

EMPIRICAL ARTICLE

Sincere or motivated? Partisan bias in advice-taking

Yunhao Zhang  and David G. Rand

MIT Sloan, Cambridge, MA, USA

Corresponding author: Yunhao Zhang; Email: zyhjerry@mit.edu

Received: 1 April 2023; **Revised:** 30 June 2023; **Accepted:** 16 July 2023

Keywords: partisan bias; political motivated reasoning; advice-taking

Abstract

Political divisions have become a central feature of modern life. Here, we ask whether these divisions affect advice-taking from co- and counter-partisans in a nonpolitical context. In an incentivized task assessing the accuracy of nonpolitical news headlines, we find partisan bias in advice-taking: Democratic participants are less swayed by (accurate) information that comes from Republicans compared to the same information from Democrats (Republican participants display no such bias). We then adjudicate between two possible mechanisms for this biased advice-taking: a preference-based account, where participants are motivated to take less advice from counter-partisans because doing so is unpleasant; versus a belief-based account, where participants sincerely believe co-partisans are more competent at the task (even though this belief is incorrect). To do so, we examine the impact of a substantial increase in the stakes, which should increase accuracy motivations (and thereby reduce the relative impact of partisan motivations). We find that increasing the stakes does not reduce biased advice-taking, hence no evidence to support the bias is driven by preference. Consistent with the belief-based account, we find that Democratic participants (incorrectly) believe their co-partisans are better at the task, and this incorrect belief is much less severe among Republican participants. Further supporting the notion that the stated beliefs are sincere, raising the stakes of the belief elicitation of relative partisan competence does not affect the stated beliefs. Finally, participants—instead of ignoring the feedback—actually substantially update in favor of their counter-partisans given feedback that suggests counter-partisans are competent.

1. Introduction

We live in an era of great political division (Bertrand & Kamenica, 2018; Evans & Fu, 2018; Finkel et al., 2020; Iyengar, 2021; Prior, 2013). A large body of work has studied the impact of partisanship on information processing in political contexts, such as political news consumption (Bail et al., 2018; Broockman & Kalla, 2022; Cinelli et al., 2021; Levy, 2021; Peterson et al., 2021), political factual beliefs (Bullock et al., 2015; Khanna & Sood, 2018; Peterson and Iyengar, 2021), and policy preferences (Hawkins & Nosek, 2012; Kahan, 2016). Across these varied political contexts, researchers regularly observe substantial effects of shared partisanship, wherein participants are more receptive to political information that is itself ideologically congenial or that comes from congenial sources. Yet politics occupies only a small fraction of most people's lives. What effect do political identities have on interactions outside the domain of politics? For example, a great deal of the news that people consume and use to inform their daily lives is nonpolitical, such as news related to products, health, and entertainment. Here, we investigate the potential for such information frictions between counter-partisans: Are people less receptive to nonpolitical information from counter-partisans? And if so, is this friction caused by preference-based motivated reasoning or by an accuracy-driven motive that

discounts information from a source sincerely perceived as being less competent? In addressing these questions, we make two main contributions.

First, we demonstrate that in the context of assessing the accuracy of nonpolitical news headlines, participants are indeed less swayed by the same (accurate) information when it comes from a counter-partisan compared to from a co-partisan—although this effect was only obtained for Democratic participants and not Republican participants. Most related to the current paper, Marks et al. (2019) found that people are more likely to solicit advice from co-partisans. They further conducted an analysis on the subsequent advice-taking behavior after selection. Since they do not experimentally vary whether participants received information from co-partisans versus counter-partisans, their results cannot assess the causal effect of advisor identity on advice-taking due to selection bias. To the best of our knowledge, we are the first to properly explore the effect of shared partisanship on advice-taking¹ in a nonpolitical context.²

Second, and more importantly, we differentiate between two alternative explanations for this partisan bias in advice-taking: one based on preferences and the other based on beliefs.³ By the ‘preference-based’ account, people are motivated reasoners who dislike taking advice from counter-partisans relative to co-partisans, regardless of their perceived competence of the advisor. For instance, even if subjects ultimately believe that co-partisans and counter-partisans are equally competent, they may choose to discount counter-partisan opinions—and therefore, intentionally sacrifice their accuracy and monetary payoff—to avoid the disutility that comes from siding with counter-partisans (e.g., agreeing to a counter-partisan opinion makes one feel uncomfortable even though one knows the opinion is accurate). By the ‘beliefs-based’ account,⁴ subjects take advice based on their honest assessment of advisor competence and are otherwise perfectly happy to take advice from counter-partisans. Any discounting of counter-partisan advice is subjectively optimal for maximizing accuracy and the monetary payoff, because the advisees sincerely believe that counter-partisans are less competent than co-partisans. Distinguishing between these two accounts is crucial for understanding how to mitigate the observed bias because different mechanisms require different solutions. However, as pointed out by recent literature (Baron & Jost, 2019; Little, 2021; Tappin et al., 2020a,b), past literature tends to ascribe all biased information processing and source selection to (preference-based) political motivated reasoning, without further clarification of the nature of the motivation.⁵

How, then, can these two accounts be distinguished?⁶ We propose raising the stakes on accurate advice-taking (e.g., high incentive on the updated response). Based on the aforementioned definitions, there is a natural tradeoff between the two accounts: while the financial incentive creates incentives for belief accuracy, the preference-based account sometimes creates a disincentive that distorts one’s advice-taking. Therefore, when the stake on accurate advice-taking is sufficiently high, accurate advice-taking based on true beliefs becomes dominant because intentionally distorting the updated response

¹Following the tradition in the advice-taking literature (see, e.g., Bonaccio & Dalal, 2006 for a review), we refer to this paradigm of belief-updating given a numerical signal from others as ‘advice-taking’. Other fields may call it ‘belief-updating given a signal’ or ‘reaction to (social) influence’.

²Prior work in Psychology has examined how advice-taking is affected by factors, such as gender (Conlon et al., 2021), similarity in preferences (Yaniv et al., 2011), general demographic information (Gino et al., 2009), or mental states (Faraji-Rad et al., 2015). Prior work in Political Science has examined how partisanship affects which information sources people seek out (Bail et al., 2018; Levy, 2021) and tend to believe (Bullock & Lenz, 2019), but it has largely focused on political information and politicized topics, rather than nonpolitical contexts.

³The preference-based account we articulate here is related to ‘taste-based’ discrimination, whereby the bias is rooted in preferences driven by animus (Coffman et al., 2021; Phelps, 1972). The belief-based account, however, differs importantly from statistical discrimination, which is rooted in rational beliefs about average group competence.

⁴Motivated beliefs or beliefs driven by expressive responding (Bromberg-Martin & Sharot, 2020; Golman et al., 2017; Kunda, 1990; Sharot, 2011; Zimmermann, 2020), are also possible and they belong to ‘preference-based’ account. The design of our secondary experiment allows us to examine whether beliefs are sincerely-held and reported.

⁵People could hold a sincere belief of a co-partisan being more competent (even though such a belief is incorrect). Then favoring the co-party is rational and not due to any emotional desires to discriminate against counter-partisans.

⁶Little (2021) has pointed out that it is hard to use departure from a Bayesian benchmark to identify political motivated reasoning because of difficulty with establishing and interpreting a Bayesian benchmark.

due to ‘preference’ would be very costly. The underlying assumption is that the gain for getting the right answer (e.g., utility gain from higher expected monetary payoff) is enough to offset the negative utility from heeding the counter-partisan—leading to less discounting of counter-partisan advice relative to the co-partisan advice. In contrast, if discrimination against counter-partisan advice is driven by the sincere belief that counter-partisans are less competent, then increasing the task’s monetary payoff will not relatively reduce the discounting of counter-partisan advice. The rationale is that higher stakes would not affect one’s perception of advisor competence if the beliefs are sincere in the first place. Very importantly, one may be concerned that higher stakes increase cognitive efforts (Enke et al., 2023) or lower ego-centric discounting, which might lead to more advice-taking. Nevertheless, effects like these should be independent of the advisor’s identity, and thus do not affect the interaction of incentive and advisor identity (e.g., these effects are canceled out in the interaction, see Section 8 of the Supplementary Material for an illustration in the regression format). Therefore, only when participants take more advice from their counter-partisans when the stakes become higher, but do not do so to the same extent facing a co-partisan advisor, we will have clear evidence for the presence of the preference-based account.

Our initial expectation was that biased advice-taking is driven by the preference-based account, as implied by literature on political motivated reasoning (Iyengar, 2021; Iyengar et al., 2019; Kahan, 2016; Khanna & Sood, 2018; Lord et al., 1979; Peterson and Iyengar, 2021). Therefore, we expected raising the stakes to reduce partisan bias in advice-taking. Contrary to this expectation and the preference-based account, however, we did *not* find a significant effect of raising the stakes on biased advice-taking. This suggests that Democrats are less swayed by the advice of counter-partisans to a similar extent in the low-stake and high-stake conditions (and Republicans were equally receptive to advice from co-partisans and counter-partisans regardless of stakes). Consistent with the beliefs-based account, when we later elicit subjects’ beliefs about the relative performance of Democratic and Republican subjects, Democrats believe that counter-partisans give worse advice than co-partisans—and this biased perception is not affected by varying incentives. Consistent with their lack of bias in the main task, Republicans showed little difference in beliefs about the competence of co- versus counter-partisans. Finally, in an auxiliary follow-up experiment aiming at providing additional support for our main results, we find that Democrats’ incorrect beliefs are not deeply entrenched: when we provide feedback about the performance of co-partisan and counter-partisans, they update their priors accordingly. This suggests that the discounting of counter-partisan advice may be mitigated if subjects are exposed to information that counter-partisans are actually competent.

2. Experimental design

In this section, we describe the task, procedures, and conditions of our experiment. Full experimental materials are provided in Supplementary Material. Our study design is pre-registered⁷ at <https://aspredicted.org/r8yu2.pdf>.

2.1. Participants

Our goal was to recruit an online sample of 1,600 left-leaning Democrats and 1,600 right-leaning Republicans. To identify partisans, we first used the pre-screening options provided by Prolific and CloudResearch to recruit subjects who identified as either Democrats or Republicans. We then did an initial screening based on demographic questions at the outset of the study, only allowing subjects who self-identified in our survey as Democratic and left-leaning, or as Republican and right-leaning, to proceed. We began by recruiting Amazon Mechanical Turk workers using CloudResearch, and then after $N = 2,342$ subjects had been recruited, switched to Prolific as CloudResearch was unable to provide more subjects.⁸

⁷Some of the analysis departs from our pre-registration as suggested by the editor. Nevertheless, all results are qualitatively the same if we follow our pre-registered analysis.

⁸The data collection for right-leaning Republicans was unexpectedly slow after we had recruited 2,200 subjects—fewer than 20 qualified subjects per day on CloudResearch. To achieve our pre-registered sample size, we turned to Prolific. To avoid double

A total of 3,206 subjects (1,601 Republicans and 1,605 Democrats) completed our study from January 12, 2021, to February 9, 2021, and were included in the analysis. The average time to complete the survey was 5.8 min. The average total earnings per subject was \$2.16 (\$0.80 fixed payment plus an average of \$1.36 additional performance-based bonus); the total cost of our study was \$9,915.84. Among the 1,601 Republicans, the average age is about 41 years old, and about 51% are female. Among the 1,605 Democrats, the average age is about 37 years old, and about 58% are female.

2.2. Procedure

After signing a consent form, subjects provided their worker ID and answered basic demographic questions of political party, political orientation, age, gender, education, income, and voting behavior in the 2020 U.S. general election. Then they completed the news rating task, followed by the partisan competency prior elicitation task, a manipulation check, and then the partisan competency belief-updating task. Finally, subjects answered a few exploratory questions before finishing the study.

2.2.1. Main task: News rating task

To assess the effect of partisanship on advice-taking in an (almost) non-politicized and (almost) nonpolitical context with practical relevance, we use a news rating task. We selected a set of 10 nonpolitical news headlines from Allen et al. (2021). Each headline's veracity had been rated by three professional fact-checkers using a seven-point scale from 1 = 'definitely false' to 7 = 'definitely true'. We take the average of the three fact-checker ratings, rounded to the nearest integer, as ground truth. We provide the list of headlines and ground-truth veracity ratings in Section 2 of the Supplementary Material. To ensure the headlines are nonpolitical, we pre-tested these headlines and found that Democrats and Republicans did not have any sizable difference in attitudes toward the headlines in their initial ratings.⁹ To measure the extent to which subjects are influenced by information from others—our key outcome—we introduce our procedure using one of the news headlines as an illustration.

First, subjects see the headline "*No One Should Be Doing Keto Diet' Says Leading Cardiologist // Dr. Kim Williams says the science behind the fad diet is 'wrong'*" and provide their initial veracity rating on a seven-point scale. In addition, the initial response is always incentivized with a low monetary bonus, which is specified as "*This question is worth \$0.01*". Next, subjects proceed to the next page and are shown a rating (which we will refer to as the 'influence') from another participant (which we will refer to as the 'advisor'). To assess the impact of partisanship on advice-taking, and to test the preference-based account of partisan bias, we vary the advisor's political identity and the incentive level for the subject's final rating. Subjects see a message that says "*Another [Republican participant who voted for Donald Trump in the 2020 election] / [Democratic participant who voted for Joe Biden in the 2020 election] gave a rating of X. Please rate the story again.*" The influence X is set to be the average fact-checker rating (truth) of each headline.¹⁰ In addition, we also specify the bonus on the final response with an additional message: "*This question is worth \$Y*", where Y is either \$10 (high-incentive condition) or \$0.01 (low-incentive condition). Subjects then rate the headline's veracity a second time (the final rating) on the same seven-point scale on the same page.

Both the initial and the final ratings are incentivized using a quadratic scoring rule against the fact-checker average rating: $\max\{\text{Maximum Bonus} - 2*|\text{Rating} - \text{Truth}|^2, 0\}$; one of the two ratings (initial or final) is chosen at random for additional bonus payment.¹¹ For example, if the maximum bonus is

entry to the survey, we asked subjects from Prolific to (optionally) provide their MTurk ID if they have one. We excluded $N = 33$ Prolific entries which had MTurk IDs that appeared in the previous CloudResearch sample. However, all results are robust if we perform the analysis using only MTurk subjects.

⁹Indeed, examining our data with 3,200 subjects finds no significant differences in initial veracity ratings for Democrats versus Republicans on 8 out of 10 headlines (see Table 1 in Section 2 of the Supplementary Material). All our main results are robust to excluding the two headlines where there were significant partisan differences in initial ratings (see Section 7 of the Supplementary Material).

¹⁰ $X = 5$ for the example headline.

¹¹The minimum additional bonus is \$0 by default as no negative bonus is allowed on the online survey platforms.

\$10 and the truth is '5', a rating with an absolute error of 1 corresponds to an \$8 bonus; a rating with an absolute error of 2 corresponds to a \$2 bonus; a rating with an absolute error of 3 or larger corresponds to \$0 bonus. Therefore, subjects are incentivized to provide a rating as close as possible to the fact-checker average and have no incentive to hedge their answers.

To summarize, Democratic (Republican) subjects will view the Democratic (Republican) advisor as a co-partisan. While the initial rating always uses a \$0.01 maximum bonus (low incentive), subjects are randomly assigned to have their final rating worth either a \$0.01 (low incentive) or \$10 (high incentive) maximum bonus. Thus, we use a 2×2 between-subject design: [co-partisan advisor, counter-partisan advisor] \times [low incentive for final rating, high incentive for final rating]. As described earlier, if preference-based motivated reasoning exists, the coefficient of this interaction should be positive. Critically, the actual influence shown to subjects for each news headline does not vary across conditions, but is always set to be the fact-checker average rating. Thus, subjects would maximize their payoffs by changing their final rating to match the influence. By holding the influence constant, our design allows us to investigate how subjects respond to the same piece of high-quality advice under different advisor identities and incentive levels. To avoid deception, we ran a large pilot that included at least one Democrat voting for Biden and Republican voting for Trump who also gave the fact-checker average rating for each headline, and reported the ratings of these specific pilot subjects.

2.2.2. Secondary task 1: Partisan competency prior belief elicitation task

In addition to observing if and how subjects update their responses in the main news sharing task, we also directly examined their beliefs about the relative competence of co- versus counter-partisans. Specifically, we asked whether, out of 100 pairs of Republican and Democratic voters from a previous study, they thought the Republicans or Democrats performed better overall, as well as the probability (between 50% and 100%) that their response to this question about relative performance was correct (see Figure 1.4 in Section 1 of the Supplementary Material for a screenshot).

To investigate potential expressive responding or motivated beliefs about partisan competence, we also randomly vary the incentive level on belief accuracy. For example, when participants are answering “Which party do you think had a better performance in the news tasks” and “What’s the probability your answer is correct”, they are told “A correct guess and an accurate probability is worth \$Z”, where Z is either \$10 (high incentive) or \$0.01 (low incentive). The value of Z is re-randomized here and is independent of condition in the previous news rating task. Subjects’ answers are scored using the brier scoring rule: $\max\{\text{Maximum Bonus} - 10 \times |\text{probability} - \text{correct probability}|^2, 0\}$. Subjects in the high-incentive condition have a maximum bonus amount of \$10 while subjects in the low-incentive condition have a maximum bonus amount of \$0.01. To construct the prior belief measure, we use the subject’s stated probability (between 50% and 100%) if they believed the other party would do better, and 100% minus their stated probability if they believed their own party would do better. Since the 100 Republicans were marginally more accurate in our pilot study, the response that maximizes the bonus is *Republican* and 100%.

2.2.3. Secondary task 2: Partisan competency belief-updating task

In another secondary task, we assess the extent to which these beliefs about the relative competency of co- versus counter-partisans are malleable. After subjects state their priors, we provide them with performance feedback: “We select 20 pairs from the 100 previous comparisons. In 17 out of 20 pairs, [Republicans] / [Democrats] had more accurate news judgments than [Democrats] / [Republicans]”. In other words, we give subjects a noisy signal which suggests the group of 100 co-partisans or 100 counter-partisans might have performed better. Since the feedback partially informs the subjects about the relative competence of 100 Republicans and 100 Democrats, subjects should update their prior belief given the feedback. Subjects are then given a manipulation check to see whether they read and understand the provided feedback. Only those who passed the manipulation check (over 97% of

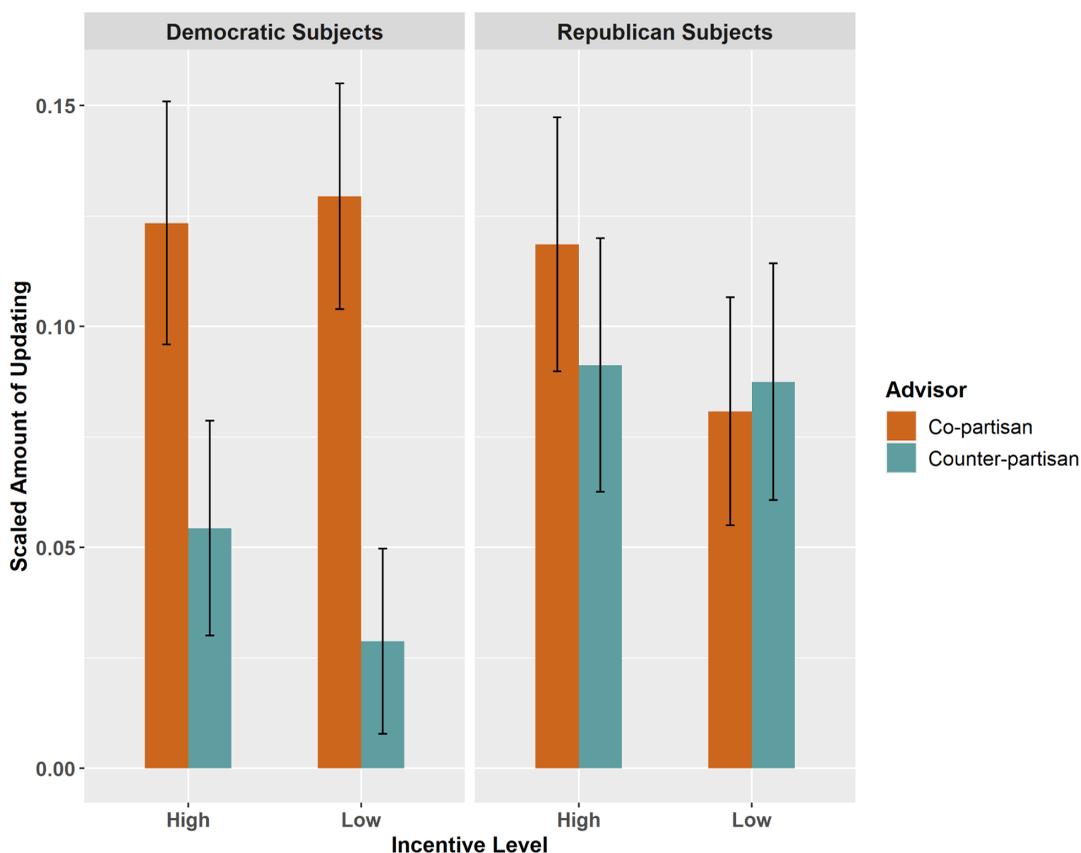


Figure 1. Shown is the Democratic subjects' (left panel) and Republican subjects' (right panel) average scaled amount of updating in the news rating task $[(\text{initial distance} - \text{updated distance}) / (\text{initial distance} + 1)]$, as a function of whether the advisor is a co-partisan (green) or counter-partisan (orange), and whether the incentive for the final rating is low (\$0.01) versus high (\$10). Error bars indicate 95% confidence intervals. We remove subjects whose initial distance is 0 (i.e., no room for updating).

participants) were allowed to continue¹². Finally, we ask the subjects to answer the same two belief elicitation questions as in the prior belief elicitation task. The incentive is \$0.01 for all subjects.

2.3. Design recap

To summarize, our design helps to distinguish between the preference-based and belief-based accounts of biased partisan advice-taking in the following ways. First, the 2×2 design in our News Veracity Task tests a key prediction of the preference-based account: to the extent that preferences are driving the discounting of counter-partisan influence, increasing the incentive for accurate responding should reduce preference-driven bias (i.e., interaction between high incentive and counter-partisan advisor identity should be positive.)¹³ Second, the Partisan Competence Belief Elicitation tests a key prediction of the belief-based account: that subjects will believe that co-partisans are more competent than counter-partisans. This task also allows us to test whether such a belief, if held, is sincere versus

¹²All results are robust if we include the 89 subjects who fail this attention check.

¹³A higher incentive may offset the disutility of following counter-partisan advice, which leads to greater updating specifically when the advisor is a counter-partisan.

motivated (by examining whether it is affected by raised incentives). Finally, the partisan competence belief-updating task assesses the malleability of subjects' beliefs. If partisan bias is purely based on sincere (but incorrect) beliefs, subjects should be willing to correct their beliefs when informed that counter-partisans are competent. Conversely, if subjects are motivated to believe that counter-partisans are inferior, they should be resistant to such information, especially given the low stakes.

3. Results

3.1. Biased partisan advice-taking

We begin by testing for partisan bias in advice-taking in the news rating task, and asking how such bias varies based on the incentive level. Specifically, as suggested by a reviewer, we construct our dependent variable as '(initial distance – updated distance) / (initial distance + 1)'.¹⁴ For example, imagine a subject with an updated rating of 5 given advice 5, her scaled amount of updating is 0.67 if her initial rating is 7 and 0.5 if her initial rating is 6. This measure better accounts for the distance between the initial rating and the advice with a correction for continuity. If we were to use the weight on advice measure¹⁵ (Logg et al., 2019), then the measure in the latter case would be 100%, where in fact, the advice might only have a small influence (e.g., the advice changes one's belief of the correct rating being 5 from 10% to 60%). Reassuringly, our results are qualitatively the same whether we use weight on advice or absolute amount of updating¹⁶ as the dependent variable (see Section 3 of the Supplementary Material).

For our main analysis, we regress the dependent variable on the following independent variables (all predictors in this paper are centered): an advisor-type indicator (1 = counter-partisan advisor, –1 = co-partisan advisor), an incentive-level indicator (1 = high incentive, –1 = low incentive, the subject's degree of right-wing extremity¹⁷, and all two- and three-way interactions. We remove subjects whose initial distance is 0 (about 13.8% of subjects' initial ratings are the same as the influence). Reassuringly, our results (including results from using the alternative dependent variables) are qualitatively the same when we include subjects whose initial distance is 0 (see Section 3 of the Supplementary Material).

We find a significant negative coefficient on the advisor-type indicator ($b = -0.017$, $p = 0.001$)¹⁸, such that in the low-incentive condition, subjects update less when the influence comes from a counter-partisan advisor.¹⁹ This establishes the basic phenomenon of interest: biased advice-taking based on partisanship exists, even though the task domain is largely nonpolitical. We also observe a significant two-way interaction between the advisor-type indicator and subject's degree of right-wing extremity ($b = 0.006$, $p = 0.0002$) such that in the low-incentive condition, Democratic subjects displayed significantly larger biased advice-taking than Republican subjects. In fact, when analyzing the two partisan subgroups separately, while the coefficient on the advisor-type indicator is highly significant for Democrats ($b = -0.04$, $p < 0.0001$), it is non-significant for Republican subjects ($b = -0.005$, $p = 0.46$). To contextualize the magnitude of the effect for Democrats in the low-incentive condition, the (scaled) amount of updating based on a Republican advisor was 68.9% smaller than for a Democratic advisor.

Next, we examine the effect of raising the stakes. Although the coefficient on the incentive-level indicator is not significant ($b = 0.008$, $p = 0.14$), it being positive suggests that subjects are 'updating more' when facing higher monetary incentive. Furthermore, the coefficient is significantly positive ($b = 0.02$, $p = 0.006$) and positive ($b = 0.045$, $p = 0.06$) if we use weight on advice and

¹⁴Initial (updated) distance means distance between the initial (updated) rating and the advice.

¹⁵The amount of update normalized by the distance between the initial answer and advice.

¹⁶The absolute amount one updates in the direction of the advice.

¹⁷Constructed as political leaning minus 4 for Democratic subjects and political leaning minus 6 for Republican subjects. Therefore, values of –3(3), –2(2), –1(1) correspond to extremely left (right), left (right), and moderately left (right) leaning. Our results are robust if we use binary classification of subject's partisanship as the independent variable.

¹⁸ b is the original regression coefficient.

¹⁹This coefficient is negative in 8 out of 10 headlines when we run this regression on each news headline separately.

absolute amount of updating—instead of the scaled amount of updating—as the dependent variable.²⁰ These results indicate that subjects respond to co-partisan influence more when the stakes are higher. This may serve as a manipulation check which suggests that the \$10 high incentive can affect subjects' decision-making.

Next, we examine the key test for the preference-based account of biased advice-taking: Does increasing the incentive reduce the discounting of counter-partisan advice relative to co-partisan advice? We first look at results pooling all Democratic and Republican subjects. We do not find a positive interaction between the advisor-type indicator and the incentive-level indicator ($b = -0.015$, $p = 0.84$).²¹ Nor do we find a significant three-way interaction ($b = 0.001$, $p = 0.60$). Furthermore, even when only examining Democrats (the subgroup who showed counter-partisan bias in the low-incentive condition), we continue to find no significant interaction between the advisor-type indicator and the incentive-level indicator ($b = 0.004$, $p = 0.51$).²² The analytical results above are consistent with a visual inspection of Figure 1 (i.e., biased advice-taking among Democrats but not Republicans; although raising the stakes increase advice-taking, it does not reduce the gap between advice-taking given co-partisan versus counter-partisan advice.).

This lack of effect of raising the stakes on the discounting of counter-partisan advice is inconsistent with the preference-based argument. If a preference to avoid counter-partisan advice sits in opposition to an economic motive for accuracy, increasing the accuracy motive would have decreased the discounting of counter-partisan advice. Thus, our empirical results from the news rating task do not support the preference-based account as a major driver of the biased advice-taking observed among Democrats.

3.2. Prior beliefs about partisan competence

We now turn to the partisan competency prior belief elicitation task to test key predictions of the belief-based account of biased advice-taking. Given our main results, the beliefs-based account would predict that, on average, Democratic subjects should believe that their co-partisans are more competent than counter-partisans, but such belief should be much weaker (or absent) among Republican subjects. Furthermore, if such beliefs are motivated, then raising the stakes on the belief elicitation should reduce the gap between co-partisans and counter-partisans: the higher incentive for stating an accurate prior should push subjects to reveal their suppressed prior belief (as well as mitigating any expressive responding). Conversely, if the beliefs are sincere, then raising the stakes should have no effect.

Consistent with the results in the main task, Figure 2 shows that although both Democratic and Republican subjects expected co-partisans to perform better on the task than counter-partisans (i.e., the mean belief of co-partisan being better than counter-partisan is greater than 50%), this belief is much stronger among Democrats than among Republicans.²³ Furthermore, the belief of prior partisan competence was not affected by the incentive level, suggesting sincere rather than motivated beliefs.

To examine the results analytically, we regress subjects' estimate of the probability that their co-partisans outperform the counter-partisan as the dependent variable on an incentive-level indicator (1 = high incentive on prior belief task, -1 = low incentive on prior belief task), subject's degree of right-wing extremism, and their interaction. First, we find a significant negative effect of the subject's degree of right-wing extremism, such that being more right-leaning is associated with less co-partisan favoritism in the prior belief task ($b = -1.89$, $p < 0.0001$). This is consistent with Republicans showing less partisan bias in the news rating task as depicted in Figure 2. Second, we find no significant effect of incentive level ($b = -0.03$, $p = 0.93$), which is consistent with the beliefs being sincere rather than motivated. We also note that believing one's co-partisans are more competent, although apparently

²⁰See Table 3 in Section 3 of the Supplementary Material.

²¹In addition, we find no significant effect on the incentive-level indicator ($b = -0.003$, $p = 0.62$); or significant interaction with subject's degree of right-wing extremism ($b = 0.001$, $p = 0.50$).

²²This interaction is also not significant when examining only Republican subjects ($b = 0.01$, $p = 0.15$).

²³See Section 4 of the Supplementary Material for the histograms of prior beliefs.

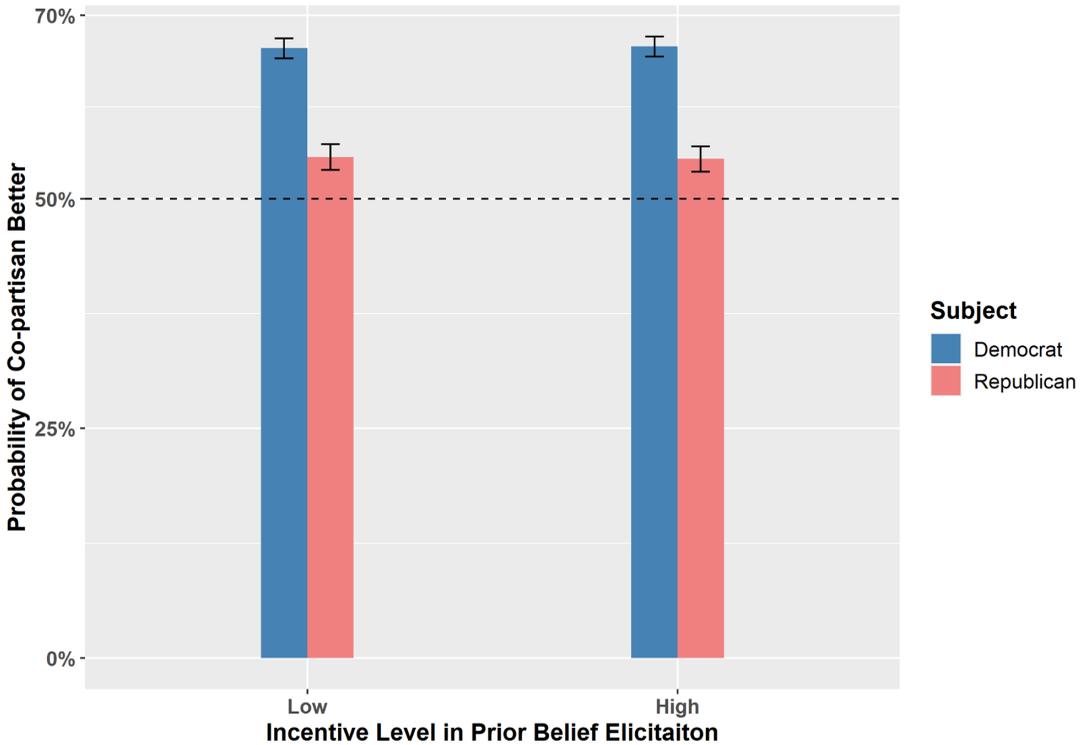


Figure 2. Shown is Democratic subjects' (blue bars) and Republican subjects' (red bars) prior probability of co-partisans being more accurate than counter-partisans, across levels of the incentive (low = \$0.01 and high = \$10) for the prior belief elicitation task. Error bars indicate 95% confidence intervals.

sincere, is not actually correct within our sample: as described in Methods, the headlines were selected to be non-partisan, and in terms of actual accuracy of the initial ratings, there was almost no difference between Democrats and Republicans.²⁴

3.3. Can feedback change beliefs about partisan competence?

Finally, we examine how subjects update their beliefs about co-partisan and counter-partisan competence when told that out of a subset of 20 Democrat–Republican pairs, either the Democrats or Republicans performed better. Given that the results of the previous task suggest that the discounting of counter-partisan advice is due to sincere (but incorrect) beliefs, we would expect participants to be willing to update their beliefs in the direction of the feedback they are provided.

Indeed, Figure 3 shows that regardless of the incentive on the prior belief elicitation, both Democratic and Republican subjects updated in the direction of the feedback. We also run a regression with the amount of change in the probability of co-partisan being better as the dependent variable, and independent variables of the feedback condition (1 = counter-partisan-better, -1 = co-partisan-better), the prior incentive level dummy (1 = high incentive on prior belief, -1 = low incentive on prior belief), subject's degree of right-wing extremity, and all two-way and three-way interactions.

Most importantly, we find a significant negative effect of counter-partisan-better feedback ($b = -16.3$, $p < 0.001$), such that subjects update very differently for the two types of feedback.

²⁴Across the 10 news headlines, the average absolute errors of initial news rating judgment are 2.0522 and 2.0520, respectively, for Republican and Democratic subjects.

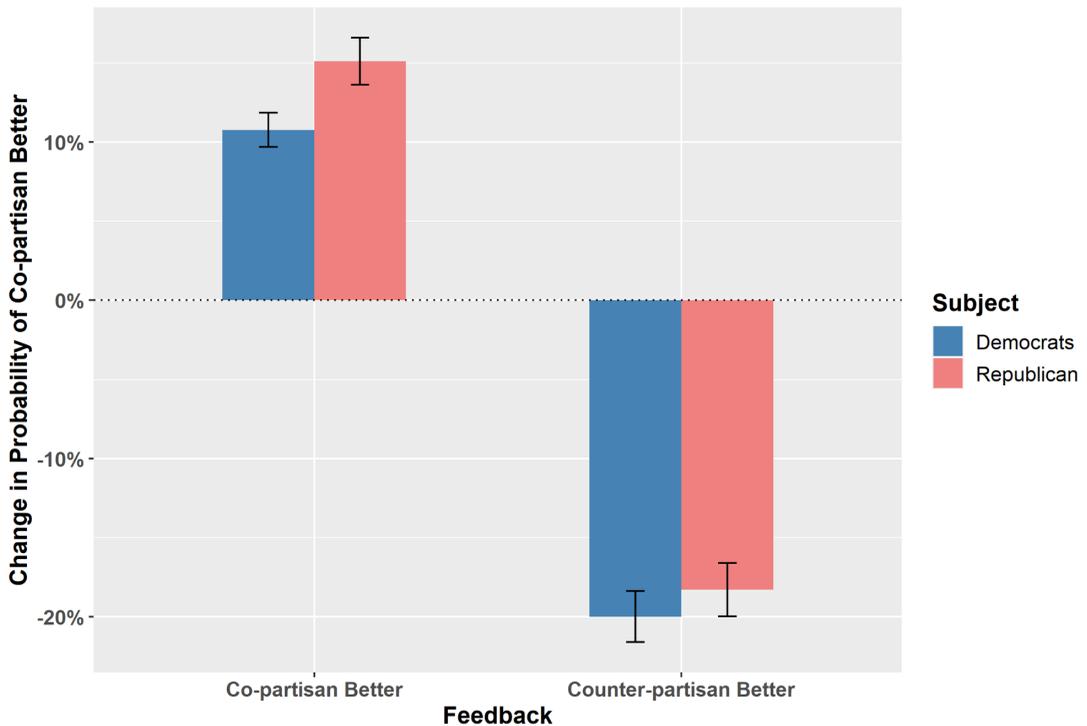


Figure 3. Shown is Democratic subjects' (blue bars) and Republican subjects' (red bars) amount of change from prior to posterior belief (posterior minus prior) when provided with feedback, based on whether the feedback indicated that co-partisans or counter-partisans performed better. Error bars indicate 95% confidence intervals.

Specifically, if subjects receive feedback that co-partisans did better, they update in favor of their own party by about 12%; if subjects receive feedback that counter-partisans did better, they update in favor of their counterpart by about 19%. Thus, rather than ignoring uncongenial feedback, people are willing to change their beliefs accordingly. Nevertheless, our results do not directly speak to the extent to which motivated reasoning matters to the formation of prior belief about partisan competence, because we cannot properly compare updating given different signals using the current experimental set-up.²⁵ We also find no significant interaction between feedback type and incentive level of prior belief elicitation ($b = -0.23$, $p = 0.59$). This suggests that 'forcing' subjects to state a sincere prior belief with higher stakes has no impact on how they subsequently process the feedback (e.g., subjects do not have the urge to state a posterior belief that favors their co-partisans), which is further suggestive evidence for subjects operating under their sincere beliefs at the moment of the advice-taking task.

4. Discussion

Here, we have shown that partisan bias can have spill-over effects in a nonpolitical domain: Democrats in our experiment were more prone to being influenced by information from co-partisans than counter-partisans, even though the information provided was exactly the same and the context had nothing to do

²⁵ Whether subjects are updating more or less than would be rational under Bayesian inference depends on their certainty about their stated prior belief of competence, which is difficult to measure; this makes clean comparison to a Bayesian benchmark in tasks such as ours infeasible, see Section 5 of the Supplementary Material for a detailed explanation.

with politics. Beyond documenting this phenomenon, our design allows us to make a further theoretical contribution by distinguishing whether the biased advice-taking is driven by a preference or is rational, based on sincere beliefs.

Our experimental results do not support the preference-based account. We do not find that incentives for accuracy reduce Democrats' rejection of information from counter-partisans relative to co-partisans. Conversely—and consistent with the belief-based account—we do find that Democratic subjects have strong incorrect perceptions that co-partisans are more competent at the (nonpolitical) news rating task. Furthermore, we find that this perception that co-partisans are more competent is not reduced by incentives, and is correctable by receiving (uncongenial) feedback. Together, these findings suggest that people are 'rational'—in the sense of following their sincerely held (albeit incorrect within the task context) beliefs about partisan competence—when processing information from party members rather than acting against their beliefs for partisan motivations.

Our observations in favor of the belief-based account over the preference-based account of biased advice-taking—although perhaps surprising given the current received wisdom about partisanship—resonate with recent work from political psychology that questions the pervasiveness of motivated reasoning during partisan interactions and emphasizes the importance of differences in prior beliefs (Bago et al., 2020; Tappin et al., 2020a,b). This alternative portrait of partisan discrimination suggests a surprisingly straightforward way for the society to mitigate the problem of ignoring information from counter-partisans: although motivated reasoning may still have a role in biased information search and prior belief formation, simply providing people with information indicating the competency of counter-partisans is at least successful in shifting beliefs. Nevertheless, future work may explore the extent to which shifted beliefs correspond to shifted behaviors.

More broadly, our results shed light on the causes of echo chambers. Instead of preference-based motivated reasoning, one important yet somewhat overlooked fundamental cause of the problem is that people may simply discount out-party information because of a perception that the out-group is less competent. By this account, echo chambers and polarization can also occur because people do not have the opportunity to learn positive information about counter-partisans. Counterfactually, if people were exposed to more positive information about counter-partisans—for example, by traditional media outlets, or social media ranking algorithms, and so forth—people's perceptions may change, and political polarization may be mitigated. Most importantly, recent evidence from the field is supportive of this claim. For example, Levy (2021) shows that exposure to counter-attitudinal news decreases negative attitudes toward an opposing political party. Broockman and Kalla (2022) shows that incentivizing regular Fox News viewers to watch news from CNN has a substantial effect on those viewers' overall political stance (e.g., political factual beliefs, attitudes, views, etc.) Kalla and Broockman (2022) shows that learning the perspectives of voters who have an opposing view on political issues reduces through in-depth two-way conversions in political campaigns substantially reduces the affective polarization of the political activists who initiate the campaigns.

Finally, there are several limitations of our study that it is important to bear in mind. First, consistent with most research in management, economics, and psychology, our experiments use convenience samples from MTurk and Prolific. It is now known that these online panels tend to be more left-leaning. Although our analysis has taken subjects' political leaning into account, it is possible that the results are different if we use a subject pool that is representative of the population. For example, Republicans and Democrats in our sample are equally competent in evaluating nonpolitical news, which may or may not be true in the population. Future work should see how our findings, and in particular, the surprising partisan asymmetries we observed, generalize to more representative samples. Second, some of our key results are null results. Of course, lack of evidence for an effect should not be confused with evidence for a lack of effect; but our sample size of 3,206 subjects is quite large, and follow-up experiments all clearly provide support for a lack of effect. Third, our identification assumes that the incentive in our high-incentive conditions (\$10) is high enough for subjects to disregard the preference-based account (if it exists). Although \$10 is quite large for online studies, it might not be large enough. Thus, it is possible that we would see some evidence of incentives reducing partisan motivations if the stakes

were higher. Relatedly, one might be worried about whether \$0.01 is already enough to eliminate the preference-based account. This is less concerning because if partisan preference can be neutralized by \$0.01, it probably does not pose a real threat to our society. Fourth, one may be concerned about experimenter demand effects, social desirability, or ‘expressive responding’. However, if such effects were indeed driving our results, we would expect raising the stakes to mitigate these effects (Berinsky, 2018)—yet, we find no effect of incentive when eliciting prior belief of partisan competence. Finally, our experiment examined advice-taking in only one nonpolitical task (news evaluation). Future work should investigate how our findings generalize across other tasks. Last but not least, future work may also adopt the concept behind our experimental design to examine the role of preference-based reasoning at different stages of belief formation and decision-making.

5. Conclusion

In sum, we find evidence, even in a nonpolitical task, of biased advice-taking that discriminates against counter-partisan opinions. To understand whether this biased advice-taking behavior is driven by partisan preferences or sincere beliefs, we developed a novel experimental design that uses a high incentive on task accuracy to reduce the impact of partisan preferences. Our findings suggest that the biased advice-taking observed among Democrats is driven by (incorrectly) believing that co-partisans are more competent at the task, rather than a desire or motivation to intentionally discount information from counter-partisans, or due to the formation of motivated beliefs. Furthermore, we find that subjects’ prior belief of partisan competence is subject to change after being given performance feedback. Taken together, these results suggest that partisan bias when processing nonpolitical information may be due to people not being exposed to the news that depicts positive images of counter-partisans, resulting in the formation of an inaccurate belief that counter-partisans are incompetent. If people somehow have more exposure to positive news of counter-partisans, our results suggest that they are indeed willing to correct their beliefs—which may possibly boost the communication across party lines and mitigate political polarization. More generally, the experimental design we propose here can be used to study whether a wide variety of other types of identity-based discriminatory behaviors (e.g., based on race or gender) are sincere or motivated.

Supplementary material. The supplementary material for this article can be found at <https://doi.org/10.1017/jdm.2023.28>.

Acknowledgments. We thank Jon Baron, Frank Schilbach, Dmitry Taubinsky, Andrew T. Little, Adam Berinsky, Tali Sharot, Steven Sloman, Drazen Prelec, Duncan Simester, Katie Coffman, Benjamin Enke, Sean McKay, Elitza Ambrus as well as various seminar and conference participants for their constructive feedback on the paper. We also thank Professor Jonathan Baron for helping us identify and correct an error in the Supplementary Material. We thank the John Templeton Foundation and the Alfred P. Sloan Foundation for financial support. For any questions or comments on the draft, please contact zyhjerry@mit.edu.

Funding statement. The authors do not have financial support to report.

Competing interest. The authors do not have competing interests to disclose.

References

- Allen, J., Arechar, A. A., Pennycook, G., & Rand, D. G. (2021). Scaling up fact-checking using the wisdom of crowds. *Science Advances*, 7(36), eabf4393.
- Bago, B., Rand, D. G., & Pennycook, G. (2020). Fake news, fast and slow: Deliberation reduces belief in false (but not true) news headlines. *Journal of Experimental Psychology: General*, 149(8), 1608.
- Bail, C. A., Argyle, L. P., Brown, T. W., Bumpus, J. P., Chen, H., Hunzaker, M. B. F., Lee, J., Mann, M., Merhout, F., & Volfovsky, A. (2018). Exposure to opposing views on social media can increase political polarization. *Proceedings of the National Academy of Sciences*, 115, 37, 9216–9221.
- Baron, J., & Jost, J. T. (2019). False equivalence: Are liberals and conservatives in the United States equally biased? *Perspectives on Psychological Science*, 14(2), 292–303.

- Berinsky, A. J. (2018). Telling the truth about believing the lies? Evidence for the limited prevalence of expressive survey responding. *Journal of Politics*, *80*(1), 211–224.
- Bertrand, M., & Kamenica, E. (2018). *Coming apart? Cultural distances in the United States over time* (Vol. w24771). Cambridge, MA: National Bureau of Economic Research.
- Bonaccio, S., & Dalal, R. S. (2006). Advice taking and decision-making: An integrative literature review, and implications for the organizational sciences. *Organizational Behavior and Human Decision Processes*, *101*(2), 127–151.
- Bromberg-Martin, E. S., & Sharot, T. (2020). The value of beliefs. *Neuron*, *106*(4), 561–565.
- Broockman, D., & Kalla, J. (2022). The manifold effects of partisan media on viewers' beliefs and attitudes: A field experiment with Fox News viewers. *OSF Preprints*, *1*.
- Bullock, J. G., Gerber, A. S., Hill, S. J., & Huber, G. A. (2015). Partisan bias in factual beliefs about politics. *Quarterly Journal of Political Science*, *10*(4), 519–578.
- Bullock, J. G., & Lenz, G. (2019). Partisan bias in surveys. *Annual Review of Political Science*, *22*, 325–342.
- Cinelli, M., Morales, G. D. F., Galeazzi, A., Quattrociochi, W., & Starnini, M. (2021). The echo chamber effect on social media. *Proceedings of the National Academy of Sciences*, *118*(9), e2023301118.
- Coffman, K. B., Exley, C. L., & Niederle, M. (2021). The role of beliefs in driving gender discrimination. *Management Science*, *67*(6), 3551–3569.
- Conlon, J. J., Mani, M., Rao, G., Ridley, M. W., & Schilbach, F. (2021). *Learning in the household* (Vol. w28844). Cambridge, MA: National Bureau of Economic Research.
- Enke, B., Gneezy, U., Hall, B., Martin, D., Nelidov, V., Offerman, T., & van de Ven, J. (2023). Cognitive biases: Mistakes or missing stakes? *Review of Economics and Statistics*, *105*(4), 818–832.
- Evans, T., & Fu, F. (2018). Opinion formation on dynamic networks: identifying conditions for the emergence of partisan echo chambers. *Royal Society Open Science*, *5*(10), 181122.
- Faraji-Rad, A., Samuelsen, B. M., & Warlop, L. (2015). On the persuasiveness of similar others: The role of mentalizing and the feeling of certainty. *Journal of Consumer Research*, *42*(3), 458–471.
- Finkel, E. J., Bail, C. A., Cikara, M., Ditto, P. H., Iyengar, S., Klar, S., Mason, L., et al. (2020). Political sectarianism in America. *Science*, *370*(6516), 533–536.
- Gino, F., Shang, J., & Croson, R. (2009). The impact of information from similar or different advisors on judgment. *Organizational Behavior and Human Decision Processes*, *108*(2), 287–302.
- Golman, R., Hagmann, D., & Loewenstein, G. (2017). Information avoidance. *Journal of Economic Literature*, *55*(1), 96–135.
- Hawkins, C. B., & Nosek, B. A. (2012). Motivated independence? Implicit party identity predicts political judgments among self-proclaimed independents. *Personality and Social Psychology Bulletin*, *38*(11), 1437–1452.
- Iyengar, S. (2021). The polarization of American politics. In M. Hannon, & J. de Ridder (Eds.), *The Routledge handbook of political epistemology* (pp. 90–100). London/New York City: Routledge.
- Iyengar, S., Lelkes, Y., Levendusky, M., Malhotra, N., & Westwood, S. J. (2019). The origins and consequences of affective polarization in the United States. *Annual Review of Political Science*, *22*, 129–146.
- Kahan, D. M. (2016). The politically motivated reasoning paradigm, Part 1: What politically motivated reasoning is and how to measure it. In R. A. Scott, & S. M. Kosslyn (Eds.), *Emerging trends in the social and behavioral sciences*. <https://doi.org/10.1002/9781118900772.etrds0417>.
- Kalla, J. L., & Broockman, D. E. (2022). Voter outreach campaigns can reduce affective polarization among implementing political activists: Evidence from inside three campaigns. *American Political Science Review*, *116*, 1516–1522.
- Khanna, K., & Sood, G. (2018). Motivated responding in studies of factual learning. *Political Behavior*, *40*(1), 79–101.
- Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, *108*(3), 480.
- Levy, R. (2021). Social media, news consumption, and polarization: Evidence from a field experiment. *American Economic Review*, *111*(3), 831–870.
- Little, A. T. (2021). Directional motives and different priors are observationally equivalent.
- Logg, J. M., Minson, J. A., & Moore, D. A. (2019). Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes*, *151*, 90–103.
- 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.
- Marks, J., Copland, E., Loh, E., Sunstein, C. R., & Sharot, T. (2019). Epistemic spillovers: Learning others' political views reduces the ability to assess and use their expertise in nonpolitical domains. *Cognition*, *188*, 74–84.
- Peterson, E., Goel, S., & Iyengar, S. (2021). Partisan selective exposure in online news consumption: Evidence from the 2016 presidential campaign. *Political Science Research and Methods*, *9*(2), 242–258.
- Peterson, E., & Iyengar, S. (2021). Partisan gaps in political information and information-seeking behavior: Motivated reasoning or cheerleading? *American Journal of Political Science*, *65*(1), 133–147.
- Phelps, E. S. (1972). The statistical theory of racism and sexism. *American Economic Review*, *62*(4), 659–661.
- Prior, M. (2013). Media and political polarization. *Annual Review of Political Science*, *16*, 101–127.
- Sharot, T. (2011). The optimism bias. *Current Biology*, *21*(23), R941–R945.
- Tappin, B. M., Pennycook, G., & Rand, D. G. (2020a). Thinking clearly about causal inferences of politically motivated reasoning: Why paradigmatic study designs often undermine causal inference. *Current Opinion in Behavioral Sciences*, *34*, 81–87.

- Tappin, B. M., Pennycook, G., & Rand, D. G. (2020b). Rethinking the link between cognitive sophistication and politically motivated reasoning. *Journal of Experimental Psychology: General*, *150*(6), 1095–1114.
- Yaniv, I., Choshen-Hillel, S., & Milyavsky, M. (2011). Receiving advice on matters of taste: Similarity, majority influence, and taste discrimination. *Organizational Behavior and Human Decision Processes*, *115*(1), 111–120.
- Zimmermann, F. (2020). The dynamics of motivated beliefs. *American Economic Review*, *110*(2), 337–361.