading people to incorporate these facts into their belief systems. Findings from social psychology hint that new facts are easily dismissed if they increase cognitive dissonance (Festinger & Carlsmith, 1959), reduce coherence among already held beliefs (Lord et al., 1979), or counter one's political allegiance (Nyhan & Reifler, 2010). In the present studies, we were interested in exploring cognitive processes that could facilitate the incorporation of facts into people's belief systems.

A central feature of beliefs—defined as statements that individuals hold to be true (Schwitzgebel, 2010)—is their dynamic nature: Beliefs are susceptible to change (Bendixen, 2002). Prior work has identified several strategies that proved effective at changing beliefs, such as using fictional narratives (Wheeler et al., 1999), manipulating memory accessibility (Vlasceanu & Coman,

Corresponding Author:
Madalina Vlasceanu, Princeton University, Department of Psychology
E-mail: mov@princeton.edu

2018), associating beliefs with emotionally arousing images (Vlasceanu, Goebel, & Coman, 2020), emphasizing normativity cues (Vlasceanu & Coman, 2020a), and nudging accuracy goals (Pennycook et al., 2020). Here, we propose that one powerful strategy to facilitate belief change might involve updating mental models through prediction errors. This conjecture builds on the seminal finding that learning is proportional to prediction error, where prediction error is the difference between the prediction one makes about a state of the world and the actual outcome (Rescorla & Wagner, 1972). Because expectations are based on the agent's model of the world, when predictions are validated, they reinforce the model of the world they stemmed from, and when they are invalidated, the model gets updated accordingly (den Ouden et al., 2012). Generating predictions is, arguably, a ubiquitous process implemented by the cognitive system that has adaptive consequences for the organism (Bar, 2009).

In the current investigation, we were interested in whether prediction errors may have a similar effect on belief change and whether this effect might be modulated by motivational factors that involve political ideology. Is the influence of prediction errors on belief update a general process, or are there partisan biases in the way prediction errors impact beliefs? On the one hand, in support of prediction errors as a general mechanism, they have been found to impact a wide range of cognitive processes, including perception (de Lange et al., 2018), action (Bestmann et al., 2008), memory (Erickson & Desimone, 1999), language (Kutas & Hillyard, 1980), cognitive control (Alexander & Brown, 2011), and decision-making (Greve et al., 2017). On the other hand, past research shows that cognitive processing can vary as a function of political ideology (Buechner et al., 2021). For instance, motivations to reach particular conclusions have been shown to affect information processing (Nyhan & Reifler, 2010). This suggests that there might be meaningful motivational differences between liberals and conservatives that could affect the relation between prediction errors and belief update (Ditto et al., 2019; Haidt et al., 2009). One way in which ideological biases could influence the belief-updating process may involve a reduced susceptibility to changing ideologically consistent beliefs as a function of prediction errors (as opposed to ideologically inconsistent or neutral beliefs). That is, people might be entrenched with respect to their party's beliefs but flexibly updating other beliefs (Toner et al., 2013). Another way in which ideological biases might impact belief update could involve a differentiation between liberals and conservatives; conservatives might be more resistant to change than liberals, as has already been shown (Jost et al., 2003; White et al., 2020). This would predict that conservatives might be less likely than liberals to change their beliefs according to the prediction errors they make, regardless of the ideological nature of those beliefs. Yet another possibility is that belief updating could be dynamically dependent on environmental factors involving uncertainty and political-identity threat (Haas & Cunningham, 2014). The more one is uncertain and threatened, the more resistant one is to changing one's beliefs.

Statement of Relevance

Misinformation spread is among the top threats facing the world today. The current unprecedented level of exposure to false information leads people to confidently hold false beliefs. As a result, policymakers face an important challenge to design campaigns guided by empirical research to combat and prevent misinformation. One of the first steps in compiling empirically grounded recommendations is large-scale testing of interventions aimed at changing ideologically charged false beliefs. In this study, we reveal such an intervention: Engaging in prediction regarding surprising belief-related evidence and making large errors followed by immediate feedback led to the successful updating of the corresponding beliefs. This effect held across ideological boundaries, making it a viable strategy for reducing ideologically charged misinformation.

We explored the relationship between prediction error and belief update in an experiment in which participants either passively viewed or actively made predictions based on belief-associated statistical evidence. We hypothesized that there would be a positive linear relationship between prediction-error size and belief update. We also hypothesized that making large prediction errors would lead to more evidence incorporation and belief change than not engaging in prediction. We did not have a priori hypotheses regarding how participant and belief ideology would interact with the effect of prediction error on belief update.

Study 1

Method

Open-science practices. We preregistered the study's experimental design and hypotheses on AsPredicted (https://aspredicted.org/zu4iq.pdf). In addition, the stimuli, pilot-study results, and data for the main study can be found on the study's OSF page (https://osf.io/aur2t). The data-analysis code (in Python) can be accessed as a Jupyter notebook at https://github.com/mvlasceanu/PredictionBelief.

Participants. We estimated that to obtain a power of .80 to detect a moderate effect size (Cohen's *d*) of 0.3 for two between-subjects comparisons, a sample of 704 participants would be needed. Participants were recruited on Cloud Research (https://www.cloudresearch.com/), an Internet-based research platform similar to Amazon Mechanical Turk but with more intensive participant-pool checks. Participants were compensated at the platform's standard rate (Litman et al., 2017). In total, we recruited 945 participants, of which 241 were excluded from the analysis on the basis of preregistered criteria (i.e., attention checks). We stopped data collection as soon as we reached the preregistered sample size of 704 valid participants (age: $M = 50.32$ years, $SD = 16.51$; 67.7% women). Of these, 352 participants self-identified as Democrats and 352 as Republicans. Each Democrat was assigned to either the experimental condition ($n = 176$) or the control condition ($n = 176$), and the Republicans were also evenly split between the two conditions. The study protocol was approved by the Princeton University Institutional Review Board.

Stimulus materials. We undertook preliminary studies to develop a set of 36 statements (see https://osf.io/nry4z/). These statements were equally split into 12 neutral statements (e.g., "Shark attack rates are similar for men and women"), 12 Democratic statements (e.g., "The US has loose gun laws"), and 12 Republican statements (e.g., "A large proportion of immigrants in the US is not in the workforce"). The 36-statement set was selected from a larger initial set of 48 statements that we pretested on an independent sample of Cloud Research participants ($N = 50$; age: $M = 41.94$ years, $SD = 15.83$; 62% women). In the pilot study, we first measured the believability of each statement with the question "How accurate do you think this statement is?" on a scale ranging from 0, *extremely inaccurate*, to 100, *extremely accurate*. We selected the final set of 36 statements so that each neutral statement was equally believed by the Democratic and Republican participants, each Democratic statement was believed significantly more by Democrats than by Republicans, and each Republican statement was believed more by Republicans than by Democrats. Overall, the neutral statements were equally endorsed by Democrats ($M = 62.83$, $SD = 9.35$) and Republicans ($M = 59.58$, $SD = 15.76$), $p = .546$; the Democratic statements were endorsed significantly more by Democrats ($M = 70.78$, $SD = 7.89$) than by Republicans ($M = 49.68$, $SD = 16.15$), $p < .001$; and the Republican statements were endorsed significantly more by Republicans ($M = 61.74$, $SD = 12.37$) than by Democrats ($M = 48.36$, $SD = 10.85$), $p < .01$ (see https://osf.io/nry4z/).

We also developed a set of 36 facts that provide evidence either in support or against the 36 statements (see https://osf.io/nry4z/). For example, for the statement "Very few Americans identify as vegetarian," the evidence in support was "5% of Americans are vegetarian," and for the statement "Many American adults exercise on a daily basis," the evidence against was "5% of Americans participate in 30 minutes of physical activity every day." These factual statistics were selected from a larger set of 48 accurate facts that we found in scientific articles or official polls and were pretested on the same sample of participants as in the pilot study. Statistics were selected to match on how strongly each piece of evidence would influence each associated statement (e.g., "How likely is this piece of evidence to influence your support for this statement?" on a scale from 1, *not at all*, to 5, *a great deal*). The 36 facts were selected so that for Democrats, the neutral ($M = 3.05$, $SD = 0.54$), Democratic ($M = 3.24$, $SD = 0.54$), and Republican ($M = 3.05$, $SD = 0.25$) facts did not significantly differ on how strongly Democrats thought they influence the statements; likewise, for Republicans, the neutral ($M = 3.56$, $SD = 0.41$), Democratic ($M = 3.57$, $SD = 0.21$), and Republican ($M = 3.43$, $SD = 0.42$) facts did not significantly differ from each other on the evidence-strength dimension (see https://osf.io/nry4z/).

In addition, a set of 36 scale-based estimation questions was constructed to be used as part of the evaluation phase. These questions were created by rephrasing the facts constructed as evidence. For example, for the fact "In the US, 3 of the 50 states require a permit to purchase a rifle," the corresponding question was "How many of the 50 states require a permit to purchase a rifle?" Each question had 12 potential answers, linearly increasing on the 12-item scale, and one of the 12 was correct. Across the 36 questions, the correct answer had an equal chance of being in any of the 12 scale positions from 1 to 12. This prevented forming probability estimates for the most likely positions on the scale to contain the correct answer.

We measured *resistance to change*, a construct that has been found to differentiate between liberals and conservatives (Jost et al., 2003). The measure was adapted from the Willingness to Compromise Scale (Wee, 2013) and was computed as the average response to the three-item scale ("I would stick to my beliefs even when others might think that they are not reasonable"; "Reality constraints should not stand in the way of one's beliefs"; and "Once I believe in something, no piece of evidence would change my mind"). All questions were rated on a scale from 1, *strongly disagree*, to 5, *strongly agree*. Thus, higher scores indicate more resistance to change, and lower scores indicate less resistance to change.

Finally, we measured participants' strength of identification with their selected political party with the question "How strongly do you identify with the party you just selected?" on a scale from 1, *not at all*, to 5, *a great deal*. We also measured their support for the current president with the question "How would you qualify president Donald Trump's performance in office for the past 3 years?" on a scale from 1, *awful*, to 7, *excellent*. We used both of these questions as measures of political polarity, which is of interest given prior work indicating cognitive failures in people holding radical beliefs (Rollwage et al., 2018) and cognitive inflexibility in people holding extreme partisan identities (Zmigrod et al., 2020). We note that 92.18% of the sample indicated that they are registered to vote for the party they identified with.

Design and procedure. The data for this study were collected between October 10, 2019, and October 14, 2019. Participants were told that they would participate in an experiment about how people evaluate information encountered on the Internet and were directed to the survey on the Qualtrics platform (https://www.qualtrics.com/platform/). After providing informed consent, participants were directed to the pretest phase, where they were instructed to answer questions about information encountered on the Internet, which meant rating a set of 36 statements (one on each page) by indicating the degree to which they believed each statement was accurate (i.e., "How accurate do you think this statement is?" from 1, *extremely inaccurate*, to 100, *extremely accurate*).

Then, in the evidence phase, each participant was randomly assigned to one of two between-subjects conditions: prediction condition and control condition. Participants in the control condition were shown a series of 36 facts that provided direct evidence either in favor of or against the set of 36 beliefs. Instead of simply being exposed to the facts, participants in the prediction condition were asked to predict the correct answers to questions equivalent in content to the 36 facts used in the control condition. After choosing an answer, participants were immediately given feedback (i.e., the correct answer). In both conditions, the evidence was presented one on each page and in a random order. Then, in a posttest believability phase, participants were instructed to rate the believability of the initial 36 statements again. Finally, participants were asked to complete the Resistance to Change Scale and a series of demographic measures including their strength of party affiliation and support for President Trump, after which they were debriefed.

Analysis and coding. We operationalized *rational belief update* as a belief change from the pretest phase to the posttest phase in the direction corresponding to incorporating the available evidence. Critically, whether this update corresponds to increasing or decreasing beliefs depends on (a) counterbalanced features of the stimuli and (b) observed features of participants' predictions. The first counterbalanced feature of the stimuli is that one half of the presented evidence supports half of the beliefs—in this case, the rational update is to increase one's belief from pretest to posttest. The other half of the evidence refutes the other half of the beliefs—in this case, the rational update is to decrease one's belief from pretest to posttest. Using this setup, we ensured that participants could not trivially infer that "correct" updates must necessarily occur in one direction. This is the only variable necessary to compute rational belief update in the control condition. For example, for the belief "Very few Americans identify as vegetarian" with the corresponding supporting piece of evidence "5% of Americans identify as vegetarian," the rational update is to increase believability from pretest to posttest. Conversely, for the belief "Many American adults exercise on a daily basis" with the corresponding piece of evidence arguing against it being "5% of Americans exercise on a daily basis," the rational update is to decrease believability from pretest to posttest.

Rational belief update in the prediction condition has two additional variables that determine in which direction update is rational for each belief: the magnitude of the correct answer (high or low on the scale) and the sign of the prediction error (positive or negative relative to the correct answer). The magnitude of the correct answer was counterbalanced across the stimuli so that the "surprise" was possible in both directions. For example, the answer to the question "How many child deaths worldwide is pneumonia responsible for every year?" is of high magnitude (i.e., 1 million deaths; the alternative answers of lower magnitudes are situated below the correct answer on the scale), whereas the answer to the question "What percentage of people who collapse on the street fully recover from receiving CPR?" is of low magnitude (i.e., 2%; the alternative answers of higher magnitudes are situated above the correct answer on the scale).

The last variable needed to determine the direction of rational update is the prediction-error sign. Prediction error is defined as the difference between the selected answer by the participant and the correct answer for that question. If a participant selects an answer of higher magnitude than the correct answer, then the prediction error for that item will have a positive sign, whereas if they select a lower answer than

Table 1. Direction of Rational Belief Update in the Prediction Condition of Study 1

Evidence type and answer type	Negative prediction error (−1)	Positive prediction error (+1)
Evidence in support (+1)		
Correct answer is high (−1)	Increase (+1)	Decrease (−1)
Correct answer is low (+1)	Decrease (−1)	Increase (+1)
Evidence against (−1)		
Correct answer is high (−1)	Decrease (−1)	Increase (+1)
Correct answer is low (+1)	Increase (+1)	Decrease (−1)

the correct one, the prediction error for that item will have a negative sign. For example, if the participant selects 20% on the question "How many Americans identify as vegetarian?" given that the correct answer is 5%, the prediction error will be positive. If, however, the participant selects 2%, the prediction error will be negative.

These three variables (evidence in support of vs. evidence against, high correct answer vs. low correct answer, and positive prediction error vs. negative prediction error) determined the direction of rational update in the prediction condition (Table 1).

An example of a trial in which the evidence is in support and the correct answer is high is the belief "Pneumonia is dangerous for children" with the associated evidence "Worldwide, 1 million children die of pneumonia each year." In this case, if the participant chooses any value lower than 1 million (prediction error is negative), they will realize that the number of deaths is higher than they thought, which should increase their belief in the danger of pneumonia. Thus, the rational update is to increase the belief that "pneumonia is dangerous for children" (Table 1). If, however, the participant chooses a value higher than 1 million (prediction error is positive), they will realize that the number of deaths is lower than they thought, which should decrease their belief in the danger of pneumonia. Thus, the rational update is to decrease the belief that "pneumonia is dangerous for children." Analogous logic applies to each of the other eight combinations of the variables determining rational update in the prediction condition (Table 1).

Analytically, we can implement these features as sign changes multiplying the observed belief update: Any factor that would drive a belief increase is defined as +1 (supporting evidence, low magnitude, positive prediction error), and any factor that would drive a belief decrease is defined as −1 (refuting evidence, high magnitude, negative prediction error). The product of these three factors in the experimental condition—or the first factor in the control condition—determines on an item-by-item basis a positive or negative rational-update direction (RUD) for belief changes (ΔB):

$$\text{Evidence} = \{\text{support, against}\} \rightarrow \text{E} = \{-1, +1\}$$
$$\text{Magnitude} = \{\text{high, low}\} \rightarrow \text{M} = \{-1, +1\}$$
$$\text{sign}(\text{prediction error}) = \{-1, +1\}$$

$$\Delta \text{B}_{\text{rational}} = \text{RUD} \times \Delta \text{B}$$

$$\text{RUD}_{\text{control}} = \text{E}$$

$$\text{RUD}_{\text{prediction}} = \text{E} \times \text{M} \times \text{sign}(\text{prediction error})$$

We took the absolute value of the prediction errors to obtain an index of prediction-error size (from 0 to 11); higher scores indicate larger errors, and lower scores indicate lower errors. We then separated the 11 prediction-error sizes into three bins: no prediction error (0), small prediction error (1–5), and large prediction error (6–11). We decided to bin our data in this manner given the higher degree of interpretability of the binned prediction-error sizes as well as increased statistical power. For the sake of transparency, we present both the binned and the unbinned results, although they are equivalent.

Results

Does prediction error linearly predict rational update? To test our first hypothesis, we analyzed rational belief update as a function of prediction error in the prediction condition. We fitted a linear regression of prediction-error size against rational belief update and found that, as hypothesized, prediction-error size linearly and positively predicted rational belief update, $\beta = 2.87$, $SE = 0.11$, $t(3699) = 24.26$, $R^2 = .137$, $p < .001$ (Fig. 1a). We verified and expanded this analysis with a more rigorous linear mixed model run in *lme4* (Version 1.1.21; Bates et al., 2015) in R (Version 3.1.0; R Core Team, 2014). This model contained rational belief update as the dependent variable, prediction-error size and belief at pretest as fixed effects, and by-participant random intercepts and by-item random intercepts. We included belief at pretest as a fixed effect to control for the effect of the baseline

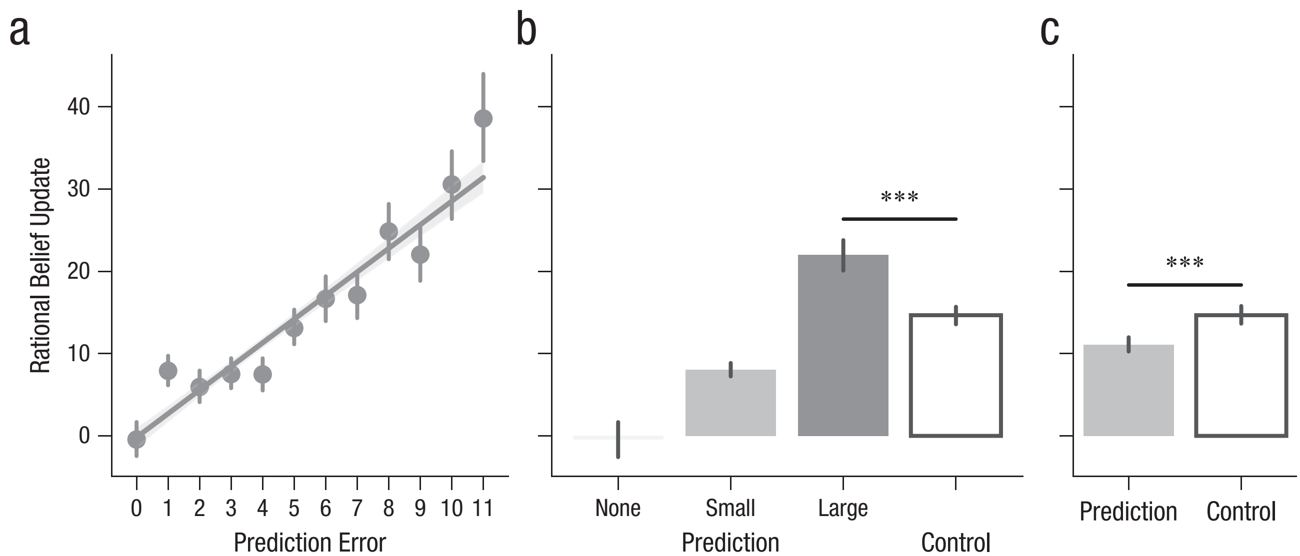

Fig. 1. Results from Study 1. Mean rational belief update (a) is shown as a function of prediction-error size. Rational belief update was calculated by subtracting pretest raw scores from posttest raw scores. Error bars represent 95% confidence intervals. The diagonal line shows the best-fitting regression, and the error band represents the 95% confidence interval. Mean rational belief update is broken down in (b) for prediction-error size (absolute value of prediction error) in each of the three prediction conditions and the control condition and in (c) for the prediction condition (collapsed) and control condition. Error bars in (b) and (c) represent ±1 *SEM*, and asterisks represent significant differences between conditions ($p < .001$).

level of belief on the degree of belief update while observing the independent effect of prediction-error size on rational belief update. Again, we found that prediction-error size linearly predicted rational belief update, $\beta = 1.58$, $SE = 0.09$, $t(12490) = 16.68$, $p < .001$.

Does more rational belief updating occur in the large-prediction-error condition compared with the control condition? We ran an independent-samples t test comparing rational belief update in the prediction and control conditions and found that items in the prediction condition were rationally updated ($M = 11.05$, $SD = 8.23$) to a lower degree than items in the control condition ($M = 14.67$, $SD = 10.28$), $t(670) = -5.14$, $p < .001$, Cohen's $d = 0.38$, 95% confidence interval (CI) = [−4.99, −2.23] (Fig. 1c). This indicates that, on average, participants' beliefs changed more when they were provided with passive evidence than when they were asked to make predictions and then provided with the correct answer.

To further explore this pattern, we assessed whether this conclusion applies independently of the size of the prediction error. Our preregistered hypothesis was that larger belief updates would occur in the large-prediction-error condition compared with the control condition. We, thus, compared the degree of rational belief update in the large-prediction-error bin of the prediction condition with the degree of rational belief update in the control condition. Consistent with our preregistered hypothesis, results of an independent-samples t test established that items in the large-prediction-error bin were rationally updated ($M = 22.00$, $SD = 17.40$) to a higher degree than items in the control condition ($M = 14.67$, $SD = 10.28$), $t(569) = 6.802$, $p < .001$, Cohen's $d = 0.51$, 95% CI = [5.21, 9.44] (Fig. 1b). Of note, the percentage of items that ended up in the three bins of the prediction condition was 9.45% in the no-prediction-error bin, 61.48% in the small-prediction-error bin, and 29.05% in the large-prediction-error bin.

Is there a partisan bias in how prediction error linearly predicts rational belief update? First, to investigate whether there was a difference in how Republicans and Democrats updated their beliefs, we turned to the prediction condition. We ran a linear mixed model with rational belief update as the dependent variable; prediction-error size, participant ideology, and belief at pretest as fixed effects; and by-participant and by-item random intercepts. The interaction between prediction error and participant ideology was not significant ($p = .18$). In other words, we did not find a difference in how Republicans and Democrats updated their beliefs as a function of prediction errors. This result was surprising for two reasons. First, participants' self-reported resistance to change was found to significantly moderate the effect of prediction error on belief update (i.e., participants who self-reported as more resistant to change were less likely to update beliefs as a function of prediction

Table 2. Rational Belief Update Predicted by a Linear Mixed Model Testing the Interaction of Prediction-Error Size With Participant Ideology and Item Ideology, Controlling for Belief at Pretest (Study 1)

Variable and item ideology	β	*SE*	*t*	*p*
Intercept	4.59	1.416	$t(45.24) = 3.24$	.002
Belief at pretest	16.8	0.255	$t(12510) = 66.00$	< .001
Democratic participants				
Neutral items	2.12	0.188	$t(6904) = 11.24$	< .001
Democratic items	1.80	0.182	$t(8156) = 9.86$	< .001
Republican items	1.29	0.181	$t(8693) = 7.13$	< .001
Republican participants				
Neutral items	2.00	0.185	$t(6870) = 10.76$	< .001
Democratic items	1.08	0.180	$t(7964) = 6.00$	< .001
Republican items	1.29	0.185	$t(8672) = 6.96$	< .001

errors), $\beta = -0.22$, $SE = 0.1$, $t(12550) = -2.12$, $p = .03$. Second, Republican participants self-reported as significantly more resistant to change ($M = 3.57$, $SD = 0.78$) than Democrats ($M = 3.31$, $SD = 0.83$), $t(700) = 4.179$, $p < .001$, Cohen's $d = 0.315$, 95% CI = [0.13, 0.37].

Furthermore, we tested for a potential ideological modulation of the effect of prediction error on rational update, this time while also taking into account the item ideology. We ran a linear mixed model testing the interaction of prediction-error size with participant ideology (Democratic and Republican) and item ideology (Democratic, Republican, neutral). The dependent variable was, again, rational belief update. We fitted prediction-error size, participant ideology, item ideology, and belief at pretest as fixed effects and included by-participant random intercepts and by-item random intercepts. The results showed that prediction-error size linearly predicted rational belief update in all six ideological conditions crossing participant ideology and item ideology (i.e., Democrats on neutral, Democratic, and Republican beliefs as well as Republicans on neutral, Democratic, and Republican beliefs; summarized in Table 2 and plotted in Fig. 2).

To exclude the possibility that a nonlinear model would better explain our data, we performed a model comparison between the linear models and alternative relationships. Because different models vary in parameter count, we compared models using the canonical Bayesian information criterion (BIC; Delattre et al., 2014; Schwarz, 1978), which normalizes data likelihoods under the models by their respective parameter counts. Lower BIC values indicate better fit. Repeating the mixed-effects model design in Table 2 (and Fig. 2) under a quadratic relationship between belief change and prediction error, we indeed found that the linear model (BIC = 123,401) was preferred to the quadratic model (BIC = 123,418).

We established that prediction errors linearly predicted rational belief update in all ideological subsamples. However, a partisan bias could still have existed in how strongly this effect manifested in these ideological subsamples. To test this possibility, we ran a linear mixed model on the prediction condition, with rational belief update as the dependent variable; prediction-error size, belief at pretest, item ideology, and participant ideology as fixed effects; and by-participant random intercepts and by-item random intercepts. We did not find a significant interaction among prediction-error size, item ideology (Democratic vs. Republican), and participant ideology (Democratic vs. Republican).

Finally, when we considered our measures of political polarization, neither strength of political-party affiliation, $\beta = 0.16$, $SE = 0.2$, $t(12474) = 0.809$, $p = .4186$, nor support for President Trump, $\beta = -0.16$, $SE = 0.15$, $t(12490) = -1.055$, $p = .2913$, significantly moderated the effect of prediction error on rational update.

Is there a partisan bias in how beliefs are rationally updated in the large-prediction-error condition compared with the control condition? To further investigate a potential ideological modulation of the uncovered effect of prediction error on rational update, we also tested whether rational update was higher in the large-prediction-error bin of the prediction condition compared with the control condition in each of the six subsamples of the data (i.e., Democrats on Democratic, Republican, and neutral items and Republicans on Democratic, Republican, and neutral items). We found that all of the independent-samples *t* tests were statistically significant for these comparisons (statistics reported in Table 3 and plotted in Fig. 3).

Given that we split the data in subsamples, we wanted to make sure that the magnitude of prediction errors did not differ between the Democratic and

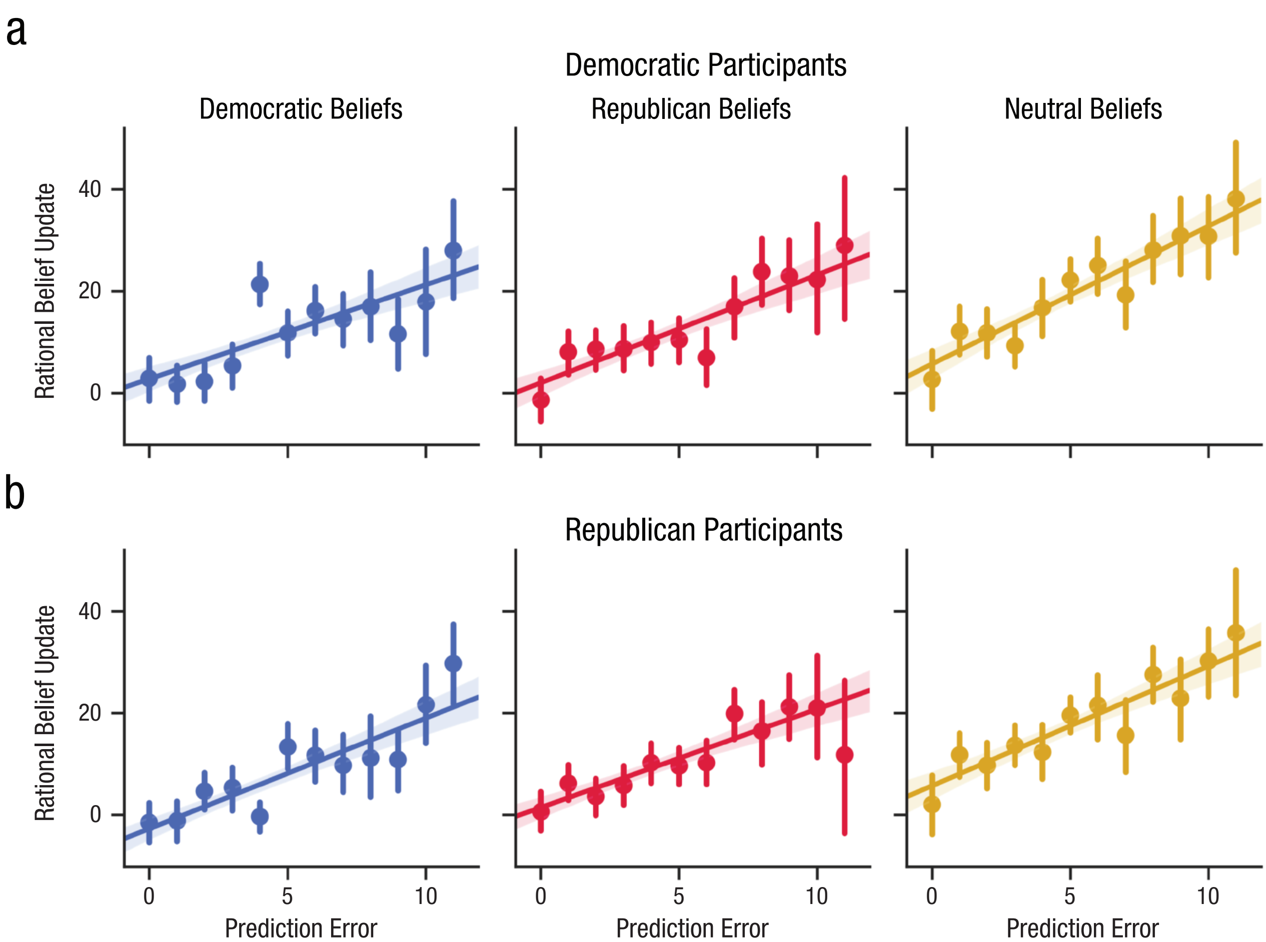

Fig. 2. Mean rational belief update by Democrats (a) and Republicans (b) as a function of prediction-error size, separately for each of the three belief ideologies (Study 1). Error bars represent 95% confidence intervals. Diagonal lines show best-fitting regressions, and error bands represents 95% confidence intervals.

Republican participants. Looking at the percentage of items that fell in each of the three prediction-error bins, we saw no evidence of such differences (Table 4).

In addition, to test for a partisan bias in the differences in belief update between the two conditions, we ran a mixed analysis of variance (ANOVA) with rational

Table 3. Difference in Rational Belief Update Between the Large-Prediction-Error and Control Conditions as a Function of All Participant Ideologies and Item Ideologies (Study 1)

Participant and item ideology	Large prediction error		Control		Between-conditions comparison			
	M	*SD*	*M*	*SD*	*t*	*p*	*d*	95% CI
Democrats								
Neutral	28.59	38.34	21.64	35.03	$t(258) = 4.45$	< .001	0.47	[5.62, 14.51]
Democratic	16.01	31.89	12.62	31.78	$t(278) = 1.98$	.048	0.21	[0.04, 7.00]
Republican	19.01	36.94	11.01	34.06	$t(240) = 4.36$	< .001	0.46	[4.96, 13.07]
Republicans								
Neutral	26.73	38.47	20.17	35.01	$t(275) = 3.81$	< .001	0.41	[3.96, 12.36]
Democratic	17.55	34.01	11.88	35.39	$t(274) = 4.31$	< .001	0.45	[4.51, 12.11]
Republican	16.22	32.86	10.69	32.35	$t(265) = 2.88$	.004	0.31	[1.70, 9.05]

Note: CI = confidence interval.

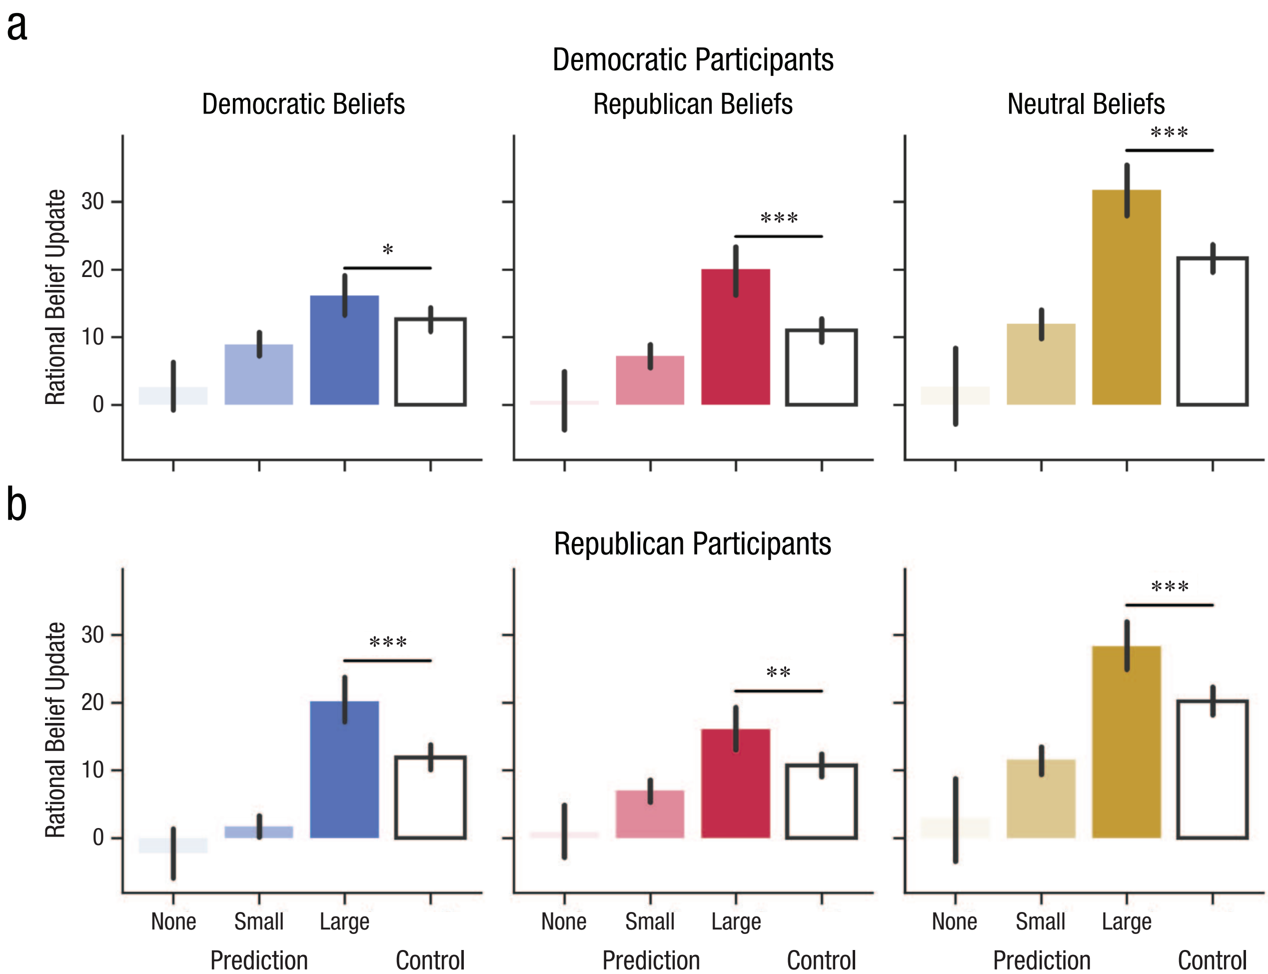

Fig. 3. Mean rational belief update by Democrats (a) and Republicans (b) as a function of prediction-error size in the three prediction conditions and the control condition (Study 1). Results are shown separately for each of the three belief ideologies. Error bars represent ±1 *SEM*, and asterisks represent significant differences between conditions (* $p < .05$, ** $p < .01$, *** $p < .001$).

belief update as the dependent variable, condition (large-prediction-error prediction vs. control) as a between-subjects variable, and congruence between participant and item ideologies as a within-subjects variable. We found a significant main effect of condition, $F(1, 1398) = 47.38$, $p < .001$, $\eta_p^2 = .03$; a significant main effect of ideology congruence, $F(1, 1398) = 3.91$, $p = .0482$; and a significant interaction of condition with ideology congruence, $F(1, 1398) = 4.89$, $p = .02717$, $\eta_p^2 = .003$, showing that, overall, participants in the large-prediction-error condition updated ideologically consistent items less than ideologically inconsistent items than did participants in the control condition. This suggests a partisan bias in rational belief update, manifested as higher rigidity for large-prediction-error beliefs in one's own ideology. To explore whether this

Table 4. Percentage of Items in the Three Bins of the Prediction Condition (Study 1)

	Neutral items		Democratic items		Republican items	
Condition	Democrats	Republicans	Democrats	Republicans	Democrats	Republicans
No prediction error	8.6%	7.6%	10.1%	10.2%	10.2%	9.9%
Small prediction error	62.8%	62.9%	58.4%	62.1%	61.4%	61.1%
Large prediction error	28.5%	29.4%	31.3%	27.5%	28.4%	29.0%

result further interacted with participant ideology, we conducted another mixed ANOVA with rational belief update as the dependent variable, condition (large-prediction-error prediction vs. control) and participant ideology (Democrats vs. Republicans) as between-subjects variables, and ideology congruence as a within-subjects variable. We did not find a significant three-way interaction, $F(1, 143) = 0.45$, $p = .501$.

Discussion

In Study 1, we found that prediction-error size linearly predicted rational belief update and that making large prediction errors led to larger belief updates than being passively exposed to evidence. These effects held for both Democrats and Republicans and for all belief types (neutral, Democratic, Republican). Despite the fact that self-reported resistance to change significantly moderated the effect of prediction error on belief update and Republicans self-reported as more resistant to change, we did not find differences in how Democrats and Republicans updated their beliefs. Finally, we found a partisan bias, which was manifested as higher rigidity for updating large-prediction-error beliefs in one's own ideology.

To assess the replicability and generalizability of these findings, we conducted a high-powered replication with a U.S. Census–matched sample.

Study 2

Method

Open-science practices. We preregistered the study's experimental design and hypotheses on an open-science platform (https://aspredicted.org/5iz27.pdf). The data for the replication study can be found on the study's OSF page (https://osf.io/aur2t). The data-analysis code (in Python) can be accessed as a Jupyter notebook at https://github.com/mvlasceanu/PredictionBelief.

Participants. For the replication, we aimed for a U.S. Census–matched sample of 1,000 participants. We recruited 1,387 Americans using the Cloud Research platform and excluded 313 of them on the basis of preregistered criteria (i.e., failed attention checks). We conducted statistical analyses on the final U.S. Census–matched sample of 1,073 participants (57% female; age: $M = 48.32$ years, $SD = 16.92$) who matched the census age, gender, race, and ethnicity quotas (Table 5). The total sample contained 552 participants who self-identified as Democrats and 521 who self-identified as Republicans. Each participant was randomly assigned to either the experimental condition (Democrats: $n = 324$, Republicans: $n = 296$) or the control condition (Democrats: $n = 228$, Republicans: $n = 225$). The study protocol was approved by the Princeton University Institutional Review Board.

We used the same stimulus materials, procedure, and coding as in Study 1. The data for Study 2 were collected between May 26, 2020, and June 4, 2020.

Results

Does prediction error linearly predict rational update? We fitted a linear regression of prediction-error size against rational belief update and replicated the result that prediction-error size linearly positively predicts rational belief update, $\beta = 2.36$, $SE = 0.09$, $t(6466) = 26.12$, $R^2 = .095$, $p < .001$ (Fig. 4a). As in Study 1, we ran a linear mixed model with rational belief update as the dependent variable, prediction-error size and belief at pretest as fixed effects, and by-participant random intercepts and by-item random intercepts. Belief at pretest was included as a fixed effect to control for the effect of baseline level of belief on degree of belief update while observing the independent effect of prediction-error size on rational belief update. We replicated the finding that prediction-error size linearly predicts rational belief update, $\beta = 1.40$, $SE = 0.07$, $t(22140) = 19.50$, $p < .001$.

Does more rational belief updating occur in the large-prediction-error condition compared with the control condition? We ran an independent-samples t test comparing rational belief update in the prediction and control conditions and replicated the result that items in the prediction condition were rationally updated ($M = 9.81$, $SD = 16.52$) to a lower degree than items in the control condition ($M = 13.49$, $SD = 9.79$), $t(825) = -5.92$, $p < .001$, Cohen's $d = 0.38$, 95% CI = [−4.32, −2.23] (Fig. 4c). To further explore this pattern as preregistered, we again assessed whether this conclusion applies independently of the size of the prediction error. With an independent-samples t test, we also replicated the result that items in the large-prediction-error bin were rationally updated ($M = 20.06$, $SD = 16.01$) to a higher degree than items in the control condition ($M = 13.49$, $SD = 9.79$), $t(1041) = 8.301$, $p < .001$, Cohen's $d = 0.47$, 95% CI = [4.90, 8.23] (Fig. 4b). Of note, the percentage of items that ended up in the three bins of the prediction condition was 9.2% in no-prediction error, 62.7% in small prediction error, and 28.1% in large prediction error.

Is there a partisan bias in how prediction error linearly predicts rational belief update? As in Study 1, to investigate whether there was a difference in how Republicans and Democrats updated their beliefs, we turned to the prediction condition. We ran a linear mixed model with rational belief update as the dependent

Table 5. Demographic Distribution of the Study 2 Sample Compared With U.S. Census Data

Variable	Census	Sample
Gender		
Male	49.4%	42.8%
Female	50.6%	57.2%
Age (years)		
18–29	22.6%	19.9%
30–39	16.8%	16.4%
40–49	16.2%	14.8%
50–59	17.8%	21.6%
60–69	14.0%	16.0%
70–99	12.4%	11.2%
Race		
Caucasian	78.8%	75.6%
African American	13.0%	9.9%
Native American	1.2%	3.4%
Asian	4.8%	4.9%
Other	2.2%	5.8%
Ethnicity		
Hispanic	16.0%	13.7%
Not Hispanic	84.0%	86.1%

variable; prediction-error size, participant ideology, and belief at pretest as fixed effects; and by-participant and by-item random intercepts. Contrary to Study 1, results showed that the interaction between prediction error and participant ideology in Study 2 reached statistical significance, $\beta = -0.57$, $SE = 0.12$, $t(22231) = -4.47$, $p < .001$, likely because of the increased sample size. This interaction indicates that Republicans updated their beliefs ($\beta = 1.11$, $SE = 0.16$) less than Democrats ($\beta = 1.68$, $SE = 0.09$) as a function of prediction errors. This result is thus consistent with the replicated findings that (a) self-reported resistance to change moderates the effect of prediction error on belief update—that is, the interaction between prediction error and resistance to change, $\beta = -0.50$, $SE = 0.08$, $t(22190) = -6.18$, $p < .001$, shows that participants who self-reported as more resistant to change were less likely to update beliefs as a function of prediction errors—and (b) Republicans self-reported as more resistant to change ($M = 3.35$, $SD = 0.87$) than Democrats ($M = 3.03$, $SD = 0.88$), $t(1062) = 5.908$, $p < .001$, Cohen's $d = 0.361$, 95% CI = [0.21, 0.42].

Furthermore, when also including item ideology (Democratic, Republican, neutral) in the model, we replicated the finding that prediction-error size linearly predicted rational belief update in all six ideological conditions crossing participant ideology and item ideology (i.e., Democrats on neutral, Democratic, and Republican beliefs as well as Republicans on neutral, Democratic, and Republican beliefs; summarized in Table 6 and plotted in Fig. 5).

As in Study 1, to exclude the possibility that a nonlinear model would better explain our data, we

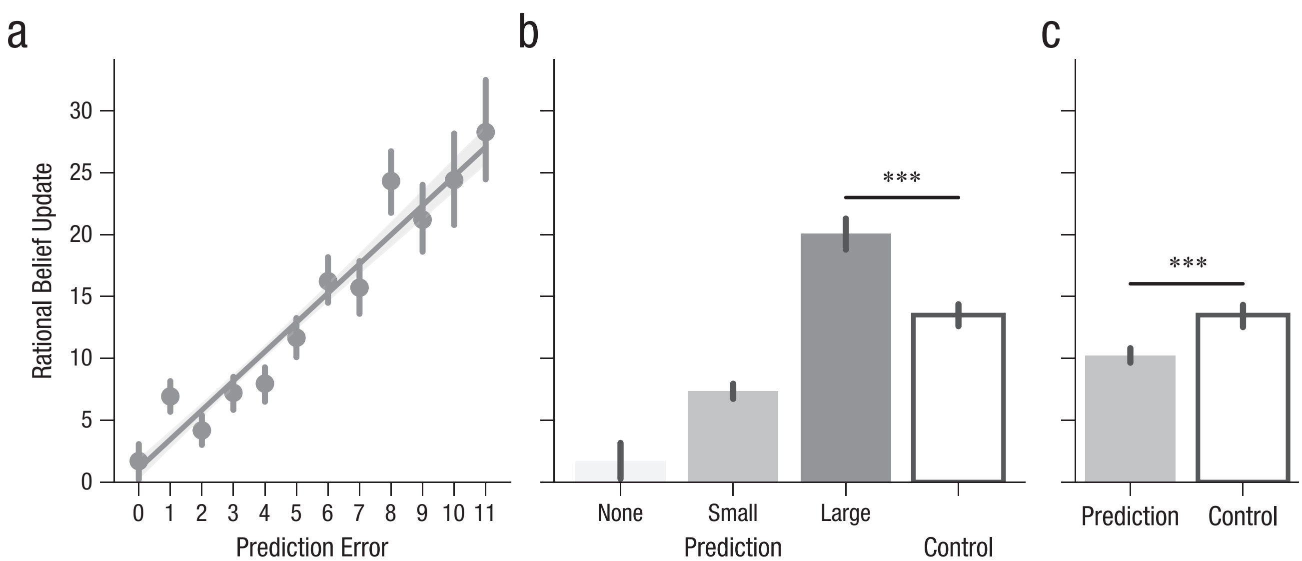

Fig. 4. Results from Study 2. Mean rational belief update (a) is shown as a function of prediction-error size. Rational belief update was calculated by subtracting pretest raw scores from posttest raw scores. Error bars represent 95% confidence intervals. The diagonal line shows the best-fitting regression, and the error band represents the 95% confidence interval. Mean rational belief update is broken down in (b) for prediction-error size (absolute value of prediction error) in each of the three prediction conditions and the control condition and in (c) for the prediction condition (collapsed) and control condition. Error bars in (b) and (c) represent ±1 *SEM*, and asterisks represent significant differences between conditions ($p < .001$).

Table 6. Rational Belief Update Predicted by a Linear Mixed Model Testing the Interaction of Prediction-Error Size With Participant Ideology and Item Ideology, Controlling for Belief at Pretest (Study 2)

Variable and item ideology	β	*SE*	*t*	*p*
Intercept	4.57	1.293	$t(40.37) = 3.53$	.001
Belief at pretest	16.4	0.191	$t(22170) = 85.97$	< .001
Democratic participants				
Neutral items	2.13	0.141	$t(14070) = 15.17$	< .001
Democratic items	1.72	0.139	$t(13920) = 12.36$	< .001
Republican items	1.35	0.138	$t(16240) = 9.76$	< .001
Republican participants				
Neutral items	1.44	0.141	$t(14310) = 10.26$	< .001
Democratic items	0.75	0.139	$t(14200) = 5.42$	< .001
Republican items	0.97	0.140	$t(16200) = 7.01$	< .001

performed a model comparison between the linear models and alternative relationships. Repeating the mixed-effects model design in Table 6 (and Fig. 5) under a quadratic relationship between belief change and prediction error, we indeed found that the linear model (BIC = 217,164) was preferred to the quadratic model (BIC = 217,188).

As before, once we had established that prediction errors linearly predicted rational belief update in all ideological subsamples of the data, we tested how

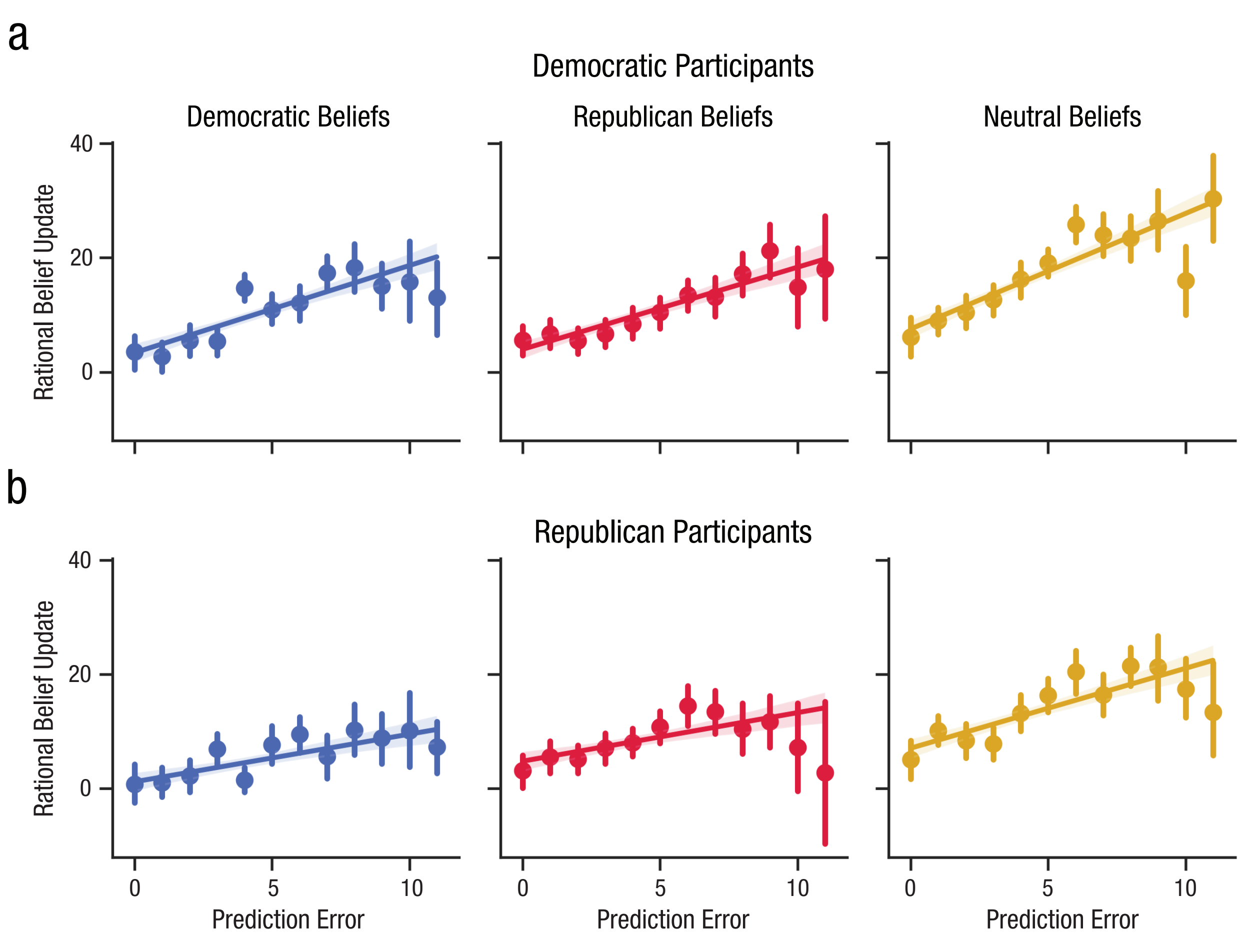

Fig. 5. Mean rational belief update by Democrats (a) and Republicans (b) as a function of prediction-error size, separately for each of the three belief ideologies (Study 2). Error bars represent 95% confidence intervals. Diagonal lines show best-fitting regressions, and error bands represents 95% confidence intervals.

Table 7. Difference in Rational Belief Update Between the Large-Prediction-Error and Control Conditions as a Function of All Participant Ideologies and Item Ideologies (Study 2)

Participant and item ideology	Large prediction error		Control		Between-conditions comparison			
	M	*SD*	*M*	*SD*	*t*	*p*	*d*	95% CI
Democrats								
Neutral	29.89	37.38	19.76	34.65	$t(526) = 6.74$	< .001	0.54	[7.81, 14.99]
Democratic	16.62	31.58	11.09	30.77	$t(539) = 5.11$	< .001	0.41	[3.88, 9.35]
Republican	17.06	34.56	10.24	32.04	$t(478) = 4.95$	< .001	0.39	[4.07, 10.41]
Republicans								
Neutral	25.14	36.26	19.00	35.37	$t(464) = 3.94$	< .001	0.32	[3.21, 10.54]
Democratic	12.57	34.15	10.63	33.62	$t(447) = 2.18$	.029	0.17	[0.08, 6.28]
Republican	14.46	35.56	10.22	31.39	$t(414) = 3.04$	.002	0.25	[1.43, 7.94]

Note: CI = confidence interval.

strongly this effect manifested in these subsamples. We ran a linear mixed model on the prediction condition, with rational belief update as the dependent variable; prediction-error size, participant ideology, item ideology, and belief at pretest as fixed effects; and by-participant random intercepts and by-item random intercepts. We replicated the lack of significance of the interaction among prediction-error size, item ideology (Democratic vs. Republican), and participant ideology (Democratic vs. Republican).

When we considered the measures of political polarization, the results were in contrast to those of Study 1. Now both strength of political-party affiliation, $\beta = -0.16$, $SE = 0.06$, $t(12402) = -2.628$, $p = .0086$, and support for President Trump, $\beta = -0.15$, $SE = 0.02$, $t(21440) = -5.767$, $p < .001$, reached statistical significance in moderating the effect of prediction error on rational update. These moderation analyses suggest that the more extreme participants were on the ideological spectrum and the more strongly they supported President Trump, the less they updated their beliefs according to prediction errors.

Is there a partisan bias in how beliefs are rationally updated in the large-prediction-error condition compared with the control condition? To further investigate a potential ideological modulation of the uncovered effect of prediction error on rational update, we again tested whether rational update was higher in the large-prediction-error bin of the prediction condition compared with the control condition in each of the six subsamples of the data. We found that all of the independent-samples *t* tests were statistically significant for these comparisons (statistics reported in Table 7 and plotted in Fig. 6).

As in Study 1, given that we split the data into subsamples, we wanted to make sure that the magnitude of Democratic and Republican participants' prediction errors did not differ. Looking at the percentage of items that fell in each of the three prediction-error bins, we saw no evidence of such differences (Table 8).

To test the partisan bias in belief update that we obtained in Study 1, according to which participants were less likely to update their ideologically consistent beliefs, we ran a mixed ANOVA with rational belief update as the dependent variable, condition (large-prediction-error prediction vs. control) as a between-subjects variable, and congruence between participant and item ideologies as a within-subjects variable. In contrast to Study 1, results did not show an interaction between condition and ideology congruence, $F(1, 28) = 0.08$, $p = .77938$. In other words, we did not observe a difference in resistance to changing one's own party's ideological beliefs compared with the other party's beliefs.

Discussion

Study 2 replicated the main findings that prediction-error size linearly predicts rational belief update and that making large prediction errors leads to a larger belief update than being passively exposed to evidence. It also replicated the result that these effects held for both Democrats and Republicans and for all belief types (neutral, Democratic, Republican). Moreover, we again found that self-reported resistance to change significantly moderated the effect of prediction error on belief update and that Republicans self-reported as significantly more resistant to change. Consistent with these effects (but in contrast to Study 1), the results of Study 2 showed that Republicans updated all beliefs less than Democrats. Notably, we no longer found the higher rigidity in updating large prediction-error beliefs in one's party ideology.

General Discussion

Changing people's beliefs is notoriously difficult. Here, in two preregistered studies—including one with a U.S.

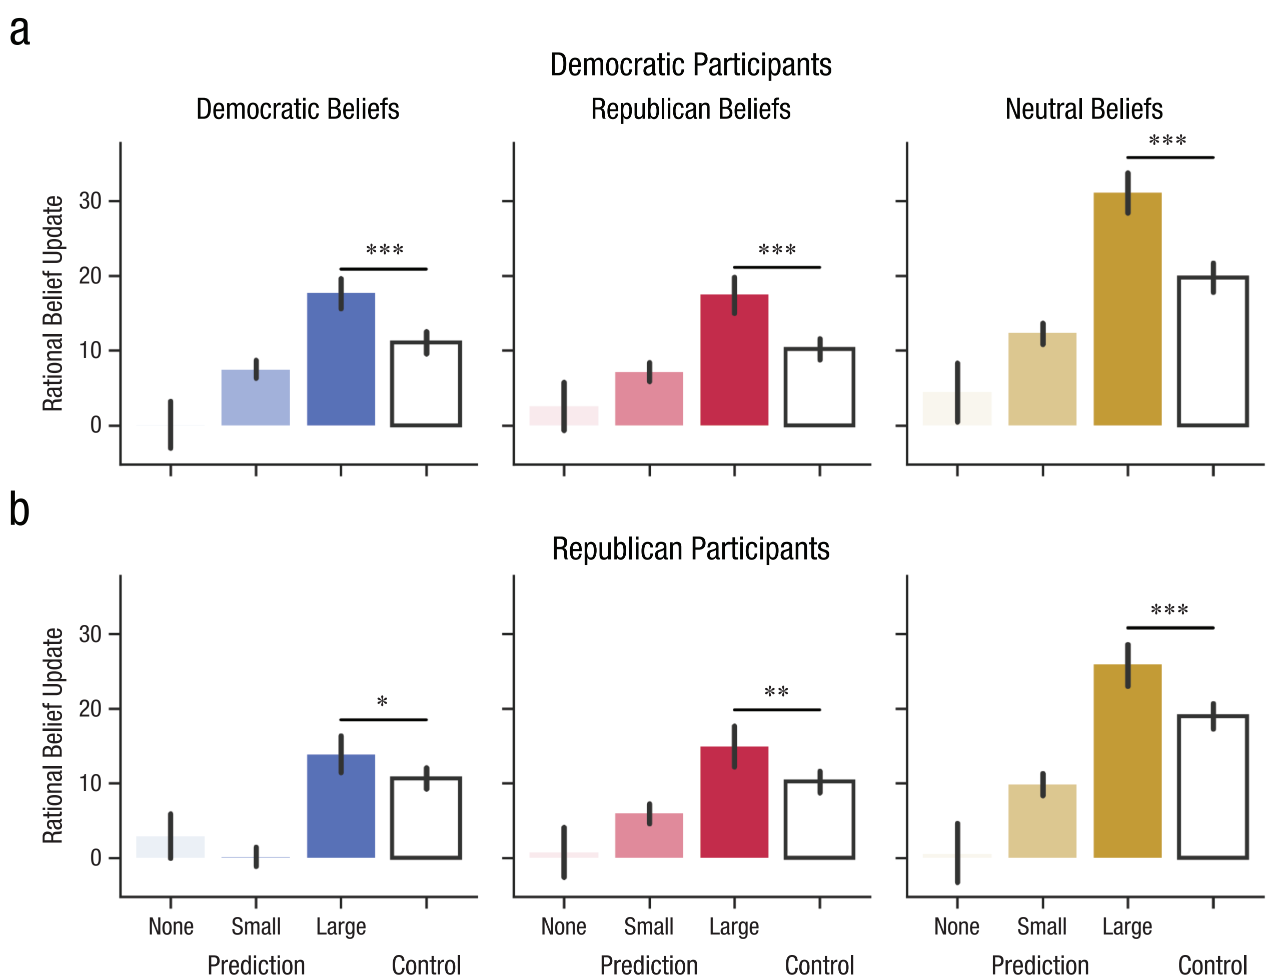

Fig. 6. Mean rational belief update by Democrats (a) and Republicans (b) as a function of prediction-error size in the three prediction conditions and the control condition (Study 2). Results are shown separately for each of the three belief ideologies. Error bars represent ±1 *SEM*, and asterisks represent significant differences between conditions (*$p < .05$, **$p < .01$, ***$p < .001$).

Census–matched sample—we found that an intervention that built on prediction errors could be successfully used, under specific circumstances, to change beliefs. Our main finding, that rational belief update is proportional to the magnitude of prediction error, aligns with the associative-learning principle that learning is proportional to prediction error (Rescorla & Wagner, 1972). We claim that our findings do not constitute a simple extension of this prior work, given that beliefs are deemed as meaningfully different from knowledge because of their associated conviction and self-referential element (Connors & Halligan, 2015). The ideological dimension of both believers and their beliefs could have amplified, attenuated, or even eliminated the

Table 8. Percentage of Items in the Three Bins of the Prediction Condition (Study 2)

	Neutral items		Democratic items		Republican items	
Condition	Democrats	Republicans	Democrats	Republicans	Democrats	Republicans
No prediction error	8.6%	9.1%	8.8%	8.6%	9.0%	9.2%
Small prediction error	62.8%	62.6%	60.9%	63.7%	63.8%	63.0%
Large prediction error	28.5%	28.2%	30.2%	27.5%	27.1%	27.7%

effect of prediction error on belief update. Yet the fact that they have not points to the effect's generalizability across the cognitive system.

Our findings also align with prior work showing that surprising information can tune knowledge, attitudes, and beliefs (Ranney & Clark, 2016). The element of surprise employed to increase acceptance of climate change, for instance, likely operates similarly to the prediction-error processes triggered in our paradigm (Ranney et al., 2001). However, our findings supplement this work in several ways. Critically, (a) we isolated the effect of prediction errors from that of evidence alone; (b) we quantified the magnitude of the surprise (i.e., prediction-error size), which we then used to predict belief update; and (c) we incorporated beliefs as well as participants from both sides of the political-ideological spectrum, which allowed the comparison of the effect's magnitude both within and across ideological boundaries.

We note an important difference between the two studies in how the effects interacted with ideology. In Study 1, we found a partisan bias in the form of higher rigidity in updating beliefs in one's own ideology (in the large-prediction-error condition compared with the control condition). In Study 2, we found that Republicans updated all beliefs less than Democrats, suggesting a partisan bias manifested as Republicans' resistance to changing all beliefs following prediction errors. At least two explanations could account for this difference in the manifestation of these partisan biases. First, the sample-size increase and the national representativeness of Study 2 may have provided the statistical power and the necessary variation to observe the true effect—Republican participants' diminished belief change based on prediction errors. Second, and perhaps more interestingly, the different sociopolitical contexts at the time of the data collection between the studies (October 2019 vs. May 2020) might have shifted the ideological bias from a symmetric effect (for both Democrats and Republicans) to the Republicans' resistance to update all beliefs. This possibility is consistent with existing work on the impact of threat and uncertainty on political beliefs (Haas & Cunningham, 2014). Although difficult to programmatically explore in a highly dynamic real-world situation (i.e., COVID-19 and nationwide antiracism protests), further research clarifying how consequential events affect belief change is certainly worth pursuing.

A reliable finding across the two studies was that the belief update in the prediction condition was, on average, significantly lower than in the control condition. We speculate that having to remember both the correct and predicted answer could create interference (when the difference between them is small), resulting in a memory decrement for the correct answer in the prediction condition compared with the control condition. This memory decrement could, in turn, lead to less rational belief update. Alternatively, there might be a cost to changing one's mind (e.g., one may appear inconsistent), so people might be willing to pay that cost only when extraordinary evidence is presented. Regardless of the mechanism, there is a pragmatic implication of the difference between the prediction and the control conditions. When addressing an audience for which one has no baseline belief information, one should simply provide accurate information. On the other hand, when one does have baseline belief information about a community, one would be well served to attack misinformation by narrowing the message to the most egregious belief violations.

Several important aspects of the belief-updating process were omitted in this study. One such factor is the credibility of the source presenting the evidence (Chung et al., 2008; Merdes et al., 2020). Future studies could explore how an information source affects the incorporation of evidence into one's belief system and how this impact might interact with ideology. For example, Democrats receiving evidence against a Democratic belief from CNN might update their belief accordingly, whereas evidence from Fox News might be completely discarded. Conversely, Republicans may be more open to evidence incorporation when watching Fox News compared with CNN (Haidt et al., 2009).

Another important extension could involve investigating the effect of conversations on prediction-based belief update and how these conversations, when they occur in larger communities, could impact collective beliefs (Vlasceanu et al., 2018). One possibility is that when given the opportunity to discuss, people would display a novelty bias and mention the evidence that is most surprising to them. Conversely, people might instead display a confirmation bias and mention the evidence they correctly predicted. Depending on what they choose to discuss, the community's collective beliefs would be shaped accordingly because previous research found that conversations influence collective beliefs (Vlasceanu & Coman, 2020b; Vlasceanu, Morais, et al., 2020). Clarifying this process would be particularly meaningful for policymakers interested in impacting communities (Dovidio & Esses, 2007).

Beyond their theoretical importance, these findings might provide useful tools in the battle against misinformation, a prominent threat facing the world today (Lewandowsky et al., 2012). For example, a third of Americans believe that global warming is a conspiracy (Jensen, 2013), and a third of American parents believe that vaccines cause autism (National Consumers League, 2014). False beliefs are dangerous when endorsed by a large proportion of people because they can shift attention and resources away from real threats, dramatically impact normative behavior, and cause suboptimal collective decisions (Kuklinski et al., 2000). Crucial steps in the misinformation-prevention battle are understanding the processes driving belief update and using that understanding to design misinformation-combating interventions. The present findings point to such interventions. For instance, our findings point to a powerful strategy that could shortcut ideological biases as an alternative to refutation, which may backfire, especially if beliefs are ideologically charged (Nyhan & Reifler, 2010). First, one needs to map the community's estimates on relevant statistics that can be used as surprising evidence. These statistics need to be carefully compiled given that people's predictions about everyday events are fairly accurate (Griffiths & Tenenbaum, 2006). After the statistics eliciting the largest misestimates are selected, these pieces of evidence need to be disseminated back to the community in a *predictions-then-feedback* format. This procedure is intensive but might have a stronger impact in diminishing misinformation than existing approaches. Conducting more empirical research to determine the stability of these findings over time, their boundary conditions, and their behavioral instantiations could offer policymakers a powerful tool to address this global epidemic.

Transparency

Action Editor: Lisa Diamond
Editor: Patricia J. Bauer
Author Contributions
M. Vlasceanu and A. Coman developed the study concept. All the authors contributed to the study design. M. Vlasceanu conducted testing and data collection. All the authors analyzed and interpreted the data. M. Vlasceanu drafted the manuscript, and A. Coman and M. J. Morais provided critical revisions. All the authors approved the final manuscript for submission.
Declaration of Conflicting Interests
The author(s) declared that there were no conflicts of interest with respect to the authorship or the publication of this article.
Open Practices
All data and stimuli have been made publicly available via OSF and can be accessed at https://osf.io/aur2t. The design and analysis plans for the study were preregistered at https://aspredicted.org/zu4iq.pdf (Study 1) and https://aspredicted.org/5iz27.pdf (Study 2). The data-analysis code (in Python) can be accessed as a Jupyter notebook at https://github.com/mvlasceanu/PredictionBelief. This article has received the badges for Open Data, Open Materials, and Preregistration. More information about the Open Practices badges can be found at http://www.psychologicalscience.org/publications/badges.

References

Alexander, W. H., & Brown, J. W. (2011). Medial prefrontal cortex as an action-outcome predictor. *Nature Neuroscience*, *14*, 1338–1344.

Bar, M. (2009). Predictions: A universal principle in the operation of the human brain. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *364*(1521), 1181–1182.

Bartels, L. M. (1996). Uninformed votes: Information effects in presidential elections. *American Journal of Political Science*, *40*(1), 194–230. https://doi.org/10.2307/2111700

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, *67*(1). https://doi.org/10.18637/jss.v067.i01

Bendixen, L. D. (2002). A process model of epistemic belief change. In B. K. Hofer & P. R. Pintrich (Eds.), *Personal epistemology: The psychology of beliefs about knowledge and knowing* (pp. 191–208). Erlbaum.

Bestmann, S., Harrison, L. M., Blankenburg, F., Mars, R. B., Haggard, P., Friston, K. J., & Rothwell, J. C. (2008). Influence of uncertainty and surprise on human corticospinal excitability during preparation for action. *Current Biology*, *18*, 775–780.

Buechner, B. M., Clarkson, J. J., Otto, A. S., Hirt, E. R., & Ho, M. C. (2021). Political ideology and executive functioning: The effect of conservatism and liberalism on cognitive flexibility and working memory performance. *Social Psychological and Personality Science*, *12*, 237–247. https://doi.org/10.1177/1948550620913187

Chung, S., Fink, E. L., & Kaplowitz, S. A. (2008). The comparative statics and dynamics of beliefs: The effect of message discrepancy and source credibility. *Communication Monographs*, *75*(2), 158–189.

Connors, M. H., & Halligan, P. W. (2015). A cognitive account of belief: A tentative road map. *Frontiers in Psychology*, *5*, Article 1588. https://doi.org/10.3389/fpsyg.2014.01588

de Lange, F. P., Heilbron, M., & Kok, P. (2018). How do expectations shape perception? *Trends in Cognitive Sciences*, *22*(9), 764–779.

Delattre, M., Lavielle, M., & Poursat, M.-A. (2014). A note on BIC in mixed-effects models. *Electronic Journal of Statistics*, *8*(1), 456–475.

den Ouden, H. E., Kok, P., & de Lange, F. P. (2012). How prediction errors shape perception, attention, and motivation. *Frontiers in Psychology*, *3*, Article 548. https://doi.org/10.3389/fpsyg.2012.00548

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.

Dovidio, J. F., & Esses, V. M. (2007). Psychological research and public policy: Bridging the gap. *Social Issues and Policy Review*, *1*, 5–14.

Erickson, C. A., & Desimone, R. (1999). Responses of macaque perirhinal neurons during and after visual stimulus association learning. *The Journal of Neuroscience*, *19*, 10404–10416.

Festinger, L., & Carlsmith, J. M. (1959). Cognitive consequences of forced compliance. *The Journal of Abnormal and Social Psychology*, *58*(2), 203–210.

Greve, A., Cooper, E., Kaula, A., Anderson, M. C., & Henson, R. (2017). Does prediction error drive one-shot declarative learning? *Journal of Memory and Language*, *94*, 149–165.

Griffiths, T. L., & Tenenbaum, J. B. (2006). Optimal predictions in everyday cognition. *Psychological Science*, *17*(9), 767–773.

Haas, I. J., & Cunningham, W. A. (2014). The uncertainty paradox: Perceived threat moderates the effect of uncertainty on political tolerance. *Political Psychology*, *35*(2), 291–302.

Haidt, J., Graham, J., & Joseph, C. (2009). Above and below left-right: Ideological narratives and moral foundations. *Psychological Inquiry*, *20*, 110–119.

Hume, D. (1993). *An abstract of a treatise on human nature*. Hackett. (Original work published 1740)

Jensen, T. (2013, April 2). *Democrats and Republicans differ on conspiracy theory beliefs*. Public Policy Polling. https://www.publicpolicypolling.com/polls/democrats-and-republicans-differ-on-conspiracy-theory-beliefs/

Jost, J. T., Glaser, J., Kruglanski, A. W., & Sulloway, F. J. (2003). Political conservatism as motivated social cognition. *Psychological Bulletin*, *129*(3), 339–375.

Kuklinski, J. H., Quirk, P. J., Jerit, J., Schwieder, D., & Rich, R. F. (2000). Misinformation and the currency of democratic citizenship. *Journal of Politics*, *62*(3), 790–816.

Kutas, M., & Hillyard, S. A. (1980). Reading senseless sentences: Brain potentials reflect semantic incongruity. *Science*, *207*, 203–205.

Lewandowsky, S., Ecker, U. K., Seifert, C. M., Schwarz, N., & Cook, J. (2012). Misinformation and its correction: Continued influence and successful debiasing. *Psychological Science in the Public Interest*, *13*(3), 106–131.

Litman, L., Robinson, J., & Abberbock, T. (2017). TurkPrime.com: A versatile crowdsourcing data acquisition platform for the behavioral sciences. *Behavior Research Methods*, *49*, 433–442.

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.

Merdes, C., von Sydow, M., & Hahn, U. (2020). Formal models of source reliability. *Synthese*. Advance online publication. https://doi.org/10.1007/s11229-020-02595-2

National Consumers League. (2014, April). *Survey: One third of American parents mistakenly link vaccines to autism*. http://www.nclnet.org/survey_one_third_of_american_parents_mistakenly_link_vaccines_to_autism

Nyhan, B., & Reifler, J. (2010). When corrections fail: The persistence of political misperceptions. *Political Behavior*, *32*(2), 303–330.

Pennycook, G., McPhetres, J., Zhang, Y., Lu, J. G., & Rand, D. G. (2020). Fighting COVID-19 misinformation on social media: Experimental evidence for a scalable accuracy-nudge intervention. *Psychological Science*, *31*, 770–780.

Ranney, M. A., Cheng, F., Garcia de Osuna, J., & Nelson, J. (2001). *Numerically driven inferencing: A new paradigm for examining judgments, decisions, and policies involving base rates* [Paper presentation]. Annual Meeting of the Society for Judgment and Decision Making, Orlando, FL, United States.

Ranney, M. A., & Clark, D. (2016). Climate change conceptual change: Scientific information can transform attitudes. *Topics in Cognitive Science*, *8*(1), 49–75.

R Core Team. (2014). *R: A language and environment for statistical computing* (Version 3.1.0) [Computer software]. https://www.r-project.org/

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: The effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning, 2: Current research and theory* (pp. 64–69). Appleton-Century-Crofts.

Rollwage, M., Dolan, R. J., & Fleming, S. M. (2018). Metacognitive failure as a feature of those holding radical beliefs. *Current Biology*, *28*(24), 4014–4021.

Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, *6*(2), 461–464.

Schwitzgebel, E. (2010). Acting contrary to our professed beliefs or the gulf between occurrent judgment and dispositional belief. *Pacific Philosophical Quarterly*, *91*(4), 531–553.

Shearer, E., & Gottfried, J. (2017, September 7). *News use across social media platforms 2017*. Pew Research Center. https://www.journalism.org/2017/09/07/news-use-across-social-media-platforms-2017/

Shu, K., Sliva, A., Wang, S., Tang, J., & Liu, H. (2017). Fake news detection on social media: A data mining perspective. *ACM SIGKDD Explorations Newsletter*, *19*(1), 22–36.

Toner, K., Leary, M. R., Asher, M. W., & Jongman-Sereno, K. P. (2013). Feeling superior is a bipartisan issue: Extremity (not direction) of political views predicts perceived belief superiority. *Psychological Science*, *24*(12), 2454–2462.

Vlasceanu, M., & Coman, A. (2018). Mnemonic accessibility affects statement believability: The effect of listening to others selectively practicing beliefs. *Cognition*, *180*, 238–245.

Vlasceanu, M., & Coman, A. (2020a). *The impact of social norms on belief update*. PsyArXiv. https://doi.org/10.31234/osf.io/gsem6

Vlasceanu, M., & Coman, A. (2020b). *Network structure impacts the synchronization of collective beliefs*. PsyArXiv. https://doi.org/10.31234/osf.io/7rq4g

Vlasceanu, M., Enz, K., & Coman, A. (2018). Cognition in a social context: A social-interactionist approach to emergent phenomena. *Current Directions in Psychological Science*, *27*(5), 369–377.

Vlasceanu, M., Goebel, J., & Coman, A. (2020). The emotion-induced belief amplification effect. In S. Denison., M. Mack, Y. Xu, & B.C. Armstrong (Eds.), *Proceedings of the 42nd Annual Conference of the Cognitive Science Society* (pp. 417–422). Cognitive Science Society.

Vlasceanu, M., Morais, M. J., Duker, A., & Coman, A. (2020). The synchronization of collective beliefs: From dyadic interactions to network convergence. *Journal of Experimental Psychology: Applied*, *26*, 453–464. https://doi.org/10.1037/xap0000265

Wee, S. (2013). Development and initial validation of the Willingness to Compromise Scale. *Journal of Career Assessment*, *21*(4), 487–501.

Wheeler, C., Green, M. C., & Brock, T. C. (1999). Fictional narratives change beliefs: Replications of Prentice, Gerrig, and Bailis (1997) with mixed corroboration. *Psychonomic Bulletin & Review*, *6*(1), 136–141.

White, K. R., Kinney, D., Danek, R. H., Smith, B., & Harben, C. (2020). The Resistance to Change-Beliefs Scale: Validation of a new measure of conservative ideology. *Personality and Social Psychology Bulletin*, *46*(1), 20–35.

Zmigrod, L., Rentfrow, P. J., & Robbins, T. W. (2020). The partisan mind: Is extreme political partisanship related to cognitive inflexibility? *Journal of Experimental Psychology: General*, *149*(3), 407–418.