Biased belief updating in causal reasoning about COVID-19

Leo Gugerty, Michael Shreeves, Nathan Dumessa

Clemson University

in press Journal of Experimental Psychology: Applied as of 6/3/2021

Author Note

Leo Gugerty, Department of Psychology, Clemson University

Michael Shreeves, Arizona State University at Lake Havasu

Nathan Dumessa, Department of Psychology, Clemson University

Initial results from Study 1 were presented at the 61st Annual Meeting of the Psychonomic Society, Poster 2147. Studies presented in this paper were preregistered with the Open Science Foundation, available here: <https://osf.io/fc65u> and <https://osf.io/u3cdk> and <https://osf.io/25tqb>

Correspondence concerning this article should be addressed to Leo Gugerty, Department of Psychology, 418 Brackett Hall, Clemson University, Clemson, SC 29634-0745

Contact: GUGERTY@clemson.edu

This project was partially supported by funding from the Clemson University Creative Inquiry program. We thank Neil Bramley for providing information and analytical functions for strong-uniform priors and Shane Donohue, Tayler Smith and Spencer Westcott of the Clemson University Creative Inquiry Program for help in conducting the studies.

Abstract

In three experiments based on 977 participants, we investigated whether people would show belief bias by letting their prior beliefs on politically charged topics unduly influence their reasoning when updating beliefs based on evidence. Participants saw data from fictional studies and made judgments of how strongly COVID-19 mitigation measures influenced the number of COVID-19 cases (political problems) or a medicine influenced number of headaches (neutral problems). We predicted that liberals would overestimate and conservatives would underestimate causal strength on political problems relative to neutral problems. In Experiments 1 and 2, liberals showed this overestimation bias. Surprisingly, college-student conservatives in Experiment 2 showed the same overestimation as liberals. These findings made sense because all three groups who overestimated the strength of mitigation measures held prior beliefs that strongly favored use of these measures. In Experiment 3, conservatives' judgments of the strength of mitigation measures after seeing evidence increased as their degree of prior support for these measures increased. Furthermore, conservatives who strongly *opposed* the use of mitigation measures *underestimated* causal strength in the political problems. These results suggest that belief bias is driven more by specific beliefs relevant to the reasoning context than to general attitudinal factors like political ideology.

Biased belief updating in causal reasoning about COVID-19

In his prescient essay “In front of your nose,” George Orwell (1946) noted that “we are all capable of believing things which we *know* to be untrue, and then, when we are finally proved wrong, impudently twisting the facts so as to show that we were right. ... It is especially in our political thinking that these vices flourish.” Orwell’s idea that people “twist” evidence that conflicts with a prior belief in order to maintain that belief presaged what psychologists currently call belief bias (Evans et al., 1983) or the prior belief effect (Edwards & Smith, 1996). In belief bias, reasoners’ judgment about the soundness of an argument is unduly influenced by their prior belief in the argument’s conclusion. In line with Orwell’s emphasis on the domain of politics, research has shown belief bias on political topics (Edwards & Smith, 1996; Gampa et al., 2018). Importantly, research has shown that a correlate of biased evaluations of argument soundness is that people maintain or even increase the strength of their beliefs after being exposed to evidence that conflicts with those beliefs (Lord et al., 1979; Nyhan et al., 2013). These polarization or backfire effects are a striking example of people sub-optimally updating their beliefs after seeing relevant evidence.

In the current article, we consider how political ideology—and prior beliefs related to one’s ideology—might bias causal reasoning about COVID-19, focusing on the effects of mask wearing and social distancing policies on the spread of the disease. Polls in the United States have shown wide ideological differences such that, compared to conservatives and Republicans, liberals and Democrats are much more likely to believe that COVID-19 is a serious threat or risk and that people should engage in mitigation measures such as social distancing, wearing masks and getting vaccinated (Malloy & Schwartz, 2020; The Marist Poll, 2021; Miller, 2020; Pew Research Center, 2020). The Pew poll suggested that party differences were a much stronger

explanation for differences in beliefs about threat from COVID-19 than race, geography, gender or age. Also, conservatives in the US trust science less than liberals (Gauchat, 2012), which would make conservatives less likely to believe and follow scientists' recommendations about COVID-19. In our first two experiments, we investigated whether differences in political ideology influenced peoples' causal reasoning. The findings of these experiments led us to the more nuanced view that peoples' specific prior beliefs about a political topic may influence their causal reasoning on that topic more directly than their general attitude about politics, as represented by their ideology. In our third experiment, we investigated how specific prior beliefs influence causal reasoning when political ideology is held constant.

Two important causal reasoning tasks where people update their beliefs based on evidence are learning the causal *structure* in a domain (i.e., what causal links are present; Griffiths & Tenenbaum, 2005) and learning the *strength* of causal relationships (Lu et al., 2008). In prior research on how ideology influences causal reasoning, Kahan & Corbin (2016) found that liberals agreed much more than conservatives with the idea that human activity causes global warming even when participants holding both ideologies were high in openness to changing their beliefs based on evidence. In this task, people are making a causal structure judgment.

In Kahan et al. (2013), participants saw how often an effect was present vs. absent when a cause was present vs. absent and judged whether the cause was generative (increasing the frequency of the effect) or preventive (decreasing the frequency). In political problems, the causal factor was banning (or not banning) concealed carry of handguns and the effect was number of subsequent crimes. In neutral problems, the cause was using (or not using) a skin cream and the effect was the number of cases of rash. Correctly answering this question required

quantitative reasoning skill (i.e., to calculate proportions from count data and compare the proportion of effects present when the cause was present vs. absent). However, even for participants high in numeracy, when an incorrect answer based on a quick non-quantitative heuristic agreed with their political biases but the correct, calculated answer disagreed with them, both conservatives and liberals answered incorrectly, as if they accepted the quick, heuristic answer. In contrast, when the heuristic answer disagreed with their biases and the quantitative answer agreed with them, high-numeracy conservatives and liberals were able to use their quantitative skill to calculate the correct answer. Here, participants engaged in effortful calculation to protect pre-existing beliefs from conflicting evidence rather than, as in many dual processing theories (e.g., Evans & Stanovich, 2013), to correct preliminary answers.

In the current studies, we focused on a different aspect of causal reasoning—judging causal strength. Our participants saw contingency data regarding a cause-effect relationship that was framed in a political or a neutral context and then judged the strength of the relationship. This allowed us to detect biasing effects by comparing performance on the political and neutral problems.

Little research has been done on how political biases influence causal strength judgments, which is our focus. It is important to study how biases affect causal strength judgments because causal strength (or effect size) information is critical for everyday decisions. Governments will not spend billions to curb greenhouse-gas emissions unless the strength of the relationship between these emissions and global warming is high and the environmental and economic effects of global warming are strong. Freimuth et al. (2017) showed that the perceived effectiveness (strength) of the flu vaccine was a significant predictor of trust in the vaccine. Finally, social science and medical researchers focus on effect size (i.e., causal strength) in addition to

statistical significance when interpreting results, because statistical significance can be influenced by research design factors (e.g., sample size) whereas causal strength is not influenced by sample size (Kelley & Preacher, 2012) and better represents the underlying importance of a causal factor when making decisions. For example, COVID-19 vaccines will not be approved by governments unless they meet standards for effectiveness (effect size) and confidence (significance or confidence interval) (Zaks, 2020).

The task used in all experiments is shown in Figure 1, which shows a political problem. In each political problem, participants saw fictional contingency data on the number of people with and without COVID-19 in a city before and after strict social distancing policies (including mask wearing) were ended. Participants judged “how strong the effect of ending social distancing was” by entering a number between 0 (which meant ending social distancing policies did not increase the number of people who got COVID-19 at all) and 100 (which meant ending social distancing policies increased the number of people who got COVID-19 as much as possible). In Experiment 2, participants also judged their confidence in the strength judgment. This task is based on Buehner et al. (2003).

Theoretical Background

The theoretical basis for our hypotheses rests on causal strength models at the rational or computational level of explanation (Anderson, 1991; Marr, 1982). Rational models help researchers evaluate whether individuals make optimal use of information when performing cognitive tasks and understand how individuals fall short of optimal. These models have been used in empirical psychological studies of a variety of types of causal reasoning, including judging confidence in a causal conclusion (Griffiths & Tenenbaum, 2005) and diagnostic reasoning (Shreeves et al., 2020).

Figure 1. Display for political causal reasoning problem.

Lu et al. (2008) developed a rational model of causal strength judgments based on Cheng's (1997) causal power theory. When applied to generative causes, which are used exclusively in the current research, their model represents a situation where candidate cause C might produce effect E with strength s_C and a single composite background cause B might produce E with strength s_B , i.e., a causal structure of $C \rightarrow E \leftarrow B$. Lu et al.'s model uses Bayesian parameter estimation (Kruschke, 2013) to predict strength—i.e., to estimate values for s_C and s_B —based on evidence (in this case, contingency data) and assumptions about an individual's prior beliefs about strength in the domain in question. Thus, the model represents how individuals update their prior beliefs about the strength of a cause based on new evidence.

The evidence that is input to the model is the contingency data showing how often E is present and absent when C is present and absent; i.e., $N(e^+|c^+)$, $N(e^-|c^+)$, $N(e^+|c^-)$, and $N(e^-|c^-)$, where e and c refer to the effect and candidate cause, respectively, and $+$ and $-$ refer to present and absent, respectively. For Figure 1, $N(e^+|c^+) = 40$, $N(e^-|c^+) = 24$, $N(e^+|c^-) = 16$, and $N(e^-|c^-) = 48$. The input parameters and outputs of Bayesian parameter estimation models consist of probability distributions instead of point estimates. The other key input to the

model is the joint distribution $P(s_B, s_C)$, which represents an individual's prior beliefs about the probabilities of various values of s_B and s_C . The model outputs the posterior distribution of the strength of cause C in causal structure $C \rightarrow E \leftarrow B$, which is computed using Bayes Theorem,

$$P(s_C|D) = \int_0^1 \frac{P(D|s_B, s_C)P(s_B, s_C)}{P(D)} ds_B, \quad (1)$$

where D is the contingency data, $P(D|s_B, s_C)$ is the likelihood term, and $P(D)$ is a normalization term representing the probability of the contingency data occurring. The likelihood term is a function of the likelihoods of the effect in the four cells of the contingency dataset, i.e., $P(e^+|c^+)$, $P(e^-|c^+)$, $P(e^+|c^-)$, and $P(e^-|c^-)$.

Lu et al. use Markov Chain Monte Carlo simulation to calculate a posterior strength distribution for a candidate cause given a particular contingency dataset and assumptions about priors. Figure 2 shows the distributions output by the model for the contingency data in Figure 1 for three prior distributions of causal strength. These prior distributions will be discussed in the following section. Although Bayesian parameter estimation models estimate a distribution instead of a point estimate for an output parameter, researchers usually choose a point estimate from the distribution in order to understand a model's predictions. We followed Lu et al. (2008) in using the mean of the posterior strength distribution for a point estimate (see Figure 2).

In Lu et al.'s model, the important characteristics of the contingency data are the total sample size and the four probabilities comprising the likelihood term. In Figure 1, $P(e^+|c^-)$ is .25, $P(e^+|c^+)$ is .625, and the total sample size is 64. According to Cheng (1997), the normative causal strength for a generative cause is,

$$s_C = (P(e^+|c^+) - P(e^+|c^-)) / (P(e^-|c^-)). \quad (2)$$

In other words, causal strength is the proportion of *new* cases of the effect (i.e., in the cause-present but not the cause-absent condition) relative to the proportion of effect-absent cases in the cause-absent condition. The denominator in this equation refers to an important aspect of sample size that is relevant to computing causal strength, as the cases that do not show the effect when the cause is absent are the only ones that can change to show the effect when the cause becomes present. We call the cell referenced in the denominator the functional sample size, $N(e^-|c^-)$. Liljeholm and Cheng (2008) called this the virtual sample size. Proportional functional sample size, $P(e^-|c^-)$, is .75 in Figure 1. In the current experiments, we defined nominal causal strength via Eq. 2, which follows Cheng's (1997) causal power theory. Nominal strength is .50 in Figure 1. Appendix A shows examples of problems with nominal strength of .25 and .75. To construct the strength judgment problems for participants, we varied nominal causal strength, proportional functional sample size, and total sample size.

How priors influence strength updating. The literature reviewed above on ideological differences in beliefs and behaviors about COVID-19 suggests how conservatives' and liberals' prior beliefs differ. Conservatives are less likely than liberals to believe that COVID-19 is risky, less likely to endorse and practice mitigation measures, and less likely to believe scientists' recommendations about these measures (Gauchat, 2012; Malloy & Schwartz, 2020; The Marist Poll, 2021; Miller, 2020; Pew Research Center, 2020). These findings suggest that, prior to observing evidence, conservatives generally believe that the causal strength of social distancing and mask wearing (the manipulation in our political fictional studies) is low compared to liberals. Note that these differences between conservatives' and liberals' prior beliefs do not necessarily mean that conservatives' prior beliefs about the strength of mitigation measures are low in an absolute sense.

Figure 2. Posterior distributions of the causal strength of the candidate cause for the problem in Figure 1, $P(s_c|D)$, given three prior distributions: strong-uniform (SU), a 20; strong-uniform, a 5, weak-uniform (WU), a 10. Vertical lines (and numbers) show the mean of each distribution.

Researchers have used a variety of probability distributions to represent peoples' prior beliefs about causal strength in rational strength models. Also, because judging causal strength requires comparing the likelihood of the effect occurring when the cause is present and absent, these models must assume a prior strength distribution for the candidate cause (cause-present condition) and for the background causes that produce the effect in the cause-absent or “control” condition.

While many rational model developers choose priors based on theoretical concerns, Yeung and Griffiths (2015) conducted important research that provided empirical evidence about peoples' prior causal strength distributions in a variety of domains. Their study showed that—even without any biasing influences—people generally have strong priors for candidate causes that vary in strength depending on the domain and express no preference regarding whether

background causes are strong or weak (which can be represented by a uniform distribution).

They called this strong-uniform combination “empirical priors.” Bramley et al. (2017) simulated Yeung and Griffith’s empirical priors using analytic functions instead of the raw empirical data.

Lu et al. (2008) represented peoples’ priors about causal strength by sparse-and-strong priors, which assumes that people believe that both the candidate and background cause are strong. Following Yeung and Griffiths (2015), we assumed that people generally hold strong-uniform priors and simulated these with an exponential distribution that was based on the distribution Lu et al. used to represent sparse-and-strong priors. Eq. 3, which expands the prior term in Lu et al.’s strength model (Eq.1), represents the probability of the possible s_B, s_C pairs as s_C and s_B each range from their minimum value of 0 to the maximum of 1 under strong-uniform priors.

$$P(s_B, s_C) \propto e^{-a(1-s_C)} + e^{-a(1+s_C)} \quad (3)$$

When individuals hold strong priors for the strength of the candidate cause (s_C), they assume that s_C is more likely to be high than low. The parameter a controls the degree to which individuals favor high over low values of s_C .

For our neutral problems (i.e., the effect of a mineral supplement on headaches), we used an a of 5 as this yielded a distribution similar to the empirical priors’ distribution for a similar type of problem, the effect of an allergy medicine on hormonal imbalance, in Yeung and Griffiths (2015). Figure 3 shows the distribution of s_C under strong-uniform and weak-uniform priors with various a values. The vertical axis shows the probability of each possible s_C value (in relative, not absolute, terms). For strong-uniform priors with a of 5, the most likely strength value for the candidate cause (s_C) is 1. There is a small probability that s_C is .5 and near 0 probability that s_C is 0. Someone whose priors are represented by this distribution is moderately

biased to believe that the candidate cause is strong and not weak. Because s_B follows a uniform distribution, the probability of any particular s_C value is the same for all values of s_B between 0 and 1 that are paired with that s_C value; e.g., for $a = 5$, $P(s_C = 1.0)$ is 5 relative probability units for all values of s_B where $s_C = 1.0$. Thus, if you imagine the s_B axis extending into the depth plane, each line in Figure 3 is the edge of a concave ramp.

Figure 3. Probability distributions for strong-uniform (SU) priors with α of 20 (for liberals on political problems) and α of 5 (for neutral problems) and for weak-uniform (WU) priors with α of 10 (for conservatives on political problems). Probability values on the vertical axis are in relative, not absolute terms.

In order to represent the liberal expectation that ending mitigation measures is an especially strong cause, we wanted strength priors that were even more biased towards strong causes than when a is 5. The domain with the strongest empirical priors in Yeung and Griffiths study (physical causes) had an a of around 20. Therefore, we used an a of 20 to represent the hypothesized liberal bias towards very strong priors on political problems (see Figure 3). For this distribution, the relative probability of the candidate cause having strength of 1 is much higher than for strong-uniform priors with a of 5 and the probability for strength values below .75 is near 0. Someone holding the strength priors in this distribution (a 20) assumes that the candidate

cause is highly likely to be extremely strong and has almost no chance of being moderately strong or weak.

When a is 0, individuals see all values of s_C as equally likely, which would be represented by a flat line in Figure 3. In contrast to liberals, we assumed that conservatives would hold weak priors for political problems; i.e., they would think that ending mitigation measures is likely to have very little effect on COVID-19 cases. Thus, we modified the equation for strong-uniform priors to create weak-uniform priors,

$$P(s_B, s_C) \propto e^{-as_C} + e^{-a(2-s_C)} . \quad (4)$$

When individuals hold weak priors for s_C , they assume that s_C is more likely to be low than high and the parameter a controls the degree to which they favor low over high values of s_C . Weak-uniform priors are also shown in Figure 3. Since we changed the strength parameter, a , by 15 (from 5 to 20) in going from the strong-uniform causes for the neutral problems to the extremely strong causes for liberals, we changed a by 15 units (from 5 in the strong dimension to 10 in the weak dimension) in going from the a for neutral problems to weak-uniform priors. For weak priors, the most probable value of s_C is 0 and values of s_C above .5 have probability 0.

Since the contingency data is the same for each posterior distribution in Figure 2, the differences in these distributions is solely due to the different strength priors. For the strength priors that we assumed to be unbiased (strong-uniform, a 5), the mean of the posterior distribution closely approximates nominal strength, which is .50 for this problem. However, the posterior means are overestimated for strong-uniform priors with a of 20 and underestimated for weak-uniform priors. Appendix A shows that these three strength priors yield the same pattern of over- and underestimation for other problems. Figure 4 shows the predictions of the Lu et al. (2008) strength model for each of these three prior distributions for the 10 contingency datasets

used to create the problems in Experiment 1, which used five nominal causal strengths and two proportional functional sample sizes (including the problem in Figure 2). Each data point in the graph is averaged over the model predictions for two proportional functional sample sizes. The figure predicts that on political problems liberals will overestimate strength relative to neutral problems, while conservatives will underestimate strength.¹

Figure 4. Predicted strength judgments for strength models with three priors.

Empirical research has shown that peoples' causal strength judgments are very sensitive to nominal causal strength, especially when using graphical problem displays like in Figure 1 (Liljeholm & Cheng, 2009; Lu et al., 2008). Although causal strength (effect size) is not influenced by total or functional sample size in statistical treatments (Kelly & Preacher, 2012)

¹ The slope of these lines varies with the assumed α parameters used for the prior distributions; also, the line for weak-uniform priors is curvilinear. Therefore, we only use these predicted strengths to justify our hypotheses about main effects of ideology and problem type, not interactions of these variables with nominal strength.

these studies showed that people do not treat strength as independent of sample size. Instead, they increase their strength judgments as total and functional sample size increase. Liljeholm and Cheng suggest that this may be due to people conflating confidence in a causal conclusion, which is influenced by sample size, with strength. In statistical terms, these people conflated statistical significance with effect size.

Overview of the Experiments

In Experiment 1, participants made causal strength judgments for 10 contingency datasets, with 5 nominal strengths (.125, .25, .50, .75, .875) and 2 proportional functional sample sizes (.50, .75). Sample size was always 64 cases in both the cause-absent and cause-present condition. In Experiment 2, participants made causal strength judgments for 12 contingency datasets, with 3 nominal strengths (.25, .50, .75), 2 proportional functional sample sizes (.50, .75) and 2 total sample sizes (16 or 64). In Experiment 2 only, participants followed their strength judgments with a judgment of confidence in the strength judgment. In both experiments, about half of the participants saw political problems and half saw neutral problems and participants consisted of liberals and conservatives. Thus, we varied the characteristics of the evidence and varied degree of biasing (by varying type of problem and measuring variation in a factor that was expected to influence prior beliefs, ideology). This allowed us to see how two key factors in Bayesian belief updating—evidence and prior beliefs—influenced strength and confidence judgments. We also measured individual differences in other factors that could affect participants' strength judgments, i.e., numeracy (Peters et al., 2007) scientific reasoning ability (Drummond & Fischhoff, 2017) and attitudes towards COVID-19 (see Appendix B).

In Experiment 1 and 2, we found that even among conservatives, most participants supported the use of mitigation measures like masks and social distancing and very few

participants held the anti-mask and anti-social-distancing beliefs that are thought to characterize conservatives (Miller, 2020) and that we needed to test some of our hypotheses. Therefore, in Experiment 3, we tested only conservatives and attempted to recruit participants with a wide range of beliefs about COVID-19 mitigation measures. The rationale and hypotheses for Experiment 3 will be explained more fully in the introduction to that study. Each of these studies was preregistered and data analyses were conducted in accordance with the preregistration.

Hypotheses for Experiments 1 and 2 Regarding Causal Strength Judgments

Effects of prior beliefs. We assume that when people are asked to make a causal strength judgment based on evidence, they update their prior beliefs about the causal strength based on the evidence. Thus, their posterior beliefs are influenced by characteristics of the evidence and their prior beliefs about the strength of the causal factor. Based on evidence that liberals' prior beliefs about COVID-19 mitigation measures are more positive than conservatives (Malloy & Schwartz, 2020; Miller, 2020; Pew Research Center, 2020) and the predictions of the Lu et al. (2008) strength model concerning how these different prior beliefs influence strength judgments, we hypothesized that liberals' causal strength judgments will be overestimated on political (COVID-19) problems relative to neutral (headache) problem whereas conservatives' causal strength judgments will be underestimated on political problems relative to neutral problems. It is important to note that these hypotheses are based on the assumption that liberals' prior beliefs will strongly support the use of COVID-19 mitigation measures and that conservatives' priors will strongly oppose these measures, as represented by strong-uniform priors with α of 20 and weak-uniform priors, respectively, in Figure 3.

In addition to these main effects, we hypothesized interactions such that, for liberals, the degree of overestimation will be greater when nominal strength is low than when it is high;

while, for conservatives, the degree of underestimation will be smaller when nominal strength is low than when it is high. We predicted these interactions with nominal strength under the assumption that the biasing effects of prior beliefs will be strongest when the evidence conflicts most with those beliefs and weakest when the evidence strongly supports those beliefs.

Because of the non-political context, we hypothesized that there will be little or no difference between liberals' and conservatives' strength judgments on neutral problems. Given this, the hypotheses concerning liberals' overestimation and conservatives' underestimation of causal strength imply that liberals' causal strength judgments will be greater than those of conservatives for political problems.

Effects of the evidence. We also hypothesized that participants' causal strength judgments will increase with nominal causal strength. This is supported by Lu et al.'s (2008) strength model and empirical findings (Liljeholm & Cheng, 2009).

Hypotheses for Experiment 2 Regarding Confidence Judgments

Griffiths and Tenenbaum's (2005) causal support model used Bayesian parameter estimation to predict how reasoners judge their confidence that a particular causal relationship exists based on the same type of contingency data used in Lu et al.'s strength model. They note that this model is analogous to a statistical significance decision. The causal support model predicts that confidence should increase strongly with proportional functional sample size and total sample size, which is consistent with the idea that confidence in causal conclusions should be strongly influenced by sample size. When Liljeholm and Cheng (2009) had participants view contingency data like in the current study and judge their confidence that a causal relationship existed, participants' confidence judgments were strongly influenced by both total and functional sample size.

We developed a model that predicts peoples' confidence in their causal strength judgments by modifying Lu et al.'s (2008) causal strength model to also calculate the precision of the model's strength estimates. This model is analogous to calculating a confidence interval around a strength estimate. Interestingly, the predictions of our precision-based model and the causal support model were almost perfectly correlated (r about .98 after a nonlinear transformation), which suggests that, theoretically, judging confidence in a causal structure decision (whether a causal link exists) is the same as judging confidence in a causal strength estimate. See Appendix C for more details. The equivalence between these two types of confidence judgments is important for our study, because, since our participants made causal strength but not structure judgments, we asked them about confidence in their strength judgments.

Effects of prior beliefs. Belief bias in causal reasoning means that participant's causal judgments are more influenced by prior beliefs and less influenced by evidence than is appropriate under normative models. Therefore, given our hypotheses that liberals and conservatives will show belief bias on political problems, we also hypothesized that all participants' confidence judgments would be higher on political than neutral problems. This hypothesis follows from the idea that when people let their prior beliefs have excessive weight in causal reasoning, this involves being overconfident in those prior beliefs. We also hypothesized that the increase in judged confidence with proportional functional sample size and total sample size would be stronger for neutral than political problems. This follows from the idea that when people are underweighting the evidence, characteristics of the evidence like functional and total sample size will have less influence on their confidence judgments.

Effects of the evidence. We hypothesized that participants' confidence judgments will increase with proportional functional sample size and with total sample size. This hypothesis is based on the confidence models and empirical findings discussed above.

We now describe the three experiments. Data in Experiment 1 were collected between July 16-18, 2020 using a sample recruited from Mechanical Turk (Clifford et al., 2015). The pre-screening capabilities of Mechanical Turk were used to screen out political moderates and people accessing the data collection program (using Qualtrics) from outside the United States. Data in Experiment 2 were collected between October 26 to November 12, 2020 using a sample of undergraduates from Clemson University, in South Carolina, USA. Because Experiments 1 and 2 are similar, we describe their methods and results together. Data in Experiment 3 were collected between February 17, 2021 and March 4, 2021 using a sample of United States citizens recruited from Mechanical Turk and Prolific (Palan & Schitter, 2018).

Experiment 1 occurred during the second peak of COVID-19 cases in the US with 65,417 new cases per day (CDC, 2020). This was after stay-at-home orders were relaxed in May and June (Secon, 2020). Experiment 2 occurred at the beginning of the third peak of cases in the US, with 98,146 new cases per day on average (CDC, 2020). Of the 440 participants in Study 2, 108 (25%) completed the survey on or after November 3, the date of the US presidential election. Experiment 3 occurred as the third peak of US cases was ending and leveling off, with 66,895 new cases per day on average (CDC, 2021).

Methods for Experiments 1 and 2

Participants

We conducted a power analysis based on the most restrictive hypothesis, i.e. the interaction between political ideology and nominal causal strength. No prior data was available

to estimate the effect size for this interaction. Therefore, we set a conservative effect size of $f=0.175$. A GPower analysis for ANOVA (fixed effects, special, main effects and interactions) with $\alpha = .05$, power = .80, numerator df = 1 and 4 groups led to a required total sample size of 260, i.e., 65 per group. In Study 1, we planned to over-collect by 15% to deal with MTurk data quality issues, resulting in a target sample size of 300. For Study 2, we planned to over-collect by 36% to account for moderates based on a prior study of participants from the university in this study, resulting in a target sample size of 406.

In Study 1, 302 online participants were recruited through Amazon Mechanical Turk. Of these, 17 were identified as political moderates and excluded (see Measures) and 1 participant was excluded due to a technical issue, resulting in a total sample size of 284. (See Table 1 for demographics.) MTurk participants were compensated at a rate of \$7.50 per hour. Participants were required to be at least 18 years of age and US residents. TurkPrime by CloudResearch was used to pre-screen participants as either Liberals (self-reporting as Very Liberal or Liberal) or Conservatives (self-reporting as Very Conservative or Conservative). TurkPrime was also used to exclude participants from suspicious IP addresses, (i.e., out-of-country, server farms, or a universal ban list). Study 2 included 443 Clemson University students who were compensated with course credit. These participants could not be pre-screened for ideology and 139 identified as political moderates. Three participants were excluded due to taking hours-long breaks during the 10 minute causal reasoning task, resulting in a final sample size of 301 liberals or conservatives. All participants were treated in accord with APA ethical guidelines.

Design

Experiment 1 followed a problem type (political, neutral; between subjects) by nominal causal strength (.125, .25, .50, .75, .875; within-subjects) by proportional functional sample size

(.50, .75; within-subjects) mixed factorial experimental design. Political ideology (liberal, conservative) was included as a measured between-subjects variable. Problems were either political (increase in COVID-19 cumulative cases after ending social distancing) or neutral (increase in headaches after taking a mineral supplement). This manipulation was implemented via instructions preceding the fictional study problems and a graphical representation of study results in the problem displays (see Figure 1). Pre-screened liberals and conservatives were recruited separately, and participants within these groups were randomly assigned to either neutral or political problems using Qualtrics' randomization function.

Table 1. Demographics for Experiment 1 and 2. SDs in parentheses.

Variable	Experiment 1						
	Liberals			Conservatives			
	Total or Average	Political problems	Neutral problems	Total or Average	Political problems	Neutral problems	
N	150*	73	77	134	71	63	
N female	78	37	41	62	35	27	
N male	69	34	35	72	36	36	
Age	37.7 (12.5)	38.0 (12.9)	37.4 (12.1)	43.5 (13.5)	42.1 (12.6)	45.1 (14.3)	
	Experiment 2						
	N	117	65	52	184	95	89*
	N female	97	56	41	113	54	59
	N male	20	9	11	70	41	29
	Age	18.8 (1.9)	18.7 (2.1)	18.8 (1.7)	18.8 (2.3)	18.9 (3.1)	18.6 (0.9)

* A few participants selected "Prefer not to say" for gender.

Experiment 2 followed a problem type (political, neutral; between subjects) by nominal causal strength (.25, .50, .75; within-subjects) by proportional functional sample size (.50, .75; within-subjects) by total sample size (1600, 6400; within-subjects) mixed factorial experimental design. Political ideology (liberal, conservative) was included as a measured between-subjects variable. All participants were randomly assigned to complete political or neutral problems using Qualtrics' randomization feature.

Measures

Political Ideology. Political ideology was measured using a composite of two 5-point scales: political viewpoint (“Which of these options comes closest to your political views”, 1= Very Liberal, 5 = Very Conservative) and Democratic-Republican party alignment (“Which of these options is closest to the political party you identify with?” 1 = “Strong Democrat”, 5 = “Strong Republican.”) For both scales, responses 1 and 2 were coded as liberal, 3 as moderate, and 4 and 5 as conservative. Political ideology was coded as Liberal if the political viewpoint and party questions were coded as both liberal or one liberal and one moderate; Moderate if the political viewpoint and party questions were coded as both moderate or one liberal and one conservative; and Conservative if the political viewpoint and party questions were coded as both conservative or one conservative and one moderate.

Causal Reasoning Task. Appendix D shows participant instructions for the political and neutral conditions of the causal reasoning task. (We changed the instructions after Experiment 1 to make them more similar between the political and neutral conditions. So, instructions for both Experiments are included.) Figure 1 shows an example of a political problem with a causal strength question. (The bulleted list explaining the social distancing policies in Figure 1 was only included in Experiment 2.) Appendix D shows a neutral problem. An example strength question for neutral problems is “How **strong** was the effect of **taking Mineral J**? Enter a number from 0 to 100 where: **0** means taking this mineral **DID NOT INCREASE** the number of people who got headaches **at all**; and **100** means taking this mineral **INCREASED** the number of people who got headaches **as much as possible**.”

The confidence question (which was only given in Experiment 2) was “Based on the results of this study, **how confident** are you that your answer to the “how strong” question would

be **similar** if another study were conducted in the **same city** [OR for the same mineral] with a **different** group of randomly chosen participants? Enter a number between 0 and 100 where: **0** means **not at all confident**; and **100** means **extremely confident**.”

In Experiment 1, for each problem type, 10 problems were created by a factorial combination of 5 nominal causal strengths and 2 proportional functional sample sizes. In Experiment 2, for each problem type, 12 problems were created by a factorial combination of 3 nominal causal strengths, 2 total sample sizes and 2 proportional functional sample sizes.

Numeracy. Participants completed a 13-item numeracy measure adapted from Peters et al. (2007).

Scientific Reasoning. In Experiment 2, participants completed the 11 item Scientific Reasoning Scale described in Drummond and Fischhoff (2017).

COVID-19 Beliefs & Behaviors. In Experiment 1, participants completed a 23-question survey with categorical or yes-no answers to assess beliefs and intended risk behaviors related to COVID-19. (See Appendix B for full survey items.) This included two measures of emotional risk perception (“How concerned are you about the COVID-19 outbreak in your area?”, 1 = “unsure”, 5 = “extremely concerned”; “How worried are you about getting sick with COVID-19?”, 1 = “unsure”, 5 = “extremely worried”) and three measures of quantitative risk perception, including the likelihood of hospitalization for someone your age, the likelihood of dying in hospital for someone your age, and how many people would catch the virus from contact with one contagious person.

To measure intended risk behaviors, participants self-reported their likelihood to enact four risk-preventive activities recommended by the Centers for Disease Control when they go out in public on a 5-point scale from “extremely unlikely” to “extremely likely.” These behaviors

were: 1) washing hands with water and soap, 2) avoiding touching the face, eyes or mouth with unwashed hands, 3) maintain physical distance from others, and 4) wearing a cloth or surgical mask. Participants also reported their current frequency of activities including leaving the home for work, leaving the home to buy groceries or supplies, getting food or groceries delivered, leaving the home to socialize with others, or attending a gathering with 10 or more people. Participants self-reported the official social distancing restrictions currently active in their community.

Participants also self-reported their frequency of use of different news outlets. News data was not analyzed for this study.

Experiment 2. In Experiment 2, the Experiment 1 measures of quantitative risk perception were removed and replaced with measures of psychometric risk factors as originally identified Bostrum et al. (2020). This included three items representing threat or dread risk and 2 items representing unknown risk. One additional item representing personal moral responsibility for combating COVID-19 was also added. All items for current and intended risk behaviors were repeated from the Experiment 1 survey.

News viewership measures were removed. Instead, three items were added representing trust in various public figures regarding conflicting information about COVID-19. Participants rated Journalists, Politicians, and Scientists on a 7-point scale from “Completely untrustworthy” to “Completely trustworthy.”

Procedure (Experiment 1)

All experimental stimuli and survey items were presented through an online Qualtrics survey. Participants completed the survey remotely using their own devices, and were required to use a computer (not a smartphone) for the task. Participants were directed from Amazon MTurk

to Qualtrics and provided digital consent before moving on to experimental tasks. Sessions lasted for approximately 30 minutes and participants were compensated with \$3.75. Liberal and conservatives were separately randomly assigned to complete either political or neutral problems. Participants first received instructions for and completed the 10 causal strength judgements in a random order. After providing 10 strength ratings, participants then completed the numeracy scale, the COVID-19 beliefs and behavior survey, and provided their age and gender. Participants were then debriefed and thanked before exiting the survey.

Procedure (Experiment 2)

Participants completed a Qualtrics survey using their own (non-Smartphone) device, as in Experiment 1. Sessions lasted for approximately 45 minutes. Participants were randomly assigned to complete either political or neutral problems. After reading instructions, participants completed both strength and confidence judgements for each of the 12 problems in a random order. Participants then completed the numeracy, scientific reasoning, and COVID-19 beliefs and behaviors measures and provided demographic information before being debriefed.

Results and Discussion for Experiments 1 and 2

In Experiment 1, we screened out moderates from the Mechanical Turk sample so all participants were conservative or liberal. In Experiment 2, we could not screen on ideology, so moderates were included in the sample. As shown in Table 1, conservatives in Experiment 1 were about 6 years older than liberals, $F(1, 280) = 14.6, p < .01$, but problem type was not significantly related to age. Table 2 presents descriptive statistics for the cognitive ability tests. The data from the COVID-19 survey is presented later. The main effect of ideology on numeracy was significant in Experiment 1, $F(1, 280) = 8.1, p = .005, \eta_p^2 = .028$, but not significant in

Experiment 2, $F(1, 297) = 1.7, p = .20$. The main effect of ideology on scientific reasoning was not significant in Experiment 2, $F(1, 297) < 1$.

Table 2. Descriptive statistics for cognitive ability tests Experiment 1 and 2. SDs in parentheses.

Variable	Experiment 1					
	Liberals			Conservatives		
	Total or Average	Political problems	Neutral problems	Total or Average	Political problems	Neutral problems
Numeracy % correct	91.5 (10.8)	92.3 (9.2)	91.5 (12.3)	87.7 (14.6)	89.2 (13.1)	86.2 (16.2)
	Experiment 2					
	Numeracy % correct	86.9 (13.8)	86.9 (15.4)	87.7 (13.8)	84.6 (15.4)	85.4 (16.9)
	Scientific reasoning % correct	63.6 (22.0)	57.6 (22.7)	70.8 (18.8)	62.2 (20.1)	62.7 (20.5)
						61.7 (19.7)

Testing Causal Strength Hypotheses

Our hypotheses regarding strength judgments were tested with repeated measures ANOVAs with within-subjects factors corresponding to the characteristics of the contingency datasets (e.g., nominal strength) and between-subjects factors of ideology and problem type (political vs. neutral). Because partial eta squared (η_p^2) overestimates effect sizes in within-subjects designs, whenever possible we used generalized eta squared, η_G^2 , to estimate effect size, which provides more accurate effect size measures in these designs. Bakeman (2005) provides calculation formulas for η_G^2 for common designs. However, since current statistical software does not calculate η_G^2 , we were not able to calculate it for some of the more complex designs in our analyses that were not included in Bakeman's article. In these cases, we used η_p^2 . This is a conservative strategy for estimating effect size. Like η_p^2 , η_G^2 represents the percentage of variance accounted for by a variable. Cohen (1988) suggested that variance accounted for values of .02, .13 and .25 represent small, medium and large effect sizes, respectively.

Experiment 1. We hypothesized that judged strength would increase with nominal strength and that, on political problems, liberals would overestimate strength and conservatives would underestimate strength relative to neutral problems. In addition, we expected that nominal strength and problem type would interact such that this overestimation or underestimation would be stronger when nominal causal strength differed the most from participants ideological bias (i.e., low nominal strength for liberals and high nominal strength for conservatives. We tested these hypotheses with separate repeated measures ANOVAs for liberals and conservatives with problem type as a between-subjects factor, nominal strength and proportional functional sample size as within-subjects factors and judged strength as the dependent variable.

Figure 5 shows the effect of ideology, problem type and nominal strength on strength judgments. Table 3 shows the effect of ideology and problem type on overall strength judgments. The figure suggests that strength judgments of all participants were strongly influenced by nominal strength but only liberals showed the expected ideological bias on political problems. The ANOVA confirmed these impressions. Liberals' strength judgments were significantly higher on political than neutral problems, although this effect was small, $F(1,148) = 7.28, p < .01, \eta_G^2 = .017$. Conservatives judgments were very similar on political and neutral problems, $F(1,132) < 1, \eta_G^2 = .001$. The hypothesized interaction such that liberals would overestimate more on political problems when the evidence conflicted more with their prior beliefs (i.e., at low nominal strength), was not significant, $F(2,26) < 1, \eta_G^2 = .002$, as liberals overestimated about equally at all nominal strengths. This means that liberals showed belief bias indiscriminately, without moderating it based on the strength of the evidence. The hypothesized interaction between problem type and nominal strength was not significant for conservatives, $F(1,148) < 1, \eta_G^2 = .001$.

Figure 5. Effect of ideology and problem type on causal strength judgments in Experiment 1. Error bars are 1 SD.

Table 3. Mean strength judgments by ideology and problem type. SDs are in parentheses.

	Experiment 1		Experiment 2		
	Conservatives	Liberals	Conservatives	Moderates	Liberals
Political problems	51.0 (12.8)	54.1 (7.3)	58.5 (12.0)	57.0 (10.8)	55.9 (17.1)
Neutral problems	51.8 (14.5)	50.5 (9.0)	53.4 (11.4)	52.5 (12.7)	51.9 (10.5)

The characteristics of the evidence in the fictional studies influenced participants' strength judgments much more strongly than ideological bias did. The hypothesized main effect of nominal strength was very large for liberals, $F(4,592) = 1674$, $p < .001$, $\eta_p^2 = .778$, and conservatives, $F(4,528) = 748.2$, $p < .001$, $\eta_p^2 = .601$. One notable unpredicted effect in these ANOVAs was that strength judgments increased with proportional functional sample size for

both liberals, $F(1,148) = 112.8, p < .001, \eta_G^2 = .094$, and conservatives, $F(1,132) = 35.3, p < .001, \eta_G^2 = .028$. These effects were small. The finding that proportional functional sample size influenced strength judgments, albeit to a much smaller extent than nominal strength, is consistent with Liljeholm and Cheng (2009). They explained this finding in terms of participants conflating confidence, which should increase with sample size, with strength.

In addition, because their biases were assumed to work in different directions, we hypothesized that liberals would make higher strength estimates than conservatives on political problems, but liberals and conservatives would make similar strength estimates on neutral problems. This hypothesis was not supported, as a simple effects ANOVA showed that liberals did not have significantly higher strength judgments than conservatives on political problems, $F(1, 142) = 2.65, p = .11, \eta_p^2 = .018$. Also, a problem type by ideology by nominal strength by proportional functional sample size repeated measures ANOVA showed that the problem type by ideology interaction was not significant, $F(1, 280) = 2.88, p = .09, \eta_p^2 = .010$. Because conservatives were older than liberals, age was a covariate in these two analyses.

Experiment 2. Figure 6 shows the effect of ideology, problem type and nominal strength on strength judgments for the college students in Experiment 2. Table 3 shows the effect of ideology and problem type on overall strength judgments. The figure suggests that, as in Experiment 1, liberals overestimated causal strength, especially when nominal strength was low; but, contrary to our hypothesis, conservatives also overestimated strength.

Simple effects ANOVAs for each ideological group confirmed these impressions. For liberals, the predicted main effect of higher strength judgments on political than neutral problems was not significant, $F(1,115) = 2.14, p = .15, \eta_G^2 = .022$. However, the predicted interaction such that liberals would overestimate strength on political problems mainly when the evidence

conflicted most with their prior beliefs (i.e., at low nominal strength) was significant, $F(2,230) = 6.63, p < .01, \eta_G^2 = .017$. For conservatives, the main effect of overestimating strength was significant, $F(1,182) = 8.73, p < .01, \eta_G^2 = .031$; but participants did not moderate their belief bias based on the strength of the evidence, as the problem type by nominal strength interaction was not significant, $F(2,364) < 1, \eta_G^2 < .001$. Given the unexpected finding that conservatives showed the same bias towards overestimating strength on political problems as liberals, our hypothesis that liberals would have greater strength judgments than conservatives on political problems was not supported, $F(1,158) < 1, \eta_p^2 < .001$.

Given that both liberals and conservatives showed the same bias, we also looked at the data from moderates, although we had no hypotheses about this group. They showed the same overestimation (see Figure 6). Moderates showed the same pattern as conservatives, with a significant main effect, $F(1,137) = 4.96, p = .03, \eta_G^2 = .027$, and a non-significant interaction, $F(2,274) < 1, \eta_G^2 < .001$.

The strength judgments of liberals, moderates and conservatives were affected similarly by the characteristics of the evidence. All showed very large main effects of nominal strength (Liberals: $F(2,230) = 667.7, p < .001, \eta_G^2 = .630$; Moderates: $F(2,274) = 961.5, p < .001, \eta_G^2 = .638$; Conservatives: $F(2,364) = 1155, p < .001, \eta_G^2 = .611$), medium effects of proportional functional sample size (Liberals: $F(1,115) = 93.3, p < .001, \eta_G^2 = .126$; Moderates: $F(1,137) = 199.8, p < .001, \eta_G^2 = .197$; Conservatives: $F(1,182) = 145.8, p < .001, \eta_G^2 = .106$), and small or negligible effects of total sample size (Liberals: $F(1,115) = 22.8, p < .001, \eta_G^2 = .017$; Moderates: $F(1,137) = 13.1, p < .001, \eta_G^2 = .010$; Conservatives: $F(1,182) = 31.4, p < .001, \eta_G^2 = .020$).

Figure 6. Effect of ideology and problem type on causal strength judgments in Experiment 2. Error bars are 1 SD. $N = 117$ liberals, 139 moderates, and 184 conservatives.

Summary. In Experiments 1 and 2, liberals give higher strength judgments when thinking about the effects of social distancing policies on COVID-19 than when thinking about politically neutral problems, which supported our hypothesis. However, the hypothesis that conservatives would underestimate strength on political problems was not supported in Experiment 1 and contradicted in Experiment 2. Thus, we found belief bias but, in one sample, not the ideological bias that we expected. Why did this happen? We think that an answer to this question follows from the idea that people are not biased by ideology per se but by specific task-relevant prior beliefs that may be associated with their ideology (Ajzen & Fishbein, 2005; Gampa et al., 2018). Therefore, we will look to the COVID-19 attitude survey to see what information it provides.

We also found that strength judgments were influenced very strongly by nominal strength ($\eta_G^2 > .6$), moderately by proportional functional sample size ($\eta_G^2 .10$ to $.20$), and less by total sample size ($\eta_G^2 .01$ to $.02$). We said in the introduction that rational models can help determine how close participants reasoning comes to optimal or normative. The fact that participants' strength judgments was influenced with a small to medium effect size by aspects of the sample size is suboptimal. This is like thinking that a vaccine is more effective if it has been tested on more participants.

Testing Hypotheses about Confidence

We only collected confidence data in Experiment 2. Because belief bias may involve overconfidence, we hypothesized that, irrespective of ideology and whether belief bias involves overestimation or underestimation of strength, confidence judgments would be higher on political than neutral problems. Since both liberals and conservatives did show belief bias for their strength judgments, we had a good situation to test this hypothesis. Using the liberals and

conservatives, we conducted a repeated measures ANOVA with problem type and ideology as between-subjects factors and nominal strength, proportional functional sample size and total sample size as within-subjects factors and confidence judgments as the dependent variable. The main effect of problem type was not significant, $F(1, 297) = 3.35, p = .068, \eta_p^2 = .011$. However, there was a significant interaction such that confidence was higher on political than neutral problems only at low and moderate nominal strength, as shown in Figure 7, $F(2, 594) = 6.66, p < .01, \eta_p^2 = .022$. Thus, the hypothesis was supported at low and moderate nominal strength but not at high nominal strength. Recall that most participants in this study showed bias by overestimating causal strength even when the evidence from the fictional studies (nominal strength) agreed with their biases by showing very high strength. Seen in this light, this confidence interaction suggests that when people make biased strength judgments when this is not necessary, they have less confidence in their biased judgment.

Figure 7. Interaction of problem type and nominal causal strength on confidence. Error bars are 1 SD.

Since confidence should increase with sample size, we hypothesized that participants' confidence judgments would increase with proportional functional sample size and with total sample size. These hypotheses were supported. For proportional functional sample size, $F(1, 297) = 20.5, p < .001$. For total sample size, $F(1, 297) = 51.3, p < .001$. Finally, we hypothesized an interaction such that the increase in judged confidence with proportional functional sample size and total sample size would be stronger for neutral than political problems. These hypotheses were not supported. For the first interaction, $F < 1$. For the second, $F(1, 297) = 1.09, p < .01$.

COVID-19 Attitude Survey

In both experiments, the COVID-19 survey had sub-scales (with 2-5 items) of: *emotional concern*, *risk perception*, *mitigation behaviors* and *risky behaviors*. In Experiment 2, we used Bostrum et al.'s (2020) psychometric model of risk perception to redefine our risk perception questions as *threat or dread risk* and add two new sub-scales related to risk perception, *uncertainty risk* and *moral responsibility*. We also added a question on likelihood of taking a COVID-19 vaccine and 3 questions on *trust in information sources* in Experiment 2. This meant that the items in only four sub-scales (*emotional concern*, *risk perception*, *mitigation behaviors* and *risky behaviors*) were common to the two surveys. In Experiment 1, the 10 items that were analyzed had a Cronbach's alpha of .49. The changes we made in this survey for Experiment 2 were beneficial, as Cronbach's alpha in the new 18-item survey was .76.

Table 4 shows the mean survey responses for participants on the sub-scales and some individual items, broken down by ideological group. For some sub-scales, we have included individual items that inform our discussion. The effect size and significance information show that in Experiment 1, conservatives showed less emotional concern and mitigation behavior and

Table 4. Responses on the COVID-19 surveys averaged over the items in sub-scales. In parentheses, on left is the number of items in the sub-scale; on right is the response scale from low to high. *d* refers to Cohen's *d* for the difference score Liberals – Conservatives.

	Experiment 1			Experiment 2			Expt. 3	
	Lib.	Con.	<i>d</i>	Lib.	Con.	<i>d</i>	Mod.	Con.
N participants	150	134		117	184		139	253
Emotional concern (2, 1-4)	3.06	2.40	0.76*					
Emotional concern (2, 1-5)				3.19	2.17	1.20*	2.38	2.32
Quantitative risk (2, 0-100)	21.8	20.6	0.06					
Threat or dread risk (3, 1-5)				3.10	2.27	1.20*	2.59	2.58
Moral responsibility (1, 1-5)				3.67	3.09	0.58*	3.22	3.02
Uncertainty risk (2, 1-5)				2.87	2.75	0.17	2.88	2.95
scientists knowledgeable				2.70	2.58	0.16	2.77	2.79
you informed				3.03	2.91	0.13	2.99	3.12
Mitig. behaviors (4, 1-5)	4.68	4.33	0.55*					
Mitig. behaviors (4, 1-7)				6.17	5.38	0.75*	5.78	5.48
Take vaccine (1, 1-7)				5.24	4.46	0.38*	4.85	3.68
Risky behaviors (2, 1-7)	1.50	1.98	-0.49*	4.13	4.98	-0.80*	4.50	2.93
Trust (3, 1-7)				4.20	3.61	0.85*	4.00	3.16
scientists				6.26	5.37	0.91*	5.97	4.60
journalists				4.03	2.92	0.69*	3.56	2.58
politicians				2.31	2.54	-0.20	2.45	2.29

* significant ($p < .001$) from t-test

more risky behavior regarding COVID-19 than liberals. In Experiment 2, compared to liberals, conservatives showed less emotional concern, lower threat or dread risk, less responsibility, less mitigation behavior (including intending to take a vaccine), less trust of scientists and journalists and more risky behavior. Most of these effects were medium or large. These findings fit with prior research on ideological differences in COVID-19 attitudes (Malloy & Schwartz, 2020; Miller, 2020; Pew Research Center, 2020). One notable sub-scale that did not show ideological

differences was uncertainty risk. Both liberals and conservatives perceived fairly small difference between their own and scientists' knowledge of COVID-19 (with negligible effect sizes, $d < .20$). Table E1 and E2 in Appendix E shows the survey responses for individual items.

Three aspects of these survey data help to explain our unexpected findings that conservatives in Experiment 2 showed bias in strength judgments that was expected of liberals. First, the conservative students in Experiment 2 showed high trust in scientists (5.4 on a 1-low to 7-high scale). Although this level of trust was less than that shown by liberals, it fit more closely with the liberal than the conservative ideological position (Gauchat, 2012). Second, conservatives and liberals in Experiment 2 considered themselves about equally well informed about COVID-19.

Most importantly, in both experiments conservatives rated themselves above the scale midpoint in terms of intention to use COVID-19 mitigation measures (in Experiment 1, 4.3 on a 1-5 scale with 5 = most likely; in Experiment 2, 5.3 on a 1-7 scale with 7 = most likely). Although conservatives endorsed mitigation behaviors to a lesser extent than liberals in both experiments, their behavioral intention regarding COVID-19 mitigation measures were closer to the liberal than the expected conservative position on this issue. This is important because the survey questions about mitigation measures (which included questions about masks and social distancing) can be taken as measuring participants prior belief about the manipulation in the fictional political studies in the causal reasoning task, which involved the ending of social distancing policies. Participants in both experiments were instructed that these policies included physical distancing, wearing masks, stay-at-home orders and closing businesses and schools (and in Experiment 2 these policies were on every display of the causal reasoning task). Recall that our hypothesis that conservatives would underestimate strength on political problems was based

on the assumption that they would strongly oppose mitigation measures. This assumption was not met in either study. These findings suggest that conservatives in Experiment 2 performed like liberals on the causal reasoning task because their prior beliefs about COVID-19 mitigation measures were similar to liberals' beliefs.

Discussion of Experiments 1 and 2

In two experiments with 724 participants and different samples, we found consistent evidence that participants' strength judgments were biased when faced with the politically charged COVID-19 problems. Four of the five ideological groups tested in the two studies—the liberals in Experiment 1 and all three ideological groups in Experiment 2—showed belief bias in which they overestimated strength on political problems relative to neutral problems. These belief bias effects were small in size, accounting for 1.7 to 3.1% of the variance in strength judgments. The only sample that did not show belief bias was conservatives in Experiment 1. However, the conservatives in Experiment 2 showed the overestimation bias predicted for liberals, which was the opposite of our prediction. Thus, our hypotheses about ideological belief bias were supported for liberals but not for conservatives.

Our post-hoc explanation for the counter-ideological belief bias shown by conservatives in Experiment 2 was that these participants did not show the strong opposition to COVID-19 mitigation measures that was expected based on their ideology. These findings suggest that conservatives failed to underestimate strength on political problems in Experiment 1 and overestimated it in Experiment 2 because their performance on the causal reasoning problems was influenced by their specific prior beliefs relevant to the COVID-19 reasoning problems, not by their ideology. This idea is consistent with evidence that general attitudes (e.g., towards policies and institutions) influence specific behaviors to a lesser extent than specific beliefs

related to the behavior (Ajzen & Fishbein, 2005). It is also consistent with our Bayesian strength model, which predicted belief bias based on specific prior beliefs about the strength of a candidate cause, not based on ideology.

There is an important theoretical takeaway from the finding that some conservatives showed belief bias that did not fit their ideology. Gampa et al. (2019) provided evidence that both conservatives and liberals were likely to rate arguments where the conclusion agreed with their ideology as sound and those where the conclusion disagreed with their ideology as unsound. Importantly, Gampa et al. concluded that the bias effects they found were not based on ideology *per se* but by whether participants' prior belief about the specific conclusion of an argument (e.g., tax increases harm the economy) agreed or disagreed with their ideology. A good example of this is Edwards and Smith's (1996) study of belief bias using realistic arguments that might be aligned with political ideology (e.g., the death penalty, abortion). Instead of measuring participants' ideology, the researchers used a pretest months before the main experiment to measure participants' strength of belief in the arguments' conclusions. When participants with a strong belief for or against a specific conclusion later evaluated arguments for those conclusions, they showed belief bias consistent with their measured prior belief. Edwards and Smith's findings support our argument that the conservatives in Experiment 2 exhibited bias that went against their ideology but did *not* show bias that went against their measured prior beliefs. Our argument here, which is similar to that of Pennycook (2020), is that ideology is a high-level variable that is sometimes correlated with belief bias—as in Kahan and Corbin (2016) and Kahan et al. (2013)—but these correlations are driven by specific beliefs associated with an ideology, not by the ideology *per se*.

This perspective suggests that a better way to frame our belief-bias hypotheses would be to suggest that people who strongly support the use of mitigation measures (rather than liberals) will overestimate strength on political (COVID-19) problems relative to neutral (non-political) problems; whereas people who weakly support (i.e., oppose) the use of mitigation measures (rather than conservatives) will underestimate strength on political problems relative to neutral problems. We tested this hypothesis in Experiment 3. We also tested the interaction hypotheses that bias will be greatest when the evidence in the fictional studies conflicts most strongly with prior belief. Thus, for strong supporters of mitigation measures, the degree of overestimation will be greater when nominal strength is low than when it is high; while, for weak supporters, the degree of underestimation will be greater when nominal strength is high than when it is low.

In order to test these hypotheses, we needed to sample participants who strongly opposed the use of mitigation measures in Experiment 3. Because participants with these beliefs are more likely to be conservative than liberal and because our hypotheses about liberals were supported in Experiments 1 and 2, we sampled only conservatives in Experiment 3. In the first two studies, our samples of college students and Mechanical Turk workers, who generally have high education levels (Clifford et al., 2015), were probably not representative of the population of the US. We conjectured that a sample that spanned wider education levels would find more variability in prior beliefs about mitigation measures.

Evidence on this question is mixed. One study found that education level was not correlated with people's perceptions of the effectiveness of COVID-19 mitigation measures (including mask wearing and social distancing), but people with more education were more likely to take a COVID-19 vaccine (Head et al., 2020; Kasting et al., 2020). A poll found that college graduates were more likely to take a COVID-19 vaccine than those with less education

(The Marist Poll, 2021). In Experiment 3, we required that 56% of participants have a high school degree as their highest level of education. We also stratified by age within both education levels.

Experiment 3 used the same causal reasoning problems as Experiment 2, except that the confidence question was not asked, i.e., participants only gave strength judgments. Participants also completed the numeracy scale and the COVID-19 survey.

Experiment 3

Methods

Participants. Experiment 3 tested similar hypotheses regarding participants' strength estimates as Experiments 1 and 2. The difference was that conservative vs. liberal ideology was replaced with weak vs. strong endorsement of mitigation measures. Thus, the same power analysis as for the first two experiments specified that we needed 130 participants to test each of the two simple effects hypotheses—for the weak and strong mitigation groups. As noted above, our main goal in Experiment 3 was to recruit participants who weakly endorsed mitigation measures in order to test our hypothesis that these participants would underestimate strength on the causal reasoning task. Therefore, we recruited only conservatives and recruited samples that were stratified by education level (highest degree of high school or less vs. 2-year college or above) and age (18-29, 30-54, 55 and above). Mechanical Turk indicated that participants in the low education group were scarce; so, we also recruited from Prolific. We set recruitment goals of 150 low education and 120 high education participants, with an equal number of participants in each age group within those education levels.

The Mechanical Turk participants were run first and participants from Experiment 1 could not participate. The Prolific participants were run a few days later. In Experiment 3, 122

participants were recruited through Mechanical Turk and 148 through Prolific. Of these, 16 were identified as political moderates (using the same definition as in Experiments 1 and 2) and excluded and 1 Prolific participant was excluded who reported participating in the same study on Mechanical Turk. The final sample size was 253. Participants in the lower and higher education groups were compensated at a rate of \$9.00 and \$7.50 per hour, respectively. All participants were treated in accord with APA ethical guidelines.

We defined weak belief in mitigation behaviors as responses of less than 4 (“neither likely nor unlikely”) on a 7-point scale from extremely unlikely (1) to extremely likely (7) for a question about the likelihood of wearing masks *and* a question about maintaining physical distance from other people on the COVID-19 survey. Strong belief in mitigation measures was defined as responses of greater than 4 on both questions. All other participants were defined as having moderate belief in mitigation measures. As Table 6 shows, recruiting more low education participants did not yield many participants with weak mitigation beliefs, as there were only 28 participants (11.1% of 253) in this group. There were 167 who had strong and 58 who had moderate beliefs about mitigation measures. In Experiment 2, only 3.4% of 440 college student were in the weak mitigation beliefs group. In Experiment 1, where we did not recruit based on education level or age, only 1.8% of 284 Mechanical Turk workers were in this group. Although 28 participants in the weak mitigation belief group was far below the sample size determined by our power analysis (130), the fact that we found so few people holding these beliefs after testing 977 participants in three experiments demonstrates that continuing to recruit participants until we achieved our sample goals was not practical. Therefore, we conducted our analyses with this inadequate sample size.

Table 6. Sample size, demographics and numeracy data for Experiment 3 for the weak and strong mitigation beliefs group. Numbers in parentheses are percentages for the sample sizes for the education groups and SDs for age.

Variable	Weak mitigation beliefs			Strong mitigation beliefs		
	Total or Average	Political problems	Neutral problems	Total or Average	Political problems	Neutral problems
N	28 (100%)	11 (100%)	17 (100%)	167(100%)	86 (100%)	81 (100%)
N low ed.	20 (71%)	7 (63%)	13 (76%)	88 (53%)	47 (53%)	41 (51%)
N high ed.	8 (29%)	4 (37%)	4 (24%)	79 (47%)	39 (47%)	40 (49%)
Age	42.6 (13.5)	45.5 (14.8)	40.6 (12.8)	43.1 (17.3)	44.0 (16.9)	42.1 (17.7)
N female	21	7	14	76	37	39
N male	7	4	3	91	49	42
Numeracy (%correct)	82.7 (17.0)	86.7 (12.0)	80.1 (19.4)	84.5 (16.0)	84.9 (16.0)	84.1 (16.0)

Design. We used a problem type (political, neutral; between subjects) by nominal causal strength (.25, .50, .75; within-subjects) by proportional functional sample size (.50, .75; within-subjects) by total sample size (1600, 6400; within-subjects) mixed factorial experimental design. Weak vs. strong beliefs in COVID-19 mitigation measures was included as a measured between-subjects variable. Within each of the six education by age recruitment groups, participants were randomly assigned to complete political or neutral problems using Qualtrics' randomization feature.

Measures. The numeracy scale was the same as in Experiment 1 and 2. The causal reasoning task was the same as in Experiment 2 except that the confidence questions were dropped. The COVID-19 survey was similar to Experiment 3 except for minor rewording of two questions.

Results and Discussion

As shown in Table 6, the mitigation belief and problem type groups differed very little in age and in numeracy. The weak mitigation belief group had a higher percentage of participants

with lower education levels (High School degree or less), 71.4%, compared to the strong mitigation belief group, 52.7%, although this difference only approached significance, $X^2 = 3.41$, $p = .065$. These data suggest that our strategy of looking for participants with weak mitigation beliefs among people with lower education levels was partially successful.

Testing Causal Strength Hypotheses. Figure 8 shows the data relevant to the hypotheses that participants with weak beliefs in COVID-19 mitigation measures would underestimate strength on the COVID-19 problems relative to the mineral problems, while those with strong beliefs would overestimate strength on the political problems. For the group with weak mitigation beliefs, the mean strength judgment for political problems was 42.9 (SD 11.4), which is less than the mean for neutral problems 52.4 (SD 11.5). A significant main effect of problem type supported the underestimation hypothesis for this group, $F(1,26) = 4.55$, $p = .042$, $\eta^2_G = .114$. However, because of the low sample size, the 95% confidence interval for the difference of 9.44 between these strength means was wide, [0.35, 18.54].

Because the low sample size for the weak mitigation group presented more opportunity for statistical error, we obtained an alternative estimate of the effect of political vs. neutral problems on strength estimates by using Bayesian parameter estimation, which uses effect sizes and estimates of confidence (credible intervals) to understand whether data support hypotheses (Kruschke & Liddell, 2018). A Bayesian t-test comparing the COVID and mineral problems for the weak mitigation group showed that the median effect size (Cohen's d) of the problem type effect was 0.63 and the 95% credible interval for d was [-0.06, 1.43]. In addition, since beliefs about COVID-19 mitigation measures should have no influence on strength judgments on neutral problems, we reasoned that there is no reason to use only the 17 participants in the weak mitigation belief group who completed neutral problems to represent participants' causal

strength judgments on these problems. Therefore, we conducted another Bayesian t-test where we ameliorated the sample size problem for the weak mitigation group by using the data for all participants on neutral problems (N 131). As expected, the mean strength judgment on neutral problems for this larger sample, 52.5 (*SD* 13.0), was very close to the 52.4 mean for the smaller sample with weak mitigation beliefs. For the second t-test, the median d was 0.61 and the 95% credible interval for d was [0.04, 1.23].

Figure 8. Effect of belief mitigation group and problem type on causal strength judgments in Experiment 3. Error bars are 1 SD.

These effect sizes from both the frequentist and Bayesian analyses are near medium or greater than medium, according to Cohen (1988). The Bayesian analyses reinforced the conclusions from the frequentist analysis regarding the most likely size of this effect and also the considerable uncertainty about the effect size estimate. This effect demonstrates that participants who weakly endorse mitigation beliefs underestimated causal strength. Furthermore, these effect size estimates are larger than the corresponding estimates for the strength overestimation effects

in Experiments 1 and 2. This underestimation finding is important because—taken together with the overestimation effects in our previous experiment—it more fully supports the predictions of the Bayesian updating model that belief bias can either increase or decrease strength estimates depending on the specific contents of participants' prior beliefs. Also, our finding that political biases can influence causal reasoning in opposite directions is similar to Kahan and Corbin's (2016) finding that ideological biases can increase or decrease belief in human-caused global warming. However, because of the low sample size, we cannot have high confidence in the exact size of the underestimation effect.

The hypothesized interaction such that people with weak mitigation beliefs would underestimate more when the evidence conflicted more strongly with their prior beliefs (i.e., at high nominal strength), was not significant, $F(2,26) < 1$, $\eta_G^2 = .004$, as participants underestimated about equally at all nominal strengths. Thus, as with the conservatives who overestimated strength in Experiment 2, participants did not moderate their belief bias based on the strength of the evidence, but instead gave biased judgments regardless of nominal strength.

Turning to the strong mitigation group, the predicted main effect of problem type—which involved overestimating strength—was not significant, $F(1,26) = 1.61$, $p = .207$, $\eta_G^2 = .008$. According to a Bayesian t-test, the median effect size of the problem type effect was $d = 0.18$, 95% credible interval $[-0.11, 0.48]$. The hypothesized interaction such that people with strong mitigation beliefs would overestimate more when the evidence conflicted most with their beliefs (at low nominal strengths) was not significant, $F(2,26) = 2.23$, $p = .119$, $\eta_G^2 = .003$.

Figure 9 shows the data relevant to the hypothesis that peoples' strength judgments should be influenced by their mitigation beliefs on political problems—with stronger mitigation beliefs associated with higher strength judgments—but not on neutral problems. We planned to

test this hypothesis with a repeated measures ANOVA with two between-subjects factors, mitigation belief group and problem type, and three within-subjects factors, nominal strength, proportional functional sample size, and total sample size. Because education level was correlated with mitigation belief group, we added it to the ANOVA. The interaction of mitigation beliefs and problem type was significant, $F(1,187) = 4.08$, $p = .045$, $\eta_p^2 = .021$, supporting this hypothesis. Education level did not show a significant main effect and did not interact with either of the other two between-subjects factors.

Because the participants in Experiment 3 were all conservatives, these findings support the claims that belief bias is driven more by specific beliefs relevant to the domain of reasoning than by political ideology, which have been made by Ajzen and Fishbein (2005) and Pennycook (2020). The conservatives in Experiment 3 had varied beliefs about COVID-19 mitigation measures and the evidence in Figure 9 demonstrates that participants' strength judgments on the COVID-19 problems increased with the strength of their beliefs in mitigation measures as predicted by the Bayesian belief updating models.

Figure 9. Mean judged causal strength by mitigation belief group and problem type. Error bars equal 1 SD.

These findings also suggest that, when investigating how political ideology may bias causal reasoning, ideological groups—the focus of Experiments 1 and 2—should not be treated as homogenous. Instead these groups should be decomposed based on their specific beliefs related to the reasoning task. The solid black and gray lines in Figure 10 show what the data for Experiment 3 look like if we ignore this injunction and present the data for all 253 conservative participants as a group, as in Experiments 1 and 2. These data are similar to the data for conservatives in Experiment 1 in that participants' strength judgments were very similar for political and neutral problems. In Experiment 1, we interpreted this finding as suggesting that conservatives did not show belief bias. However, Figure 10 also shows the data when the conservative ideological group is decomposed into participants who held strong and weak beliefs about mitigation measures (these are the same data as in Figure 8). Figure 10 suggests that we might find different belief bias effects in Experiments 1 and 2 if the ideological groups were decomposed based on participants' mitigation beliefs. This issue will be discussed in the General Discussion.

The lack of belief bias effects for the overall group of conservative participants in Experiment 3 (Figure 10) is similar to the findings for conservatives in Experiment 1 (Figure 5) but not Experiment 2, where conservatives overestimated strength on political problems like liberals did (Figure 6). To understand the reasons for this dissociation, we turn to the COVID-19 survey and discussion of demographic differences between these groups.

COVID-19 Attitude Survey. Table 4 presents the mean responses for Experiment 3 participants on the COVID-19 survey items. Because the conservative participant groups in Experiments 2 and 3 performed differently on the causal reasoning task, we compare the survey responses of those two groups here. All group differences mentioned in this section are

significant by a t-test, with alpha set to .005 to control for multiple tests. The average degree of support for four COVID-19 mitigation behaviors in Experiment 3 was 5.48 (above the midpoint of the 7 point scale), which was similar to Experiment 2 (5.38). Thus, conservatives in Experiment 3 were similar to all the ideological groups in the current studies in their positive attitude towards mitigation measures.

Figure 10. Effect of nominal strength on neutral and political problems for all participants in Experiment 3 and on political problems for the weak and the strong mitigation belief groups. Error bars are 1 SD.

Of particular importance in understanding the differences in performance by conservatives on the causal reasoning task in Experiments 2 and 3 is the finding that conservatives in Experiment 3 showed significantly lower overall trust of scientists, journalists and politicians as sources of information about COVID-19 (3.16, 7 point scale) compared to conservatives in Experiment 2 (3.61, Cohen's $d = 0.38$). Importantly, trust in scientists was quite high in the Experiment 2 conservatives (5.37) and significantly lower in Experiment 3 (4.60, $d =$

0.55). These differences in trust may be due to age differences, as the college students in Experiment 2 (mean age of conservatives 18.8) were much younger than the conservatives in Experiment 3 (mean age 43.0). For the Experiment 3 conservatives, a regression with age and education level as predictors showed that trust in journalists decreased significantly with age, $t = 2.94$, $Beta = -.182$, $p = .004$. This negative age correlation approached significance for combined trust in all three information sources, $t = 2.49$, $Beta = -.155$, $p = .013$. Education level was not related to trust in either regression.

In the introduction, we argued that conservatives would underestimate the strength of COVID-19 mitigation measures because the lower trust in science that characterizes their political ideology discourages them from accepting scientists' recommendations about mitigation measures. Thus, the conservative students in Experiment 2 went against the grain of their ideology with regard to trust in scientists and other information sources, while the conservatives in Experiment 3 did this to a significantly lesser extent. Perhaps the greater trust in scientists among the college-student conservatives in Experiment 2 compared to the older conservatives in Experiment 3 can help explain why the former group overestimated causal strength on the political problems but the latter group did not. The fact that the counter-ideological beliefs about science held by the college students were associated with causal strength judgments that showed a counter-ideological "liberal" belief bias provides further support for the claim that specific beliefs influence causal reasoning more than general attitudes.

General Discussion

Effects of Belief Bias on Causal Strength Judgments

When public figures argue that masks are not protecting you from COVID-19 (Dick, 2020) or that the ice caps are not melting to the degree predicted by global warming models

(@realDonaldTrump, 2013), they are making statements about causal strength, i.e., effect size, in order to make a political point. In this project, we investigated how belief bias, which may be grounded in ideological affiliations, influences peoples' causal strength judgments. As these examples and empirical research (Freimuth et al., 2017) demonstrate, peoples' beliefs about the strength of a cause can have an important influence on their decisions and actions. However, to our knowledge, little previous research has focused on how belief bias influences strength judgments. In much of the research on belief bias, peoples' judgments about the soundness of an argument are influenced by their prior beliefs about the argument's conclusion. For example, in Gampa et al. (2019) prior beliefs consisted of directional causal statements like "Tax increases harm the economy." In the current studies, peoples' judgments about the strength of a cause after seeing relevant evidence are assumed to be influenced by their prior beliefs about the strength of the cause, e.g., beliefs like "social distancing has no (or a strong) effect on COVID-19 cases."

We based our hypotheses about this question on rational Bayesian models of causal strength judgments (Lu et al., 2008). In particular, we adapted empirical work demonstrating that people tend to assume that most causes are strong (i.e., strong-uniform priors; Bramley et al., 2017; Yeung & Griffiths, 2015) by suggesting that, under the influence of belief bias, people may be predisposed to see some causes as even stronger than usual and other causes as weak (weak-uniform priors). The idea of weak priors for causal strength was the basis for our hypothesis that conservatives will underestimate strength when thinking about the effect of masks and social distancing on COVID-19 cases. To our knowledge, the idea of weak priors has not been discussed in the literature. We feel that using quantitative models to inform our predictions about belief bias adds needed precision to theories about belief bias.

In Experiments 1 and 2, where participants were sampled from fairly different populations, we found evidence supporting our hypothesis that liberals would show belief bias by overestimating causal strength when their prior political views caused them to view a causal factor (ending COVID-19 mitigation measures) as strong. However, conservatives in Experiment 1 did not show any belief bias and conservatives in Experiment 2 showed the overestimation bias expected for liberals. Experiment 3 provided evidence that peoples' specific prior beliefs about the domain of a causal reasoning problem influence belief bias when ideology was held constant. Conservatives' judgments of the strength of mitigation measures after seeing evidence increased as their degree of prior support for these measures increased. Furthermore, conservatives who strongly opposed mitigation measures underestimated causal strength on COVID-19 problems relative to neutral problems. This finding of belief bias involving underestimation is the first piece of evidence in the current studies to support Bayesian models assuming that people's prior beliefs about causal strength sometimes follow a weak-uniform distribution. However, this finding must be viewed as preliminary since it is based on only 28 participants.

Most of the belief bias effects in the current experiments were small. Effect sizes for the overestimation bias shown by liberals, moderates and conservatives in Experiments 1 and 2 ranged from 1.7 to 3.1% of variance accounted for. However, in Experiment 3, the effect size for the underestimation bias shown by conservatives who weakly supported mitigation measures was larger, accounting for 11.4% of variance. We note that other studies of belief bias have shown small effect sizes. For example, Gampa et al. (2019) found that belief bias effects driven by ideology and specific political beliefs were statistically significant but small.

Our belief bias findings were consistent for liberals, as both of the liberal groups we tested (Experiments 1 and 2) overestimated strength on political relative to neutral problems as

predicted. The one group of political moderates we studied (in Experiment 2) also overestimated strength. The belief bias effects for conservatives (considered as a group) were inconsistent and did not support our hypotheses regarding ideologically-driven belief bias, as conservatives either showed no belief bias in Experiments 1 and 3 or overestimated strength in Experiment 2.

The findings of Experiment 3 may provide an explanation for both the consistent belief bias effects for liberals and moderates and the mixed findings for conservatives. This experiment provided evidence that, instead of participants' general ideological attitudes influencing their judgments about the causal strength of mitigation measures, their specific beliefs about COVID-19 mitigation measures influenced these strength judgments. This supports empirical findings and theoretical claims in Ajzen and Fishbein (2005), Pennycook (2020) and Gampa et al. (2019). To make this point more clearly, Table 7 shows, for each experiment, the percentage of participants within each political ideology who showed weak vs. strong endorsement of mitigation measures. Importantly, these percentages are shown only for participants who saw the political, COVID-19 causal reasoning problems, as these are the participants whose causal strength judgments could be influenced by beliefs about COVID-19 mitigation measures.

Table 7. For each experiment and ideological group, the percentage of participants who saw political problems who were in the weak vs. strong mitigation belief groups.

	Experiment 1		Experiment 2		Experiment 3	
ideology	weak	strong	weak	strong	weak	strong
conservative	0	71.8	2.1	67.4	9.0	70.5
moderate			0	85.9		
liberal	0	94.5	3.1	90.8		

The table suggests that the reason liberal and moderate participants in Experiments 1 and 2 overestimated strength is that most of them (85-95%) strongly endorsed mitigation measures and almost none of them (0-4%) weakly endorsed these measures. Thus, for these groups, the

effect for the ideological group may be due to specific mitigation beliefs, since ideology and beliefs were strongly confounded. The mitigation beliefs of conservatives show more variance than liberals or moderates, with 67-72% in the strong endorsement group and 0-9% in the weak group. The conservatives in Experiment 3, where the mitigation beliefs varied the most, was the only sample where we had (barely) enough participants to test our hypotheses about causal strength judgments increasing with the degree of endorsement of mitigation measures. We found support for this hypothesis. However, despite showing greater variance in beliefs about mitigation measures than liberals or moderates, these conservatives still strongly supported these measures. This may explain why—as a group— conservatives in Experiment 3 did not show the ideological bias predicted in Experiments 1 and 2, i.e., underestimation of causal strength.

One remaining question concerns why conservatives (considered as a group) in Experiments 1 and 3 showed no evidence of belief bias while conservatives in Experiment 2 showed significant and counter-ideological belief bias. In the discussion of Experiment 3, we suggested that this unexpected bias effect may have occurred because the Experiment 2 participants trusted scientists more and were younger than Experiment 3 participants. Research suggests that trust in science plays a key role in people's beliefs about COVID-19 mitigation measures. In a correlational study conducted in the United States, Plohl and Musil (2020) found that political ideology had an indirect effect on compliance with COVID-19 mitigation measures mediated by trust in science but no direct effect. In addition, education had no influence on mitigation compliance when trust in science was in the model. In a similar study in Germany, trust in science and age were the two strongest predictors of compliance with COVID-19 mitigation measures in a model that included other demographic factors, perceptions of COVID-

19 risk and trust in politics (Dohle et al., 2020). In both studies, trust in science was positively correlated with compliance with mitigation measures.

Consistent with these findings, in Experiment 3, trust in scientists was strongly correlated with the 4-item subscale in the COVID-19 survey measuring endorsement of mitigation measures, $r = .498, p < .001$. For the Experiment 2 conservatives, where trust in science was higher and had less variability, trust in science still showed moderate correlations with the mitigation endorsement subscale, $r = .229, p = .002$. Thus, the greater trust in science among the college-student conservatives in Experiment 2 compared to the conservatives in Experiment 3 could have led the former group to have greater trust in the public recommendations of scientists that COVID-19 mitigation measures are a salutary response to the pandemic. This high trust in science, which is more characteristic of liberals than conservatives, could have biased these conservative students to show “liberal” overestimation bias in their judgments about the causal strength of mitigation measures.

This analysis suggests that the causes of belief bias are multifaceted and complex. In Figure 11, we present a hypothesized causal network that could be explored in further research. Some of the links in this network—between ideology, trust in science and endorsement of mitigation measures—are supported by prior research (e.g., Dohle et al., 2020; Gauchat, 2012; Plohl & Musil, 2020). All of these links are supported by significant correlations from the current experiments. In addition to the links just discussed, Experiments 1 and 2 documented the negative correlation between conservatism and endorsement of mitigation measures. This link is dashed to represent Plohl and Musil’s finding that ideology influenced compliance with mitigation measures only indirectly, through trust in science. Education and endorsement of mitigation measures were correlated in Experiment 3 (see Table 6). Finally, the belief bias

findings in Experiment 3 support the link between endorsement of mitigation measures and belief bias.

This hypothesized causal network emphasizes the idea that ideology is a general attitude that influences behavior indirectly—through the influence of more specific attitudes and beliefs like trust in science and beliefs about mitigation measures—rather than directly. In particular, we argued above that the belief bias effects in the three current experiments are better interpreted as driven by participants’ beliefs about mitigation measures than by ideology. Experiment 3, which showed that beliefs about mitigation measures influenced causal strength judgments when ideology was held constant, provides direct evidence for the claim that there should not be a direct link from ideology to belief bias in the causal network.

Figure 11. Hypothesized causal network showing influences on belief bias.

Effects of the Evidence on Causal Strength Judgments

In contrast to the small to medium belief bias effects we found, the characteristics of the evidence (i.e., the contingency data in the fictional studies) influenced causal strength much more strongly. Nominal strength accounted for more than 60% of variance in judged strength (a very large effect), while proportional functional sample size accounted for 10-20% of the variance. These findings mean that belief bias does not operate by disabling evaluation of the evidence. This is consistent with the Bayesian belief updating models discussed in the

introduction. According to these models, belief updating is influenced by both the evidence and prior beliefs in a multiplicative fashion. The extent to which the evidence vs. prior beliefs influence posterior beliefs is determined by the relative strength of these two components. For example, someone with an extremely strong prior belief that approached certainty would reach a conclusion supporting their belief even after evaluating moderately strong evidence that disconfirms their belief. It seems that, in our experiments, participants prior beliefs were not strong enough to overwhelm the impact of the evidence.

The very strong influence of characteristics of the evidence such as nominal causal strength is not surprising given the way the causal reasoning problems were displayed in our experiments. The contingency information needed to infer causal strength was displayed graphically and all the instances in a contingency dataset remained visible until the participant answers. Research demonstrates that people reason better with graphical than numerical displays of quantitative information (Garcia-Retamero & Hoffrage, 2013). From our theoretical perspective of Bayesian updating based on evidence, the graphical displays make the characteristics of the evidence that normatively should influence updating of causal strength judgments highly salient. This seems to explain why effect sizes for the within-subjects manipulation of nominal causal strength were very large. Nevertheless, belief bias explained additional variance in causal strength judgments after controlling for nominal strength.

How Beliefs about COVID-19 Vary with Ideology

One notable and unexpected finding of these studies was the high frequency of beliefs favoring use of COVID-19 mitigation measures and the rarity of beliefs opposing these measures, even among conservatives. In Experiments 1 and 2, which included liberals, moderates and conservatives, 1.8% and 3.4% of participants weakly endorsed both mask wearing

and social distancing. In Experiment 3, which included conservatives (only) and lower education levels than the other studies, 11.1% were in this category. In a study of attitudes towards wearing masks during the COVID-19 pandemic conducted in the United States and Canada, Taylor et al. (2021) found that 16% of participants did not wear masks and that not wearing masks was associated with political conservatism. These findings are consistent with our data based on behavioral intentions regarding two mitigation behaviors. Despite the media attention about strong anti-mask and anti-social-distancing attitudes among some people (Richardson, 2020; Szekares, 2020), it seems that attitudes towards mitigation measures are largely positive and compliance with using these measures is high in the United States. As educators associated with Clemson University, where Experiment 2 was conducted, we applaud these pro-public-health and pro-science attitudes among Clemson students of all political ideologies.

Applications

Since we found consistent belief bias involving overestimating causal strength in Experiments 1 and 2, we remind readers that overestimating the strength of a cause can be just as harmful as underestimating its effects. For example, people who believe that the COVID-19 or flu vaccine guarantees them immunity may take excessive health risks and become infected. Thus, the current research can help people evaluate health risks by increasing awareness of their bias blind spots (Pronin et al., 2002). We emphasize here that whether a particular scientific question becomes politicized and leads to bias depends on factors that may differ between and within countries.

We chose ending COVID-19 mitigation measures as the manipulation in the political fictional studies because, when we were planning Experiment 2 in June of 2020, the first surge of COVID-19 cases had diminished and many local governments in the US were doing just that,

e.g., removing mandates for social distancing and stay-at-home orders. This process was called “opening up” by politicians and journalists. Unfortunately, the following sequence has occurred in the US between July 2020 and January 2021 (IHME COVID-19 Forecasting Team, 2021) and in India between October 2020 and April 2021 (Mallapaty, 2021)—a surge leads to mitigation measure mandates being implemented, which are removed when the surge has only partially diminished, which leads to another surge. Thus, the current studies shed light on how politically-driven beliefs may bias the public’s understanding of the critical public health idea that premature removal of mitigation measures can have deleterious consequences.

One advantage of rational models is that researchers can determine the extent to which individuals’ reasoning is normative. We found that the sample size information in the fictional studies (e.g., proportional functional sample size) had a moderately strong influence on participant’s causal strength judgments, which suggests that participants conflated confidence with causal strength. This behavior is suboptimal, as increasing sample size should lead to higher confidence in causal conclusions but not higher causal strength judgments. The COVID-19 pandemic has created many situations where citizens have paid attention to what particular scientific studies have revealed about the effects of interventions including masks, social distancing, COVID-19 treatments or vaccines. For the public to understand these studies well, they need to understand the difference between causal strength and confidence in a study’s conclusions. For example, news reports of recent results of phase-3 clinical trials of COVID-19 vaccines (e.g., Grady, 2020) highlight the causal strength (effectiveness) of the vaccine and present information about sample size and statistical significance, but sometimes fail to clarify the degree of confidence in the effectiveness estimate (e.g., by using confidence intervals).

Educating journalists and the public about the distinction between causal strength and confidence would be a good place to start.

Since we asked participants in Experiment 2 and 3 how likely they were to take a COVID-19 vaccine, our findings are relevant to the question of how political ideology and associated beliefs influence vaccination intentions. Our vaccine intention data replicated our findings for other COVID-19 mitigation measures regarding the importance of trust in science. Although liberal college students in Experiment 2 were more likely to take a COVID-19 vaccine than conservatives, trust in science also had a strong positive correlation with vaccination intentions in Experiment 2 and for the conservatives in Experiment 3. These findings suggest that reducing strong vaccine resistance requires addressing mistrust of science.

Limitations

An obvious limitation is our very small sample size of participants who weakly endorsed mitigation measures in Experiment 3. However, since we found very few participants who held such beliefs after testing 977 participants, we are unclear how to rectify this. The prevalence of positive attitudes towards COVID-19 mitigation measures that we found is a good sign for humanity but, given our research question, a bad situation for researchers.

Conclusion

In conclusion, we found evidence in three experiments that participants made biased causal strength judgments when reasoning about evidence in the context of a politically-charged topic, COVID-19. The evidence from these experiments suggests that belief bias is driven more by specific beliefs relevant to the reasoning context than by general attitudinal factors like political ideology. This timely information could improve the ability of public health officials to advocate for preventive health interventions, including vaccines.

References

- Ajzen, I., & Fishbein, M. (2005). *The Influence of Attitudes on Behavior*. In D. Albarracín, B. T. Johnson, & M. P. Zanna (Eds.), *The handbook of attitudes* (p. 173–221). Lawrence Erlbaum Associates Publishers.
- Akesson, J., Ashworth-Hayes, S., Hahn, R., Metcalfe, R. D., & Rasooly, I. (2020). *Fatalism, beliefs, and behaviors during the COVID-19 pandemic* (No. w27245). National Bureau of Economic Research.
- Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review*, 98(3), 409. <https://doi.org/10.1037/0033-295X.98.3.409>
- Bakeman, R. (2005). Recommended effect size statistics for repeated measures designs. *Behavior research methods*, 37(3), 379-384. <https://doi.org/10.3758/BF03192707>
- Bostrom, A., Böhm, G., O'Connor, R. E., Hanss, D., Bodi-Fernandez, O., & Halder, P. (2020). Comparative risk science for the coronavirus pandemic. *Journal of Risk Research*. <https://doi.org/10.1080/13669877.2020.1756384>
- Bramley, N. R., Dayan, P., Griffiths, T. L., & Lagnado, D. A. (2017). Formalizing Neurath's ship: Approximate algorithms for online causal learning. *Psychological Review*, 124(3), 301–338. <https://doi.org/10.1037/rev0000061>
- Buehner, M. J., Cheng, P. W., & Clifford, D. (2003). From covariation to causation: A test of the assumption of causal power. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29(6), 1119–1140. <https://doi.org/10.1037/0278-7393.29.6.1119>
- Centers for Disease Control (2020, Dec. 1). *United States COVID-19 cases and deaths by state over time*. [Data set]. <https://data.cdc.gov/Case-Surveillance/United-States-COVID-19-Cases-and-Deaths-by-State-o/9mfq-cb36>

- Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, 104(2), 367-405. <https://doi.org/10.1037/0033-295X.104.2.367>
- Clifford, S., Jewell, R. M., & Waggoner, P. D. (2015). Are samples drawn from Mechanical Turk valid for research on political ideology?. *Research & Politics*, 2(4), 1-9. <https://doi.org/10.1177/2053168015622072>
- Dick, A. (2020. Nov. 18). *New medical study: Your face mask is not protecting you*. Ron Paul Institute for Peace and Prosperity. <https://www.ronpaulinstitute.org/archives/peace-and-prosperity/2020/november/18/new-medical-study-your-face-mask-is-not-protecting-you/>
- Dohle, S., Wingen, T., & Schreiber, M. (2020). Acceptance and adoption of protective measures during the COVID-19 pandemic: The role of trust in politics and trust in science. *Social Psychological Bulletin*, 15(4), e4315. <https://doi.org/10.32872/spb.4315>
- Drummond, C., & Fischhoff, B. (2017). Development and validation of the scientific reasoning scale. *Journal of Behavioral Decision Making*, 30(1), 26-38. <https://doi.org/10.1002/bdm.1906>
- Evans, J. St. B. T., Barston, J. & Pollard, P. (1983) On the conflict between logic and belief in syllogistic reasoning. *Memory and Cognition* 11:295–306. <https://doi.org/10.3758/BF03196976>
- Edwards, K., & Smith, E. E. (1996). A disconfirmation bias in the evaluation of arguments. *Journal of Personality and Social Psychology*, 71(1), 5.
- Evans, J. S. B., & Stanovich, K. E. (2013). Dual-process theories of higher cognition: Advancing the debate. *Perspectives on Psychological Science*, 8(3), 223-241. <https://doi.org/10.1177/1745691612460685>

- Fischhoff, B., Slovic, P., Lichtenstein, S., Read, S., & Combs, B. (1978). How safe is safe enough? A psychometric study of attitudes towards technological risks and benefits. *Policy Sciences*, 9(2), 127-152. <https://doi.org/10.1007/BF00143739>
- Freimuth, V. S., Jamison, A. M., An, J., Hancock, G. R., & Quinn, S. C. (2017). Determinants of trust in the flu vaccine for African Americans and Whites. *Social Science & Medicine*, 193, 70-79. <https://doi.org/10.1016/j.socscimed.2017.10.001>
- Gampa, A., Wojcik, S. P., Motyl, M., Nosek, B. A., & Ditto, P. H. (2019). (Ideo) logical reasoning: Ideology impairs sound reasoning. *Social Psychological and Personality Science*, 10(8), 1075-1083. <https://doi.org/10.1177/1948550619829059>
- Garcia-Retamero, & Hoffrage, U. (2013). Visual representation of statistical information improves diagnostic inferences in doctors and their patients. *Social Science & Medicine*, 83, 27-33. <http://dx.doi.org/10.1016/j.socscimed.2013.01.034>
- Gauchat, G. (2012). Politicization of science in the public sphere: A study of public trust in the United States, 1974 to 2010. *American Sociological Review*, 77(2), 167-187. <https://doi.org/10.1177/0003122412438225>
- Grady, D. (2020, November 16). Early Data Show Moderna's Coronavirus Vaccine Is 94.5% Effective. *The New York Times*. <https://www.nytimes.com/2020/11/16/health/Covid-moderna-vaccine.html?searchResultPosition=10>
- Griffiths, T. L., & Tenenbaum, J. B. (2005). Structure and strength in causal induction. *Cognitive Psychology*, 51(4), 334-384. <https://doi.org/10.1016/j.cogpsych.2005.05.004/>
- Gugerty, L. G., & Shreeves, M. (2020, Nov. 20). *Belief bias in causal learning*. [Poster] 61st Annual Meeting of the Psychonomic Society, Virtual

- Head, K. J., Kasting, M. L., Sturm, L. A., Hartsock, J. A., & Zimet, G. D. (2020). A national survey assessing SARS-CoV-2 vaccination intentions: Implications for future public health communication efforts. *Science Communication*, 42(5), 698-723.
<https://doi.org/10.1177/1075547020960463>
- Hoffrage, U. & Gigerenzer, G. (1998). *Academic Medicine*, 73(5), 538-540.
- IHME COVID-19 Forecasting Team (2021). Modeling COVID-19 scenarios for the United States. *Nature Medicine*, 27, 94-105. <https://doi.org/10.1038/s41591-020-1132-9>
- Kahan, D. M., & Corbin, J. C. (2016). A note on the perverse effects of actively open-minded thinking on climate-change polarization. *Research & Politics*, 3(4), 1-5.
<https://doi.org/10.1177/2053168016676705>
- Kahan, D. M., Peters, E., Dawson, E., & Slovic, P. (2013). Motivated numeracy and enlightened self-government. *Behavioural Public Policy*, 1, 54-86. <https://doi.org/10.1017/bpp.2016.2>
- Kasting, M. L., Head, K. J., Hartsock, J. A., Sturm, L., & Zimet, G. D. (2020). Public perceptions of the effectiveness of recommended non-pharmaceutical intervention behaviors to mitigate the spread of SARS-CoV-2. *PloS one*, 15(11), e0241662.
<https://doi.org/10.1371/journal.pone.0241662>
- Kelley, K., & Preacher, K. J. (2012). On effect size. *Psychological Methods*, 17(2), 137-152.
<https://doi.org/10.1037/a0028086>
- Kruschke, J. K. (2013). Bayesian estimation supersedes the t test. *Journal of Experimental Psychology: General*, 142(2), 573–603. <https://doi.org/10.1037/a0029146>
- Kruschke, J. K. & Liddell, T. (2018). The Bayesian New Statistics: Hypothesis testing, estimation, meta-analysis, and power analysis from a Bayesian perspective. *Psychonomic Bulletin & Review*, 25, 178–206. DOI 10.3758/s13423-016-1221-4

- Liljeholm, M., & Cheng, P. W. (2009). The influence of virtual sample size on confidence and causal-strength judgments. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 35(1), 157-172. <https://doi.org/10.1037/a0013972>
- Lord, C. G., Ross, L., & Lepper, M. R. (1979). Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology*, 37(11), 2098-2109. <https://doi.org/10.1037/0022-3514.37.11.2098>
- Lu, H., Yuille, A. L., Liljeholm, M., Cheng, P. W., & Holyoak, K. J. (2008). Bayesian generic priors for causal learning. *Psychological Review*, 115(4), 955-984.
<https://doi.org/10.1037/a0013256>
- Mallapaty, S. (2021). India's massive COVID surge puzzles scientists. *Nature*, 592, 667-668
<https://doi.org/10.1038/d41586-021-01059-y>
- Malloy, T., & Schwartz, D. (2020). Biden crushes "Sanders in democratic race; Quinnipiac University national poll finds more disapprove of Trump's response to coronavirus." March 9, 2020 [Survey Report]. Quinnipiac University Poll. Retrieved from
<https://poll.qu.edu/national/release-detail?ReleaseID=3657>
- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. San Francisco: W. H. Freeman.
- Miller, S. V. (2020). "Palmetto Poll Finds Broad Concern and Major Partisan, Racial/Ethnic Differences in South Carolina About COVID-19." April 21-29, 2020. [Survey report] Clemson, SC: Clemson University Palmetto Poll. Retrieved from
<http://newsstand.clemson.edu/wp-content/uploads/2020/05/palmetto-poll-covid-effects.pdf>
- Nyhan, B., Reifler, J., Richey, S. & Freed, G. (2013). Effective messages in vaccine promotion: A randomized trial. *Pediatrics*, 113(4), e835-e842.

- Orwell, G. (1946, March 22). In front of your nose. *Tribune*.
- Palan, S. & Schitter, C. (2018). Prolific.ac—A subject pool for online experiments. *Journal of Behavioral and Experimental Finance*, 17, 22–27. <https://doi.org/10.1016/j.jbef.2017.12.004>
- Pennycook, G. (2020). Belief bias and its significance for modern social science. *Psychological Inquiry*, 31(1), 57-60. <https://doi.org/10.1080/1047840X.2020.1722577>
- Peters, E., Dieckmann, N., Dixon, A., Hibbard, J. & Mertz, C. K. (2007). Less Is More in Presenting Quality Information to Consumers. *Medical Care Research Review*, 64, 169-190. doi: 10.1177/10775587070640020301
- Pew Research Center (2020). “Republicans, Democrats Move Even Further Apart in Coronavirus Concerns.” June 25, 2020 [Survey report]. Retrieved from <https://www.pewresearch.org/politics/2020/06/25/republicans-democrats-move-even-further-apart-in-coronavirus-concerns/>
- Plohl, N., & Musil, B. (2021). Modeling compliance with COVID-19 prevention guidelines: The critical role of trust in science. *Psychology, Health & Medicine*, 26(1), 1–12. <https://doi.org/10.1080/13548506.2020.1772988>
- Pronin, E., Lin, D. Y., & Ross, L. (2002). The bias blind spot: Perceptions of bias in self versus others. *Personality and Social Psychology Bulletin*, 28, 369–381.
- @realDonaldTrump. (2013, April 24). [Tweet] "@Jenil93_: @realDonaldTrump No melting of ice caps in the degree that global warming predicted. On the contrary increased ice in Antarctica <https://twitter.com/realDonaldTrump/status/326972596581191680>
- Richardson, V. (2020, July 20). Mask madness: Coronavirus confrontations escalate across U.S. Washington Times. <https://www.washingtontimes.com/news/2020/jul/20/coronavirus-mask-confrontations-escalateacross-us/>. [cited 2020 Jul 20].

- Secon, H. (2020, June 3). *An interactive map of the US cities and states still under lockdown — and those that are reopening*. Business Insider. <https://www.businessinsider.com/us-map-stay-at-home-orders-lockdowns-2020-3>
- Shreeves, M., Gugerty, L. & Moore, D. (2020). Individual Differences in Strategy Use and Performance during Fault Diagnosis. *Cognitive Research: Principles and Implications*. Online First Publication October 2020. <https://doi.org/10.1186/s41235-020-00250-5>
- Szekeres, J. (2020, Sep 27) Fight breaks out on bus in Surrey after passenger refuses to wear mask. News 1130. <https://www.citynews1130.com/2020/09/27/violent-fight-surrey-board-translink-bus/>. [cited 2020 Sept 27].
- Taylor, S., & Asmundson, G. J. G. (2021). Negative attitudes about facemasks during the COVID-19 pandemic: The dual importance of perceived ineffectiveness and psychological reactance. *PLOS ONE*, 16(2), e0246317. <https://doi.org/10.1371/journal.pone.0246317>
- The Marist Poll. (2021). NPR/PBS NewsHour/Marist Poll Results: the Biden Administration & COVID-19 (March 2021). Marist College Institute for Public Opinion. Downloaded March 31, 2021 from <http://maristpoll.marist.edu/npr-pbs-newshour-marist-poll-results-the-biden-administration-covid-19/#sthash.4d62OVNx.mOyoADL4.dpbs>
- Yeung, S., & Griffiths, T. L. (2015). Identifying expectations about the strength of causal relationships. *Cognitive Psychology*, 76, 1-29. <https://doi.org/10.1016/j.cogpsych.2014.11.001>
- Zaks, Tal. (2020). *A Phase 3, randomized, stratified, observer-blind, placebo-controlled study to evaluate the efficacy, safety, and immunogenicity of mRNA-1273 SARS-CoV-2 vaccine in adults aged 18 years and older*. (Protocol mRNA-1273-P301, Amendment 3). ModernaTX,

Inc. Retrieved from <https://www.modernatx.com/sites/default/files/mRNA-1273-P301-Protocol.pdf>

Appendix A: Posterior Distributions from Strength Model

The figure below shows posterior distributions of the causal strength of the candidate cause, $P(s_c|D)$, based on Lu et al.'s (2008) Bayesian causal strength model for three of the causal strength problems used in the current studies. Vertical lines (and numbers) show the mean of each posterior distribution. For each problem, three posterior distributions are shown, one for each of three prior distributions of causal strength: strong-uniform, $a = 20$; strong-uniform, $a = 5$; weak-uniform, $a = 10$. The problem associated with each graph is shown on the right. Problems are defined in terms of how often the effect is present and absent when the candidate cause is present and absent (E = effect; C = candidate cause; + = present; - = absent). Nominal causal strength is based on Eq. 2 in the text.

	E +	E -
C -	16	48
C +	52	12

nominal strength: .75

	E +	E -
C -	16	48
C +	40	24

nominal strength: .50

	E +	E -
C -	16	48
C +	28	36

nominal strength: .25

Appendix B: COVID-19 Surveys for Experiment 1 and 2

Experiment 1

Instructions:

Your next task consists of **about 20 questions** about **Coronavirus**.

When people become infected with **Coronavirus**, it causes a disease called **COVID-19**.

Questions:

1. How concerned are you about the COVID-19 outbreak in your area?

- ☐ unsure (14)
- ☐ not concerned at all (15)
- ☐ a little concerned (16)
- ☐ moderately concerned (17)
- ☐ extremely concerned (18)

2. How worried are you about getting sick with COVID-19?

- ☐ unsure (1)
- ☐ not worried at all (2)
- ☐ a little worried (3)
- ☐ moderately worried (4)
- ☐ extremely worried (5)

3. How **likely is it that someone your age will be hospitalized** if they become infected with Coronavirus? Enter a number from 0 to 100 where 0 = no chance, 50 = a 50% chance, and 100 = absolutely certain.

4. How **likely is it that someone your age would die if they are hospitalized** as a result of Coronavirus? Enter a number from 0 to 100 where 0 = no chance, 50 = a 50% chance, and 100 = absolutely certain.

5. On average, **how many people** do you think will catch the Coronavirus from **direct contact** with one contagious person? **Only consider** cases transmitted by coughing, sneezing, breathing, touch, or other **direct contact** with the first contagious person.
Enter a number

6. When any shelter-in-place or shutdown order related to COVID-19 was in effect (either in the past or currently), did you have to **leave the home to work** for any reason? (For example, if you are considered an “essential worker” or need to continue working for financial reasons.)

- ☐ No (1)
- ☐ Yes (2)
-

7. If and when you go out in public, how likely are you to do each of the following activities because of Coronavirus? *[all items use same response scale]*

Wash your hands with water and soap more than you usually would.

- ☐ extremely unlikely (1)
- ☐ unlikely (2)
- ☐ neither likely or unlikely (3)
- ☐ likely (4)
- ☐ extremely likely (5)

Avoid touching your face, eyes, or mouth with unwashed hands.
Maintain physical distance from other people.
Wear a cloth or surgical mask.

8. These questions have to do with how often **you or someone you live with** engages in various activities **at the current time**. *[all items use same response scale]*

How often do you leave the home for work?

- ☐ never (1)

- ☐ once a month (2)
- ☐ once every few weeks (3)
- ☐ once a week (4)
- ☐ a few times per week (5)
- ☐ daily (6)
- ☐ multiple times per day (7)

How often do you or someone you live with leave the home to buy groceries or supplies?

How often do you or someone you live with get food or groceries delivered?

How often do you leave the home to socialize with other people?

How often do you attend a gathering with 10 or more people?

9. How often do you get news from the following sources? *[all items use same response scale]*

CNN

Fox News

MSNBC

NPR

Breitbart News

PBS

- ☐ not at all (1)
- ☐ rarely (2)
- ☐ occasionally (3)
- ☐ frequently (4)
- ☐ very frequently (5)
- ☐ have not heard of this source (6)

Other (fill in, optional)

10. Is your state or community **currently** under any shelter-in-place or shutdown order related to COVID-19? (Only allowed to leave home for essential activities)

- ☐ No (1)
- ☐ Yes (2)

11. Please select **ALL** restrictions that apply to your state or community **at the current time**. If none apply, select "no restrictions."

- ☐ You are under a strict shelter-in-place order. (Only allowed to leave home for essential activities) (1)
- ☐ Only essential businesses are open. (2)
- ☐ Essential businesses and some non-essential businesses are open. (For example, some states have re-opened theaters or gyms.) (3)
- ☐ People are required to wear cloth or surgical masks in public. (4)
- ☐ Other (5) _____
- ☐ No restrictions (7)
-

Experiment 2

Instructions: Your next task consists of **21 questions** about **Coronavirus** and **COVID-19**. When people become infected with **Coronavirus**, it causes a disease called **COVID-19**.

Questions:

How concerned are you about the COVID-19 outbreak in your area?

- ☐ not concerned at all
- ☐ a little concerned
- ☐ moderately concerned

- ☐ very concerned
- ☐ extremely concerned

How worried are you about getting sick with COVID-19?

- ☐ not worried at all
- ☐ a little worried
- ☐ moderately worried
- ☐ very worried
- ☐ extremely worried

How much does the thought of getting COVID-19 fill you with dread? Getting COVID is ...

- ☐ not dreadful at all
- ☐ a little dreadful
- ☐ moderately dreadful
- ☐ very dreadful
- ☐ extremely dreadful

How serious a threat is COVID-19 to peoples' health? (same response scale as below)

How serious a threat is COVID-19 to your health?

- ☐ no threat at all
- ☐ a little serious
- ☐ moderately serious
- ☐ very serious

- ☐ extremely serious

How much can you personally contribute to mitigating COVID-19?

- ☐ can do nothing personally
- ☐ can do a little bit personally
- ☐ can do a moderate amount personally
- ☐ can do a lot personally
- ☐ can do a great deal personally

How well is COVID-19 currently understood by science?

- ☐ not at all understood
- ☐ a little bit understood
- ☐ moderately well understood
- ☐ very well understood
- ☐ extremely well understood

How well informed do you feel about COVID-19?

- ☐ not at all informed
- ☐ a little bit informed
- ☐ moderately well informed
- ☐ very well informed
- ☐ extremely well informed

If and when you go out in public, how likely are you to do each of the following activities **because of Coronavirus?**

Wash your hands with water and soap more than you usually would

Avoid touching your face, eyes, or mouth with unwashed hands

Maintain physical distance from other people

Wear a cloth or surgical mask

- ☐ extremely unlikely
- ☐ unlikely
- ☐ somewhat unlikely
- ☐ neither likely nor unlikely
- ☐ somewhat likely
- ☐ likely
- ☐ extremely likely

If a COVID-19 vaccine that was approved by the US government as safe and effective were available today, would you take this vaccine?

- ☐ extremely unlikely
- ☐ unlikely
- ☐ somewhat unlikely
- ☐ neither likely nor unlikely
- ☐ somewhat likely

- ☐ likely
- ☐ extremely likely

These questions have to do with how often **you** engage in various activities **at the current time**.

How often do you leave the home for work?

How often do you leave the home to socialize with other people?

How often do you attend a gathering with 10 or more people?

- ☐ never
- ☐ once a month
- ☐ once every few weeks
- ☐ once a week
- ☐ a few times per week
- ☐ daily
- ☐ multiple times per day

Journalists, politicians, and scientists sometimes give factual information and advice **about COVID-19** that conflicts with information and advice from other sources. If you want to find out **what is going on with COVID-19 or get advice on how to deal with COVID-19**, indicate **how trustworthy** you believe each of these groups to be.

Journalists

Politicians

Scientists

- ☐ Completely untrustworthy
- ☐ Untrustworthy
- ☐ A little bit untrustworthy
- ☐ Neither trustworthy nor untrustworthy
- ☐ A little bit trustworthy
- ☐ Trustworthy
- ☐ Completely trustworthy

In the media and in politicians' statements, you can hear very different opinions about how **serious a threat the COVID-19** pandemic presents to the **health** of people in our country, ranging from a very serious health threat that requires a major effort to manage it to a relatively minor one. You also hear very different opinions about **measures to slow the spread of COVID-19**, like social distancing, masks or vaccines. In your opinion, how serious a health threat is presented by COVID-19? Also, what is your opinion about measures to slow the spread of COVID-19?

Please select **ALL** restrictions that apply **at the current time** to **the community where you live**. You **may select more than one**. If none apply, select "no restrictions."

- ☐ You are under a strict shelter-in-place order. (Only allowed to leave home for essential activities)
- ☐ Only essential businesses are open.
- ☐ Essential businesses and some non-essential businesses are open. (For example, some states have re-opened theaters or gyms.)
- ☐ People are required to wear cloth or surgical masks in public.

☐

Other _____

☐

No restrictions

Appendix C: Correlation between judgments of confidence in causal structure and strength

Conceptually, the Griffiths and Tenenbaum's (2005) causal support model and our precision model (discussed in the Hypotheses section) are different, since they estimate confidence in causal structure and strength judgments, respectively. However, we ran both the causal support and precision models on 64 contingency datasets based on a wide range of causal powers (.25, .5, .75, 1), proportional virtual sample size (.25, .5, .75, 1) and total sample sizes (32, 48, 64, 128). The output of the causal support model, *logBayesFactor*, had a linear correlation of 0.83 with the output of the precision model, *logPrecision*. However, the model outputs had a curvilinear (power-function) relationship and the shape of this relationship varied with causal power. For causal powers of 1 and .75, the correlation of *logPrecision* with *logBayesFactor^{exp}* was maximized at 0.99 and 0.98, respectively, when *exp*=0.0001, i.e., when $\log Precision \approx \log (\log BayesFactor)$. For power of .5, *r* was 0.97 when *exp*=0.35. For power of .25, *r* was 0.97 when *exp*=0.42. Thus, confidence in causal structure is mathematically very similar to confidence in causal strength.

Appendix D: Instructions and Example Problems for Causal Reasoning Task

Example Neutral Problem

The text for the strength and confidence questions is presented in the text.

This problem has a nominal strength of .25, a proportional sample size of .5 and a total sample size of 64.

Instructions

Below is the first part of the instructions for political and neutral problems in both experiments. After this introductory text, participants saw diagrams of the causal reasoning problems with explanations of their elements and an example of the strength and (in Experiment 2) confidence questions.

Experiment 1 – Political problems

Here is some info you will need to complete your first task:

When people become infected with **Coronavirus**, it causes a disease called **COVID-19**.

At certain times, many states and cities in the US have put in place **strict social distancing policies** in order to slow the spread of COVID-19.

In some locations, these policies have included some or all of the following requirements.

Strict social distancing policies

- keeping at least **6 feet away** from other people in public places

- **wearing masks** when less than 6 feet away from people in public
- **closing many businesses and schools**, so that only essential workers are working at their normal workplace
- **shelter in place**, where people cannot leave their neighborhood except for essential trips like going to the grocery store or doctor's office

In answering the following questions, imagine that a group of researchers is studying whether or not **ending these strict social distancing policies** (in other words, **opening up**) **increases the number of people who get diagnosed with COVID-19**.

These researchers have conducted studies in **10 small cities**.

In each of these cities, imagine that **all of the strict social distancing policies** described above were put into effect within a few days after COVID-19 cases started to be diagnosed in the city.

These policies were in effect for a few months and were **followed by most people** in the city.

Then, in each city, **all** of these policies were suddenly **ended** (no longer in effect).

In **each city**, the researchers randomly chose a **group of 6400 people** to participate in the study.

- **JUST BEFORE the social distancing policies were ended**, the researchers measured **how many people in this group had been diagnosed with COVID-19** during the few months since COVID-19 cases started to be diagnosed in that city. In other words, they measured the **total, cumulative number of COVID-19 diagnoses** up to that point in time.
- **A FEW MONTHS AFTER the social distancing policies were ended**, the researchers again measured how many people in this group had been diagnosed with COVID-19 since COVID-19 cases started to be diagnosed in that city. In other, they measured the **total, cumulative number of COVID-19 diagnoses** up to that later point in time.

You will see what the researchers found from each of the 10 studies. After seeing the results of each study, you will answer a question about what you think the results mean. Here is an **example problem** showing what the results of a study will look like.

Experiment 1 – Neutral problems

Here is some info you will need to complete your first task:

A drug company is developing an **allergy medicine** and wants to find out if the medicine might have a side effect of **headaches**.

The scientists believe that the different **minerals** that make up the allergy medicine may influence headache differently, so they are testing each of the **10 minerals** in the medicine.

In answering the following questions, imagine that researchers are studying whether or not **taking each mineral increases the number of people who get headaches**.

Each mineral is tested in a research study run in a different lab.

In **each study**, the researchers from that lab randomly chose a **group of 6400 people** from the local area to participate in the study.

- **BEFORE the people in a study took the medicine**, the researchers measured **how many of them had headaches**.
- **AFTER the people in a study took the medicine**, the researchers measured **how many of them had headaches**.

You will see what the researchers found from each of the 10 studies. After seeing the results of each study, you will answer a question about what you think the results mean. Here is an **example problem** showing what the results of a study will look like.

Experiment 2 – Political problems

Here is some info you will need to complete your first task:

When people become infected with **Coronavirus**, it causes a disease called **COVID-19**.

At certain times, many states and cities in the US have put in place **strict social distancing policies** in order to slow the spread of COVID-19. In some locations, these policies have included some or all of the following requirements.

Strict social distancing policies

- keeping at least **6 feet away** from other people in public places
- **wearing masks** when less than 6 feet away from people in public
- **closing many businesses and schools**, so that only essential workers are working at their normal workplace
- **stay-at-home** (also called **shelter-in-place**) **orders**, where people cannot leave their neighborhood except for essential trips like going to the grocery store or doctor's office

In answering the following questions, imagine that a group of researchers wants to find out if **ending these strict social distancing policies** (in other words, **opening up**) might have an effect on **the number of people who get diagnosed with COVID-19**.

These researchers believe that **ending social distancing policies** may influence **the number of COVID-19 diagnoses differently in different US cities**, so they ran studies in **12 small or large cities in various parts of the US**. The researchers are studying whether or not **ending strict social distancing policies in a particular city** increased **the number of people diagnosed with COVID-19 in that city**.

In each of these cities, imagine that **all of the strict social distancing policies** described above were put into effect within a few days after COVID-19 cases started to be diagnosed in the city. These policies were in effect for a few months and were **followed by most people** in the city. **Then**, in each city, **all** of these policies were suddenly **ended** (no longer in

effect).

In **each city**, the researchers randomly chose a **group of 1600 or 6400 people** to participate in the study.

- **JUST BEFORE the social distancing policies were ended**, the researchers measured **how many people in this group had been diagnosed with COVID-19** during the few months since COVID-19 cases started to be diagnosed in that city. In other words, they measured the **total, cumulative number of COVID-19 diagnoses** up to that point in time.
- **A FEW MONTHS AFTER the social distancing policies were ended**, the researchers again measured how many people in this group had been diagnosed with COVID-19 since COVID-19 cases started to be diagnosed in that city. In other words, they measured the **total, cumulative number of COVID-19 diagnoses** up to that **later** point in time.

You will see what the researchers found from each of the 12 studies. After seeing the results of each study, you will answer some questions about what you think the results mean. Here is an **example problem** showing what the results of a study will look like.

Experiment 2 – neutral problems

Here is some info you will need to complete your first task:

In answering the following questions, imagine that a drug company is developing a new allergy medicine and wants to find out if the **allergy medicine** might have a side effect of **headaches**.

Researchers at the drug company believe that the **different minerals** that make up the allergy medicine may influence **headache differently**, so they ran studies testing each of the **12 minerals** in the medicine. The researchers are studying whether or not **taking each mineral** increased **the number of people who got headaches**.

In **each study**, the researchers randomly chose a **group of 1600 or 6400 people** to participate in the study.

- **BEFORE the people in the study took the mineral**, the researchers measured **how many of them had headaches**.
- **AFTER the people in that study took the mineral**, the researchers measured **how many of them had headaches**.

You will see what the researchers found from each of the 12 studies. After seeing the results of each study, you will answer some questions about what you think the results mean. Here is an **example problem** showing what the results of a study will look like.

Appendix E: COVID-19 Survey Data

Table E1. Responses on the individual items of the COVID-19 survey for Experiment 1. The response scale from low to high is in parentheses. d refers to Cohen's d for the difference score Liberals – Conservatives.

Sub-scales	Items	Liberals	Conservatives	d
Emotional concern (1-4)	concern	3.23	2.54	0.75
	worry	2.89	2.25	0.69
Risk perception (0-100)	risk hospitalized	27.63	25.74	0.08
	risk die	15.97	15.54	0.02
Safety behavior (1-7)	wash hands	4.67	4.44	0.30
	no touch	4.47	4.25	0.25
	distance	4.76	4.50	0.36
	wear mask	4.83	4.13	0.78
Risky behavior (1-7)	socialize	1.88	2.34	-0.34
	large gatherings	1.12	1.63	-0.53
Frequency of consuming media (1-5)	CNN	3.29	1.84	1.21
	Fox	1.44	3.17	-1.51
	MSNBC	2.59	1.69	0.80
	NPR	2.89	1.83	0.78
	Breitbart	1.25	1.95	-0.58
	PBS	2.30	1.81	0.44

Table E2. Responses on the individual items of the COVID-19 survey for Experiment 2. The response scale from low to high is in parentheses. *d* refers to Cohen's *d* for the difference score Liberals – Conservatives.

Sub-scales	Items	Liberals	Conserv- atives	<i>d</i>	Moderates
Emotional concern (1-5)	concern	3.50	2.37	1.24*	2.63
	worry	2.88	1.96	0.92*	2.14
Threat or dread risk (1-5)	dread	3.02	2.15	0.82*	2.36
	threat to community	3.97	2.96	1.21*	3.45
	threat to you	2.32	1.69	0.72*	1.95
Responsibility(1-5)	moral responsibility	3.67	3.09	0.58*	3.22
Unknown risk (1-5)	scientists knowledgeable	2.70	2.58	0.16	2.77
	you are informed	3.03	2.91	0.13	2.99
Safety behaviors (1-7)	wash hands	6.19	5.67	0.41*	5.96
	no touch	5.53	4.86	0.44*	5.23
	distance	6.16	4.94	0.82*	5.50
	wear mask	6.80	6.05	0.70*	6.45
	take vaccine	5.24	4.46	0.38*	4.85
Risky behaviors (1-7)	socialize	4.89	5.63	-0.70*	5.35
	large gatherings	3.37	4.33	-0.70*	3.65
Trust (1-7)	Journalists	4.03	2.92	0.83*	3.56
	politicians	2.31	2.54	-0.20	2.45
	scientists	6.26	5.37	0.91*	5.97

* significant ($p < .05$) from t-test