

Supplementary Information for

Social Learning and Partisan Bias in the Interpretation of Climate Trends

This PDF file includes:

Supplementary Materials and Methods

Supplementary Discussion

Supplementary Figures

Supplementary References

Supplementary Materials and Methods

Experiment Design. In each trial of this study, subjects were randomized into one of four experimental conditions: *i*) a control group consisting of participants with the same political ideology, *ii*) a social network composed of an equal number of conservative and liberal subjects, where subjects were shown the average answer of their network neighbors as a social signal, *iii*) a social network composed of an equal number of conservative and liberal subjects, where subjects were exposed to the logos of the Democratic and Republican parties, along with the average answer of their network peers, and *iv*) a network with an equal number of conservative and liberal subjects, where participants were exposed to the political ideology of their peers, along with their average answer. In each trial, each condition contained 40 individuals, such that each experimental trial contained 200 individuals. We conducted 12 independent trials of this design.

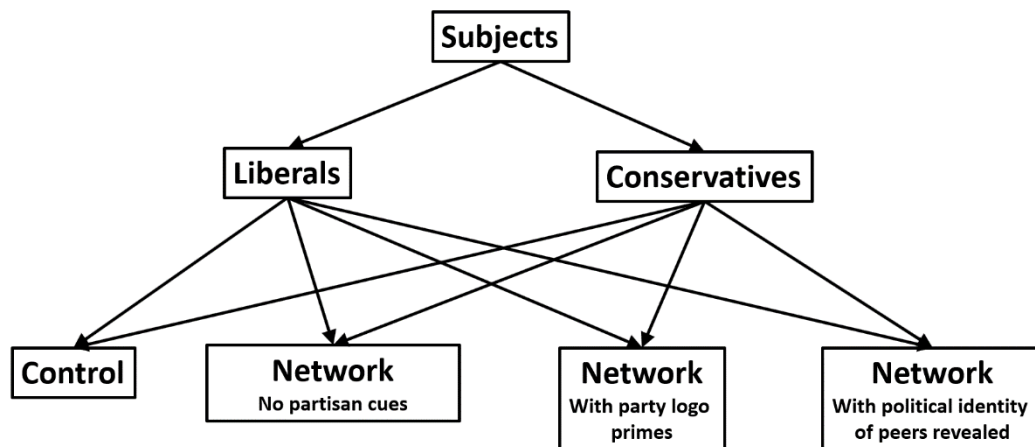


Fig. S1. A schematic presentation of the experimental design.

When subjects entered the network conditions, they were randomly assigned to one node in the network, and they maintained this position throughout the experiment. The network conditions employed a random network where every node had the same number of connections.

We constructed a network with 4 edges per node, and we employed the same network topology across all network conditions to minimize variance. We used random decentralized networks because previous experiments illustrate that this topology is the most reliable for generating social learning in online collective intelligence tasks (1). Centralized networks have been shown to improve collective intelligence only when the central, more influential nodes are more accurate. However, if the central nodes are not more accurate than the network, then their increased influence can bias the answers of the network toward their more inaccurate estimates. For this study, decentralized networks provided a way to prevent the partisan biases of individuals from spreading throughout the whole network. Decentralized networks give every participant access to the collective intelligence of their group, where exposure to the answers of peers can present individuals with opposing views and thereby create the opportunity to overcome biases.

Every condition contained an equal number of liberals and conservatives. We chose to sort subjects based on political ideology because political ideology has been shown to be a highly salient dimension for partisan bias in the domain of climate change (2, 3), and because people have been found to strongly define their sense of political identity on the basis of political ideology (4–6). We chose to use the logos for the Republican and Democratic parties as a minimal prime, because recent studies (7–9) show that these logos are highly effective at priming implicit partisan bias based both on party membership and political ideology.

Data collection for this design took place across three experimental periods. In each period, subjects were randomized into either a control condition or a network condition. In the first period, participants were randomized into either *i*) a control group consisting of participants with the same political ideology, or *ii*) a social network composed of an equal number of conservative and liberal subjects, in which subjects were shown only the mean of their network neighbors' estimates. In

the second period, participants were randomized into either *i*) a control group consisting of participants with the same political ideology, or *ii*) a network with an equal number of conservative and liberal subjects, where participants were exposed to the average estimates of their network neighbors, along with their political ideology. In the third period, participants were randomized into either *i*) a control group of participants with the same political ideology, or *ii*) a social network composed of an equal number of conservative and liberal subjects, where subjects were exposed to the logos of the Democratic and Republican parties below the mean of their neighbors' estimates. In each experimental period, 4 independent groups were gathered for both the liberal and conservative control conditions, and 12 independent groups were gathered for the network condition.

There were no significant baseline differences in the average trial-level estimates of conservatives across experimental conditions ($n=53$, $P=0.54$, Kruskal-Wallis H test), nor among liberals across experimental periods ($n=53$, $P=0.92$, Kruskal-Wallis H test), including the data from the robustness tests on echo chamber effects. In the main analyses of this study, we compare 12 independent control groups for both liberals and conservatives, with 12 independent networks without partisan priming, 12 independent networks with the political identities of network peers revealed, and 12 independent networks with partisan priming. We further conducted trials on structured echo chamber effects, in which 5 independent networks were filled with only conservatives, and 5 independent networks were filled with only liberals.

Subjects were presented with NASA's graph (Fig. 1) and asked to forecast the amount of Arctic Sea ice in 2025. To identify bias in the participants' interpretation of NASA's climate data, we consider whether participants' estimates correspond to the correct trend –i.e., above or below

the final point on the graph – as identified by climate change experts at NASA. In this case, the correct trend is downward from the endpoint in the graph (10,11).

Subject Experience. To isolate the effect of social influence, the interface was identical across the network and control conditions. When interpreting the graph, participants were given two opportunities to revise their answer. The only difference between conditions was that, in the network conditions, subjects were shown peer information while they made their revision, whereas controls were not. In networks without partisan cues, subjects were shown the average answers of their network neighbors, and then they were permitted to revise their responses. In networks with partisan priming, subjects were not only shown the average answers of their network neighbors, but they were also shown the party logos of the Democratic and Republican parties on round 2 and 3. In networks with political identity markers, subjects were not only shown the average answers of their network neighbors, but they were also shown the usernames and political ideology of their network neighbors (see Fig. S4). The same standardized usernames were shown to all players so that identity signaling was restricted to political ideology. The usernames shown were designed to be neutral, so that the choice of username did not reflect additional identity signaling, such as names that evoke identity characteristics related to gender or race. Also, political ideology did not dictate the order of the names displayed below the group average. This order was determined randomly for each player to avoid order effects (See Fig. S2).

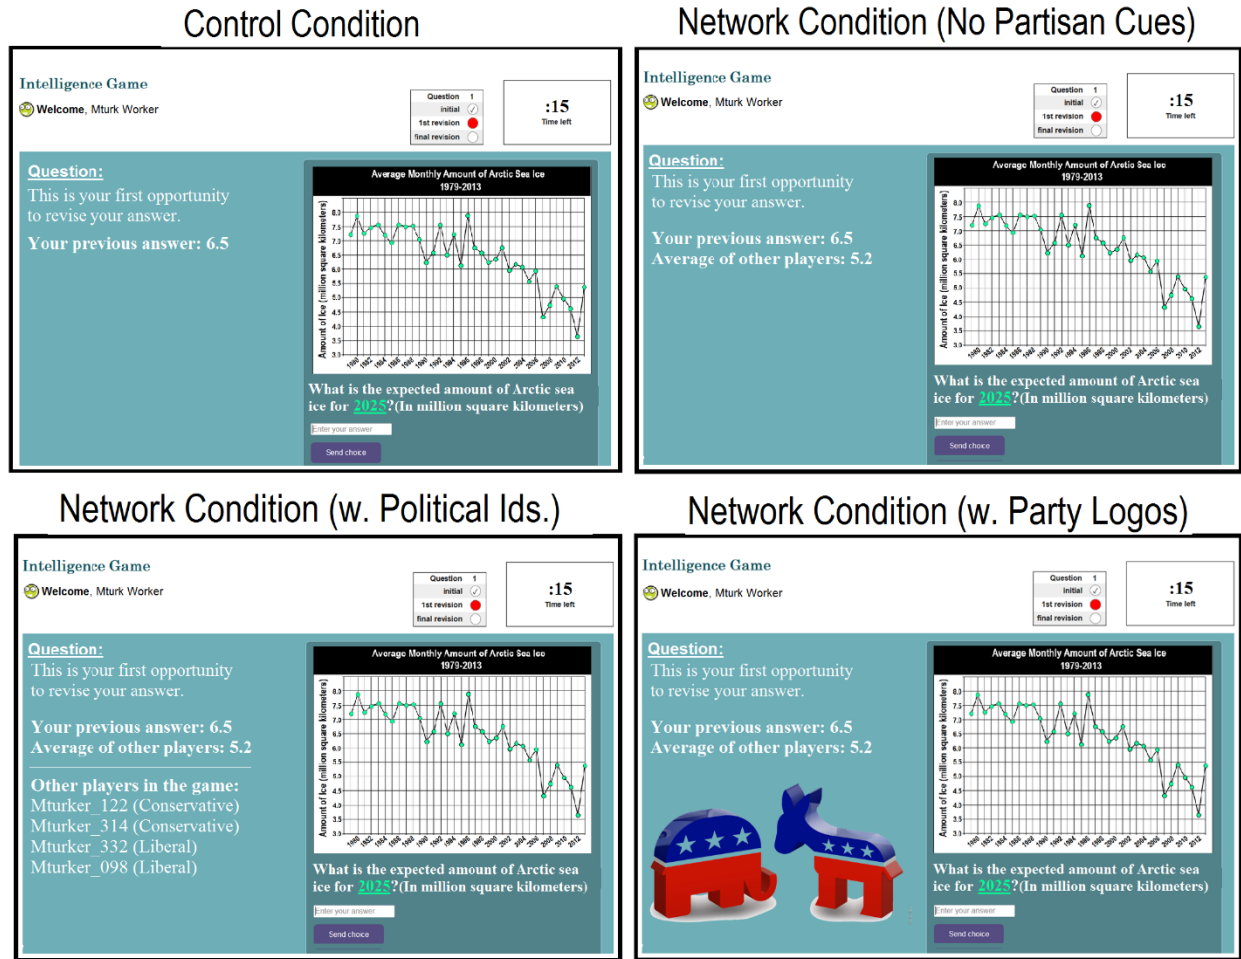


Fig. S2. Screenshots of subject experience, by condition. w., with; Ids., identities.

Subjects had one minute for each response, and the experiment lasted for five minutes. Subjects were paid based on their error as a percent of the true value. The true value of 4.04 was determined using NASA's projections (10,11). Answers which were exactly correct earned the maximum payout (\$2.50). Answers which were within 1% of true value received \$1.25; within 10%, \$1.00; within 15%, \$0.75; within 20%, \$0.35; within 30%, \$0.25; within 90%, \$0.15. Answers more than 90% from the true value did not earn payment. This payment schedule was not observable to participants, who were only told "The more accurate your answers, the more you win!" Subjects were not informed about their accuracy and their payments until after the experiment. Also,

subjects were not told that the true information was determined by NASA, in order to avoid known biases associated with the organizational sources of information (10). Designing the interface in this fashion allowed us to isolate the effects of networked information exchange on the capacity for subjects to accurately interpret trends in the graph.

Subject Recruitment. Our participants were recruited over the Amazon Turk platform to be players in an “Intelligence Game.” We recruited participants based on their self-identified political ideology. We only included subjects who identified as conservative or liberal. Subjects who chose not to identify with one of these categories were not invited to the experiment. Upon arriving at the study website, participants viewed instructions on how to play the intelligence game, and they waited while other subjects arrived. When a sufficient number of subjects arrived, all subjects were randomized to a condition and the trial would begin. Data were collected over a 12-month period, from March 5th, 2017 to March 31st, 2018. 2,400 participants in total were recruited from the World Wide Web to participate in this study. Exactly half of the participants were self-identified as liberals, and the other half were self-identified as conservatives. A quarter of the participants (i.e. 600 subjects) were randomly sorted into each of the main conditions: the conservative control condition, the liberal control condition, the bipartisan network condition without partisan cues, the bipartisan network condition with party logo primes, and the bipartisan network condition with the political identities of network peers revealed.

Statistical Information. Each trial of this study is independent from one another, and each condition within each trial is independent from one another. To compare the trend accuracy of participants across experimental conditions, we first compute the percentage of participants in each

experimental condition within each trial who estimated the correct trend. This approach produces 12 independent, group-level measures of accuracy for each experimental condition. To compare experimental conditions across trials, we use the nonparametric Wilcoxon test. Because our data collection took place over four experimental periods, all comparisons between conditions use the unpaired Wilcoxon rank sum test. All comparisons using the Wilcoxon test are two-tailed to not only test for the hypothesis that social learning increases accuracy, relative to controls, but also for the possibility that it decreases accuracy, as studies on biased assimilation suggest (12–14).

To examine improvements in trend prediction at the individual-level, we used McNemar's exact test of symmetry, which is designed for paired within-subject comparisons concerning changes in binary responses before and after a mediating event. Specifically, we used McNemar's exact test to examine whether the proportion of subjects who changed their trend in the correct direction was symmetrical to the proportion of subjects who changed their trend in the incorrect direction, within the control condition. A significant p-value indicates that, among those subjects who changed their trend predictions within the control condition, significantly more revised their trend in the correct direction. Consistent with our trial-level results, our individual-level analyses showed that improvements to trend accuracy in the control condition were small among both liberals (4.7 percentage points, $n=446$, $P<0.01$, McNemar's exact test) and conservatives (3.1 percentage points, $n=440$, $P=0.08$, McNemar's exact test), where the observed improvements among conservatives failed to reach strong statistical significance.

Social Learning in Point Estimate Predictions. Accuracy can also be measured by the absolute distance between an individual's point estimate prediction and the 'true' amount of Arctic Sea ice predicted by NASA. In this case, collective intelligence processes can be observed in terms of

whether individual and collective error increases as a result of individual learning versus social influence, where the error of point estimates pertains to the absolute distance between an estimate and the ‘true’ answer. The following analysis compares changes in point estimate error, across control groups, bipartisan networks without partisan cues, bipartisan networks with minimal priming via party logos, and bipartisan networks with political identities revealed. To average across point estimates in a representative manner, all estimates that were greater than 100 million kilometers of Arctic Sea ice were removed from this analysis. These estimates are likely the result of misinterpreting the units of the graph, and they make up only 1% of the data. Removing them from the analyses has no effect on the main results for trend accuracy. Fig. S3 shows, for each condition, the round by round changes in the average estimates for conservative participants who estimated the incorrect trend at round 1.

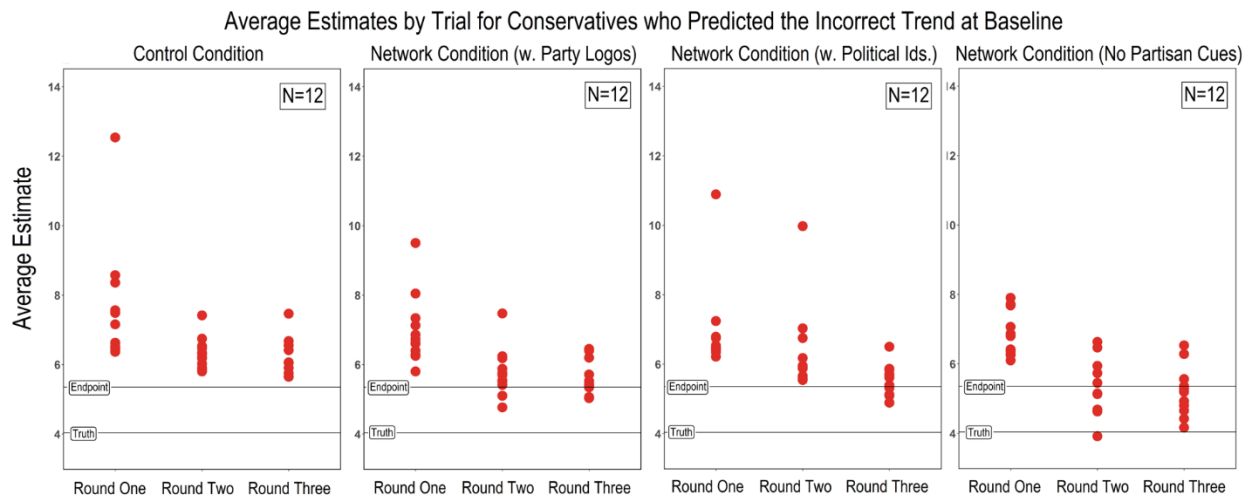


Fig. S3. In the network conditions, conservatives are more likely to revise their answers in the correct direction with a magnitude that correctly changes the direction of their predicted trend. In the bipartisan networks without partisan priming, conservatives are more likely to correctly update the direction of their trend. Note that for this figure, one of the groups is not shown at round 2 for the conservatives in networks without partisan priming because this group had an average estimate of 20.3 due to one subject’s especially high answer. Because this answer was not high at round 1

or 3 it is likely the result of data entry error on the behalf of that subject. Removing this data point allowed clearer visualization of the changes in trends. w., with; Ids., identities.

At baseline, there was no significant difference in the average point estimate error among conservatives across conditions ($n=48$, $P=0.43$, Kruskal-Wallis H test). Average point estimate error significantly decreased for conservatives in the bipartisan networks with priming via party logos (by 0.56 million square kilometers, $n=12$, $P<0.01$, Wilcoxon signed rank test), networks with exposure to the political identities of peers (by 0.57 million square kilometers, $n=12$, $P<0.01$, Wilcoxon signed rank test), and networks without partisan cues (by 0.57 million square kilometers, $n=12$, $P<0.01$, Wilcoxon signed rank test), but not in the control condition (by 0.33 million square kilometers, $n=12$, $P=0.10$, Wilcoxon signed rank test). We find that the average point estimate error decreased significantly more among conservatives in networks without partisan priming than conservatives in control groups ($n=24$, $P=0.05$, Wilcoxon rank sum test). However, the average point estimate error did not decrease more between conservatives in networks with political identities revealed than conservatives in control groups ($n=24$, $P=0.12$, Wilcoxon rank sum test), nor between conservatives in networks with party logo primes than conservatives in control groups ($n=24$, $P=0.12$, Wilcoxon rank sum test).

Fig. S3 displays the average change in response for conservatives who estimated the incorrect trend at round 1, across conditions. We found that at round 1 there was no significant difference in the response distribution between initially biased conservatives across experimental conditions ($n=48$, $P=0.60$, Kruskal-Wallis H test). By round 3, we found no significant difference in the extent to which conservatives lowered their estimates between the networks with party logo priming and the control condition ($n=24$, $P=0.31$, Wilcoxon rank sum test), nor between the networks with identities revealed and the control condition ($n=24$, $P=0.58$, Wilcoxon rank sum

test). By contrast, we found that conservatives in the bipartisan network condition without partisan priming lowered their estimates significantly more than conservatives in the control condition (by 0.45 million square kilometers, $n=24$, $P=0.05$, Wilcoxon rank sum test). These results indicate that social learning at the level of point estimates was particularly present among conservatives in the bipartisan networks without partisan cues, while the improvements observed in the bipartisan networks with partisan cues were indistinguishable from individual learning in the control condition.

Liberals also exhibited significant social learning at the level of point estimates as a result of exposure to opposing interpretations. It is important to note that liberals increased in their average absolute error in only one of the twelve instances in the networks without partisan cues; however, this increase in error (by 0.88 million square kilometers) skews the mean of error for the liberals in this condition, so when needed, the following results also report the median.

The average absolute point estimate error significantly decreased for liberals in the identity marker networks (by 0.58 million square kilometers, $n=12$, $P<0.01$, Wilcoxon signed rank test), the networks with party logos (by 0.47 million square kilometers, $n=12$, $P=0.01$, Wilcoxon signed rank test), and the networks without partisan cues (by 0.26 million square kilometers [mean], by 0.35 million square kilometers [median], $n=12$, $P=0.03$, Wilcoxon signed rank test), but not in the control condition (by 0.18 million square kilometers, $n=12$, $P=0.10$, Wilcoxon signed rank test). As a result, liberals decreased their absolute error more than control groups in the networks with identity markers (by 0.40 million square kilometers [mean], by 0.42 million square kilometers [median], $n=24$, $P<0.01$, Wilcoxon rank sum test) and networks without partisan cues (by 0.08 million square kilometers [mean], by 0.26 million square kilometers [median], $n=24$, $P=0.04$, Wilcoxon rank sum test), but not in the networks with party logo priming (by 0.29 million square

kilometers [mean], by 0.15 million square kilometers [median], $n=24$, $P=0.11$, Wilcoxon rank sum test).

Testing Social Learning in Politically Homogenous Networks. As a robustness test, we ran 5 additional trials where subjects were randomized into one of two experimental conditions: 1) a social network composed solely of conservative participants, and 2) a social network composed solely of liberal participants. These networks used a random decentralized topology of 40 unique people where, identical to the bipartisan network conditions, every node had 4 connections. Participants were randomly assigned to a node within a given network, and they maintained this position throughout the experiment. These robustness trials were motivated by the theory of echo chambers (15), which holds that interaction within politically homogenous groups can amplify the pre-existing biases shared by members of that homogeneous group (16, 17). Echo chambers are especially widespread and problematic in the polarized discussion of climate change (18–20). We tested whether participants in homogenous networks without partisan priming would exhibit improvements in their trend predictions similar to subjects in the bipartisan networks without partisan priming, or whether their baseline bias would remain unchanged.

Using only the responses of participants in the all-conservative or all-liberal networks, we replicated the result in Fig. 1, which shows that conservatives are significantly less likely to interpret NASA's graph correctly. At round 1, 59.2% of conservatives ($n=184$) estimated the correct trend, compared to 74.8% of liberals ($n=183$) ($P<0.01$, Chi-squared test). When combined with the data from main text, we find that altogether 60.6% of conservatives ($n=1,262$) estimated the correct trend, compared to 74.1% of liberals ($n=1,269$) ($P<0.001$, Chi-squared test).

To test whether the politically homogenous networks exhibited social learning, we compared their accuracy throughout the experiment to the accuracy of the same political group in the control condition. To draw these comparisons, we computed the percent of subjects within each group who predicted the correct trend, and then we used the Wilcoxon rank sum test to detect whether group-level measures of trend accuracy in the homogeneous networks were significantly greater than group-level measures of trend accuracy in control groups. When including the homogenous networks, we continue to find that there are no significant baseline differences in trial-level trend accuracy among conservatives across all experimental conditions ($n=53$, $P=0.54$, Kruskal-Wallis H test), nor among liberals across all experimental conditions ($n=53$, $P=0.92$, Kruskal-Wallis H test).

At round 1, the trend accuracy of liberals in the politically homogenous networks was not significantly different than the trend accuracy of liberals in the control condition ($n=17$, $P=0.67$, Wilcoxon rank sum test). However, by round 3, 10.2 percent more liberals (in percentage points) predicted the correct trend in the echo chamber conditions than liberals in control groups ($n=17$, $P<0.01$, Wilcoxon rank sum test). By contrast, the trend accuracy of conservatives in politically homogenous networks was not significantly different than conservatives in control groups at round 1 ($n=17$, $P=0.83$, Wilcoxon rank sum test) or at round 3 ($n=17$, $P=0.26$, Wilcoxon rank sum test). In these robustness trials, liberals showed signs of social learning in politically homogenous networks ($n=5$, $P=0.057$, Wilcoxon signed rank test), but conservatives in politically homogeneous networks did not ($n=5$, $P=0.125$, Wilcoxon signed rank test).

Lastly, we compared members of each political group within homogenous networks and within bipartisan networks, each without partisan priming. We find that liberal subjects in echo chambers and in bipartisan networks without priming were not significantly different in their

capacity to interpret climate data, either at round 1 ($n=17$, $P=0.87$, Wilcoxon rank sum test) or round 3 ($n=17$, $P=0.87$, Wilcoxon rank sum test). However, while conservatives in homogeneous and bipartisan networks without priming were not significantly different in their accuracy at round 1 ($n=17$, $P=0.39$, Wilcoxon rank sum test), we find that conservatives were significantly more accurate in bipartisan networks (88.3%) than conservatives in echo chambers (71.3%) by round 3 (by 17 percentage points, $n=17$, $P=0.02$, Wilcoxon rank sum test).

In these robustness trials, liberal participants, who are systematically more accurate in the task at baseline, appeared to exhibit social learning in both politically homogeneous and bipartisan networks, in the absence of partisan priming. However, while conservative participants exhibited social learning in the bipartisan networks without priming, they did not show significant improvement in the politically homogeneous networks without priming. These robustness trials provide preliminary evidence that communication in conservative echo chambers can fail to eliminate bias, while bipartisan networks without partisan priming can improve accuracy among liberals and conservatives. Altogether, these robustness trials suggest that the political composition of networks can alter the effects of belief exchange on social learning and partisan bias.

Testing the Effects of Bipartisan Identity Signaling when Subjects are Shown the Individual Answers of their Network Peers. We conducted an additional robustness test to examine whether revealing the political ideologies of peers would reduce social learning when subjects were exposed to the individual answers of their network peers. An additional 677 subjects were recruited from Amazon's Mechanical Turk to participate in this study. 289 subjects were self-identified conservatives, and 388 were self-identified liberals. All subjects provided their political ideology and informed consent as part of their registration. Subjects were randomized into one of two

conditions: *i*) a condition showing the four individual answers of bot confederates, as well as their mean, and *ii*) an identical condition that showed a political ideology beside each of the confederate answers shown in condition *i*).

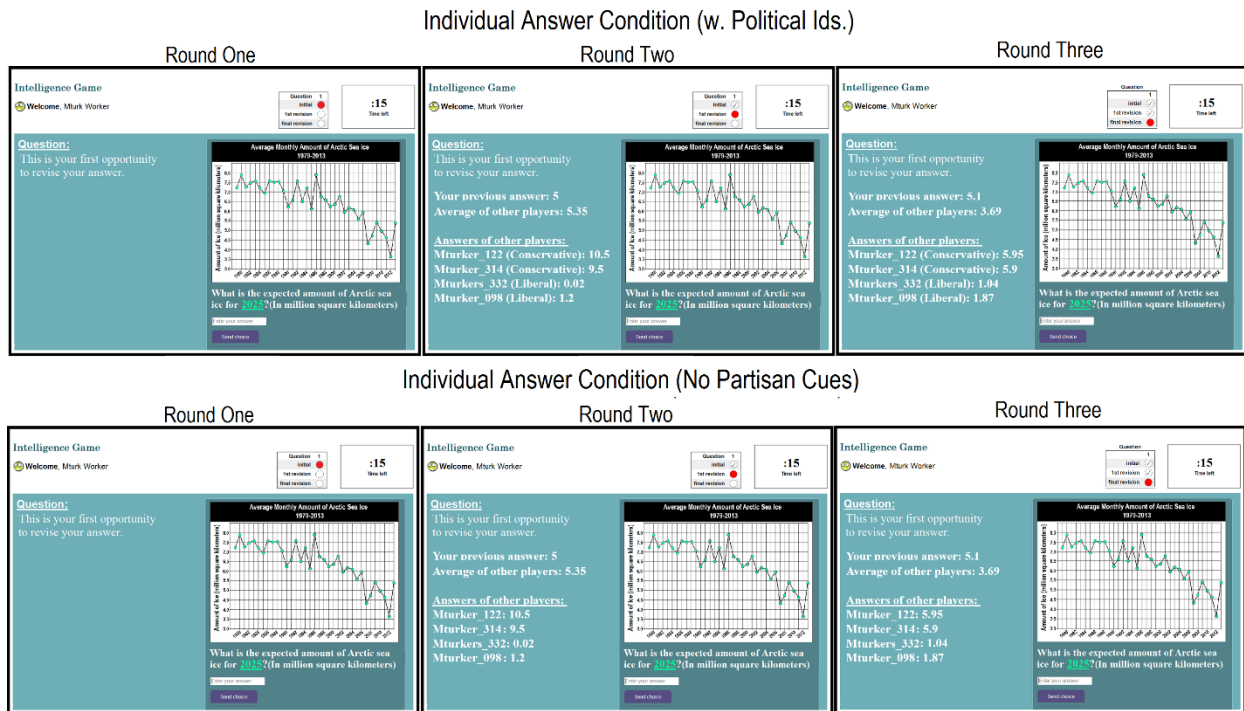


Fig. S4. Screenshot of subject experience in the individual answer condition with and without the political identities of peers revealed. w., with; Ids., identities.

The confederate answers were chosen such that two conservative answers were above the endpoint, and two liberal answers were below the endpoint, where the answers of these political groups were equally distant from the endpoint. The exact same answers were shown in the same order between condition *i*) and condition *ii*). Furthermore, the political identity markers associated with the individual answers were held constant and shown in an identical order in all runs of condition *ii*). The only difference between condition *i*) and condition *ii*) was whether or not the confederate answers were associated with a political ideology congruent with the trend bias

indicated by the answer (see Fig. S4). For this analysis, subjects were removed if they did not successfully provide an answer at both round 1 and round 3.

To choose the confederate answers for the conservative bots, we began by excluding all round 1 responses by conservatives that were either in the correct trend direction or more than 20 million square kilometers away from the endpoint. From the remaining data, we selected two responses from the ten most incorrect answers among conservatives (see Fig. S4). For the liberal confederates, we constructed two responses that were as distant from the endpoint as the two responses for the conservative bots, except in the correct trend direction. To generate the revised responses for round 2, we constructed an OLS model that regressed the round 1 responses of players, as well as the average of each player's network peers, on their round 2 responses, controlling for the distance between each player's round 1 response and their neighborhood mean, since this distance has been shown to moderate the magnitude of revision (1). A separate OLS model was generated for both the conservative and liberal subjects. We generated the round 2 responses for the confederate bots by using these regressions to predict their revisions based on their round 1 responses as determined by the above procedure.

Using only the responses of participants in the individual answer conditions with and without identity markers, we replicated the result in Fig. 1, which shows that conservatives are significantly less likely to interpret NASA's graph correctly. At round 1, 51.2% of conservatives ($n=287$) estimated the correct trend, compared to 73.3% of liberals ($n=388$) ($P<0.01$, Chi-squared test).

In support of our hypotheses, a significant proportion of conservatives corrected the direction of their trend predictions in the individual answer condition without partisan cues ($n=149$, $P=0.02$, McNemar's exact test). However, this learning effect was significantly dampened among

conservatives in the individual answer condition with political identity markers ($n=140$, $P=0.21$, McNemar's exact test). Furthermore, as predicted, liberal subjects significantly improved their trend accuracy in the individual answer condition with ($n=188$, $P<0.01$, McNemar's exact test) and without ($n=200$, $P<0.01$, McNemar's exact test) the political identities of peers revealed.

Subject Demographics. After providing their political affiliation, subjects were linked to an additional survey that invited them to share more demographic information about themselves. A total of 964 subjects (379 conservatives; 578 liberals) volunteered information about their gender and age. This information indicated that our sample of conservatives and liberals was well-balanced along these demographic variables. As such, this data adds confidence that the baseline differences in how conservatives and liberals interpreted NASA's graph were due to their political affiliations and not to other demographic traits that correlate with political alignment.

Among conservatives, 49.5% identified their gender as 'Male', 50.5% identified as 'Female', and 0% identified as 'Non-binary'. A similar distribution was observed among Liberals. Among liberals, 48.1% identified as 'Male', 50.8% identified as 'Female', and 1% identified as 'Non-binary'. We also find similarities in the age distribution across political groups. Among conservatives, 5% identified as between 18-24 years old, 32.8% identified as between 25-34 years old, 32.3% identified as between 35-44 years old, 15.2% identified as between 45-54 years old, 8% identified as between 55-64 years old, and 5% identified as above 65 years old. Among liberals, 13.1% identified as between 18-24 years old, 42.6% identified as between 25-34 years old, 25.8% identified as between 35-44 years old, 10.7% identified as between 45-54 years old, 6.3% identified as between 55-64 years old, and 1.1% identified as above 65 years old.

This data is consistent with broad demographic surveys of the Mturk population, which indicate that Mturk workers are likely to be below 50 years of age, without significant differences in the demographic composition of liberal and conservative workers (25). While we did not collect demographic data on education levels and ethnicity, broad demographic surveys of the Mturk population suggest that Mturkers are more educated than the people studied by the Cooperative Congressional Election Survey (21), with a significant proportion possessing an undergraduate degree. The same survey also found that there were no substantial differences in the levels of education of younger and older Mturk workers (21). This survey study indicates that the Mturk population is Caucasian by majority, though this study maintains that Mturk provides better representation of other ethnicities (e.g. self-identified Hispanics) than the Cooperative Congressional Election Survey (25). A recent national survey shows that American's support of climate change policy is fairly equally distributed across ethnicities (22), suggesting that the ethnic imbalance within the Mturk population is unlikely to alter responses to NASA's graph.

Subject Retention. To calculate attrition rates, we counted the number of participants who entered a response at both round 1 and round 3. Among the players who entered the platform, there was a high retention rate, with 95.1% of all subjects entering a response at both the first and final round. There was no significant difference in retention rates across control conditions, with 93.5% of subjects finishing the study in the liberal control condition, and 94.2% of subjects finishing in the conservative control condition ($P = 0.94$, Chi-squared test). In the bipartisan network conditions, 94.9% of subjects completed the task, which was not significantly different from the retention rates in the liberal control condition ($P=0.82$, Chi-squared test) nor in the conservative control condition ($P=0.94$, Chi-squared test). In our main analyses, we only used data from users who successfully

input an answer at both round 1 and 3, so that we could measure how their answers changed as a result of the experimental manipulations.

Ensuring Data Quality. We took several precautions to ensure that subjects did not violate the design of the experiment. Such precautions can be more difficult in online experiments because researchers may have less control over the behavior of subjects than in traditional laboratory settings. We took several steps to ensure that each participant was a unique user. To prevent individuals from participating in the study multiple times, we designed the system so that if a user tried to use a second browser tab to simultaneously participate, the system would produce an error, and only allow one active browser tab to communicate on the same computer. Additionally, we required users to enter their Mturk ID before playing the game, and all payments were sent to participants using these unique IDs through Mturk's API. We chose to recruit through Amazon Turk in part because Amazon provides strong safeguards for ensuring that each registered user is unique and associated with a single account consisting of verified personal information. The interface was simple and was explained with a set of instruction pictures that all users viewed as they waited for the game to start, so there was little reason to believe there was any skill or learning that could occur from having played similar games before in other experiments (1).

Supplementary Discussion

Limitations. In this study, we observed comparable group-level behavioral outcomes in networks with political identity-signaling and networks with minimal partisan priming. Both methods for increasing the salience of partisanship reduced social learning and maintained belief polarization, where these effects were especially strong through exposure to party logos, potentially because of

their visual and cultural salience (7, 8). A limitation of current theories of identity-effects is that much of the data supporting these theories derives from individual-level surveys, where the psychological effects of identity signaling in bipartisan communication networks cannot be directly measured (4, 5, 9, 23). Some studies suggest that identity-effects and partisan priming operate via different cognitive processes (4, 5, 9, 23,24), while others suggest that partisan primes activate implicit bias by reminding subjects of their own political identity (25, 26). The results of our study are compatible with many hypothesized psychological processes and do not arbitrate between competing psychological theories; our results concern how the exchange of information in social networks mediates the collective integration of beliefs. Our findings thus suggest the need for future work that integrates psychological measures into collective intelligence experiments to examine how individual-level cognitive processes are mediated by social network dynamics.

An important limitation of our study not discussed in the main text is that we were unable to identify any long-term impact of our intervention on the climate change beliefs of subjects in this study. Our primary aim was to test whether the structured network dynamics of social learning can reduce bias in people's interpretations of climate data known to evoke motivated reasoning (10,11). While our approach is based on providing a concrete demonstration that a network-based intervention can bolster effective scientific communication, we expect that valuable extensions of our design could be used to address longer term changes in attitudes. We anticipate that future studies will be able to explore the effects of collective intelligence interventions on subjects' long-term beliefs and their resulting impact on behavior change through peer reinforcement (27).

Another notable scope condition for this study is that we used a single validated question due to its established tendency to elicit endpoint bias (10, 11); this provided a reliable set of conditions for observing motivated reasoning. More work is needed to test how partisan cues

influence collective intelligence across different questions and topic domains. Moreover, it will be useful to investigate the effects of different network topologies on social learning in politicized contexts. Our study used a single decentralized network topology that has been shown to reliably enable social learning (1). Keeping this topology constant across experimental conditions allowed us to reduce variance and causally isolate the effects of partisan cues on social learning. Recent work suggests that centralized network topologies are also capable of facilitating social learning, though less reliably (1). This suggests that the effects of belief exchange in bipartisan networks may generalize to other topologies, though further investigation is required.

The results of our study are also constrained by the subject pool characteristics of participants over Mturk. As our sample indicates, Mturk workers are more likely to be young. While surveys also show that Mturkers are more likely to be left-leaning, we controlled for this by blocking each experimental condition into equally sized groups of conservative and liberal Mturkers (21). It may not be possible to treat our sample of liberals and conservatives as fully representative of all liberals and conservatives throughout the US. Yet even with the left-leaning bias of Mturk workers, we detected a substantial baseline difference in partisan bias between conservative and liberal participants in our task. It is possible that if our study included a more representative sample of conservatives and liberals, we may detect an even more salient difference in baseline bias between these groups. Meanwhile, the Mturk user population is rapidly growing, and recent research indicates that Mturk workers are often more attentive than traditional samples of college students, thus constituting a viable study population (28).

Supplementary References

1. Becker J, Brackbill D, Centola D (2017) Network dynamics of social influence in the wisdom of crowds. *Proc Natl Acad Sci USA* 114(26): 5070-5076.
2. McCright A, Dunlap R (2010) Anti-reflexivity: The American conservative movement's success in undermining climate science and policy. *Theory, Culture & Society* 27: 100–133.
3. McCright A, Dunlap R (2011) The politicization of climate change and polarization in the American public's view of global warming, 2001–2010. *The Sociological Quarterly* 52: 155–194.
4. Mason L (2018) Ideologues without issues: The polarizing consequences of ideological identities. *Public Opinion Quarterly* 82(1): 280–301. <https://doi.org/10.1093/poq/nfy005>.
5. Devine C (2014) Ideological social identity: Psychological attachment to ideological in-groups as a political phenomenon and a behavioral influence. *Political Behavior* 37:1–27.
6. Levendusky M (2009) *The partisan sort: How Liberals became Democrats and Conservatives became Republicans*. (University of Chicago Press, Chicago).
7. Iyengar S, Westwood S (2014) Fear and loathing across party lines: New evidence on group polarization. *American Journal of Political Science* 59(3): 690–707.
8. Slothus R, de Vreese C (2010) Political parties, motivated reasoning, and issue framing effects. *The Journal of Politics* 72(3): 630-645.
9. Malka A, Lelkes Y (2010) More than ideology: Conservative–Liberal identity and receptivity to political cues. *Social Justice Research* 23(2–3): 156–88.
10. Hardy B, Jamieson K (2017) Overcoming endpoint bias in climate change communication: The case of Arctic Sea ice trends. *Environmental Communication* 11(2): 205 – 217.
11. Jamieson K, Hardy B (2014) Leveraging scientific credibility about Arctic Sea ice trends in a polarized political environment. *Proc Natl Acad Sci USA* 111(4): 13598–13605.

12. Lord C, Ross L, Lepper M (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.
13. Sunstein C (2009) *Going to extremes* (Oxford University Press, Oxford).
14. Hart S, Nisbet E (2012) Boomerang effects in science communication: How motivated reasoning and identity cues amplify opinion polarization about climate mitigation policies. *Communication Research* 39(6): 701–723.
15. Jamieson K, Cappella J (2008) *Echo chamber: Rush Limbaugh and the conservative media establishment*. (Oxford University Press, Oxford).
16. Myers D, Bishop G (1970) Discussion effects on racial attitudes. *Science* 169(3947): 778-779.
17. Myers D, Bishop G (1971) Enhancement of dominant attitudes in group discussion. *Journal of Personality and Social Psychology* 20(3): 386-391.
18. Farrell J (2015) Politics: Echo chambers and false certainty. *Nature Climate Change* 5: 719-720.
19. Jasny L, Waggle J, Fisher D (2015) An empirical examination of echo chambers in US climate policy networks. *Nature Climate Change* 5: 782-786.
20. Williams H, McMurray J, Kurz T, Lambert H (2015) Network analysis reveals open forums and echo chambers in social media discussions of climate change. *Glob. Environ. Change* 32: 126-138.
21. Huff C, Tingley D (2015) “Who are these people?” Evaluating the demographic characteristics and political preferences of Mturk survey respondents. *Research and Politics* 2(3): 1–12.

22. Leiserowitz A, Akerlof K (2010) Race, ethnicity and public responses to climate change. Yale University and George Mason University. New Haven, CT: Yale Project on Climate Change. Working paper.
23. Iyengar S, Sood G, Lelkes Y (2012) Affect, not ideology: A social identity perspective on polarization. *Public Opinion Quarterly* 76(3): 405-431
24. Otten S, Moskowitz G (2000) Evidence for implicit evaluative in-group bias: Affect-biased spontaneous trait inference in a minimal group paradigm. *Journal of Experimental Social Psychology* 36(1): 77–89.
25. Price V (1989) Social identification and public opinion: Effects of communicating group conflict. *The Public Opinion Quarterly* 53(2): 197–224.
26. Boutyline A, Vaisey S (2017) Belief network analysis: A relational approach to understanding the structure of attitudes. *American Journal of Sociology* 122(5): 1371-1447.
27. Zhang J, Brackbill D, Yang S, Centola D (2015) Efficacy and causal mechanism of an online social media intervention to increase physical activity: Results of a randomized controlled trial. *Preventive Medicine Reports* 2: 651-657.
28. Hauser D, Schwarz N (2016) Attentive Turkers: Mturk participants perform better on online attention checks than do subject pool participants. *Behavioral Research* 48: 400-407.