Journal of Experimental Psychology: Applied

Interactive Crowdsourcing to Fact-Check Politicians

Santos Espina Mairal, Florencia Bustos, Guillermo Solovey, and Joaquín Navajas

Online First Publication, August 31, 2023. https://dx.doi.org/10.1037/xap0000492

AMERICAN PSYCHOLOGICAL ASSOCIATION

Journal of Experimental Psychology: Applied

ISSN: 1076-898X

https://doi.org/10.1037/xap0000492

Interactive Crowdsourcing to Fact-Check Politicians

Santos Espina Mairal[1], Florencia Bustos[1], Guillermo Solovey[2, 3], and Joaquín Navajas[1, 3, 4]

[1] Laboratorio de Neurociencia, Universidad Torcuato Di Tella

[2] Instituto de Cálculo, Facultad de Ciencias Exactas y Naturales, Universidad de Buenos Aires–Consejo Nacional de Investigaciones Científicas y Técnicas

[3] Consejo Nacional de Investigaciones Científicas y Técnicas

[4] Escuela de Negocios, Universidad Torcuato Di Tella

The discourse of political leaders often contains false information that can misguide the public. Fact-checking agencies around the world try to reduce the negative influence of politicians by verifying their words. However, these agencies face a problem of scalability and require innovative solutions to deal with their growing amount of work. While the previous studies have shown that crowdsourcing is a promising approach to fact-check news in a scalable manner, it remains unclear whether crowdsourced judgements are useful to verify the speech of politicians. This article fills that gap by studying the effect of social influence on the accuracy of collective judgements about the veracity of political speech. In this work, we performed two experiments (Study 1: $N = 180$; Study 2: $N = 240$) where participants judged the veracity of 20 politically balanced phrases. Then, they were exposed to social information from politically homogeneous or heterogeneous participants. Finally, they provided revised individual judgements. We found that only heterogeneous social influence increased the accuracy of participants compared to a control condition. Overall, our results uncover the effect of social influence on the accuracy of collective judgements about the veracity of political speech and show how interactive crowdsourcing strategies can help fact-checking agencies.

Public Significance Statement

This article studies the effect of social influence on the accuracy of crowdsourcing strategies to fact-check statements made by politicians. We found that only exposing individuals to the judgements of people supporting the opposite political party improves their performance at the individual and collective level.

Keywords: social influence, crowdsourcing, fact-checking, interactive, political speech

Supplemental materials: https://doi.org/10.1037/xap0000492.supp

Santos Espina Mairal https://orcid.org/0000-0002-8724-0618

Guillermo Solovey https://orcid.org/0000-0003-4093-2649

Joaquín Navajas https://orcid.org/0000-0001-8765-037X

This work was supported by the James McDonnell Foundation twenty-first Century Science Initiative in Understanding Human Cognition—Scholar Award (Grant 220020334) and by a Sponsored Research Agreement between Meta and Fundación Universidad Torcuato Di Tella (Grant INB2376941). The authors have no known conflict of interest to disclose.

Santos Espina Mairal played an equal role in conceptualization, data curation, formal analysis, investigation, methodology, software, visualization, writing–original draft, and writing–review and editing. Florencia Bustos played an equal role in conceptualization and investigation. Guillermo Solovey played an equal role in conceptualization, investigation, methodology, project administration, supervision, writing–original draft, and writing–review and editing. Joaquín Navajas played an equal role in conceptualization, funding acquisition, investigation, methodology, project administration, resources, supervision, writing–original draft, and writing–review and editing.

An early version of this article was posted to a preprint archive and is available at https://psyarxiv.com/u9eyw.

The data are available at https://osf.io/ufs5c/.

The experimental materials are available at https://osf.io/ufs5c/.

The preregistered design and analysis plan is accessible at https://aspredicted.org/p6ma7.pdf.

Correspondence concerning this article should be addressed to Santos Espina Mairal, Laboratorio de Neurociencia, Universidad Torcuato Di Tella, Avenida Figueroa Alcorta 7350, Buenos Aires C1428, Argentina. Email: santosespinamairal@gmail.com

The spread of false information is a major concern in recent years, posing tangible risks to public health (Burel et al., 2020; Burki, 2019), democratic life (Frau-Meigs, 2018), and the fight against climate change (van der Linden et al., 2017). In response to this threat, a vast number of fact-checking organizations aimed at increasing the quality of information in public debates have surged around the World. Despite their enormous efforts, these agencies are very overloaded and they cannot keep pace with the amount of false information that they need to process (Burel et al., 2020). For this reason, finding novel solutions to scale and reduce the time-intensive work of fact-checking organizations has become an urgent issue in social science (Lazer et al., 2018).

On the basis of the observation that combining several layperson estimates about factual issues can outperform expert judgements, a phenomenon popularly known as "the wisdom of crowds" (Larrick et al., 2012; Surowiecki, 2004), recent research tested the reliability of crowdsourcing as a potential tool to assist fact checkers (Allen et al., 2021; Pennycook & Rand, 2019a; Resnick et al., 2021). For example, one study found that averaging 16 laypeople ratings about the truthfulness of news articles led to more accurate judgements than the ones made by three qualified journalists (Resnick et al., 2021). Similarly, Allen et al. (2021) observed that a crowd of 26 lay raters predicted expert judgements with substantial accuracy, suggesting that crowdsourcing may become a promising avenue to boost fact-checking scalability.

However, one limitation of the previous research is that it only focused on fake news, while false information can take different forms. For example, a very demanding endeavor of fact-checking agencies is to check politicians for false statements. Far from being a secondary activity, checking the discourse of political leaders is critical, as it is known to directly influence public behavior. Just to give one clear example, research has found that behavioral metrics of social distancing in Brazil severely reduced right after the president inaccurately minimized the mortality of COVID-19 (Ajzenman et al., 2020). The first goal of the present work is to test whether crowdsourcing strategies can be useful to fact-check statements made by politicians.

A main challenge in this goal is that laypeople's judgements might be subject to partisan effects. In the context of fake news, partisanship has been previously shown to be a strong predictor of people's beliefs about false information and their willingness to subsequently share it in social media (Faragó et al., 2019; Osmundsen et al., 2021; Pennycook & Rand, 2021; Pereira et al., 2018; Vegetti & Mancosu, 2020). However, given that this research was tested with fake news specifically, it is still unknown whether and how partisan biases influence the perceived accuracy of claims made by politicians. The previous studies have found that leader statements can be interpreted as party cues which may awaken tribal motives (Levy Yeyati et al., 2020), and that support for political figures can remain unchanged even after their claims have been fact checked as false (Swire-Thompson et al., 2020). Therefore, we hypothesized that statements made by leaders from a supported party (which we here refer to as "concordant" statements) will be more likely to be classified as "true" compared to statements made by leaders from the opposite party (which we call "discordant" statements).

In principle, the abovementioned hypothesis, if confirmed, would set a clear limitation on the applicability of crowdsourcing to fact-check politicians. For this reason, a second goal of the present work is to investigate whether partisan biases can be reduced in social interactive settings and to test whether social influence that involves diverse opinions can increase the accuracy of aggregated judgements. The previous studies examining the impact of social influence on the wisdom of crowds have also provided mixed evidence. While part of the literature has found that it sometimes reduces the diversity of opinions and degrades the accuracy of crowdsourced judgements (Frey & van de Rijt, 2021; Lorenz et al., 2011; Mavrodiev & Schweitzer, 2021), several studies reported remarkably positive effects (Becker et al., 2019; Guilbeault et al., 2018; Jayles et al., 2017; Navajas et al., 2018; Pescetelli et al., 2021). In the political domain, a large-scale observational study on Wikipedia has found evidence that articles edited by "heterogeneous" collaborators who support opposite political parties are of higher quality compared to articles edited by "homogenous" collaborators who share the same political affiliation (Shi et al., 2019). This study suggests that interacting with politically diverse individuals may reduce partisan biases by breaking filter bubbles and promoting informational flow (Cinelli et al., 2021; Rhodes, 2022; Sunstein, 2009). This is consistent with findings relating homogeneous social settings with polarization of political stances, and heterogeneous interaction resulting in individuals becoming more accuracy-based (Klar, 2014). Although previous literature highlights how group diversity and dissent can improve decision making (Schulz-Hardt et al., 2006), other studies have shown that the benefit of heterogeneous social influence is absent in small groups (Pescetelli et al., 2021), and yet another study has found the accuracy of the wisdom of crowds is robust to social interaction within politically homogeneous networks (Becker et al., 2019). Therefore, on the basis of these disparate findings, here, we considered the two possibilities and examined the effect of social influence in both politically homogeneous and heterogeneous settings.

To summarize, in this work, we empirically tested whether crowdsourced judgements about the veracity of claims made by politicians become more accurate after social influence from politically heterogeneous or homogeneous individuals. Given previous studies suggesting that analytical thinking, quantified by scores obtained in the Cognitive Reflection Test (CRT; Frederick, 2005), plays an important role in shaping people's ability to detect misinformation (Pennycook & Rand, 2019b), here, we also sought to evaluate whether the effect of social influence persists after controlling for individual differences in CRT scores. We present results from two experiments where participants either interacted in dyads (Study 1) or received feedback from supporters of the same or opposite political party (Study 2). To anticipate our results, we found that social influence improved crowdsourced judgements to fact-check politicians but only if participants interacted with or received feedback from politically heterogenous, but not homogenous, individuals.

Study 1

Method

Transparency and Openness

All data and codes necessary to reproduce our findings are publicly available at the Open Science Framework and can be accessed at https://osf.io/ufs5c/ (Espina Mairal et al., 2023). This study's design and its analysis were not preregistered. An early version of this article (Espina Mairal et al., 2022) was posted to a preprint archive and is available at https://psyarxiv.com/u9eyw.

Sample size was determined based on a power analysis assuming a medium-sized effect (Cohen's $d = .5$) and an unpaired t test between the experimental and control conditions. The selected sample size yields over 80% power with a two-tailed 5% significance level.

Ethics

The study has been approved by the Ethics Committee of "Centro de Educación Médica e Investigaciones Clínicas" (protocol ID 435—Version 5) and was performed in line with the principles of the Declaration of Helsinki. Participants provided informed consent and were paid a flat fee of 400 Argentine Pesos (roughly 4 United States dollars at the time) for completing the experiment. On top of that, we incentivized accuracy by providing an additional bonus of the same amount to the 10% best-performing participants at the task. These monetary compensations were informed before the experiment.

Participants

This experiment was performed in Argentina, a politically polarized country (Freira et al., 2021; Navajas et al., 2019; Zimmerman et al., 2022) which is largely underrepresented in the study of misinformation. We recruited $N = 180$ Argentinian participants (56% female, aged 26.3 ± 8.4 years old, 22% having completed university education) through student mailing lists at Universidad Torcuato Di Tella and Universidad de Buenos Aires. Potential participants completed an online sign-up form that included several political identity questions. We only recruited participants with a defined political identity, i.e., those who (a) reported positive affect (categorically, yes/no) for Mauricio Macri (former President of Argentina, and the main political leader of the center-right party *Juntos por el Cambio*) and negative affect for Cristina Fernández de Kirchner (current Vice-President, former President, and the main political leader if the center-left party *Frente de Todos*), or vice versa; and simultaneously (b) stated they would vote for the corresponding party in a hypothetical Presidential election taking place the following week. For clarity purposes, within this report we will refer to those participants who identified with Mauricio Macri and *Juntos por el Cambio* as right-wing individuals and those displaying a preference for Cristina Fernandez de Kirchner and the *Frente de Todos* party as left-wing individuals.

Design

Previous to the experiment, we randomly paired participants into dyads and created three experimental conditions. In the "heterogeneous" condition ($n = 30$ dyads), 60 participants (30 left-wing and 30 right-wing individuals) were matched with someone supporting the opposite political party. In the "homogeneous" condition ($n = 30$ dyads), 60 participants (30 left-wing and 30 right-wing individuals) were matched with someone supporting the same political party (15 dyads were composed by two left-wing individuals and 15 dyads by two right-wing individuals). Last, the remaining 60 participants (30 left-wing and 30 right-wing individuals) were assigned to a Control condition in which no dyads were formed. Participants were blind to this assignment throughout the experiment.

Statements

We selected a corpus of 20 claims (Supplemental Table S1) made by Argentinian politicians that had already been classified as either true or false by Chequeado (https://chequeado.com/), the only fact-checking agency in Argentina affiliated with the International Fact-Checking Network (https://ifcncodeofprinciples.poynter.org/). Half of these statements were made by left-wing politicians, and the remaining half by right-wing politicians. Moreover, half were classified as true, and half were classified as false by Chequeado. The corpus consisted of five true left-wing phrases, five false left-wing phrases, five true right-wing phrases, and five false right-wing phrases. This balanced design implied that each participant, upon presentation of the whole corpus, would be exposed to the same number of concordant as well as discordant phrases, and simultaneously to an equal number of true and false phrases. Participants were unaware of this balanced design.

Procedure

Before the experiment, participants first completed a form consisting of a series of demographic items and a CRT test (Frederick, 2005; Primi et al., 2016). The experiment consisted of a three-stage procedure (Figure 1). The procedure is identical in structure to the one implemented in two previous studies looking at the effect of deliberation on the wisdom of crowds (Navajas et al., 2018), the probability of reaching consensus in polarized moral issues (Navajas et al., 2019), and the effect of diversity on herding behavior (Navajas et al., 2022). In Stage 1, participants were instructed to individually classify each one of the 20 statements as true or false, as well as to provide a measure of confidence in their reported answer, in a scale ranging from 1 (*Low Confidence*) to 5 (*High Confidence*). Each phrase was read out by the experimenter and presented for a duration of 10 s. The only information displayed to participants was the statement itself, the name of the politician as well as a brief description of his/her public position, and the approximate date when the statement was made. For example, the information displayed would read: *Alberto Fernández, the president of Argentina, said:* "we have started the largest vaccination campaign in Argentinian history." *He said this in March, 2021. Is what he said true or false?* Supplemental Figure S1 shows how this information was displayed.

After completion of Stage 1, the experimenter showed participants the responses provided by each of them, introducing social influence. Therefore, in Stage 2, participants learned about the responses of another person and were given time to freely discuss their disagreements (1 min per disagreed statement). In the Control condition, we asked participants to privately classify as true or false a set of unrelated general knowledge statements (i.e., comparing city populations, see Supplemental Table S2) lasting approximately the same time as the one taken to complete Stage 2 in the two treatment conditions. Participants in the Control condition did not interact with each other at all. Stage 3 consisted of a re-run of Stage 1, in which participants were informed that they could revise their initial individual answers and confidence ratings. Timing and format were identical to those of Stage 1.

Figure 1
Experimental Procedure

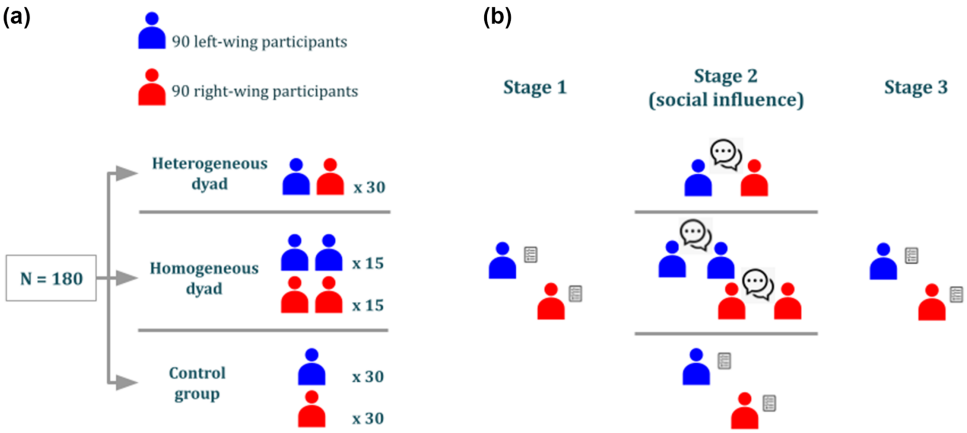

Note. (a) Participants were paired in dyads of similar or opposed partisan affinity, depending on the experimental condition. (b) The experiment had a 3-stage structure: after individually classifying 20 statements as true or false, participants were exposed to social influence, followed by another individual stage in which they could revise their initial responses. See the online article for the color version of this figure.

Data Collection

Participants were invited to attend an online meeting hosted in *Zoom* (https://zoom.us/), where they were assigned to separate breakout rooms with their corresponding pair and one research assistant (experimenter). Microphones were active and cameras turned off for the complete duration of the experiment. Participants were instructed to write down their answers on article and then type them into a preassigned, individual Google Sheet (https://docs.google.com/spreadsheets) file sent to them (Supplemental Figure S2).

This manual upload of individual responses happened after Stage 1 was completed and after Stage 3 completion. The experimenter saved a backup copy of each file once participants finished typing their responses to avoid post hoc editing. The research assistant had access to a Google Sheet fed by both individual Sheets, enabling them to simultaneously show both members of the dyad their responses during Stage 2. A total of 10 experimental sessions were carried out between April and May 2020. Research assistants were blind to condition throughout the experiment.

Data Analysis

Reported statistical analyses were performed in Matlab R2018b. Full specification and details on the mixed-effects logit model reported in the Results section can be found in Supplemental Table S3. Figures were generated with Python Jupyter Notebooks in Google Colab (https://colab.research.google.com/) and Matlab.

For aggregating individual judgements into collective estimates (see Results section), we randomly sampled n individual responses (with replacement) for a given politician statement from each experimental group's pool, where n thus indicates *group size*. A group-level response was constructed based on the simple majority of the individual responses: If the majority of the sampled responses for a given group classified the statement as *true* (*false*), the overall group response was *true* (*false*). This collective estimate could in turn be correct or not, upon comparison with the ground truth (assumed to be the fact-checking agency's classification). Running this procedure for each n across the total number of politician statements (20) yielded a score (the number of correct answers divided by the number of phrases, 20) for each crowd size. We iterated this sequence 1,000 times and calculated the mean score and the standard error of the mean for each n.

Results

Initial Individual Responses

We initially set out to determine whether participants could discriminate between true and false statements made by politicians. For the purpose of this work, we considered a decision to be "correct" if it matched the classification of the fact-checking agency. Prior to social influence, participants classified correctly, on average, approximately 12 out of 20 statements (Figure 2a), leading to a better-than-chance classification performance ($d = 0.87$, $t_{179} = 11.6$, $p < .001$; two-tailed t tests are employed from now, with Cohen's d as measure for effect size). There were no differences in initial accuracy between experimental groups, $\eta^2 = 0.01$, $F(2, 177) = 0.59$, $p = .55$; partial η-squared is employed for effect size in analysis of variance tests from now on, and performance was approximately equal across politically concordant and discordant statements ($d = -0.1$, $t_{358} = -0.92$, $p = .36$). Thus, in our study, we found no evidence of participants being better when classifying concordant information, a pattern that has been documented regarding truth discernment in previous studies focused on fake news (Pennycook & Rand, 2021).

No differences between groups were found regarding CRT scores, $\eta^2 = 0.005$, $F(2, 177) = 0.48$, $p = .62$, and education levels, $\eta^2 = 0.004$, $F(2, 177) = 0.35$, $p = .71$, and no correlation between initial accuracy and CRT performance was found at the individual level ($r = 0.03$, $p = .66$).

In addition, the scope of this investigation involved studying the effect of partisanship on participants' judgements. If partisanship played a role, we would expect that people would be more likely to

Figure 2
Stage 1 Responses

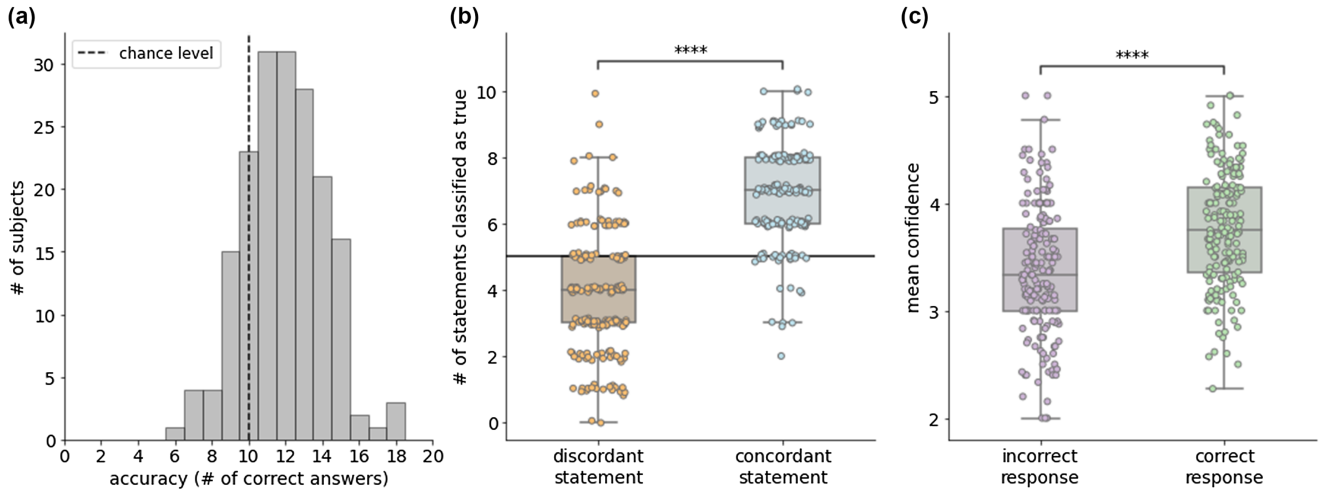

Note. (a) Accuracy of initial individual responses. (b) Number of statements classified as true by partisan concordance. Theoretically unbiased participants would result in no significant differences between both conditions. Each point marks the count of answers by participant for each condition. (c) Mean declared confidence assigned to responses by their accuracy, revealing participants' introspective access. Each point represents mean confidence by participant for each condition. Significance of the corresponding *t* test is depicted at 0.01% (****). See the online article for the color version of this figure.

believe in statements made by politicians of the same political party they support compared to statements made by politicians of the opposite party. Given that each participant was exposed to the same number of concordant and discordant statements, a participant with no partisan bias should classify as "true" the same number of phrases in both conditions. However, in Stage 1, participants tended to classify politically concordant statements as "true" more frequently than discordant ones ($d = 1.22$, $t_{179} = 16.4$, $p < .001$), suggesting that partisanship predicts beliefs about the veracity of statements made by politicians (Figure 2b). Defining *bias* as the difference between the number of "true" answers assigned to concordant and discordant statements, we found that, on average, participants classified an excess of approximately three statements as "true" ($M \pm SD = 2.99 \pm 2.45$) comparing both conditions. (Notice that bias could take a maximum value of 10 in this work). We observed no pretreatment differences in bias between experimental conditions, $\eta^2 = 0.023$, $F(2, 177) = 2.05$, $p = .13$.

To test whether participants had introspective access into their competence as raters (Figure 2c), we examined confidence ratings. Individuals assigned higher confidence values in a 1 to 5 scale to correct responses than to incorrect ones (mean confidence in correct trials, $M \pm SD = 3.77 \pm 0.54$; mean confidence in incorrect trials, $M \pm SD = 3.36 \pm 0.61$; $d = 0.7$, $t_{179} = 11.3$, $p < .001$), revealing they had—to some extent—insight into their performance, without significant pretreatment differences displayed by experimental conditions, $\eta^2 = 0.002$, $F(2, 177) = 0.22$, $p = .81$.

Revised Individual Responses Following Social Influence

The second main goal of this study was to test the effect of social influence on the accuracy of beliefs (Figure 3a). Comparing revised individual answers across different treatments, we found a medium-sized effect between the heterogenous and control conditions ($d = 0.53$, $t_{118} = 2.93$, $p = .004$) and no significant difference between the homogenous and control condition ($d = 0.02$, $t_{118} = 0.09$, $p = .93$).

We then studied the change in accuracy across different conditions. As a benchmark, individuals in the Control condition showed a small, albeit significant, increase in performance ($d = 0.39$, $t_{59} = 3.01$, $p = .004$). We found that participants in the Homogeneous condition did not significantly alter their performance compared to Stage 1 ($d = 0.09$, $t_{59} = 0.71$, $p = .48$), and this change was statistically indistinguishable to the one observed by the individuals lacking social influence ($d = 0.23$, $t_{118} = 1.25$, $p = .21$). By contrast, Heterogeneous dyad members showed an increase in performance compared to Stage 1 ($d = 0.58$, $t_{59} = 4.5$, $p < .001$) which was significantly larger in magnitude to the one observed in participants in the Control condition ($d = 0.44$, $t_{118} = 2.40$, $p = .018$).

Although we did not find evidence for a different improvement in accuracy between members of the Homogeneous and Control conditions, we performed an equivalence test to evaluate if the difference fell within a small interval against the null hypothesis of a significant difference between groups (Lakens et al., 2018). We ran a two one-sided tests procedure for equivalence assuming that a change in one phrase is the minimum practical difference in accuracy. This test rejected the null hypothesis of a meaningful difference between both conditions (95% CI [−0.12, 0.85], $p = .017$).

To test the robustness of our results, we ran a multivariate analysis on the probability to provide a correct answer in Stage 3 with dummy variables coding for the two treatments with social influence (Figure 3b). In this mixed-effects model, we controlled for a series of variables such as age, gender, education, party affiliation, CRT performance, and whether the statement was debated or not, and added random effects at the individual level and at the phrase level. This analysis provided evidence that the positive effect of

Figure 3
Treatment Effect

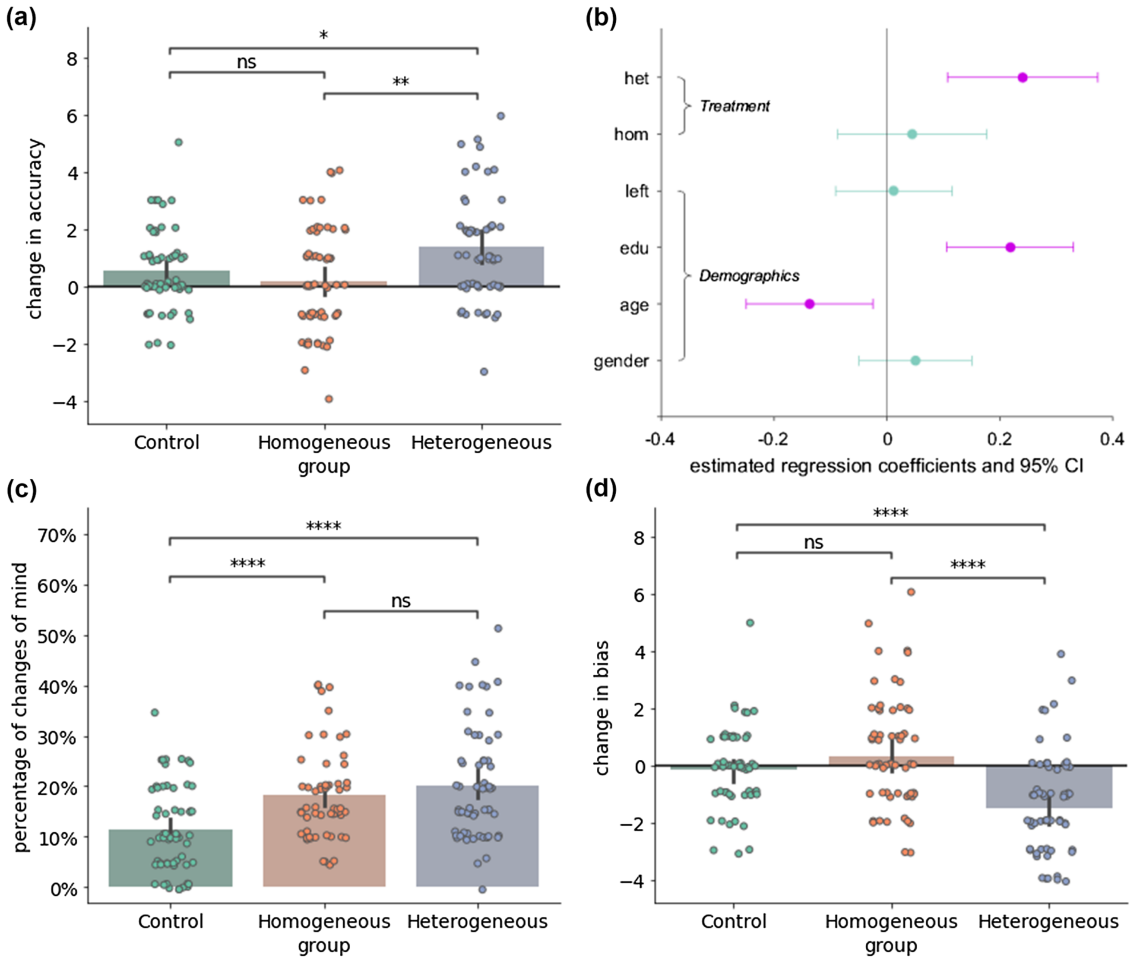

Note. (a) Mean and standard error of the mean of change in the number of individual correct answers between Stage 1 and Stage 3, by group. Outliers are omitted in the plots but included in the analysis. (b) Coefficients and CIs of predictors resulting from running a Generalized Linear Mixed-Effects Logistic regression on the probability of giving a correct answer in Stage 3. Abbreviated explanatory variables on display are: *het* and *hom*, dummies that distinguish heterogeneous dyad members from homogeneous dyad members; *left*, a dummy for left-wing participants; *edu*, the number of education years attained. (c) Mean and standard error of the mean of response revisions between Stage 1 and Stage 3, by group. (d) Mean and standard error of the mean of change in individual bias between Stage 1 and Stage 3, by group. Outliers are omitted in the plots but included in the analysis. Significance of the corresponding *t* test is depicted at 5% (*), 1% (**), 0.01% (****), or nonsignificant (ns). See the online article for the color version of this figure.

heterogeneous social influence is robust under this specification (for full model specification and results, refer to Supplemental Table S3).

To better understand the causes of this increase in accuracy, we then studied whether social influence decreased the partisan bias observed in Stage 1. We found that Heterogeneous and Homogeneous dyad members did not differ in the number of decisions that they revised (Figure 3c, $d = 0.21$, $t_{118} = 1.14$, $p = .26$, comparing both conditions). However, Heterogeneous dyad members tended to revise their answers in opposition to their partisan stance resulting in a significant reduction in bias (Figure 3d, $d = -0.74$, $t_{59} = -5.70$, $p < .001$). Hence, by switching ideologically congruent answers for incongruent ones (i.e., revising against their partisanship), they reduced the magnitude of the initial partisan bias by approximately 50%. Meanwhile, Homogeneous dyad members ($d = 0.15$, $t_{59} = 1.16$, $p = .25$) and individuals in the Control condition ($d = -0.11$, $t_{59} = -0.84$, $p = .40$) did not significantly change their initial bias across stages. Although social influence increased the overall likelihood to revise answers ($d = 0.88$, $t_{118} = 4.84$, $p < .001$ and $d = 0.76$, $t_{118} = 4.20$, $p < .001$ comparing Control participants to Heterogeneous and Homogeneous ones, respectively), heterogeneous communication resulted in a reduction of preexisting partisan bias, while homogeneous interaction failed to do so.

In the results presented so far, we used the percentage of correct answers as the main indicator of performance. However, an increase in accuracy could be due to a better classification between true and false statements (truth discernment) or a change in the general

tendency to believe that statements are true (overall belief). To better understand the roots of the effect of social influence on performance, we analyzed our results under the signal detection theory framework (Batailler et al., 2022; Pennycook & Rand, 2021). We follow the standard calculation of truth discernment as the difference between a participant's belief in true statements (hit rate) and belief in false statements (false alarm rate). Overall belief is defined as the average belief in both true and false statements.

Our main results regarding the effect of social influence hold when examining truth discernment instead of accuracy. We found no initial differences between groups regarding truth discernment, $\eta^2 = 0.006$, $F(2, 177) = 0.56$, $p = .57$, nor overall belief, $\eta^2 = 0.015$, $F(2, 177) = 1.37$, $p = .26$. Posttreatment, members of Heterogeneous dyads had a higher truth discernment compared to Homogeneous dyad members ($d = 0.53$, $t_{118} = 2.89$, $p = .005$) and Control group members ($d = 0.5$, $t_{118} = 2.71$, $p = .008$), whereas no differences in overall belief among groups were detected, $\eta^2 = 0.003$, $F(2, 177) = 0.22$, $p = .80$. This means heterogeneous social influence increased the ability of participants to distinguish between true from false statements without changing their propensity to believe in politicians in general.

Crowd Performance: Aggregating Individual Knowledge

Our main goal involved studying crowdsourcing as a potential tool for fact-checking politicians. By aggregating individual revised responses (see Data Analysis section), we found that harnessing the wisdom of crowds after social influence has an amplifying effect on collective accuracy for all groups as crowd size increases (Figure 4a). In particular, Heterogeneous crowds (formed with Stage 3 answers of individuals in the Heterogeneous condition) outperformed crowds formed by individuals in the Control and Homogeneous conditions. A multivariate linear regression showed that, controlling for crowd size ($\beta = 0.313 \pm 0.037$, $t = 8.36$, $p < .001$), Heterogeneous crowds performed better than Control crowds ($\beta = 0.442 \pm 0.043$, $t = 10.2$, $p < .001$), whereas Homogeneous crowds did significantly worse ($\beta = -0.575 \pm 0.043$, $t = -13.3$, $p < .001$). Overall, we observed that crowds formed by individuals who interacted in the Heterogeneous condition correctly classified an average of 15 out of 20 statements, a performance which is comparable with the one obtained by the previous studies on crowdsourcing and misinformation (Allen et al., 2021).

To further understand the effect of social influence on crowd accuracy, we defined a variable called "crowd score change" as the increase in crowd accuracy from Stage 1 to Stage 3 for each possible group size, and then averaged this variable across all group sizes (Figure 4b). We found that the crowd score change was significantly larger than zero for both the Control ($d = 6.41$, $t_{29} = 35.1$, $p < .001$) and the Heterogeneous conditions ($d = 5.51$, $t_{29} = 30.2$, $p < .001$) and significantly smaller than zero for the Homogeneous condition ($d = -1.73$, $t_{29} = 9.48$, $p < .001$). The crowd score change in the Heterogeneous condition was larger than the one observed for the Control condition ($d = 2.49$, $t_{58} = 9.66$, $p < .001$), suggesting that social influence had a positive effect on crowd accuracy. Instead, crowds in the Homogeneous condition showed a significantly lower crowd score change compared to the Control condition ($d = 8.2$, $t_{58} = 31.8$, $p < .001$).

The fact that Heterogeneous crowd accuracy stabilized slightly above 75% deserves further attention. In principle, this could reflect that crowds consistently misclassified some of the phrases (i.e., that the crowd correctly classified 75% of the phrases across all iterations) or, alternatively, it could mean that the crowdsourcing approach developed here produced inconsistent classifications across all phrases (i.e., that the crowd correctly classified all phrases across 75% of the iterations), or a combination of both. To

Figure 4
Wisdom of Crowds

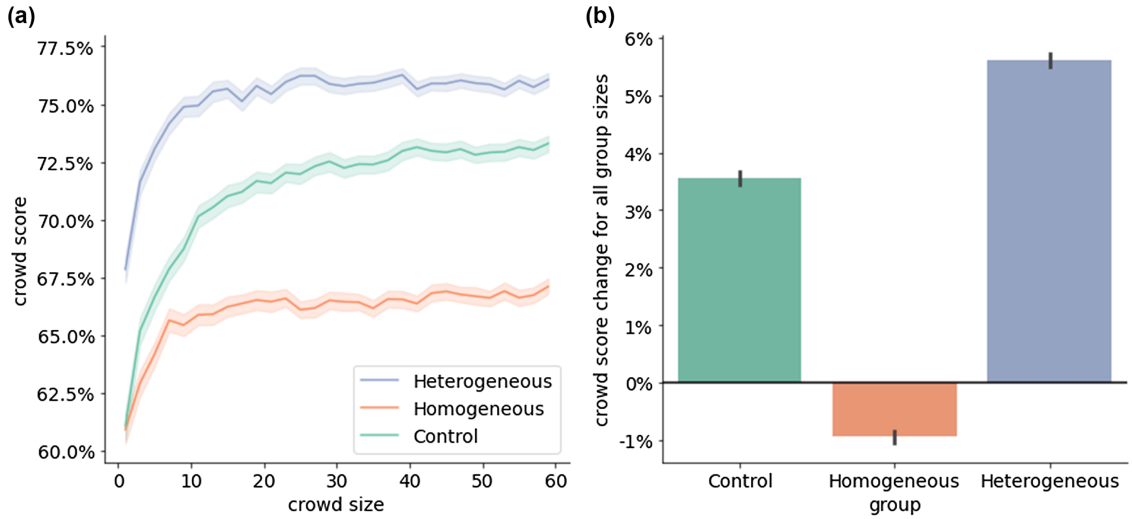

Note. (a) Accuracy of the crowd by group and by crowd size, using Stage 3 responses. Score is defined as the number of correct responses as a % of total statements (20). Mean (solid line) and standard error of the mean (shaded area) across 1,000 iterations. (b) Change in crowd accuracy between Stage 1 (initial answers) and Stage 3 (revised answers). Mean (bars) and standard error of the mean (error bars) across 1,000 iterations and for all group sizes. See the online article for the color version of this figure.

study this issue, we computed the crowd score for each phrase individually (Supplemental Figure S3). Our results provide evidence that most crowd classifications were mostly internally consistent and that in only one phrase there was some ambivalence in the crowd recommendation.

Discussion

In this study, we found individuals can distinguish true from false statements made by politicians with above-chance accuracy. We observed their performance was limited by partisan bias, which could be partially corrected by means of social interaction in dyads of opposing political affinity. In line with this observation, the accuracy of participants at fact-checking politicians increased because of politically heterogeneous communication but failed to do so if social interaction happened in dyads of same partisanship. We also found that crowdsourcing strategies based on aggregating individual fact-checks boosted performance and that the effect of social influence on collective accuracy parallels the one found on individual accuracy.

However, one main limitation of this study points toward the scalability of this tool. Because dyadic interaction may not be cost-effective in crowdsourcing approaches, it remains to be tested if paradigms with minimal social feedback can attain similar results without depending on fact-checkers communicating with one another.

Another limitation in our study concerns the fact that the observed experimental effect of social interaction is conditional on its specific format: Dyadic communication between participants was time-constrained and limited to vocal, unstructured interaction in a virtual platform. Further research is required to test whether and how our results are sensitive to the duration and modality of communication, and how they depend on the number of group members.

In addition, the amount of social influence to which each participant was exposed to in this study was unbalanced across conditions, as Heterogeneous dyads naturally tended to disagree more often than Homogeneous dyads and thus discussed relatively more between themselves. Although we found the number of phrases discussed in dyads was not a predictor of changes in bias or accuracy, balancing exposure to social information would be a key extension of this design to further uncover the mechanisms by which the structure of social influence modulates belief in politician speech.

To replicate our main results in a more scalable format involving only minimal social feedback, we ran a second, preregistered study that also balanced the amount of heterogeneous and homogeneous social influence. This also aimed to extend our findings regarding the effect of heterogeneous social influence on fact-checking ability using a larger sample of participants.

Study 2

Method

Transparency and Openness

All data and codes necessary to reproduce our findings are publicly available at the Open Science Framework and can be accessed at https://osf.io/ufs5c/. This study's hypotheses, design, and analyses were preregistered at https://aspredicted.org/p6ma7.pdf.

To determine sample size, we ran a Monte Carlo power analysis (Zhang, 2014) using data from Study 1. We generated synthetic data sets by sampling initial and revised responses from Study 1, with replacement. For each possible sample size from $N = 30$ to $N = 450$, we conducted 10,000 simulations where we replicated our main analyses. Specifically, we ran three statistical tests. The first test compared the postinteraction accuracy of the Control condition with the Heterogeneous condition. The second test performed the same comparison for the Homogenous condition against the Heterogeneous condition. The last test compared the increase in accuracy for the two conditions with social influence against the increase in accuracy for the Control condition. In all cases, we performed unpaired two-tailed t tests with a significance level of 5%. We estimated power as the fraction of times where we rejected the null hypotheses on each particular test. Based on this analysis, we decided to collect data from $N = 240$ participants (i.e., 80 individuals per condition), yielding a statistical power of 80% for the most restrictive of our tests (Heterogeneous vs. Control), and over 95% power for the other two mentioned tests (Supplemental Figure S4).

Ethics

This study was approved by the Ethics Committee of "Centro de Educación Médica e Investigaciones Clínicas" (protocol ID 435—Version 5) and was performed in line with the principles of the Declaration of Helsinki. Participants provided informed consent and no monetary compensations were given.

Participants

We recruited $N = 240$ Argentinian participants (44% female, aged 35.9 ± 12.5 years old, 37% having completed university education) with a defined political identity through student mailing lists at Universidad Torcuato Di Tella and Universidad de Buenos Aires. The criteria for defining partisanship were identical to those involved in recruitment for Study 1. No participant in Study 2 had previously been part of Study 1. As preregistered, we also asked participants to rate their subjective affect for all politicians mentioned in our corpus of phrases, in a scale ranging from −2 (*dislike*) to 2 (*like*).

Design

Participants were randomly assigned one of three possible online forms to complete: a left-wing form, a right-wing form, or a control form, where the form type combined with the partisanship of the participant determined the experimental group he/she belonged to. Left-wing (right-wing) participants assigned to the left-wing (right-wing) form belonged to the Homogeneous condition. Contrarily, participants assigned to a form of opposing political sign, as in left-wing (right-wing) participants assigned to the right-wing (left-wing) form, belonged to the Heterogeneous condition. Control group participants were assigned the control form. Participants were blind to this arrangement throughout the experiment.

Statements

The corpus of phrases was identical to that in Study 1, and each statement was displayed in the same format.

Procedure

Before the experiment, participants completed a CRT as in Study 1. All experimental forms had a two-stage structure (Figure 5a), where the first stage was identical and involved participants classifying each of the 20 statements as true or false, as well as providing a measure of their confidence in their answer in a 1–5 scale. The information displayed about each statement was identical to the one shown in Study 1. After completion of the first stage, participants answered a simple attentional question before proceeding to the second stage.

In the second stage of the experiment, participants were given the opportunity to revise their original answers. In the two treatment conditions, they were also shown the majority answer for each statement based on the data collected in Study 1. Thus, they had a chance to change their original answers while receiving feedback in the form of "the majority answered True/False." These majority

Figure 5
Study 2 Outline and Results

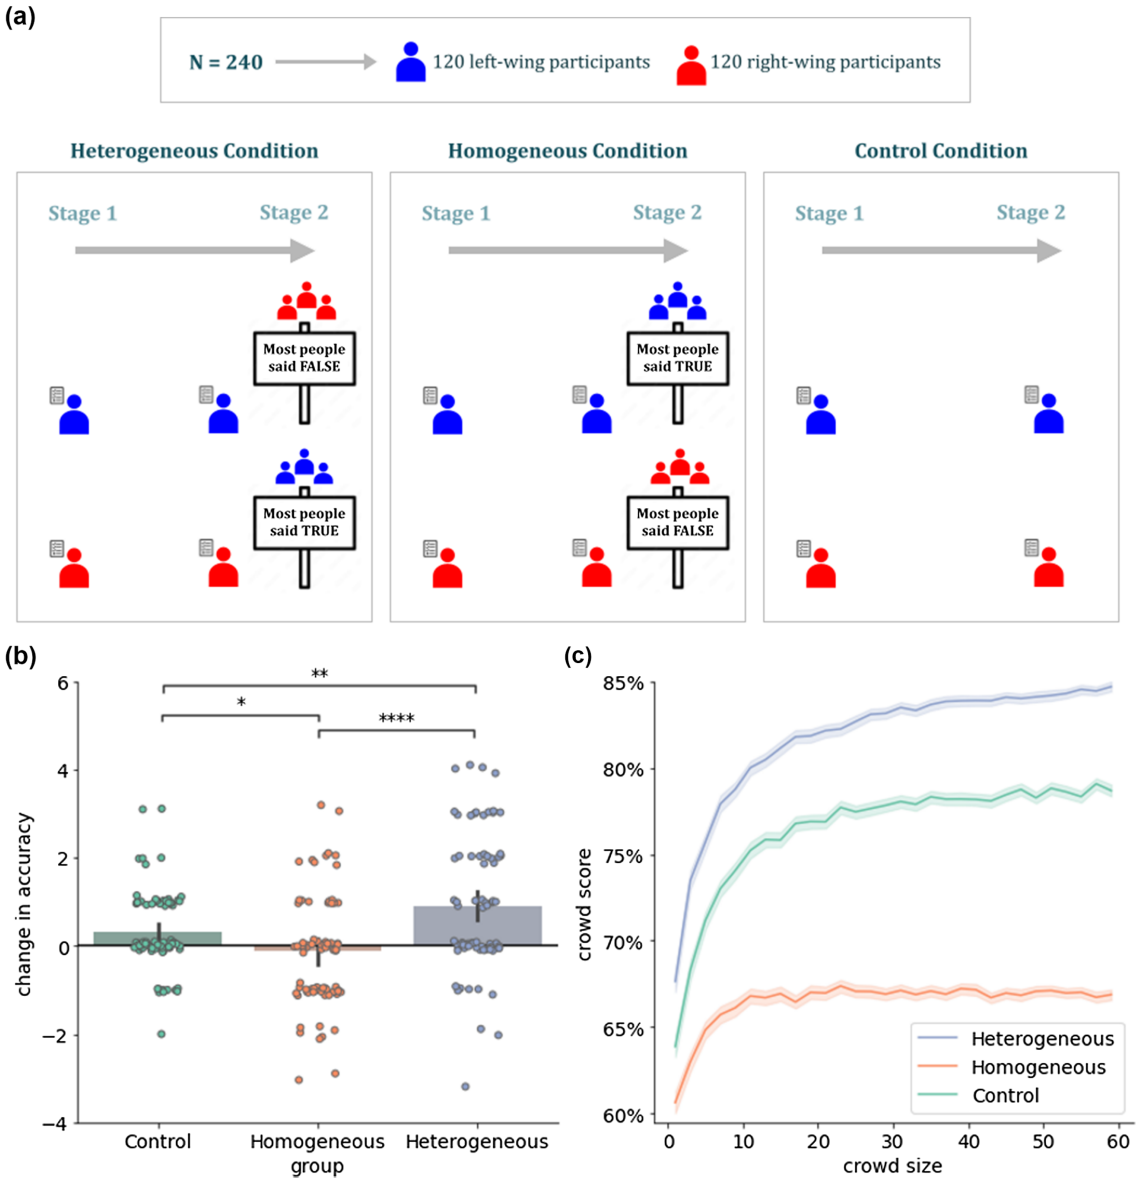

Note. (a) Experimental procedure for Study 2. Participants answered individually in both stages of the experiment, but for the Heterogeneous (Homogeneous) condition they were shown the majority answer from a group of participants of opposed (aligned) partisan affinity when given the opportunity to revise their initial responses. (b) Mean and standard error of the mean of change in the number of individual correct answers between Stage 1 and Stage 2, by group. Outliers are omitted in the plots but included in the analysis. (c) Accuracy of the crowd by group and by crowd size, using Stage 2 responses. Score is defined as the number of correct responses as a % of total statements (20). Mean (solid line) and standard error of the mean (shaded area) across 1,000 iterations. Significance of the corresponding *t* test is depicted at 5% (*), 1% (**), or 0.01% (***). See the online article for the color version of this figure.

votes were drawn from the left-wing or right-wing participants recruited in Study 1. Participants in the Heterogeneous condition were shown the majority answer given by individuals with opposed political sign, and participants in the Homogeneous condition were exposed to the majority answer participants of the same partisanship. A snapshot example of how this information was displayed can be found in Supplemental Figure S5, and the full list of majority votes for each statement given by left-wing and right-wing participants from Study 1 is listed on Supplemental Table S4. Control condition participants did not receive any kind of social information and simply had a re-run of the first stage.

Data Collection

Data were collected using Google Forms (https://docs.google.com/forms) during the first three weeks of March 2023.

Data Analysis

Reported statistical analyses were performed in Matlab R2018b, and figures were generated with Python Jupyter Notebooks in Google Colab (https://colab.research.google.com/).

Regarding aggregation of individual responses into collective ones, the sampling procedure was done exactly as described in the Materials and Method section of Study 1 (see Data Analysis subsection).

Results

Initial Individual Responses

Seeking to replicate results from Study 1, we initially asked whether performance prior to the social influence was above-chance level. We found participants correctly classified an average of 12.4 statements out of 20, which means that their accuracy was significantly better-than-chance ($d = 1.22$, $t_{239} = 18.9$, $p < .001$). There were no differences in initial performance between groups, $\eta^2 = 0.007$, $F(2, 237) = 0.82$, $p = .44$, and we found participants were able to better classify discordant statements compared to concordant ones ($d = 0.35$, $t_{478} = 3.86$, $p < .001$). No differences between groups were found regarding CRT scores, $\eta^2 = 0.001$, $F(2, 237) = 0.82$, $p = .94$, and education levels, $\eta^2 = 0.02$, $F(2, 237) = 1.86$, $p = .16$, and we found a nonsignificant positive correlation between initial accuracy and CRT performance at the individual level ($r = 0.11$, $p = .09$).

We then studied how partisanship modulated participants' beliefs. As in Study 1, we found participants were more likely to classify politically concordant statements as "true" compared to discordant ones ($d = 1.24$, $t_{239} = 19.2$, $p < .001$), specifically by an excess of slightly more than 3 statements ($M \pm SD = 3.38 \pm 2.73$). No differences in bias were observed prior to social influence between experimental conditions, $\eta^2 = 0.013$, $F(2, 237) = 1.58$, $p = .21$.

As preregistered, in Study 2, we measured the subjective affect that participants had for each political figure involved in the statements. This was done to evaluate if this variable was a driver of bias. We calculated the absolute value of the difference in affect between left and right-wing politicians and normalized this variable to range from 0 to 1, calling it affective bias: Values close to 1 indicate a strong difference in how a participant feels for politicians of one party with respect to the other, and values close to 0 imply affective ambivalence or lack of knowledge about the politicians in general. We found no significant differences in affective bias across experimental groups, $\eta^2 = 0.006$, $F(2, 237) = 0.66$, $p = .52$, and that this variable was positively correlated with detection bias at the individual level ($r = 0.61$, $p < .001$).

As in Study 1, we found participants had introspective access into their ability as raters, as they tended to assign higher confidence ratings to correct responses than to incorrect ones (mean confidence in correct trials, $M \pm SD = 3.64 \pm 0.65$; mean confidence in incorrect trials, $M \pm SD = 3.32 \pm 0.74$; $d = 0.46$, $t_{239} = 11.4$, $p < .001$), with no significant differences across groups, $\eta^2 = 0.011$, $F(2, 237) = 1.37$, $p = .26$.

Revised Individual Responses Following Social Influence

The key goal of this preregistered study was to replicate the main finding of Study 1: The distinctive effect of heterogeneous social influence on the ability of individuals to discriminate between true and false statements made by politicians. When comparing revised answers across experimental groups, we found a small-to-medium effect size in performance between the Heterogeneous and Control conditions ($d = 0.35$, $t_{158} = 2.21$, $p = .029$) and between the Control and Homogeneous conditions ($d = 0.36$, $t_{158} = 2.3$, $p = .023$).

Focusing on the change in accuracy across experimental groups (Figure 5b), a significant improvement was found in participants from the Heterogeneous ($d = 0.58$, $t_{79} = 5.2$, $p < .001$) and Control ($d = 0.36$, $t_{79} = 3.18$, $p = .002$) conditions, while this was not the case for participants in the Homogeneous group ($d = -0.09$, $t_{79} = -0.81$, $p = .42$). On the contrary, participants in the Control condition outperformed those in the Homogenous condition ($d = 0.38$, $t_{158} = 2.4$, $p = .018$). As per our preregistration, we found that participants in the Heterogeneous condition had an improvement in performance that was significantly larger than that one observed in both the Control ($d = 0.47$, $t_{158} = 2.98$, $p = .003$) and Homogeneous ($d = 0.7$, $t_{158} = 4.45$, $p < .001$) conditions.

Examining the underlying mechanisms behind the improvement in accuracy, we found no differences in the number of responses revised by participants in the Heterogeneous and Homogeneous conditions ($d = 0.006$, $t_{158} = 0.04$, $p = .97$). Nonetheless, individuals in the Heterogeneous condition were more likely to revise their answers against their partisan ideology, significantly reducing their initial bias ($d = 0.59$, $t_{79} = 5.3$, $p < .001$). On the other hand, participants in the Control condition did not change their bias when given the opportunity to revise their initial answers ($d = 0.1$, $t_{79} = 0.91$, $p = .36$), and those in the Homogeneous condition incremented their bias ($d = 0.42$, $t_{79} = 3.74$, $p < .001$). Even though social feedback increased the propensity to revise initial answers ($d = 0.57$, $t_{158} = 3.58$, $p < .001$ and $d = 0.54$, $t_{158} = 3.44$, $p < .001$ comparing Control participants to Heterogeneous and Homogeneous ones, respectively), heterogeneous social influence reduced preexisting partisan biases, whereas homogeneous feedback actually incremented it.

As in Study 1, we performed an analysis based on signal detection theory, focusing on truth discernment and overall belief rather than simply on accuracy. We found no initial differences in truth discernment, $\eta^2 = 0.005$, $F(2, 237) = 0.56$, $p = .57$, and overall belief, $\eta^2 = 0.005$, $F(2, 237) = 0.62$, $p = .54$, between experimental groups. After being given the opportunity to revise their answers, participants in the Heterogeneous condition had a higher truth

discernment when compared to Homogeneous individuals (d = 0.69, $t_{158} = 4.39$, $p < .001$) and Control group members ($d = 0.33$, $t_{158} = 2.06$, $p = .04$), although posttreatment overall belief remained similar across groups, $\eta^2 = 0.003$, $F(2, 237) = 0.34$, $p = .72$. To summarize, politically heterogeneous social influence distinctively enhanced the capacity of individuals to discriminate true from false statements without affecting their overall proclivity to believe in their veracity.

Crowd Performance: Aggregating Individual Knowledge

Our last main preregistered analyses involved replicating the findings of Study 1 about collective judgements. We found that accuracy is boosted when individual judgements were aggregated into collective estimates for all experimental groups as the size of the crowd increases (Figure 5c; see the Data Analysis Section in Study 1, for details on the aggregation procedure). Crowd judgements based on revised responses from Heterogeneous participants outperformed Homogeneous and Control crowds, as verified by running a multivariate linear regression controlling for crowd size ($\beta = 0.298 \pm 0.033$, $t = 8.99$, $p < .001$). We also found that Heterogeneous crowds achieved a higher accuracy than Control crowds ($\beta = 0.355 \pm 0.038$, $t = 9.27$, $p < .001$), while Homogeneous crowds underperformed them ($\beta = -0.673 \pm 0.038$, $t = -17.6$, $p < .001$). Altogether, crowds formed by individuals belonging to the Heterogeneous condition correctly classified an average of 17 out of 20 statements.

Studying the change in crowd accuracy between Stage 1 and Stage 2 (averaged for all group sizes) to understand how social influence modulates collective performance, we found the improvement was significantly larger than zero in Control (d = 7.56, $t_{29} = 41.4$, $p < .001$) and Heterogeneous crowds ($d = 2.88$, $t_{29} = 15.8$, $p < .001$), and significantly lower than zero in Homogeneous crowds ($d = 3.46$, $t_{29} = 19.0$, $p < .001$). The change in crowd score for the Heterogeneous condition was larger than for the Control condition ($d = 0.59$, $t_{58} = 2.28$, $p = .026$), implying heterogeneous social feedback amplified collective accuracy. Contrarily, Homogeneous crowds registered a lower improvement in performance between Stage 1 and Stage 2 when compared to the Control crowd benchmark ($d = 7.42$, $t_{58} = 28.7$, $p < .001$).

Finally, to increase the comparability of both studies presented in this work, we replicated the exact the same analyses displayed in Figures 2–4 and Supplemental Table S3, which are displayed in Supplemental Figures S6–S8 and Table S5, respectively.

Discussion

In this preregistered study, we replicated all of our main findings of Study 1. Specifically, we verified that politically heterogeneous social influence distinctively improved the ability of individuals to discriminate true from false statements relative to a control benchmark with no social feedback, and that this effect generalizes when participants' judgements are aggregated into collective estimates that outperform individual ones. Unlike Study 1, we achieved this by means of a minimal experimental procedure with no direct social interaction where participants were simply exposed to the majority answer of another group of individuals. This hints toward potential scalable solutions to fact-check politicians with minimal social feedback paradigms.

The most striking difference in results with respect to Study 1 is that Homogeneous groups significantly increased their bias after social influence. We believe this is because, unlike Study 1, participants in Study 2 received feedback for every single phrase. In Study 1, participants discussed only those phrases that they disagreed upon, and so, on average, Homogeneous dyads discussed fewer statements than Heterogeneous dyads. In Study 2, the amount of social feedback was balanced across both experimental conditions and this could explain why we observed greater bias in politically homogeneous groups. This result is consistent with the documented group polarization effect, which predicts people's attitudes or judgements become more extreme following group discussion with others holding similar opinions (Lord et al., 1979; Moscovici & Zavalloni, 1969; Myers & Lamm, 1976).

In this study, we were also able to measure participants' affective bias toward the political figures being evaluated, and we found this to be closely related to the bias stemming from their direct responses. By reciprocity, this result implies that one should be able to extract the underlying subjective partisanship of lay raters based on their direct responses, with no actual need to ask them about their partisan stance or their subjective affect for political figures.

General Discussion

We found that lay raters are able to detect false information enclosed in claims made by politicians with above-chance performance and that this ability can be exploited by crowdsourcing approaches. While their accuracy was partly bounded by preexisting partisan biases, here, we show that these biases could be significantly reduced through social communication in dyads with opposing political ideology or by means of politically heterogeneous minimal social feedback paradigms. We believe that these results are promising regarding the potential usefulness of interactive crowdsourcing approaches to check politicians and reduce the workload of fact-checking agencies.

We found that belief about the veracity of political discourse was significantly distorted by partisanship, extending previous results focused on factual information (Bullock et al., 2015; Guilbeault et al., 2018) and fake news (Faragó et al., 2019; Pennycook & Rand, 2021; Pereira et al., 2018; Vegetti & Mancosu, 2020) to a previously unexplored format. Studying the effect of social influence on preexisting partisan bias, we found it to act as a partial antidote to polarization only when happening in heterogeneous settings.

In principle, one explanation for this finding is that individuals in the Heterogeneous condition were exposed to more social information differing from their initial opinions and this could have nudged them to engage in more revisions, therefore increasing their accuracy. However, our data allow ruling out this trivial explanation. While exposure to social information differing from initial beliefs was indeed higher in the Heterogeneous condition, we found participants in both conditions to be equally likely to change their minds. Instead, the fundamental difference observed between the Heterogeneous and Homogeneous conditions was the nature of those revisions. Participants in the Homogeneous condition tended to change their mind in a direction that increased their partisan congruence, whereas participants in the Heterogeneous condition were more likely to change their mind prioritizing accuracy at the expense of political consistency (Supplemental Figure S9). This suggests that the underlying cause behind the depolarization of

preexisting biases was not the amount of exposure to incongruent social information but rather the exposure to interpartisan judgements per se. One candidate mechanism to explain changes in accuracy after cross-party social influence involves the straightforward act of realizing that someone else, who in these tasks is not necessarily perceived as an outgroup, may think negatively about ingroup political leaders. Future research should examine the psychological mechanisms underlying partisan depolarization when in heterogeneous settings.

Our findings are consistent with Klar (Klar, 2014) but differing from Becker et al. (Becker et al., 2019) in the sense that task accuracy may be sensitive to the partisan composition of interacting networks. We speculate that this discrepancy may be because Becker et al. (Becker et al., 2019) used a factual estimation task, whereas we asked participants to directly report whether or not political leaders were telling the truth. Relative to numerical tasks, whatever their political baggage, directly evaluating the trustworthiness of politicians might exacerbate the salience of partisanship in individuals, a hypothesis coherent with the observation that people do not seem to change their level of support for political leaders even after their claims have been fact checked as false (Swire-Thompson et al., 2020). In this sense, fact-checking politicians might be an endeavor vulnerable to homogeneous social influence, rather than resistant to it, because the enhancement of partisanship may trump the benefits of information exchange. Further research is needed to understand to what extent these differences in task characteristics modulate the effect of homogeneous influence on individual and collective accuracy.

An unresolved limitation of this work is that individual accuracy was possibly limited by time and information constraints determined by experimental design, especially taking into account that participants of both studies only had access to the spoken phrase and the name of the politician. Enabling lay raters with additional resources could increase overall accuracy, although future research is required to test this idea. In addition, our findings stem from a sample of participants with a defined, disclosed partisan affinity. Although extending this setup to moderate participants is challenging given the strength of political polarization in Argentina (Freira et al., 2021; Levy Yeyati et al., 2020), further efforts could be made to understand whether and how partisans perform at identifying false politician statements when compared to neutral or politically uninvolved individuals.

Overall, these results suggest that people can rate the veracity of political speech with above-chance performance and reduce their partisan biases through social influence in heterogeneous settings. Moreover, these findings are a proof of concept that crowdsourcing tools are useful to reduce the workload of agencies aiming to fact-check political statements. Further research is needed to extend this setup to diverse contexts, countries, and formats.

References

Ajzenman, N., Cavalcanti, T., & Da Mata, D. (2020). *More than words: Leaders' speech and risky behavior during a pandemic* (SSRN scholarly paper ID 3582908). Social Science Research Network. https://doi.org/10.2139/ssrn.3582908

Allen, J., Arechar, A. A., Pennycook, G., & Rand, D. G. (2021). Scaling up fact-checking using the wisdom of crowds. *Science Advances*, *7*(36), Article eabf4393. https://doi.org/10.1126/sciadv.abf4393

Batailler, C., Brannon, S. M., Teas, P. E., & Gawronski, B. (2022). A signal detection approach to understanding the identification of fake news. *Perspectives on Psychological Science*, *17*(1), 78–98. https://doi.org/10.1177/1745691620986135

Becker, J., Porter, E., & Centola, D. (2019). The wisdom of partisan crowds. *Proceedings of the National Academy of Sciences of the United States of America*, *116*(22), 10717–10722. https://doi.org/10.1073/pnas.1817195116

Bullock, J. G., Gerber, A. S., Hill, S. J., & Huber, G. A. (2015). Partisan bias in factual beliefs about politics. *Quarterly Journal of Political Science*, *10*(4), 519–578. https://doi.org/10.1561/100.00014074

Burel, G., Farrell, T., Mensio, M., Khare, P., & Alani, H. (2020). *Co-spread of misinformation and fact-checking content during the COVID-19 pandemic* [Conference session]. Proceedings of the 12th International Social Informatics Conference (SocInfo), Pisa, Italy. https://link.springer.com/chapter/10.1007/978-3-030-60975-7_3

Burki, T. (2019). Vaccine misinformation and social media. *The Lancet Digital Health*, *1*(6), e258–e259. https://doi.org/10.1016/S2589-7500(19)30136-0

Cinelli, M., De Francisci Morales, G., Galeazzi, A., Quattrociocchi, W., & Starnini, M. (2021). The echo chamber effect on social media. *Proceedings of the National Academy of Sciences of the United States of America*, *118*(9), Article e2023301118. https://doi.org/10.1073/pnas.2023301118

Espina Mairal, S., Bustos, F., Solovey, G., & Navajas, J. (2022). *Interactive crowdsourcing to fact-check politicians*. PsyArXiv. https://doi.org/10.31234/osf.io/u9eyw

Espina Mairal, S., Navajas, J., & Solovey, G. (2023). *Interactive crowdsourcing to fact-check politicians*. https://osf.io/ufs5c/

Faragó, L., Kende, A., & Kreko, P. (2019). We only believe in news that we doctored ourselves: The connection between partisanship and political fake news. *Social Psychology*, *51*(2), 77–90. https://doi.org/10.1027/1864-9335/a000391

Frau-Meigs, D. (2018). *Societal costs of fakenews in the digital single market*. European Parliament.

Frederick, S. (2005). Cognitive reflection and decision making. *The Journal of Economic Perspectives*, *19*(4), 25–42. https://doi.org/10.1257/089533005775196732

Freira, L., Sartorio, M., Boruchowicz, C., Lopez Boo, F., & Navajas, J. (2021). The interplay between partisanship, forecasted COVID-19 deaths, and support for preventive policies. *Humanities & Social Sciences Communications*, *8*(1), Article 192. https://doi.org/10.1057/s41599-021-00870-2

Frey, V., & van de Rijt, A. (2021). Social influence undermines the wisdom of the crowd in sequential decision making. *Management Science*, *67*(7), 4273–4286. https://doi.org/10.1287/mnsc.2020.3713

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 of the United States of America*, *115*(39), 9714–9719. https://doi.org/10.1073/pnas.1722664115

Jayles, B., Kim, H. R., Escobedo, R., Cezera, S., Blanchet, A., Kameda, T., Sire, C., & Theraulaz, G. (2017). How social information can improve estimation accuracy in human groups. *Proceedings of the National Academy of Sciences of the United States of America*, *114*(47), 12620–12625. https://doi.org/10.1073/pnas.1703695114

Klar, S. (2014). Partisanship in a social setting. *American Journal of Political Science*, *58*(3), 687–704. https://doi.org/10.1111/ajps.12087

Lakens, D., Scheel, A. M., & Isager, P. M. (2018). Equivalence testing for psychological research: A tutorial. *Advances in Methods and Practices in Psychological Science*, *1*(2), 259–269. https://doi.org/10.1177/2515245918770963

Larrick, R. P., Mannes, A. E., & Soll, J. B. (2012). The social psychology of the wisdom of crowds. In J. I. Krueger (Ed.), *Social judgment and decision making* (pp. 227–242). Psychology Press.

Lazer, D. M. J., Baum, M. A., Benkler, Y., Berinsky, A. J., Greenhill, K. M., Menczer, F., Metzger, M. J., Nyhan, B., Pennycook, G., Rothschild, D., Schudson, M., Sloman, S. A., Sunstein, C. R., Thorson, E. A., Watts, D. J.,

& Zittrain, J. L. (2018). The science of fake news. *Science*, *359*(6380), 1094–1096. https://doi.org/10.1126/science.aao2998

Levy Yeyati, E., Moscovich, L., & Abuin, C. (2020). Leader over policy? The scope of elite influence on policy preferences. *Political Communication*, *37*(3), 398–422. https://doi.org/10.1080/10584609.2019.1698681

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

Lorenz, J., Rauhut, H., Schweitzer, F., & Helbing, D. (2011). How social influence can undermine the wisdom of crowd effect. *Proceedings of the National Academy of Sciences of the United States of America*, *108*(22), 9020–9025. https://doi.org/10.1073/pnas.1008636108

Mavrodiev, P., & Schweitzer, F. (2021). The ambigous role of social influence on the wisdom of crowds: An analytic approach. *Physica A: Statistical Mechanics and Its Applications*, *567*, Article 125624. https://doi.org/10.1016/j.physa.2020.125624

Moscovici, S., & Zavalloni, M. (1969). The group as a polarizer of attitudes. *Journal of Personality and Social Psychology*, *12*(2), 125–135. https://doi.org/10.1037/h0027568

Myers, D. G., & Lamm, H. (1976). The group polarization phenomenon. *Psychological Bulletin*, *83*(4), 602–627. https://doi.org/10.1037/0033-2909.83.4.602

Navajas, J., Álvarez Heduan, F., Garrido, J. M., González, P. A., Garbulsky, G., Ariely, D., & Sigman, M. (2019). Reaching consensus in polarized moral debates. *Current Biology*, *29*(23), 4124–4129.e6. https://doi.org/10.1016/j.cub.2019.10.018

Navajas, J., Armand, O., Moran, R., Bahrami, B., & Deroy, O. (2022). Diversity of opinions promotes herding in uncertain crowds. *Royal Society Open Science*, *9*(6), Article 191497. https://doi.org/10.1098/rsos.191497

Navajas, J., Niella, T., Garbulsky, G., Bahrami, B., & Sigman, M. (2018). Aggregated knowledge from a small number of debates outperforms the wisdom of large crowds. *Nature Human Behaviour*, *2*(2), 126–132. https://doi.org/10.1038/s41562-017-0273-4

Osmundsen, M., Bor, A., Vahlstrup, P. B., Bechmann, A., & Petersen, M. B. (2021). Partisan polarization is the primary psychological motivation behind political fake news sharing on twitter. *The American Political Science Review*, *115*(3), 999–1015. https://doi.org/10.1017/S0003055421000290

Pennycook, G., & Rand, D. G. (2019a). Fighting misinformation on social media using crowdsourced judgments of news source quality. *Proceedings of the National Academy of Sciences of the United States of America*, *116*(7), 2521–2526. https://doi.org/10.1073/pnas.1806781116

Pennycook, G., & Rand, D. G. (2019b). Lazy, not biased: Susceptibility to partisan fake news is better explained by lack of reasoning than by motivated reasoning. *Cognition*, *188*, 39–50. https://doi.org/10.1016/j.cognition.2018.06.011

Pennycook, G., & Rand, D. G. (2021). The psychology of fake news. *Trends in Cognitive Sciences*, *25*(5), 388–402. https://doi.org/10.1016/j.tics.2021.02.007

Pereira, A., Harris, E. A., & Bavel, J. J. V. (2018). *Identity concerns drive belief: The impact of partisan identity on the belief and dissemination of true and false news*. PsyArXiv. https://doi.org/10.31234/osf.io/7vc5d

Pescetelli, N., Rutherford, A., & Rahwan, I. (2021). Modularity and composite diversity affect the collective gathering of information online. *Nature Communications*, *12*(1), Article 3195. https://doi.org/10.1038/s41467-021-23424-1

Primi, C., Morsanyi, K., Chiesi, F., Donati, M. A., & Hamilton, J. (2016). The development and testing of a new version of the cognitive reflection test applying item response theory (IRT). *Journal of Behavioral Decision Making*, *29*(5), 453–469. https://doi.org/10.1002/bdm.1883

Resnick, P., Alfayez, A., Im, J., & Gilbert, E. (2021). *Informed crowds can effectively identify misinformation*. ArXiv. https://doi.org/10.48550/arXiv.2108.07898

Rhodes, S. C. (2022). Filter bubbles, echo chambers, and fake news: How social media conditions individuals to be less critical of political misinformation. *Political Communication*, *39*(1), 1–22. https://doi.org/10.1080/10584609.2021.1910887

Schulz-Hardt, S., Brodbeck, F. C., Mojzisch, A., Kerschreiter, R., & Frey, D. (2006). Group decision making in hidden profile situations: Dissent as a facilitator for decision quality. *Journal of Personality and Social Psychology*, *91*(6), 1080–1093. https://doi.org/10.1037/0022-3514.91.6.1080

Shi, F., Teplitskiy, M., Duede, E., & Evans, J. A. (2019). The wisdom of polarized crowds. *Nature Human Behaviour*, *3*(4), 329–336. https://doi.org/10.1038/s41562-019-0541-6

Sunstein, C. R. (2009). *Going to extremes: How like minds unite and divide*. Oxford University Press.

Surowiecki, J. (2004). *The wisdom of crowds: Why the many are smarter than the few and how collective wisdom shapes business, economies, societies, and nations*. Little, Brown.

Swire-Thompson, B., Ecker, U. K. H., Lewandowsky, S., & Berinsky, A. J. (2020). They might be a liar but they're my liar: Source evaluation and the prevalence of misinformation. *Political Psychology*, *41*(1), 21–34. https://doi.org/10.1111/pops.12586

van der Linden, S., Leiserowitz, A., Rosenthal, S., & Maibach, E. (2017). Inoculating the public against misinformation about climate change. *Global Challenges*, *1*(2), Article 1600008. https://doi.org/10.1002/gch2.201600008

Vegetti, F., & Mancosu, M. (2020). The impact of political sophistication and motivated reasoning on misinformation. *Political Communication*, *37*(5), 678–695. https://doi.org/10.1080/10584609.2020.1744778

Zhang, Z. (2014). Monte Carlo based statistical power analysis for mediation models: Methods and software. *Behavior Research Methods*, *46*(4), 1184–1198. https://doi.org/10.3758/s13428-013-0424-0

Zimmerman, F., Garbulsky, G., Ariely, D., Sigman, M., & Navajas, J. (2022). Political coherence and certainty as drivers of interpersonal liking over and above similarity. *Science Advances*, *8*(6), Article eabk1909. https://doi.org/10.1126/sciadv.abk1909

Received December 21, 2022
Revision received June 21, 2023
Accepted June 28, 2023 ■