

The Wisdom of Partisan Crowds

Joshua Becker^{a,b,c,1}, Ethan Porter^d, and Damon Centola^{a,e}

^aAnnenberg School for Communication, University of Pennsylvania; ^bKellogg School of Management, Northwestern University; ^cNorthwestern Institute on Complex Systems, Northwestern University; ^dSchool of Media and Public Affairs, George Washington University; ^eSchool of Engineering, University of Pennsylvania

This manuscript was compiled on April 13, 2019

Theories in favor of deliberative democracy are based on the premise that social information processing can improve group beliefs. While research on the “wisdom of crowds” has found that information exchange can increase belief accuracy on non-controversial factual matters, theories of political polarization imply that groups will become more extreme—and less accurate—when beliefs are motivated by partisan political bias. A primary concern is that partisan biases are associated not only with more extreme beliefs, but also a diminished response to social information. While bipartisan networks containing both Democrats and Republicans are expected to promote accurate belief formation, politically homogeneous networks are expected to amplify partisan bias and reduce belief accuracy. To test whether the wisdom of crowds is robust to partisan bias, we conducted two web-based experiments in which individuals answered factual questions known to elicit partisan bias before and after observing the estimates of peers in a politically homogeneous social network. In contrast to polarization theories, we found that social information exchange in homogeneous networks not only increased accuracy but also reduced polarization. Our results help generalize collective intelligence research to political domains.

collective intelligence | polarization | networks | social influence

A major concern for democratic theorists is that citizens are simply too ignorant of basic political facts to benefit from deliberation (1), yet research on the “wisdom of crowds” (2–4) has found the aggregated beliefs of large groups can be “wise”—i.e., factually accurate—even when group members are individually inaccurate. While these statistical theories offer optimistic support for democratic principles (5, 6), normative theories of deliberative democracy remain challenged by the argument that social influence processes—in contrast with the aggregation of independent survey responses—amplify group biases (7–9).

One argument against deliberative democracy derives from a common premise in wisdom of crowds theory, which states that in order for groups to produce accurate beliefs, individuals within those groups must be statistically independent, such that their errors are uncorrelated and cancel out in aggregate (3, 10, 11). When individuals can influence each other, the dynamics of herding and groupthink are expected to undermine belief accuracy (10, 11), an argument that has raised concerns about the value of deliberative democracy (12). However, experimental research has shown that when individuals in a group can observe the beliefs of other members, information exchange can improve group accuracy even as individuals become more similar (13, 14). This effect can be explained by the observation that individuals who are more accurate revise their answers less in response to social information, thus pulling the mean belief toward the true answer (13, 15).

While such results are promising, political beliefs are shaped by cognitive biases that are not present in the nonpartisan estimation tasks (e.g., distance estimates) that have frequently

been employed in experimental studies of the wisdom of crowds (11, 13, 14). A key finding of political attitude research is that partisan bias can shape not only value statements but also beliefs about facts (16–19). Such biases persist even when survey respondents are offered a financial incentive for their accuracy (17, 20). One explanation for the emergence of partisan bias in factual beliefs is motivated reasoning (21). Motivated reasoning results from the psychological preference for cognitive consistency, which means that people will adjust their beliefs to be consistent with each other (22). This preference can affect political attitudes, such that people will adjust their beliefs about the world to support their preferences for different parties or politicians (18).

Even when inaccurate beliefs are shaped by motivated reasoning and when corrected beliefs would be less supportive of party loyalties, experimental evidence suggests that accuracy can be improved by information exposure (23). In politically heterogeneous networks containing both Democrats and Republicans, social influence has been found to improve belief accuracy and reduce partisan biases (20, 24). However, theories of political polarization suggest that homogeneous networks—containing members of only one political party—will reverse the expected learning effects of social information processing and instead amplify partisan biases (9, 25, 26).

The risk of homogeneous networks derives from the expectation that response to social information on partisan topics is correlated with belief extremity, rather than belief accuracy (25, 26). However, previous research on political polarization (9, 16, 26) has been concerned primarily with attitude differences, and has not directly examined the effect of social

Significance Statement

Normative theories of deliberative democracy are based on the premise that social information processing can improve group beliefs. Research on the “wisdom of crowds” has found that information exchange can increase belief accuracy in many cases, but theories of political polarization imply that groups will become more extreme—and less accurate—when beliefs are motivated by partisan political bias. While this risk is not expected to emerge in politically heterogeneous networks, homogeneous social networks are expected to amplify partisan bias when people communicate only with members of their own political party. However, we find that the wisdom of crowds is robust to partisan bias. Social influence not only increases accuracy but decreases polarization without between-group network ties.

J.B. and E.P. designed the experiment, analyzed the data, and wrote the paper. J.B. collected the data. J.B. and D.C. constructed the data collection tool for Experiment 1. J.B. constructed the data collection tool for Experiment 2. All authors commented on and approved the final manuscript.

The authors declare no conflict of interest

¹To whom correspondence should be addressed. E-mail: joshua.aaron.becker@gmail.com

61 influence on belief accuracy. To understand the potential effects of partisan bias on the wisdom of crowds, we first study
62 a formal model of belief formation to generate hypotheses
63 relating polarization theories to political belief accuracy. This
64 model is formally identical to that used in previous research on
65 the wisdom of crowds (13, 27), but parameterized to account
66 for a possible correlation between belief extremity and adjustment
67 to social information. Echoing previous experimental
68 findings (20), this model shows that opposing biases cancel
69 out in politically diverse bipartisan networks, leaving the
70 average belief unchanged even when bias is correlated with
71 response to social information. However, in politically homogeneous
72 "echo chamber" networks, a correlation between bias
73 and adjustment causes group beliefs to become more extreme
74 and less accurate (Fig. S4), consistent with political theories of
75 polarization (26) (see *SI Appendix* for detailed model results).
76
77 To test whether the wisdom of crowds is robust to partisan
78 bias, we conducted two web-based experiments examining
79 social influence in homogeneous social networks. Contrary
80 to predictions based on the "law of group polarization" (26)
81 we find that homogeneous social networks are not sufficient
82 to amplify partisan biases. Instead, we find that beliefs become
83 more accurate and less polarized. These results suggest
84 that prior models of the wisdom of crowds generalize to factual
85 belief formation on partisan political topics in politically
86 homogeneous networks.

87 1. Experimental Design

88 Following a pre-registered experimental design, our first experiment
89 asked subjects recruited from Amazon Mechanical Turk to answer
90 four fact-based questions (e.g., "What was the unemployment rate
91 in the last month of Barack Obama's presidential administration?").
92 Subjects were compensated for their participation according to the
93 accuracy of their final responses. The four questions used in this
94 experiment (*Materials and Methods*) were selected because they
95 showed the greatest levels of partisan bias among 25 pre-tested
96 questions.

97 Subjects were randomly assigned to either a social condition
98 or a control condition. For each question, subjects first
99 provided an independent answer ("Round 1"). In the social
100 condition, subjects were then shown the average belief of four
101 other subjects connected to them in a social network, and were
102 prompted to provide a second, revised answer ("Round 2").
103 Subjects in the social condition were then shown the average
104 revised answer of their network neighbors and were prompted
105 to provide a third and final answer ("Round 3"). In the control
106 condition, subjects were prompted to provide their answer
107 three times, but with no social information. Besides the absence
108 or presence of social information, subject experience was
109 identical in both social and control conditions. Subjects in
110 both conditions were provided 60 seconds to provide their
111 answer each round, for a total of 3 minutes per question. As
112 soon as subjects provided their response, they were advanced
113 to the next round, even if there was time remaining.

114 Each trial contained 35 subjects. For each trial in the social
115 condition, all subjects participated simultaneously. Subjects
116 in the social condition were connected to each other in random
117 networks in which each subject observed the average response
118 of four other subjects and was observed by those same four
119 subjects, forming a single connected network of 35 subjects. To
120 test whether the wisdom of crowds is robust to partisan bias

121 in politically homogeneous networks, each trial in each condition
122 consisted of either only Republicans or only Democrats. Subjects
123 in the social condition interacted anonymously and were not
124 informed that they were observing the responses by people who
125 shared their partisan preferences.

126 We controlled for question order effects by using four question
127 sets, each of which were identical except for the order in which
128 questions were presented (see *SI Appendix*). For each question
129 set, we collected data for 3 networked groups and 1 control group
130 for each political party (i.e., 4 independent groups for each party).
131 In total, we collected data for 12 networks and 4 control groups
132 for each party (1,120 subjects in total). Figure S1 (*SI Appendix*)
133 illustrates our experimental design.

134
135 The experimental questions have true answers with values ranging
136 from 4.9 to 224,600,000. In order to compare across questions,
137 we follow similar studies (11) and log-transform all responses
138 and true values prior to analysis using the natural logarithm.
139 This allows for comparison across conditions because $\log(A) - \log(B)$
140 approximates percent difference, and thus calculated errors for
141 each response are approximately equal to percent error. This also
142 accounts for the observation that estimates of this type are
143 frequently distributed log-normally (11, 28). We find that
144 alternative normalization procedures produce comparable results
145 (*SI Appendix*).

146 Because responses by individuals within a social network are
147 not independent, we measure all outcomes at the trial level.
148 To produce this metric, we first calculate the mean (logged) belief
149 of the 35 responses given for a single round of a single question
150 in a single trial. We then measure group error for each round of
151 each question as the absolute value of the arithmetic difference
152 between the mean (logged) belief and the (logged) true value.
153 We then measure the change in error for each question of each
154 trial as the arithmetic difference between the error of the mean
155 at Round 1 and the error of the mean at Round 3. This method
156 produces four measurements of change in error for each trial, i.e.
157 one for each question. We then calculate the average of this
158 value over all four questions completed by each trial to measure
159 average change in error for each trial. We thus produce 24
160 independent observations of the effect of social influence on group
161 accuracy when beliefs are motivated by partisan bias, including
162 12 independent observations of Republican networks and 12
163 independent observations of Democrat networks. In addition, we
164 produce 8 independent control observations, including 4
165 independent observations of Republican control groups and 4
166 independent observations of Democrat control groups.

167
168 We replicated this entire design in a second experiment, with
169 modifications intended to increase the effect of partisan bias
170 on responses to social information. We describe this replication
171 below.

172 Results (Experiment 1)

173 We find no evidence that social influence in homogeneous
174 networks either reduces accuracy or increases polarization on
175 factual beliefs. Instead, we find that social influence increased
176 accuracy for both Republicans and Democrats and also decreased
177 polarization despite the absence of between-group ties. We
178 begin our analysis by confirming that in Experiment 1, subjects'
179 independent beliefs demonstrated partisan bias, as expected
180 based on previous research (5, 17, 20). In Round

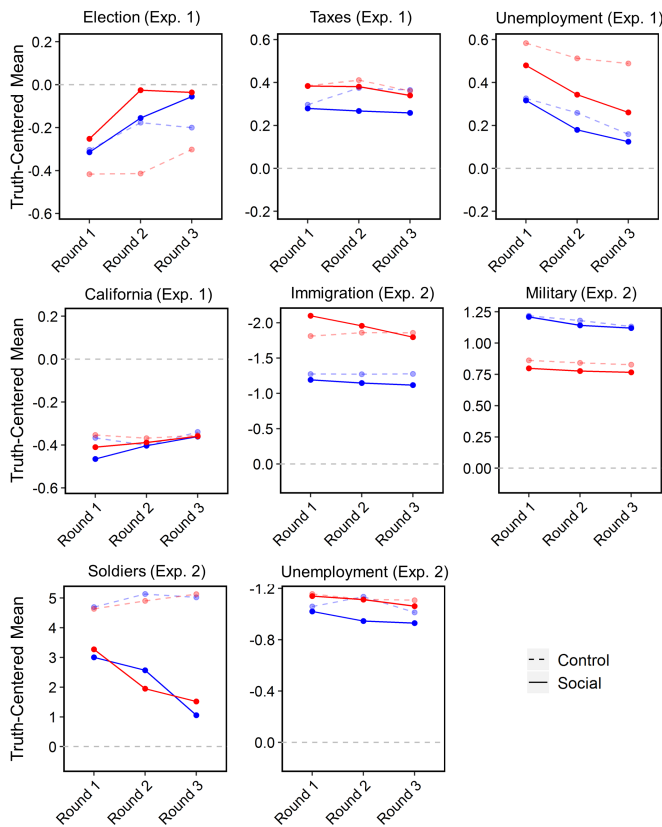


Fig. 1. Normalized, truth-centered mean at each round, averaged across 12 social trials (solid line) and 4 control trials (dashed line). Control groups show more random variation than social groups due to the smaller sample size. Each panel shows one question. Red indicates responses by Republicans, and blue indicates responses by Democrats. For questions with a negative true answer (Immigration, Unemployment) the normalization process in Experiment 2 reverses the sign, and the y-axis is inverted to show relative under- and over-estimates (e.g., subjects overestimated immigration.)

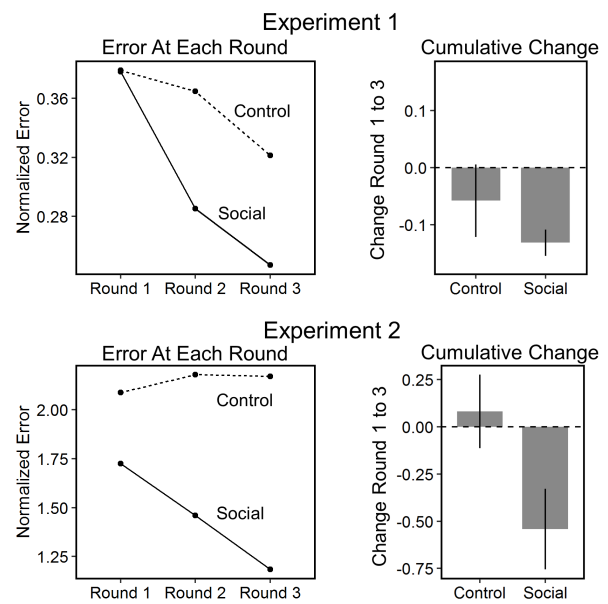


Fig. 2. LEFT: Normalized error of the mean, averaged across 24 social conditions (solid line) and 8 control conditions (dashed line) at each round of the experiment. RIGHT: Cumulative change in error from Round 1 to Round 3. Error bars display standard 95% confidence interval around the mean.

we find that error at Round 3 was significantly lower than error at Round 1 for every one of the 12 Republican trials ($P < 0.001$) as well as every one of the 12 Democrat trials ($P < 0.001$) in Experiment 1. Across both Republicans and Democrats, we find that the average error of the mean decreased by 35% from Round 1 to Round 3.

One possibility is that improvement in the social condition is due to the opportunity for subjects to revise their answers. To test whether this is the case, we compared improvement in the social condition with improvement in the control condition. Following the procedure described above, we calculate the average change in error for the 24 social network trials and the 8 control trials, shown in Figure 2. We find that error did decrease slightly in the control condition ($P < 0.15$), but that the change in the social condition was significantly greater than the control condition ($P < 0.03$), indicating that the reduction in error in homogeneous social networks cannot be explained by individual learning effects. The error of the mean in control groups decreased by only 15%, a substantially smaller change than the 35% decrease in social networks. Thus while providing individuals the opportunity to revise their answer may improve belief accuracy, these results suggest that social information processing—even in homogeneous partisan groups—can help counteract the effects of partisan bias.

Another possibility is that individuals became less accurate even as the group mean became more accurate, which would occur if individual beliefs become more widely dispersed—e.g., if moderates and extremists moved in opposite directions. To investigate this possibility, we first measure the standard deviation of responses by each of the 24 networked groups in Experiment 1 before and after information exchange, averaging across all four questions. We find that standard deviation decreased significantly from Round 1 to Round 3 in social networks ($P < 0.001$) but did not significantly change for control groups ($P = 0.25$). We find that the change in networks was

181 1 (before social influence), responses provided by Democrats
 182 were significantly different from responses provided by Repub-
 183 licans for all questions (See Fig. 1 and *SI Appendix*; $P < 0.001$
 184 for all questions except race in California, for which $P < 0.05$).

185 **Effect of Social Influence on Belief Accuracy.** To illustrate the
 186 change in beliefs for each question, Figure 1 shows the truth-
 187 centered mean of normalized beliefs (so that a negative value
 188 indicates an underestimate, and a positive value indicates
 189 an overestimate) in social conditions at each round of both
 190 experiments. The value for each data point is obtained by
 191 calculating the arithmetic difference between the mean belief
 192 and the true value at each round for each question, and then
 193 averaging this value across all 12 social network trials for each
 194 political party. In every case, the average estimate became
 195 closer to the true value after social influence.

196 To test whether this change could be explained by random
 197 fluctuation, we calculate the error for each round of each
 198 question as the absolute value of the truth-centered mean
 199 (i.e., the absolute distance from truth). We then calculate
 200 the change in error from Round 1 to Round 3, and average
 201 this value across all 4 questions to measure average change
 202 in absolute error within each trial. This analysis determines
 203 whether, on average, the group mean became closer to the true
 204 value after social influence. For those in the social condition,

240 significantly greater than change in control groups ($P < 0.001$),
241 suggesting that information exchange in homogeneous social
242 networks leads to increased similarity among group members.

243 We also directly test the effect of social influence on average
244 individual error (as opposed to the error of the average).
245 This quantity is measured by first averaging error across all
246 individuals within a group for a given question, then averaging
247 across all questions in a trial, and then averaging across all
248 24 social network trials. For Experiment 1, we find that average
249 individual error decreased in social networks ($P < 0.001$).
250 While individual error also decreased slightly in control groups
251 ($P < 0.11$), the improvement was significantly smaller in control
252 groups than social networks ($P < 0.001$), with a 7% decrease in
253 the average error of isolated individuals as compared with a
254 33% decrease in error by individuals in social networks.

255 Robustness to Partisan Priming (Experiment 2)

256 One possibility is that Experiment 1 did not fully capture
257 the effects of partisan bias. A notable observation is that
258 estimation bias—the tendency to under- or overestimate—was
259 in the same direction for both Republicans and Democrats.
260 However, nearly all of the 25 pilot questions generated bias
261 in the same direction. We also find this pattern in previous
262 research on partisan factual beliefs (17), suggesting that same-
263 direction bias is a common feature of partisan beliefs. While
264 this same-direction bias runs counter to intuitive expectations
265 about partisan polarization, it is consistent with previous
266 research on estimation bias, which shows that people have
267 a general tendency to under- or overestimate for any given
268 question (28). The belief differences between Democrats and
269 Republicans can be understood as an additional partisan bias
270 added on top of a general estimation bias.

271 Nonetheless, a limitation of Experiment 1 is that ques-
272 tions were chosen based on the numeric magnitude of bias
273 in pre-testing, and not the controversial nature of the ques-
274 tions. Moreover, the experimental interface was politically
275 neutral and did not communicate to subjects in the social
276 condition that they were in homogeneous partisan networks,
277 factors which may have prevented subjects from perceiving the
278 questions as partisan in nature. We therefore replicated our
279 initial experiment with several changes designed to increase
280 the effect of partisan bias on response to social information.

281 **Replication Methods.** Instead of choosing questions based on
282 numeric polarization in pre-testing, we selected questions based
283 on their connection to controversial policy topics. For example,
284 we asked participants about the number of illegal immigrants
285 in the U.S. at a time when illegal immigration was at the
286 center of national debate (when disagreement over “the wall”
287 with Mexico led to a U.S. government shutdown in January
288 2019). We also framed questions to emphasize change (i.e. we
289 requested numeric estimates for the magnitude and direction of
290 change) to allow for more partisan expressiveness. We re-used
291 one question from Experiment 1, asking about unemployment,
292 because that question taps into a strong policy controversy
293 (the economy) and showed the greatest partisan bias in the
294 first experiment. By re-using this question with an emphasis
295 on directional change, we expected to observe demonstration
296 of a split-direction partisan bias. Exact wording of all four
297 questions is provided in *Materials and Methods*.

298 In addition to selecting more controversial questions, we also

299 modified the experimental interface to include partisan primes
300 that have been shown in prior research (20) to enhance the
301 effects of partisan bias on social information processing. First,
302 we required all subjects to confirm their political party prior
303 to entering the experimental interface, to prime them to the
304 political nature of the study. Second, we included an image of
305 an elephant and a donkey (i.e., symbols for the Democratic and
306 Republican parties) on the experimental interface (see Fig. S3
307 in the *SI Appendix*). Third, for subjects in the social condition,
308 we indicated the party membership of other subjects in the
309 study when providing social information. Finally, subjects
310 upon recruitment were invited to participate in the “Politics
311 Challenge,” and the URL to the web platform included the
312 phrase “Politics Challenge.”

313 Questions in this second experiment allowed negative
314 answers, for which the logarithm is not defined, and so we normal-
315 ize results by dividing by the true answer, which also represents
316 percent difference. However, this method leaves our analysis
317 extremely sensitive to large values as might occur through
318 typographic error. While these extreme values do not change
319 our statistical analysis, the inclusion of all responses yields
320 implausible effect sizes. (For example, we find that error in
321 the social condition decreased by $3.6 \times 10^7\%$ while error in the
322 control groups increased by $5.3 \times 10^4\%$.) We therefore present
323 results in the main text and figures after manually removing
324 extremely large values, a process which impacts fewer than 1%
325 of responses. An analysis that includes all submitted responses
326 is provided in the *SI Appendix*.

327 **Replication Results.** As with Experiment 1, we begin our repli-
328 cation analysis by ensuring that subjects showed partisan
329 bias, finding significant differences between Republicans and
330 Democrats for all four questions ($P < 0.001$). For the question
331 on unemployment, which was re-used from Experiment 1 and
332 reframed to emphasize change, we now observe a meaningful
333 split between the two parties: a majority (54%) of Democrats
334 stated that unemployment decreased under Obama, while a
335 majority (67%) of Republicans stated the opposite. Nonethe-
336 less, the overall numeric bias was still in the same direction:
337 the mean answer for both parties was an overestimate. As
338 this example shows, divergent beliefs between Democrats and
339 Republicans can nonetheless generate numeric estimation bias
340 in the same direction.

341 Figures 1 and 2 show outcomes of the replication. We again
342 find that social influence increased the accuracy of mean beliefs
343 for both Democrats ($P < 0.03$) and Republicans ($P < 0.001$).
344 Across all trials, we found that the error of the mean decreased
345 by 31% for subjects in the social condition, approximately the
346 same effect size observed in Experiment 1. In contrast, we
347 saw a 4% increase in error for the control condition, though
348 this change was not statistically significant ($P > 0.46$). The two
349 conditions were significantly different ($P < 0.002$), indicating
350 that the benefits of social information cannot be explained by
351 individual learning effects.

352 Similar to Experiment 1, we found that standard deviation
353 decreased significantly in the social condition ($P < 0.001$), but
354 increased slightly in the control condition ($P > 0.19$) and the two
355 conditions were significantly different ($P < 0.001$). This result
356 shows that subjects became more similar over time as a result
357 of social information, indicating that social learning effects are
358 robust to explicit partisan primes. In addition to learning at
359 the group level, we found a 34% decrease in individual error for

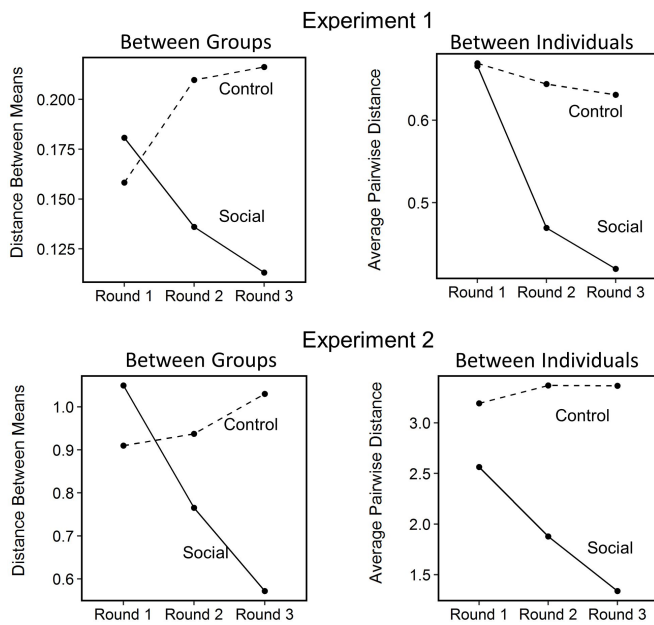


Fig. 3. Points indicate polarization at each round of the experiment for both social networks (solid line) and control groups (dashed line). LEFT: Difference in the normalized mean belief of Democrats and the normalized mean belief of Republicans. RIGHT: Average pairwise distance of normalized responses, which measures the expected difference between a randomly selected Democrat and Republican.

($P < 0.01$ for Exp. 1, $P < 0.08$ for Exp. 2).

As a second measure of polarization, Figure 3 (right) shows the average pairwise distance between individual Republicans and Democrats. This metric measures the average distance between every possible 2-person cross-party pairing, and reflects the expected distance between the belief of a randomly selected Democrat and a randomly selected Republican. This outcome can be understood as reflecting the expected distance in belief between a Democrat and a Republican who could meet by chance in a public forum. For this metric, we found that Democrats and Republicans embedded in homogeneous social networks became more similar in all 24 trials across both experiments, with a 37% decrease in average pairwise distance for Experiment 1 ($P < 0.001$) and a 48% decrease for Experiment 2. Outcomes for control groups show that this value did not change reliably in the absence of social information, showing a nominal decrease in Experiment 1 (6% change, $P > 0.12$) but a nominal increase in Experiment 2 (5% change, $P = 0.25$). Overall, decrease in average pairwise distance was significantly greater in social networks than in control groups ($P < 0.01$ for each experiment).

Discussion

We observed that the mean response to objective, fact-based questions became more accurate as a result of social influence, despite the fact that beliefs were shaped by partisan bias and individuals were embedded in politically homogeneous social networks. In contrast to theories of polarization (26), our results are consistent with the explanation that accurate individuals exert the greatest influence on factual political beliefs as predicted by prior research on the wisdom of crowds (13). In the context of growing concerns about the effects of partisan echo chambers, our results suggest that deliberative democracy may be possible even in politically segregated social networks. Homogeneous social networks, such as those we study, are not on their own sufficient to increase partisan political polarization.

This finding, however, presents a tension: information exchange can mitigate partisan bias, yet public opinion remains polarized. Although we observe decreased polarization and increased accuracy, some error remains as well as some differences between political parties. Polarization can exist despite the potential for social learning. The co-existence of polarization and social learning may be due to structural factors such as network centralization (i.e., the presence of disproportionately central individuals), which can generate and sustain belief polarization in social networks. Network centralization in general has been found to undermine the wisdom of crowds (13); and the ability to obtain central positions in social networks (e.g., through broadcast media or web-based platforms) could allow extremists to exert disproportional influence on group beliefs. In simulation (SI Appendix) we find that a correlation between belief extremity and social network centrality can cause the wisdom of crowds to fail, such that social influence simply enhances existing partisan bias, as predicted by the law of group polarization.

In considering the limitations of our study, it is important to address the generalizability of our research. One concern is that our subject population is not a nationally representative sample; Amazon Mechanical Turk (MTurk) attracts subjects who are younger and more digitally sophisticated than the gen-

360 subjects in the social conditions ($P < 0.001$) and a nominal 3%
 361 increase in individual error for control subjects ($P > 0.74$). The
 362 two conditions were significantly different ($P < 0.001$), showing
 363 that social learning is robust to partisan priming for both
 364 group-level improvement and individual improvement.

Polarization and the Wisdom of Crowds

365
 366 Results from both experiments show that the wisdom of crowds
 367 in networks is robust to political partisan bias. We find that an
 368 increase of in-group belief similarity generates improvements
 369 at both the group level and the individual level. One risk, how-
 370 ever, is that this increase of in-group similarity is accompanied
 371 by a decrease in between-group similarity, generating increased
 372 belief polarization even as groups become more accurate. To
 373 measure belief polarization, we conduct a paired analysis for
 374 each experiment matching the 12 Republican networks with
 375 the 12 Democrat networks (based on trial number, as per our
 376 pre-registered analysis) and calculating their similarity at each
 377 Round (see SI Appendix).

378 We measured polarization using two outcomes. Figure
 379 3 shows the average distance (absolute value of the
 380 arithmetic difference, see SI Appendix) between the mean nor-
 381 malized belief for Republicans and the mean normalized belief
 382 for Democrats at each round of the experiment. Among sub-
 383 jects in the social condition, the average distance between the
 384 mean belief of Democrats and the mean belief of Republicans
 385 decreased by 37% for Experiment 1 ($P < 0.01$) and 46% for
 386 Experiment 2 ($P < 0.02$). In contrast, the distance between the
 387 mean Republican and Democrat belief nominally increased
 388 for the control condition in both experiments, though the ef-
 389 fects were not statistically significant ($P < 0.13$ for Exp. 1, and
 390 $P > 0.87$ for Exp. 2). Overall, the change in polarization was
 391 significantly different between the control and social conditions

eral population (29). Subjects in our experiment may thus have relied more effectively on web search, placing less weight on social information, and so our results may be weaker than would be expected in the general population. MTurkers also tend to skew liberal, and so our sample may have underestimated initial polarization. Generally, however, analyses of political research find that research on non-representative samples such as MTurk typically replicate well on nationally representative samples (30), suggesting our experimental results are likely to replicate. A second concern about generalizability is ecological validity, i.e. whether our experiment reflects the dynamics of political belief formation more broadly. We paid subjects for accuracy, which was necessary to discourage subjects from entering nonsense answers, but political attitudes are typically formed without financial incentive. However, prior research on political beliefs has found that subjects can become more accurate even when they are not compensated for accuracy (23), suggesting that financial incentives could impact the effect sizes (17) but not the direction of belief change. Nonetheless, some empirical contexts may produce perverse incentives that drive people away from accuracy, if, for example, people are motivated to be provocative instead of accurate.

Because accuracy incentives appear necessary for the wisdom of crowds to emerge, an important direction for future work is to examine how individual motivations toward accuracy can vary across empirical settings. A single person motivated by controversy would not be likely to disrupt the wisdom of crowds (unless they hold a central network position), but an entire population motivated by controversy might meet the conditions required for the law of group polarization to hold. Under the assumption that some people are not generally motivated toward accuracy, the robustness of our findings to different empirical settings would depend on the proportion of individuals who are motivated to hold accurate beliefs and the proportion of individuals who are motivated to advance controversial views.

The primary goal of this research was to test whether the wisdom of crowds is robust to partisan bias by studying belief formation about controversial topics in politically homogeneous networks. Based on our experimental results, we reject the hypothesis that social information in politically homogeneous networks will always amplify existing biases. Rather, we find that in the networks studied here, information exchange increases belief accuracy and reduces polarization. While the wisdom of crowds may not hold in all possible empirical settings, our results open the question of when—if ever, and in what circumstances—the wisdom of partisan crowds will fail.

Materials and Methods

Subjects provided informed consent prior to entering the experimental interface. Experiment 1 was run on a custom platform and approved by University of Pennsylvania IRB, Experiment 2 was run on the Empirica.ly platform and approved by Northwestern University IRB. See *SI Appendix* for replication data and code.

Questions for Experiment 1: (1) In the 2004 election, individuals gave \$269.8 million to Republican candidate George W. Bush. How much did they give to Democratic candidate John Kerry? (Answer in millions of dollars - e.g., 1 for \$1 million.) (2) According to 2010 estimates, what percentage of people in the state of California identify as Black/African-American, Hispanic, or Asian? (Give a number 0-100) What was the U.S. unemployment rate at the end of Barack Obama's presidential administration - i.e., what percent of people were unemployed in December 2016? (Give a number 0-100)

(4) In 1980, tax revenue was 18.5% of the economy (as a proportion of GDP). What was tax revenue as a percent of the economy in 2010? (Give a number 0 to 100).

Questions for Experiment 2: (1) For every dollar the federal government spent in fiscal year 2016, about how much went to the Department of Defense (US Military)? Answer with a number between 0 and 100. (2) In 2007, it was estimated that 6.9 million unauthorized immigrants from Mexico lived in the United States. How much did this number change by 2016, before President Trump was elected? Enter a positive number if you think it increased, and a negative number if you think it decreased. Express your answer as a percent change. (3) How much did the unemployment rate in the United States change from the beginning to the end of Democratic President Barack Obama's term in office? Enter a positive number if you think it increased, and a negative number if you think it decreased. Express your answer as a percent change. (4) About how many U.S. soldiers were killed in Iraq between the invasion in 2003 and the withdrawal of troops in December 2011?

ACKNOWLEDGMENTS. This project was made possible by the Summer Institute in Computational Social Science (SICSS). We thank Matt Salganik, Chris Bail, participants at SICSS, Michael X. Delli-Carpini, and Steven Klein for helpful comments. We are grateful to Joseph "Mac" Abruzzo for help with recruitment and to Abdullah Almaatouq for support with the Empirica.ly platform.

- Somin I (2010) Deliberative democracy and political ignorance. *Crit Rev* 22(2-3):253–279.
- Pennycook G, Rand DG (2019) Fighting misinformation on social media using crowdsourced judgments of news source quality. *Proc Natl Acad Sci* 116(7):2521–2526.
- Nofer M, Hinz O (2014) Are crowds on the internet wiser than experts? The case of a stock prediction community. *Journal of Business Economics* 84(3):303–338.
- Wolf M, Krause J, Carney PA, Bogart A, Kurvers RHJM (2015) Collective intelligence meets medical decision-making: the collective outperforms the best radiologist. *PLoS ONE* 10(8):e0134269.
- Bartels LM (1996) Uninformed votes: Information effects in presidential elections. *American Journal of Political Science* pp. 194–230.
- Galton F (1907) Vox populi (The wisdom of crowds). *Nature* 75:450–51.
- Chambers S (2003) Deliberative democratic theory. *Annu Rev Polit Sci* 6(1):307–326.
- Moscovici S, Zavalloni M (1969) The group as a polarizer of attitudes. *Journal of Personality and Social Psychology* 12(2):125–135.
- Sunstein CR (2002) The law of group polarization. *J Polit Philos* 10(2):175–195.
- Surowiecki J (2004) *The wisdom of crowds*. (Anchor).
- Lorenz J, Rauhut H, Schweitzer F, Helbing D (2011) How social influence can undermine the wisdom of crowd effect. *Proc Natl Acad Sci* 108(22):9020–9025.
- Solomon M (2006) Groupthink versus the wisdom of crowds: The social epistemology of deliberation and dissent. *The Southern Journal of Philosophy* 44(S1):28–42.
- Becker J, Brackbill D, Centola D (2017) Network Dynamics of Social Influence in the Wisdom of Crowds. *Proceedings of the National Academy of Sciences* 114(26):E5070–E5076.
- Jayles B, et al. (2017) How social information can improve estimation accuracy in human groups. *Proceedings of the National Academy of Sciences* p. 201703695.
- Madirolas G, de Polavieja GG (2015) Improving collective estimations using resistance to social influence. *PLoS computational biology* 11(11):e1004594.
- Bartels LM (2002) Beyond the running tally: Partisan bias in political perceptions. *Political Behavior* 24(2):117–150.
- Bullock JG, Gerber AS, Hill SJ, Huber GA (2015) Partisan Bias in Factual Beliefs about Politics. *Quarterly Journal of Political Science* 10(4):519–578.
- Lodge M, Taber CS (2013) *The rationalizing voter*. (Cambridge University Press).
- Taber CS, Lodge M (2006) Motivated skepticism in the evaluation of political beliefs. *American Journal of Political Science* 50(3):755–769.
- Guilbeault D, Becker J, Centola D (2018) Social learning and partisan bias in the interpretation of climate trends. *Proceedings of the National Academy of Sciences* p. 201722664.
- Kunda Z (1990) The case for motivated reasoning. *Psychological bulletin* 108(3):480.
- Gawronski B (2012) Back to the future of dissonance theory: Cognitive consistency as a core motive. *Social Cognition* 30(6):652–668.
- Wood T, Porter E (2018) The elusive backfire effect: Mass attitudes' steadfast factual adherence. *Political Behavior* pp. 1–29.
- Shi F, Tepiltkiy M, Duede E, Evans JA (2019) The wisdom of polarized crowds. *Nature Human Behaviour* p. 1.
- Schkade D, Sunstein CR, Kahneman D (2000) Deliberating about dollars: The severity shift. *Columbia Law Review* pp. 1139–1175.
- Sunstein CR (2009) *Going to extremes: How like minds unite and divide*. (Oxford University Press).
- DeGroot MH (1974) Reaching a consensus. *Journal of the American Statistical Association* 69(345):118–121.
- Kao AB, et al. (2018) Counteracting estimation bias and social influence to improve the wisdom of crowds. *J R Soc Interface* 15(141).
- Munger K, Luca M, Nagler J, Tucker J (2019) Age Matters: Sampling Strategies for Studying Digital Media Effects.
- Coppock A, Leeper TJ, Mullinix KJ (2018) The Generalizability of Heterogeneous Treatment Effect Estimates Across Samples.