



Dogmatism manifests in lowered information search under uncertainty

Lion Schulz^{a,b,c,1,2} , Max Rollwage^{a,d,1} , Raymond J. Dolan^{a,d} , and Stephen M. Fleming^{a,b,d,2}

^aWellcome Centre for Human Neuroimaging, Institute of Neurology, University College London, London WC1N 3BG, United Kingdom; ^bDepartment of Experimental Psychology, University College London, London WC1H 0AP, United Kingdom; ^cDepartment of Computational Neuroscience, Max Planck Institute for Biological Cybernetics, 72076 Tübingen, Germany; and ^dMax Planck University College London Centre for Computational Psychiatry and Ageing Research, Institute of Neurology, University College London, London WC1B 5EH, United Kingdom

Edited by Susan T. Fiske, Princeton University, Princeton, NJ, and approved October 15, 2020 (received for review May 14, 2020)

When knowledge is scarce, it is adaptive to seek further information to resolve uncertainty and obtain a more accurate worldview. Biases in such information-seeking behavior can contribute to the maintenance of inaccurate views. Here, we investigate whether predispositions for uncertainty-guided information seeking relate to individual differences in dogmatism, a phenomenon linked to entrenched beliefs in political, scientific, and religious discourse. We addressed this question in a perceptual decision-making task, allowing us to rule out motivational factors and isolate the role of uncertainty. In two independent general population samples ($n = 370$ and $n = 364$), we show that more dogmatic participants are less likely to seek out new information to refine an initial perceptual decision, leading to a reduction in overall belief accuracy despite similar initial decision performance. Trial-by-trial modeling revealed that dogmatic participants placed less reliance on internal signals of uncertainty (confidence) to guide information search, rendering them less likely to seek additional information to update beliefs derived from weak or uncertain initial evidence. Together, our results highlight a cognitive mechanism that may contribute to the formation of dogmatic worldviews.

information search | dogmatism | politics | metacognition | computational modeling

A never-ending flow of informational choices is a defining feature of modernity (1). We are in charge of gathering information critical to our health (2), the survival of democracies (3), or the conservation of the planet (4). These decisions to seek information are in turn a crucial determinant of our beliefs. Cognitive science has studied information search extensively, providing us with a rich empirical and theoretical perspective on these choices (5–8). This research indicates that people prefer to seek information that confirms their beliefs and has positive valence, as when we read an additional news story about the victory of a favored political party. This type of motivated search is evident both in laboratory experiments (9–11) and in real-world data (12–14).

In contrast, normative frameworks propose that uncertainty, rather than valence, should determine where and when we should seek information (5, 7, 15). In the absence of external feedback, humans rely on internal confidence signals to guide such uncertainty-based information search. Bayes-optimal agents should seek out more information when they have lower confidence in a decision, because the likelihood of a mistake will then outweigh the cost of further information (16). Empirical data bear out such predictions (17, 18), showing that people are more likely to seek information when they express low confidence (i.e., higher uncertainty) in their decisions.

In everyday situations, both motivational influences and failures in uncertainty-guided information search can lead to biased or inaccurate beliefs, albeit via distinct mechanisms. For example, a person who does not believe in climate change is likely to show a preference for media that refutes its occurrence (19), reinforcing preexisting beliefs. Alternatively, people with doubts

about the science of global warming (20) might fail to act on this uncertainty, and as a consequence not seek out further, potentially corrective evidence. In both of these cases, an unwillingness to seek out corrective information is one potential source of dogmatism, a worldview that involves a rigid maintenance of one's beliefs (21–23) regardless of their accuracy (24). The scope of this worldview is wide-ranging and transcends specific issues and positions, affecting political (25), scientific (26), and religious debate (21, 26). Prior questionnaire-based research suggests a link between such a dogmatic style of thinking and a willingness to seek further information (21, 27). However, how motivation and uncertainty contribute to this phenomenon remains unknown.

Here, we address this question using a precise assay of uncertainty-guided information search in the context of a low-level perceptual decision-making task. Leveraging the computational precision afforded by this approach, we test (in both exploratory and replication samples) whether individual differences in a sensitivity to uncertainty explain a disposition to hold dogmatic beliefs. Our approach builds on previous research on the influence of confidence on information search (17, 28) and allows us to rule out possible motivational influences: Participants are unlikely to approach such a low-level task with vested interests or prior knowledge, and should not hold differing appraisals of the helpfulness of further information. Moreover, eliciting trial-by-trial estimates of confidence enabled us to infer how participants use uncertainty to guide their search.

Significance

Dogmatic individuals are reluctant to seek out new information to refine their views, often skewing political, scientific, and religious discourse in the process. The cognitive drivers of this reluctance are poorly understood. Here, we isolate an influence of uncertainty on information search using a low-level perceptual decision-making task. We show that people with dogmatic views are both less likely to seek information before committing to a decision and to use fluctuations in uncertainty to guide their search. Our results highlight a cognitive mechanism that may contribute to the formation of dogmatic worldviews.

Author contributions: L.S., M.R., R.J.D., and S.M.F. conceptualized the study; L.S. and M.R. developed the methodology under the supervision of S.M.F.; L.S. conducted the experiment and analyzed the data under the supervision of M.R. and S.M.F.; L.S. wrote the paper; and M.R., R.J.D., and S.M.F. revised the manuscript.

The authors declare no competing interest.

This article is a PNAS Direct Submission.


This open access article is distributed under [Creative Commons Attribution License 4.0 \(CC BY\)](https://creativecommons.org/licenses/by/4.0/).

¹L.S. and M.R. contributed equally to this work.

²To whom correspondence may be addressed. Email: lion.schulz@tuebingen.mpg.de or stephen.fleming@ucl.ac.uk.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2009641117/-DCSupplemental>.

First published November 19, 2020.



Results

We studied a sample of 370 US adults (study 1) and replicated all key findings in an independent second sample of 364 US participants (study 2). Both samples were recruited through Amazon Mechanical Turk and comprised a wide range of ages and educational backgrounds (see *Materials and Methods* and *SI Appendix* for details). Participants first completed an information-seeking task and then answered a number of questionnaires that allowed us to measure general belief rigidity and dogmatism, political beliefs, authoritarianism, and intolerance to opposing political attitudes (see *Materials and Methods* for details regarding the questionnaires). This methodology builds on our previous research as to how dogmatic individuals construct a sense of confidence (24).

Measuring Dogmatism. We derived a measure of dogmatism from a factor analysis applied to the questionnaire battery (Fig. 1). The breadth of the battery allowed us to quantitatively distinguish dogmatism from other, possibly related, constructs and study their interplay. The most parsimonious factor structure contained three factors, capturing 40% of questionnaire variance. A first factor represented individuals' position on a left-right political spectrum, and a second factor described their




Fig. 1. Liberal and conservative extremes of the political spectrum are associated with higher levels of dogmatism and belief superiority. (A) A factor analysis revealed a three-factor structure underlying responses to multiple questionnaires assessing political convictions, authoritarianism, and belief rigidity. The three factors identified 1) "political orientation" (liberal to conservative), 2) "dogmatism" representing a domain-general belief certainty, and 3) "political belief superiority" characterizing participants' confidence in specific political convictions. Item loadings for each question are presented with individual questionnaires indicated by colors. (B–D) We examined the relationships between these constructs by investigating individual scores for each factor (combined data from studies 1 and 2 are plotted). (B and C) We observed that a combined linear-quadratic model provided the best fit to the relationship between both political orientation and dogmatism as well as between political orientation and political belief superiority. (D) The relationship between dogmatism and political belief superiority was best characterized by a linear relationship (see *SI Appendix* for further details about the factor analysis).

domain-general dogmatism and rigidity of worldview (21). A third factor captured variance in beliefs as to the superiority of participants' policy preferences (a factor related to but also theoretically independent of dogmatism) (22). Whereas this last factor is specific to political policy, dogmatism itself is a broader construct that describes the general way beliefs are held and acted upon (23, 29, 30). Dogmatism thereby transcends specific political preferences as evident, for example, in a link between dogmatism and religious fundamentalism (21).

We first explored interrelationships between individuals' political orientation, dogmatism, and political belief superiority (*SI Appendix, Table S2*). We found both positive linear (study 1: $\beta_{\text{linear}} = 0.16$, $P = 0.001$; study 2: $\beta_{\text{linear}} = 0.24$, $P < 10^{-6}$) and quadratic (study 1: $\beta_{\text{quadratic}} = 0.35$, $P < 10^{-13}$; study 2: $\beta_{\text{quadratic}} = 0.37$, $P < 10^{-13}$) relationships between political orientation and dogmatism across both samples. These findings indicate that individuals on both the far left and far right of the political spectrum show enhanced dogmatism, but interestingly this increase in dogmatism is more marked for those on the far right (see Fig. 1, *Materials and Methods*, and *SI Appendix, Table S1* for further information). Conversely, a negative linear (study 1: $\beta_{\text{linear}} = -0.33$, $P < 10^{-10}$; study 2: $\beta_{\text{linear}} = -0.32$, $P < 10^{-9}$) and positive quadratic (study 1: $\beta_{\text{quadratic}} = 0.43$, $P < 10^{-20}$; study 2: $\beta_{\text{quadratic}} = 0.34$, $P < 10^{-10}$) relationship between political orientation and political belief superiority reveals that individuals on both the far left and far right show heightened beliefs in the superiority of their respective policy positions, but more so on the far left. Finally, we found a positive linear relationship between political belief superiority and dogmatism, indicating more dogmatic subjects also tended to be more confident in the superiority of their specific political convictions (study 1: $\beta = 0.26$, $P < 10^{-06}$; study 2: $\beta = 0.14$, $P = 0.006$).

Measuring Information Search. We next tested our primary hypothesis of a link between dogmatism and uncertainty-guided information search. To probe this, we presented participants with a perceptual information-seeking task (Fig. 2; see *Materials and Methods* for details) where they received monetary reward for correctly judging which of two flickering boxes contained the greater number of dots. On each trial, participants were first presented with an initial pair of boxes, made an unincentivized judgment as to which box contained more dots and simultaneously reported their confidence in this judgment. Following this initial decision, participants were asked whether they wanted to see additional evidence so as to improve their decision accuracy. If participants decided to see this additional information, they were presented with another set of flickering dots. This additional information was always helpful (the correct option continued to have a greater number of dots) and was presented at a higher stimulus strength (greater dot difference). We imposed a variable cost for seeing this information through a deduction of points, allowing concurrent assessment of participants' sensitivity to information cost. If participants decided not to see more information, they instead saw two empty black boxes for an identical duration and were not deducted any points. Finally, regardless of whether subjects decided to see additional information or not, they were asked to provide a final decision and confidence rating. To incentivize subjects to strive for the best possible overall accuracy, they were paid only for their performance on this final decision.

First, we validated that participants used the additional information adaptively (see *Materials and Methods* for details and *SI Appendix, Fig. S2* for overview). Participants chose to see additional information more often after initial mistakes (study 1: $\beta = -0.76$, $P < 10^{-69}$; study 2: $\beta = -0.77$, $P = 10^{-52}$) and were more likely to make an accurate final decision after deciding to see additional information (study 1: $\beta = 1.23$, $P < 10^{-80}$; study 2: $\beta = 1.12$, $P < 10^{-74}$). Similarly, examining individual differences




Fig. 2. Example experimental trial. Participants first had to judge whether a left or right square contained more flickering dots. They then chose whether they wanted to see a stronger, more helpful, version of this stimulus again, costing them either 5 or 20 points. After seeing either this additional stimulus or blank boxes, they again made a judgment as to which box contained more dots. We compensated participants for the accuracy of this final decision alone (100 points for a correct judgment; 0 points for an incorrect judgment). Participants rated their confidence (on a six-point scale from “sure left” to “sure right”) at both the initial and the final decision. The difficulty of the initial decision was fixed through an individually predetermined difference in dot number that resulted in ~71% accuracy. The post decision-evidence strength was yoked to this predetermined dot difference, to make the final decision easier (information on stimulus calibration is provided in *Materials and Methods* and *SI Appendix*).

in information search revealed that participants who sought additional information more often also performed better in their final decisions (study 1: $\beta = 0.68, P < 10^{-50}$; study 2: $\beta = 0.71, P < 10^{-56}$) and received a higher payoff (study 1: $\beta = 0.43, P < 10^{-17}$, study 2: $\beta = 0.49, P < 10^{-22}$). Importantly participants were also sensitive to cost, seeking information less often when it was more expensive (study 1: $\beta = -1.22, P < 10^{-38}$; study 2: $\beta = -1.34, P < 10^{-66}$). Finally, participants sought out additional evidence less often when they were more confident in their initial decision (study 1: $\beta = -2.02, P < 10^{-126}$; study 2: $\beta = -1.93, P < 10^{-130}$), demonstrating that internal signals of uncertainty were used to guide information search.

Information Search and Dogmatism. We next asked whether more dogmatic participants differed in their propensity to seek out information. To that end, we sought to explain variance in dogmatism factor scores using behavioral measures derived from the information-seeking task. In line with our hypothesis, higher levels of dogmatism were associated with a reduced willingness to seek out information (study 1: $\beta = -0.15, P = 0.005, R^2 = 0.02$; Fig. 3A). No significant relationships with initial decision accuracy (study 1: $\beta = 0.02, P = 0.72$) or overall confidence level (study 1: $\beta = -0.03, P = 0.61$) were found, and our analyses controlled for key demographic variables including age, gender, and education (see Fig. 3A and *Materials and Methods* for details). We replicated this lowered tendency for dogmatic subjects to seek out information in a second, independent sample in study 2 ($\beta = -0.10$, one-tailed $P = 0.039, R^2 = 0.01$; Fig. 3), again in the absence of differences in initial decision accuracy ($\beta = -0.09, P = 0.13$) and confidence ($\beta = -0.03, P = 0.53$). To assess the overall strength of the relationship between dogmatism and information seeking, we conducted an internal meta-analysis of this effect by pooling data from both samples (31). This revealed a consistently negative effect across both samples ($\beta = -0.12, P = 0.002, R^2 = 0.012$). Overall, our findings highlight that, in the absence of motivational factors, more dogmatic participants seek out less information before committing to a decision—even when this information would be helpful.

A key question arising from this finding is whether more dogmatic individuals’ final accuracy and payoff suffered because of their lowered information search, or whether they simply sought information more efficiently. Here, a mediation analysis (Fig. 3B; see *Materials and Methods* for details) showed that more dogmatic participants were in fact less accurate in their final decision (total effect: study 1, $\beta = -0.11, P = 0.001$; study 2, $\beta = -0.09$, one-tailed $P = 0.01$), and that this effect was fully

mediated by a lowered willingness to seek information (mediation effect: study 1: $\beta = -0.08, P = 0.005$; study 2: $\beta = -0.05$, one-tailed $P = 0.038$; corrected direct effect, study 1: $\beta = -0.03, P = 0.097$; study 2: $\beta = -0.03, P = 0.12$). To obtain a meta-analytical estimate of this mediation analysis, we again pooled data from both studies to establish that our effect was stable across conditions (total effect: $\beta = -0.098, P = 0.0001$; mediation effect: $-0.064, P = 0.002$;




Fig. 3. Dogmatism is characterized by a reduction in information search, leading to less veridical judgments. (A) Dogmatism was predicted by a reduced willingness to seek out more information before committing to a decision, controlling for several demographic and task variables. We present standardized β coefficients \pm SE of predictors for study 1 (left markers, $n = 370$) and study 2 (right markers, $n = 364$). Effects in study 2 were tested one-tailed based on the directional hypothesis derived from study 1. $*P < 0.05$, $**P < 0.01$, and $***P < 0.001$. A significant effect of age was found only in study 2, which we further discuss in *SI Appendix*. (B) A reduction in information search mediated less accurate overall final judgments in more dogmatic participants (mediation results for study 1 are presented in the figure; see main text for results from study 2).

corrected direct effect: -0.036 , $P = 0.016$). More dogmatic participants also earned less money overall, indicating their lowered information seeking did not entail any strategic benefits (study 1: $\beta = -0.24$, $P = 0.008$, $R^2 = 0.02$; study 2: $\beta = -0.21$, one-tailed $P = 0.009$, $R^2 = 0.01$; pooled internal meta-analysis: $\beta = -0.23$, $P = 0.0003$, $R^2 = 0.017$).

Trial-by-Trial Modeling of Information Search. We next sought to develop a more detailed account of how dogmatic individuals' trial-by-trial information seeking choices were informed by confidence judgments and the cost of information. This model can be expressed as a logistic regression predicting the choice to seek information:

$$P(\text{Information Seeking}) = \frac{1}{1 + \exp(-(\beta_0 + \beta_1 * \text{Confidence} + \beta_2 * \text{Cost}))}$$

The three β 's capture three independent behavioral phenomena (Fig. 4 A and B): Differences in the model's intercept, β_0 , represent a general shift in willingness to seek out information; β_1 represents how strongly participants' information-seeking choices are influenced by confidence; and β_2 indicates the influence of information cost on subjects' willingness to seek out more information.

We were interested in whether any of these parameters were associated with individual differences in dogmatism, i.e., whether dogmatism was linked to a general tendency to seek out less information (β_0), a differential influence of confidence on information seeking (β_1), or an altered sensitivity to information costs (β_2). The relation between each parameter, β , and dogmatism was thereby estimated directly within a hierarchical framework, such that each individual's parameter (indexed by i) was a function of a group mean (μ_B) and their dogmatism score (Dogmatism_i). For instance, the relation between dogmatism and a general tendency to seek out less information can be formalized as follows:

$$\beta_{0,i} = \mu_{B_0} + \rho_0 * \text{Dogmatism}_i + \varepsilon_{B_0}$$

Here, ρ_0 describes the relation between dogmatism and β_0 , while ε_{B_0} represents individual variation in this parameter that is not explained by dogmatism. If the credible interval of ρ_0 does not include zero, this indicates a significant association between dogmatism and β_0 (32). In what follows, we report 95% credible intervals.

We found more dogmatic subjects had lower values of β_0 (Fig. 4D) 95% $\text{CI}_{\rho_0} = -0.40, -0.07$), in accordance with our model-agnostic findings that dogmatic participants show lower information-seeking behavior. While we found no association between dogmatism and cost sensitivity (95% $\text{CI}_{\rho_2} = -0.06, 0.07$), dogmatism was associated with higher values of the confidence parameter β_1 (95% $\text{CI}_{\rho_1} = 0.02, 0.27$). Because β_1 values were generally negative (see Fig. 4B for distribution of individual parameter values), this positive shift suggests that the information-seeking decisions of individuals with higher levels of dogmatism were less coupled to fluctuations in subjective confidence than those of individuals with lower levels of dogmatism. In other words, participants with higher levels of dogmatism were less likely to use feelings of confidence or uncertainty to guide their search for more information. Together, this dual shift in both β_0 and β_1 parameters combines to produce marked differences in information search under low confidence (high uncertainty). On these trials, individuals with lower levels of dogmatism were more likely to (adaptively) seek out new information compared to individuals with high levels of dogmatism. In contrast, participants with both

higher and lower levels of dogmatism showed similar profiles of information-seeking behavior when they were more confident in their decision (Fig. 4 E and F).

Information Search and Other Factor Scores. Given a long-standing debate over diverging cognitive profiles of liberals and conservatives (33–36), we also investigated the relationship between information search and political orientation. Here, we found that position on the political spectrum (right vs. left) was not predicted by a willingness to seek information (study 1: $\beta = -0.07$, $P = 0.19$; study 2: $\beta = -0.07$, $P = 0.23$; SI Appendix, Fig. S3). Additionally, there was no consistent association between the extremity of political opinion, as indexed by the absolute value of the political orientation, and information-seeking behavior (study 1: $\beta = 0.07$, $P = 0.21$; study 2: $\beta = -0.13$, $P = 0.02$; SI Appendix, Fig. S3). Similarly, policy-specific political belief superiority was not related to changes in information seeking (study 1: $\beta = 0.03$, $P = 0.62$; study 2: $\beta = -0.07$, $P = 0.24$; SI Appendix, Fig. S3).

Discussion

We show that dogmatic individuals are less likely to seek out additional information before committing to a decision. By foregoing this opportunity, they in turn tend to form less accurate overall judgments. Trial-by-trial modeling revealed that two factors drove dogmatic individuals' altered information seeking: 1) a shift in the general willingness to seek information and 2) a decreased influence of confidence on information-seeking behavior. Together, these effects gave rise to a distinct pattern: Whereas dogmatism had little effect on information seeking after high confidence decisions, more dogmatic subjects were less likely (relative to moderates) to seek out additional information when they were uncertain about their decision.

A key aspect of our results is that we find this disadvantageous pattern of information seeking in a low-level perceptual decision-making task. This stands in contrast to previous studies on information seeking in the political domain that have relied on questionnaires or experimental tasks with overt political content (27, 37). By capitalizing on the neutral valence and personal irrelevance of simple dot stimuli, we could isolate uncertainty-driven information-seeking behavior from possible confounding effects of motivated reasoning. Observing such an effect in this neutral setting is consistent with a proposal that domain-general cognitive factors contribute to real-world attitudes (38–40). Nevertheless, in most real-world decision-making scenarios, it is likely that both motivational and cognitive (uncertainty-driven) effects contribute to biases in information seeking (6), and it is interesting to consider that the latter may even become magnified in the presence of affective influences.

Our trial-by-trial model fits revealed that while participants generally use internal signals of uncertainty (as assayed by confidence ratings) to guide information search, dogmatic individuals did so to a lesser extent. This points to a general alteration in the way that confidence guides actions, a process usually described as metacognitive control (41). Metacognitive control is hypothesized to not only regulate information search, but also other phenomena in which effort must be weighed against accuracy, such as cognitive offloading (42) or speed-accuracy trade-offs (43). From a theoretical perspective, metacognitive control complements metacognitive monitoring (41), which describes a process that gives rise to, and updates, representations of confidence. However, while metacognitive monitoring has received considerable attention from a neural (44, 45) and individual differences perspective (24, 46), metacognitive control processes remain underinvestigated. Such research might therefore provide fruitful for understanding the drivers of altered information search. Further work is needed also to disentangle how different models of confidence formation [such as

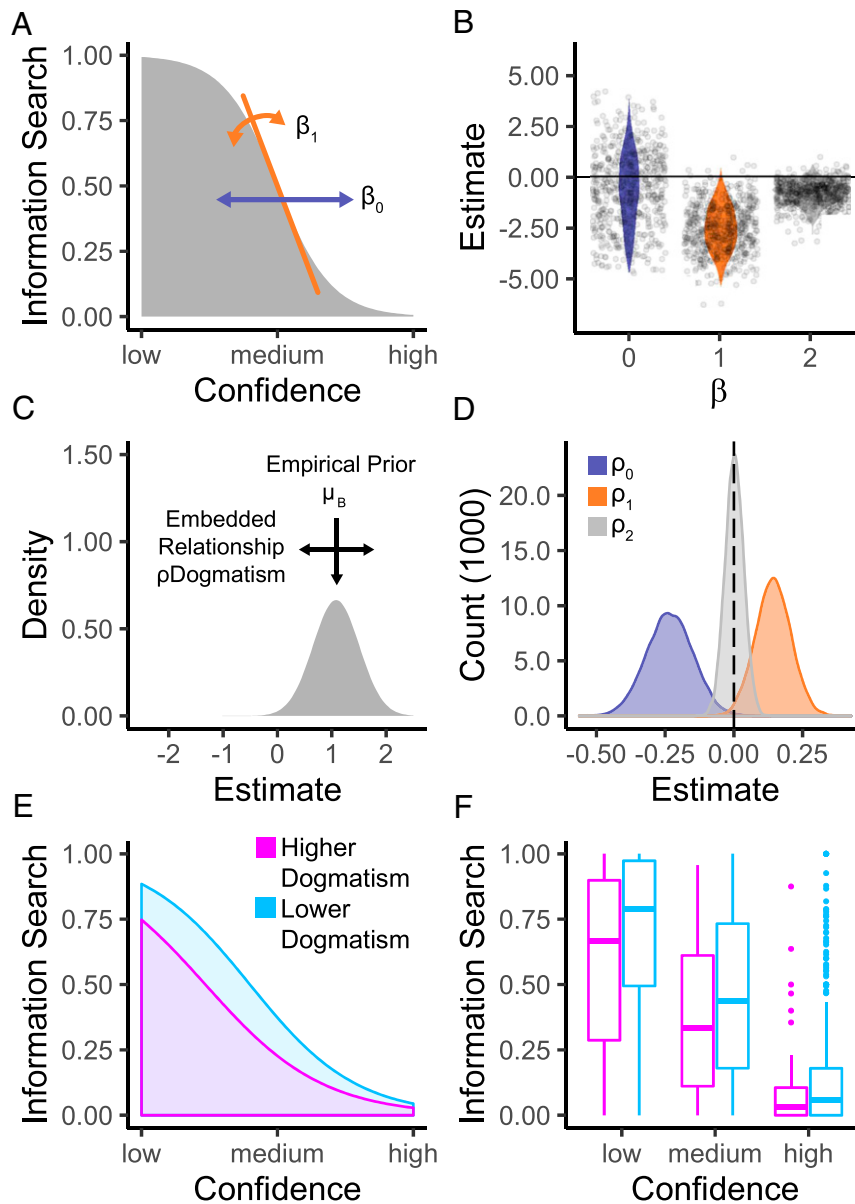


Fig. 4. Individual differences in confidence-driven information search as captured by a trial-by-trial model. (A) The decision to seek additional information was captured using a model with three parameters: an intercept β_0 , a confidence parameter β_1 , and a cost parameter β_2 (not depicted here). (B) Distribution of individual-level parameters displaying the generally negative influence of higher confidence (β_1) and higher cost (β_2) on information search. (C) We captured dogmatism-related differences in these model parameters through a hierarchical fitting procedure, whereby each parameter's empirical prior varies as a function of another set of parameters ρ that encode the influence of subjects' dogmatism scores in a hierarchical estimation scheme. (D) Posterior distribution of embedding parameters ρ encoding dogmatism-driven shifts in parameter means. We found a dogmatism-related decrease in the parameter capturing baseline information search (ρ_0) and an increase in the parameter capturing the tuning of information search to participants' confidence (ρ_1). No effect of dogmatism on the cost parameter was observed (ρ_2). The dotted vertical line represents a null effect. (E and F) Dogmatic individuals seek out less information than moderates when they are uncertain. To visualize this effect, we compared the 10% most dogmatic participants to the remainder of the sample. We plot (E) model predictions and (F) actual data (medians with upper/lower quartiles), averaged over both levels of information cost.

postdecisional or second-order architectures (47, 48)] affect both monitoring and control processes, and in turn determine the interplay of confidence and information search.

Dogmatic individuals were less likely to seek information in situations of uncertainty compared to their peers. At a single-trial level, this is consistent with basing a final judgment on less evidence, leading to less accurate judgments overall. Because uncertain decisions are also less likely to be correct, this meant dogmatic individuals were less likely to seek out contradictory evidence when they were wrong—a form of confirmation bias.

Over a longer time horizon, and in the absence of external feedback (49), such a self-reinforcing feedback loop might in turn lead dogmatic individuals to think that their initial judgments are already sufficiently optimal and that investing in acquiring more information is unnecessary. A useful extension of our work will be to investigate how dogmatic individuals manage information search in situations that span more than one trial and require iterative learning. In such scenarios, adequately managing the exploration/exploitation trade-off is central to effective learning (5, 8), such that small differences in a tendency

toward or against uncertainty-driven information search may summate and lead to skewed representations of reality.

While a psychophysical approach provides us with the precise control required to characterize dogmatic individuals' information search, our task is necessarily contrived relative to real-world decision problems. It remains unknown whether the types of search behavior observed here are representative of real-world search behavior, for instance on the internet (1). However, we can be cautiously optimistic about the generalizability of the current results, given the domain-general nature of our task and recent observations that real-life behavior adheres to cognitive models of uncertainty-based exploration (50). One difference between our paradigm and real-world decisions is the guaranteed helpfulness of future information. The calculus changes when a first source is trustworthy, but future information might be unreliable. In that case, it might be adaptive to rely more heavily on one's initial judgment, and refrain from seeking new information even when uncertain.

In sum, we highlight a generic resistance to seek out additional information in more dogmatic individuals, a difference that is most marked when initial decisions are uncertain. This is disconcerting in the current cultural landscape. While the internet has heralded access to a plethora of well-vetted information, fake news remains rife (1, 3). In such cases, the mere availability of correcting information might not be enough to prevent the formation of unsupported beliefs in dogmatic individuals, because even feelings of uncertainty would not trigger corrective information-seeking behavior. On a systemic level, such results suggest that the veracity of first contact with a news story is therefore critical (51, 52). On an individual level, instilling successful uncertainty-based search may be enabled by the extension of training of metacognitive monitoring (53) to also target metacognitive control. Finally, our research shows that psychophysical paradigms in conjunction with trial-by-trial modeling of behavior provide important tools for identifying mechanisms behind dogmatism, polarization, and their consequences (40).

Materials and Methods

Online Recruitment and Sample. Both studies were conducted online and recruited US adults through the online labor marketplace Amazon Mechanical Turk (54–56). They were approved by the Research Ethics Committee of University College London (#1260-003), and subjects gave informed consent.

In study 1, 370 subjects' data were analyzed (see *SI Appendix* for exclusion criteria). We based this sample size on previous studies conducted to detect interindividual differences in cognition across the political spectrum (24, 38) and in disorders (46). Subjects were paid a basic payment of \$1 and earned a bonus of up to \$6 based on their adequate completion of the questionnaires and their performance on the information-seeking task (see *Experimental Design, Information-seeking task* for a detailed description of the remuneration). Participants were 50% female (49.7% male, 0.3% "other/would rather not say") and the mean age was 36.62 y (SD, 11.61; range, 19 to 81 y; *SI Appendix, Fig. S1A*). In study 2 (replication), we analyzed data from 364 participants with the same payment scheme as in study 1. An a priori power analysis based on the information-seeking effect size from study 1 determined our sample size in study 2, giving us a power above 80% to detect the association between dogmatism and average information search. The sample consisted of 52% women (47% male, 1% other/"would rather not say"; mean age, 36.55 y; SD, 11.09; range, 18 to 74 y; *SI Appendix, Fig. S1A*). Participants in both studies came from a broad range of educational backgrounds, which was comparable to the general US adult population (*SI Appendix, Fig. S1B*).

Factor Analysis. There is significant debate about the exact structure of political ideology and its relationships to related constructs such as dogmatism (22, 25, 57). As used previously (24), here we administered multiple questionnaires measuring political orientation, identification with the two major US parties, the social and economic conservatism scale (58), as well as a questionnaire assessing specific policy positions and subjects' belief in the superiority of these positions (22). Additionally, right- and left-wing (59, 60) authoritarianism was assessed. Finally, participants also filled out a dogmatism questionnaire (21). We conducted a factor analysis (see *SI*

Appendix for a detailed discussion of the results) using the `fa()` function in the "Psych" R package on all 78 questionnaire items using maximum-likelihood estimation with an oblique rotation (`oblimin`). This mirrors methodology previously employed to study political beliefs (24) and mental health (46). We determined the number of factors through the Cattell–Nelson–Gorsuch test (61) where a sharp drop in the eigenvalues indicates the point at which there is little benefit to retaining additional factors. To maximize the precision of the factor loading estimates and the factor scores, we pooled the present sample with the one from Rollwage et al. (24) where subjects had completed the same questionnaire battery. This resulted in total sample of 2,135 participants for the factor analysis. We observed qualitatively similar pattern of factor loadings for both the pooled sample of 2,135 participants and the two individual samples.

Experimental Design.

Stimuli. We used the JavaScript library JsPsych (version 5.0.3) (62) to program the task and hosted the experiment on the online research platform Gorilla (63), which subjects could access through their browser. Two black squares, each 250 pixels in height and width, were presented as discrimination stimuli, with one square positioned left and one square right of center (see Fig. 3 for task overview). Each square consisted of 625 cells, randomly filled with white dots, so that one baseline square always held 313 dots and the other target square contained a greater number determined during a calibration phase (*SI Appendix*). During each dot-discrimination trial, subjects were presented with five such configurations for 150 ms each in order to create the impressions of flickering dots. Within each trial, the location of the individual dots per configuration within one square was random. However, the difference in number of dots between the target and baseline squares remained the same within each trial. The location of the target was pseudorandomized between trials.

Task and procedure. Both studies followed the same protocol and participants spent around 45 min on the experiment, which was divided into three parts. Participants first received information and reported their demographic information. Following this, they then first completed a 120-trial calibration phase to individually determine task difficulty (*SI Appendix*), identical to previous procedures (24). There, participants simply had to indicate which of the two flickering dots contained more dots by a press of the "2" or "6" key (indicating left and right) and received feedback about their correctness through a colored frame around their chosen option. This was followed by the information-seeking task (Fig. 2) in which subjects received no feedback about the correctness. The information-seeking task consisted of four blocks, each containing 25 trials. Participants then went on to fill out the aforementioned questionnaires.

Information-seeking task. Across the 100 trials of the information-seeking task, participants were presented with the stimulus strength determined in the calibration phase (study 1: mean, 73.80%; SD, 6.57%; study 2: mean, 73.67%; SD, 6.50%). As in the calibration phase, participants had to decide whether more dots were in the left or in the right box (the initial decision). Simultaneously, they indicated their confidence in this decision by pressing one of three buttons per side to indicate low, medium, or high confidence (the "1" to "3" or "5" to "7" keys in the number row). Crucially, the information-seeking task allowed participants to choose whether they wanted to see a second, additional display of the stimulus to improve the accuracy of their initial judgment. Subjects were specifically instructed about the helpful nature of this information. If they decided to see the stimulus again, the subjects saw a stronger version of the stimulus (i.e., one with a higher dot difference; *SI Appendix*). If they decided to forego the second stimulus, they were instead presented with two empty black boxes, to prevent them from artificially speeding through the task. Therefore, the only cost associated with the additional information was the deduction of points (5 points or 20 points, depending on the block). Regardless of whether subjects decided to see additional information or not, they then made another judgment (the final decision), indicating both the side they believed contained more dots and their confidence in this final decision (using the same response keys as for the initial decision). Importantly, we only incentivized the accuracy of this final decision: Subjects received 100 points for a correct and 0 points for an incorrect final decision.

Participants' bonus payment was linked to their performance in the task: They received a \$2 bonus for completing the task and an extra 4 cents for every 100 points they had earned on the task (average points-based bonus, study 1: mean, \$3.11; SD, \$0.34; study 2: mean, \$3.11; SD, \$0.35).

Statistical Analysis.

Task analysis. We conducted several analyses to ensure participants understood the task and were able to perform it adequately (see *SI Appendix, Fig.*

S1 for an overview). Within-participant effects (see *SI Appendix, Table S1* for an overview) were investigated using trial-by-trial hierarchical mixed effects models computed and analyzed in the “afex” package (64). Specifically, we constructed logistic models with binary outcomes as respective dependent variables and the corresponding predictors as fixed effects (see *SI Appendix, Table S1* for details). We included per-participant random intercepts and slopes and employed likelihood-ratio tests to obtain *P* values (65). To quantify relationships between subjects’ average information seeking and their final decision accuracy, we set up a general linear model using the `lm()` function in R. All analyses were performed separately for the two studies. **Statistical analysis.** We conducted the following regression analysis using the `lm()` function in R. All analyses were performed separately for the two studies, and effects were tested two-tailed if not stated otherwise.

- 1) To investigate the relationship between the factor scores themselves, we constructed polynomial regression models. Specifically, we built these models for each possible factor combination and compared 1) a linear fit, 2) a quadratic fit, and 3) a combined linear and quadratic fit based on their Bayesian information criterion (see *SI Appendix, Table S2* for an overview).
- 2) To investigate the relationship between information seeking and the factors observed through our questionnaire, we set up one generalized linear model per factor, explaining the respective variance in this factor score through participants’ average information seeking. Following previous work (24), we controlled for the following covariates: age, gender, education, subjects’ average performance and confidence level on the initial decision, objective stimulus strength (indicated by the logarithm of the dot difference), and performance on the stronger version of the stimulus (as recorded during the calibration phase; see *SI Appendix*). We standardized the continuous outcome and predictor variables to obtain standardized β coefficients. For significant variables of interest, we calculated R^2 values by comparing the variance explained by a full model including information seeking relative to a model excluding this predictor.
- 3) Finally, to check whether dogmatism was associated with a reduction in points earned on our task, we set up the same model used for the information-seeking analysis but replaced the information-seeking predictor with the points earned on the task.

To investigate whether dogmatism was linked to a reduction in final decision accuracy and whether this arose from a lowered propensity to seek out information, we conducted a mediation analysis. This analysis was conducted using the “mediate” package in R (66), which uses a quasi-Bayesian Monte Carlo method based on normal approximation to estimate the significance of the mediation effect (67). We again entered the covariates used for the original information-seeking analysis as control variables into all paths of the mediation analysis. To conduct an internal meta-analysis of behavioral results obtained across the two studies, we pooled the two samples as recommended by Braver et al. (31) and applied the same analysis as detailed above.

Trial-by-trial modeling. To probe the underlying mechanisms contributing to dogmatic individuals’ information search, we set up a trial-by-trial model that investigated the factors impacting an individual’s decision to seek out

more information. Specifically, we modeled the information-seeking choices as a function of the confidence level and the current information cost (see main text and *SI Appendix*).

Because classical maximum-likelihood-based methods can frequently provide noisy estimates with so few data points, we employed a hierarchical fitting procedure (68). In such a hierarchical model, individual parameters, $\beta_{0,i}$, are drawn from a group-level prior distribution. For example, for the first parameters, $\beta_{0,i}$, we can write as follows:

$$\beta_{0,i} \sim \mathcal{N}(\mu_{B_0}, \sigma_{B_0})$$

Here, μ_{B_0} represents the population mean that then informs the estimation of $\beta_{0,i}$, the individual parameters of β_0 for participant *i*, from a population distribution, $\mathcal{N}(\mu_{B_0}, \sigma_{B_0})$. Conventionally, parameters obtained through such an approach can then be correlated with an external measure of differences between individuals. However, this procedure is suboptimal because it assumes no variability in the mean of the population in the initial model fit, possibly distorting or minimizing potential relationships between the parameter and external factors (32). To maintain the advantages of hierarchical fitting while avoiding such pitfalls regarding individual differences, here we employ a procedure recently prescribed by Moutoussis et al. (32). There (Fig. 4C), the relationship between the parameters and individual differences is embedded into the estimation of the parameters themselves through the prior, so that:

$$\beta_{0,i} \sim \mathcal{N}(\mu_{B_0} + \rho_0 * \text{Dogmatism}_{i}, \sigma_{B_0})$$

To capture interindividual differences in the parameter, we allow the mean of the population distribution to vary as a function of dogmatism through the embedded parameter ρ_0 . To enable accurate hierarchical estimation, we pooled the samples from both studies and only included subjects that sought out information on at least 5% and at most on 95% of trials. In doing so, we achieved a total sample of 568 subjects. We built the model using the programming language Stan (69), which uses a form of Markov chain Monte Carlo sampling, Hamilton Monte Carlo sampling, to estimate posteriors over parameters.

Further details on model fitting and a discussion of the influence of different parameters on participant’s payoff are presented in *SI Appendix*.

Data Availability. Fully anonymized data are available from the corresponding author upon reasonable request. The final code for data analysis has been deposited in a dedicated GitHub repository (<https://github.com/metacoglab/SchulzRollwageDolanFleming>).

ACKNOWLEDGMENTS. We thank E. Schulz for comments on an earlier version of the manuscript and P. Dayan for helpful discussions on the research. The Wellcome Centre for Human Neuroimaging is supported by core funding from the Wellcome Trust (203147/Z/16/Z). L.S. was supported by a scholarship from the German Academic Scholarship Foundation. S.M.F. is supported by a Sir Henry Dale Fellowship jointly funded by the Wellcome Trust and the Royal Society (206648/Z/17/Z).

1. T. T. Hills, The dark side of information proliferation. *Perspect. Psychol. Sci.* **14**, 323–330 (2019).
2. L. Garrett, COVID-19: The medium is the message. *Lancet* **395**, 942–943 (2020).
3. D. M. J. Lazer et al., The science of fake news. *Science* **359**, 1094–1096 (2018).
4. M. C. Lemos, C. J. Kirchhoff, V. Ramprasad, Narrowing the climate information usability gap. *Nat. Clim. Chang.* **2**, 789–794 (2012).
5. E. Schulz, S. J. Gershman, The algorithmic architecture of exploration in the human brain. *Curr. Opin. Neurobiol.* **55**, 7–14 (2019).
6. T. Sharot, C. R. Sunstein, How people decide what they want to know. *Nat. Hum. Behav.* **4**, 14–19 (2020).
7. J. Gottlieb, P. Y. Oudeyer, Towards a neuroscience of active sampling and curiosity. *Nat. Rev. Neurosci.* **19**, 758–770 (2018).
8. T. T. Hills, P. M. Todd, D. Lazer, A. D. Redish, I. D. Couzin, Cognitive Search Research Group, Exploration versus exploitation in space, mind, and society. *Trends Cogn. Sci.* **19**, 46–54 (2015).
9. F. Gesiarz, D. Cahill, T. Sharot, Evidence accumulation is biased by motivation: A computational account. *PLoS Comput. Biol.* **15**, e1007089 (2019).
10. C. J. Charpentier, E. S. Bromberg-Martin, T. Sharot, Valuation of knowledge and ignorance in mesolimbic reward circuitry. *Proc. Natl. Acad. Sci. U.S.A.* **115**, E7255–E7264 (2018).
11. E. Jonas, S. Schulz-Hardt, D. Frey, N. Thelen, Confirmation bias in sequential information search after preliminary decisions: An expansion of dissonance theoretical research on selective exposure to information. *J. Pers. Soc. Psychol.* **80**, 557–571 (2001).
12. E. Bakshy, S. Messing, L. A. Adamic, Political science. Exposure to ideologically diverse news and opinion on Facebook. *Science* **348**, 1130–1132 (2015).
13. A. L. Schmidt et al., Anatomy of news consumption on Facebook. *Proc. Natl. Acad. Sci. U.S.A.* **114**, 3035–3039 (2017).
14. W. Hart et al., Feeling validated versus being correct: A meta-analysis of selective exposure to information. *Psychol. Bull.* **135**, 555–588 (2009).
15. P. Schwartenbeck et al., Computational mechanisms of curiosity and goal-directed exploration. *eLife* **8**, e41703 (2019).
16. P. Dayan, N. D. Daw, Decision theory, reinforcement learning, and the brain. *Cogn. Affect. Behav. Neurosci.* **8**, 429–453 (2008).
17. K. Desender, A. Boldt, N. Yeung, Subjective confidence predicts information seeking in decision making. *Psychol. Sci.* **29**, 761–778 (2018).
18. A. Boldt, C. Blundell, B. De Martino, Confidence modulates exploration and exploitation in value-based learning. *Neurosci. Conscious.* **2019**, niz004 (2019).
19. T. P. Newman, E. C. Nisbet, M. C. Nisbet, Climate change, cultural cognition, and media effects: Worldviews drive news selectivity, biased processing, and polarized attitudes. *Public Underst. Sci.* **27**, 985–1002 (2018).
20. H. Fischer, D. Amelung, N. Said, The accuracy of German citizens’ confidence in their climate change knowledge. *Nat. Clim. Chang.* **9**, 776–780 (2019).
21. B. Altemeyer, Dogmatic behavior among students: Testing a new measure of dogmatism. *J. Soc. Psychol.* **142**, 713–721 (2002).
22. K. Toner, M. R. Leary, M. W. Asher, K. P. Jongman-Sereno, Feeling superior is a bipartisan issue: Extremity (not direction) of political views predicts perceived belief superiority. *Psychol. Sci.* **24**, 2454–2462 (2013).

23. M. Rokeach, The nature and meaning of dogmatism. *Psychol. Rev.* **61**, 194–204 (1954).
24. M. Rollwage, R. J. Dolan, S. M. Fleming, Metacognitive failure as a feature of those holding radical beliefs. *Curr. Biol.* **28**, 4014–4021.e8 (2018).
25. J. W. van Prooijen, A. P. M. Krouwel, Extreme political beliefs predict dogmatic intolerance. *Soc. Psychol. Personal. Sci.* **8**, 292–300 (2016).
26. E. Fast, E. Horvitz, “Identifying dogmatism in social media: Signals and models” in *Proceedings of 2016 Conference Empirical Methods in Natural Language Processing*, J. Su, K. Duh, X. Carreras, Eds. (Association for Computational Linguistics, Stroudsburg, PA, 2016), pp. 690–699.
27. B. H. Long, R. C. Ziller, Dogmatism and predecisional information search. *J. Appl. Psychol.* **49**, 376–378 (1965).
28. K. Desender, P. Murphy, A. Boldt, T. Verguts, N. Yeung, A post-decisional neural marker of confidence predicts information-seeking in decision-making. *J. Neurosci.* **39**, 3309–3319 (2019).
29. R. Cross, D. A. Snow, Radicalism within the context of social movements: Processes and types. *J. Strateg. Secur.* **4**, 115–130 (2011).
30. R. Wintrobe, *Rational Extremism: The Political Economy of Radicalism* (Cambridge University Press, 2006).
31. S. L. Braver, F. J. Thoenes, R. Rosenthal, Continuously cumulating meta-analysis and replicability. *Perspect. Psychol. Sci.* **9**, 333–342 (2014).
32. M. Moutoussis, A. K. Hopkins, R. J. Dolan, Hypotheses about the relationship of cognition with psychopathology should be tested by embedding them into empirical priors. *Front. Psychol.* **9**, 2504 (2018).
33. N. J. Shook, R. H. Fazio, Political ideology, exploration of novel stimuli, and attitude formation. *J. Exp. Soc. Psychol.* **45**, 995–998 (2009).
34. J. T. Jost, Ideological asymmetries and the essence of political psychology. *Polit. Psychol.* **38**, 167–208 (2017).
35. J. Greenberg, E. Jonas, Psychological motives and political orientation—The left, the right, and the rigid: Comment on Jost et al. (2003). *Psychol. Bull.* **129**, 376–382, discussion 383–393 (2003).
36. M. E. Fiagbenu, J. Proch, T. Kessler, Of deadly beans and risky stocks: Political ideology and attitude formation via exploration depend on the nature of the attitude stimuli. *Br. J. Psychol.*, 10.1111/bjop.12430 (2019).
37. B. M. Tappin, G. Pennycook, D. G. Rand, Thinking clearly about causal inferences of politically motivated reasoning: Why paradigmatic study designs often prevent causal inference. *Curr. Opin. Behav. Sci.* **34**, 81–87 (2020).
38. L. Zmigrod, P. J. Rentfrow, T. W. Robbins, Cognitive underpinnings of nationalistic ideology in the context of Brexit. *Proc. Natl. Acad. Sci. U.S.A.* **115**, E4532–E4540 (2018).
39. L. Zmigrod, P. J. Rentfrow, T. W. Robbins, The partisan mind: Is extreme political partisanship related to cognitive inflexibility? *J. Exp. Psychol. Gen.* **149**, 407–418 (2020).
40. M. Rollwage, L. Zmigrod, L. de-Wit, R. J. Dolan, S. M. Fleming, What underlies political polarization? A manifesto for computational political psychology. *Trends Cogn. Sci.* **23**, 820–822 (2019).
41. T. O. Nelson, L. Narens, Metamemory: A theoretical framework and new findings. *Psychol. Learn. Motiv.* **26**, 125–173 (1990).
42. E. F. Risko, S. J. Gilbert, Cognitive offloading. *Trends Cogn. Sci.* **20**, 676–688 (2016).
43. K. Desender, A. Boldt, T. Verguts, T. H. Donner, Confidence predicts speed-accuracy tradeoff for subsequent decisions. *eLife* **8**, e43499 (2019).
44. S. M. Fleming, R. S. Weil, Z. Nagy, R. J. Dolan, G. Rees, Relating introspective accuracy to individual differences in brain structure. *Science* **329**, 1541–1543 (2010).
45. A. G. Vaccaro, S. M. Fleming, Thinking about thinking: A coordinate-based meta-analysis of neuroimaging studies of metacognitive judgements. *Brain Neurosci. Adv.* **2**, 2398212818810591 (2018).
46. M. Rouault, T. Seow, C. M. Gillan, S. M. Fleming, Psychiatric symptom dimensions are associated with dissociable shifts in metacognition but not task performance. *Biol. Psychiatry* **84**, 443–451 (2018).
47. S. M. Fleming, N. D. Daw, Self-evaluation of decision-making: A general Bayesian framework for metacognitive computation. *Psychol. Rev.* **124**, 91–114 (2017).
48. T. J. Pleskac, J. R. Busemeyer, Two-stage dynamic signal detection: A theory of choice, decision time, and confidence. *Psychol. Rev.* **117**, 864–901 (2010).
49. M. Rouault, P. Dayan, S. M. Fleming, Forming global estimates of self-performance from local confidence. *Nat. Commun.* **10**, 1141 (2019).
50. E. Schulz et al., Structured, uncertainty-driven exploration in real-world consumer choice. *Proc. Natl. Acad. Sci. U.S.A.* **116**, 13903–13908 (2019).
51. M. Del Vicario et al., The spreading of misinformation online. *Proc. Natl. Acad. Sci. U.S.A.* **113**, 554–559 (2016).
52. G. Pennycook, D. G. Rand, Fighting misinformation on social media using crowdsourced judgments of news source quality. *Proc. Natl. Acad. Sci. U.S.A.* **116**, 2521–2526 (2019).
53. J. Carpenter et al., Domain-general enhancements of metacognitive ability through adaptive training. *J. Exp. Psychol. Gen.* **148**, 51–64 (2019).
54. M. Buhmester, T. Kwang, S. D. Gosling, Amazon’s Mechanical Turk: A new source of inexpensive, yet high-quality data? *Perspect. Psychol. Sci.* **6**, 3–5 (2011).
55. J. J. Horton, D. G. Rand, R. J. Zeckhauser, The online laboratory: Conducting experiments in a real labor market. *Exp. Econ.* **14**, 399–425 (2011).
56. J. Chandler, P. Mueller, G. Paolacci, Nonnaïveté among Amazon mechanical turk workers: Consequences and solutions for behavioral researchers. *Behav. Res. Methods* **46**, 112–130 (2014).
57. A. Malka, Y. Lelkes, N. Holzer, “Rethinking the rigidity of the right model: Three suboptimal methodological practices and their implications” in *Frontiers of Social Psychology: Politics of Social Psychology*, J. T. Crawford, L. Jussim, Eds. (Psychology Press, 2017), pp. 116–135.
58. J. A. Everett, The 12 item social and economic conservatism scale (SECS). *PLoS One* **8**, e82131 (2013).
59. F. Funke, The dimensionality of right-wing authoritarianism: Lessons from the dilemma between theory and measurement. *Polit. Psychol.* **26**, 195–218 (2005).
60. A. Van Hiel, B. Duriez, M. Kossowska, The presence of left-wing authoritarianism in Western Europe and its relationship with conservative ideology. *Polit. Psychol.* **27**, 769–793 (2006).
61. R. L. Gorsuch, *Factor Analysis* (Lawrence Erlbaum, 1983).
62. J. R. de Leeuw, jsPsych: A JavaScript library for creating behavioral experiments in a Web browser. *Behav. Res. Methods* **47**, 1–12 (2015).
63. A. L. Anwyll-Irvine, J. Massonnié, A. Flitton, N. Kirkham, J. K. Evershed, Gorilla in our midst: An online behavioral experiment builder. *Behav. Res. Methods* **52**, 388–407 (2020).
64. H. Singmann et al, Package “afex” (2018). <https://cran.r-project.org/web/packages/afex/index.html>. Accessed 1 October 2018.
65. H. Singmann, D. Kellen, “An introduction to mixed models for experimental psychology” in *New Methods in Cognitive Psychology*, D. H. Spieler, E. Schumacher, Eds. (Psychology Press, London, 2019), pp. 4–31.
66. D. Tingley, T. Yamamoto, K. Hirose, L. Keele, K. Imai, Mediation: R package for causal mediation analysis. *J. Stat. Softw.* **59**, 1–38 (2014).
67. K. Imai, L. Keele, D. Tingley, A general approach to causal mediation analysis. *Psychol. Methods* **15**, 309–334 (2010).
68. B. Efron, *Large-Scale Inference: Empirical Bayes Methods for Estimation, Testing, and Prediction* (Cambridge University Press, 2012), vol. 1.
69. B. Carpenter et al., Stan: A probabilistic programming language. *J. Stat. Softw.* **76**, 1–32 (2017).