Group Identity and Belief Formation: Implications for Political Polarization*

Kevin Bauer Yan Chen Florian Hett Michael Kosfeld

March 7, 2025

Abstract

To evaluate the impact of group identity on belief formation, we conducted online experiments before and after the 2020 US presidential election. We elicit participants' beliefs about future unemployment statistics and provide relevant news summaries. We find that people pay money to avoid information from political outgroups and attribute lower weight to this information when updating beliefs. An intervention which unlabels information sources decreases outgroup information avoidance by 50%, an effect driven by groupish participants. A debiasing intervention equalizing instrumental values of information sources reduces only universalists' information avoidance. We establish *source utility* as a key mechanism contributing to polarization.

Keywords: group identity, information demand, information processing, political polarization

*We thank Marina Agranov, George Akerlof, Roland Bénabou, Doug Bernheim, Antonio Cabrales, Alain Cohn, Eugen Dimant, Markus Eyting, Russell Golman, YingHua He, Daniel Houser, Iris W. Hung, Matt Jackson, Navin Kartik, Victor Klockmann, Rachel Kranton, Erin Krupka, Dorothea Kübler, Steve Leider, Ro'ee Levy, Annie Liang, George Loewenstein, Yusufcan Masatlioglu, Juanjuan Meng, Muriel Niederle, David Poensgen, Tanya Rosenblat, Al Roth, Patrick Schneider, Alicia von Schenk, Ferdinand von Siemens, Zahra Sharafi, Georg Weizsäcker, Erte Xiao, Songfa Zhong, Na Zou and audiences at the Berlin Behavioral Economics Seminar, Carnegie Mellon, Cornell, Goethe University Frankfurt, Johns Hopkins, Lund, Mannheim, Max Planck Institute for Research on Collective Goods Bonn, NYU Abu Dhabi, Stanford, Universidad Carlos III de Madrid, University of Amsterdam, Copenhagen, Fribourg, Hamburg, Lüneburg, Michigan, Pennsylvania, Regensburg, Uppsala, Würzburg, Zürich, the 2021 ERINN Conference, and other conferences for their helpful comments. We thank Lucy Jiang, Madhavan Somanathan, and Paul Weingärtner for excellent research assistance. Goethe-Universität Frankfurt IRB granted this project exempt status. The research was financially supported by the Leibniz Institute for Financial Research SAFE and the University of Michigan. Bauer and Kosfeld: Faculty of Economics and Business, Goethe University Frankfurt Theodor-W.-Adorno-Platz 4, D-60323 Frankfurt am Main, Germany. Email: bauer@wiwi.uni-frankfurt.de; kosfeld@uni-frankfurt.de. Chen: School of Information, University of Michigan, 105 South State Street, Ann Arbor, MI 48109-2112. Email: yanchen@umich.edu. Hett: Johannes Gutenberg University Mainz, Chair of Digital Economics, Jakob-Welder-Weg 4, D-55128 Mainz, Germany. Email: florian.hett@uni-mainz.de.

1 Introduction

People act based on what they believe. These beliefs may include their assessment of the correctness of stated facts (Peterson and Iyengar 2021), opinions about optimal policies (Alesina et al. 2020), and moral values and norms (Andre et al. 2024). In recent years, academic and public discussions have emerged, focusing on the increasing polarization of opinions in the United States (Allcott et al. 2020), particularly emphasizing its increasingly affective nature. Surveys indicate that hostility towards political outgroups has increased sharply in recent decades (see, e.g., Figure 1) and led to increased skepticism toward policies supported by outgroups (Iyengar, Sood and Lelkes 2012). These developments have raised concerns about the erosion of constructive public discourse and the diminishing ability of democracy to operate effectively.

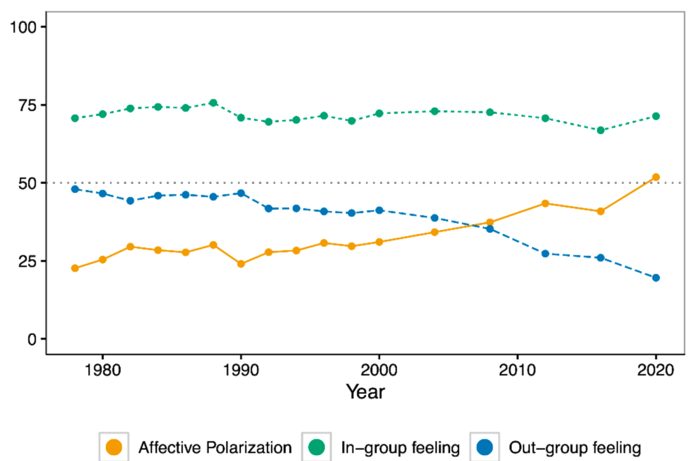

Figure 1: Affective polarization (solid yellow line) increases between 1980 and 2000: Vertical axis denotes the feeling thermometer, with 0-50 indicating cold and 50-100 indicating warm feelings. Affective polarization is the difference between in-party (dotted green line) and out-party temperature (dashed blue line). The figure suggests that Americans increasingly dislike and distrust those from the other party. This figure uses data from the American National Election Studies (Druckman and Levy 2022).

But how do divergent, or polarized, beliefs emerge? Who drives political polarization? In this paper, we analyze the role of group identity in the formation of individual beliefs. Both social psychologists (Tajfel and Turner 1979) and more recently behavioral and experimental economists have found that group identity affects individual behavior in a wide range of economic activities (Akerlof and Kranton 2000, Charness and Chen 2020, Li 2020, Shayo 2020, Chowdhury 2021). We hypothesize that group identity also affects individual belief formation, as individuals tend to selectively demand and process information in a way that is sensitive to the group associations of the information source.

Our study examines the impact of group identity on the entire cycle of individual belief formation in two waves of online experiments, using the 2020 US presidential election as our setting. To measure belief dynamics, we incentivized a representative sample of US participants to predict the trajectory of the official unemployment rate and the US public health system ranking 10 months after the election, conditional on which candidate became president. We consider two possible manifestations of intergroup preferences in belief formation: 1) people prefer information coming from political ingroup instead of outgroup sources; and 2) when updating their beliefs, people attribute a higher weight to political ingroup than outgroup information sources.

We begin by examining whether participants demand different information based on their perceived political affiliations. We provide participants with summaries of news articles on the respective topic before asking if they would like to update their initial predictions about the trajectory of policy-sensitive statistics. Although the articles cover the same facts, they differ with respect to their attributed sources, which range from left-leaning (e.g. *The New York Times*) to nonpartisan (e.g. *Nature*) to right-leaning (e.g. Fox News). We find that participants show strong demand for nonpartisan versus political outgroup information, yet no preference for political ingroup versus nonpartisan information sources. That is, almost 40% of the participants are willing to pay part of their monetary endowment to read articles from nonpartisan versus the opposing political camp. In a treatment that withholds the sources of the articles (*unlabeling treatment*), we find that participants reduce their outgroup information avoidance expenditures by 50% (from 14.9% to 7.4% of their monetary endowment).

We next examine whether participants process information differently conditional on its source. We ask them to read two news article summaries (one left-leaning and one right-leaning) that discuss the policy consequences of different presidential election outcomes. After reading the two articles, participants can update their predictions on the respective topic. Given that participants already start with a partisan gap in prior beliefs (they predict a better policy trajectory if their candidate wins), one may expect a narrowing of this gap after reading both left-leaning and right-leaning articles. This is not the case. Rather, we find that priors are sticky as they carry 62% to 80% of the weight in participants' posteriors. Investigating the mechanism behind this, we find that participants with stronger initial partisan gaps also evaluate the quality difference of ingroup and outgroup articles to be larger. Our results are consistent with the theoretical predictions of Fryer, Harms and Jackson (2019).

To isolate the underlying mechanisms driving our results and rule out alternative explanations (perceived content valence and instrumental value of information), we conducted a second wave of social learning experiments with the same participants three months later, asking them to predict whether an urn contains more green or yellow marbles. As before, we offer participants the possibility to learn from different information sources, this time from a private (statistically informative) signal or prior guesses made by either Democrats or Republicans. The results reinforce the existence of outgroup information avoidance and biased information processing. Specifically, we find that participants spend almost twice as much of their endowment to avoid guesses made by members of an out- versus ingroup and place 33% more weight on guesses from an ingroup than outgroup source.

The simplicity of the social learning tasks enables us to identify the extent to which intergroup preferences in belief formation are driven by *source utility*, that is, whether it is based on the group association of the information independent of its instrumental value and content valence. In a treatment variation consisting of truthfully informing half of the participants that there is no difference in guess accuracy between Democrats and Republicans (*debiasing treatment*), we find persistence of outgroup information avoidance and differential processing of in- versus outgroup information.

To further gauge source utility as a mechanism and to relate intergroup preferences in belief formation to group identity, we classify participants into *groupish* and *universalist* types based on their behavior in a classic bystander allocation game in randomly assigned minimal groups (Tajfel et al. 1971, Chen and Li 2009, Kranton et al. 2020) in both waves of the experiment. A participant is classified as *groupish* if they favor ingroup members in randomly assigned minimal groups, and *universalist* otherwise. These types are stable across the two waves. Furthermore, a participant's groupishness in the minimal group context strongly predicts their groupishness in the political group context. We test the moderating role of this classification in our subgroup analysis. Indeed, we find that groupish participants have a stronger partisan gap in their prior beliefs and exhibit stronger outgroup information avoidance. Further, the treatment effect of unlabeling information sources in wave 1 is entirely driven by groupish participants. In contrast, the treatment effect of debiasing in wave 2 only affects the universalists. We do not observe heterogeneities with respect to information processing, though. Documented heterogeneities in prior beliefs and information demand further support the existence of source utility as an important mechanism underlying intergroup preferences in individual belief formation.

The rest of this paper proceeds as follows. Section 2 discusses how our paper contributes to the literature on identity economics, information preferences and political polarization. In Section 3, we detail the experimental design and results from the first wave of our study. In Section 4, we present the design and results from the second wave. Section 5 provides discussions and concludes.

2 Related Literature

Our study builds on and contributes to three different streams of research within the social science literature: identity economics, information preferences, and political polarization.

First, we contribute to the **identity economics** literature pioneered by Akerlof and Kranton (2000). Research in this area has examined the various behavioral consequences of group identity, using both natural identities (Goette, Huffman and Meier 2006) and artificial ones induced in the laboratory (Eckel and Grossman 2005, Charness, Rigotti and Rustichini 2007, Chen and Li 2009). This vast and growing experimental literature on group identity primarily investigates its effects on *allocation decisions* in various games, redistribution, and labor market outcomes, as summarized in four recent surveys (Charness and Chen 2020, Li 2020, Shayo 2020, Chowdhury 2021). More recent research in identity economics expands into two new directions: the effects of identity on belief formation and heterogeneity in intergroup preferences. We discuss how our study contributes to each of these directions.

Identity and belief formation. Our study expands the research on identity economics by broadening the scope of the effects of group identity to the domain of belief formation showing that political affiliations shape outgroup information avoidance and weighting due to source utility. We also complement concurrent online experiments that study various aspects of identity and belief formation. In particular, Dekel and Shayo (2023) study social learning, conformity, and differentiation within and between groups. Their experiment design is similar to ours in that it consists of two waves of online experiments. However, the settings and timing of the waves differ in the two studies. Specifically, for the first wave focusing on belief formation in a political context, we focus on 2020 just prior to the US presidential election, whereas they focus on 2021 just after a contentious Israeli election. For the second wave focusing on social learning, we use the same set of participants in 2021 whereas they focus on the role of expertise in social learning with a new set of UK participants. Interestingly, the results of the two studies complement and strengthen each other. Both studies find that in social learning tasks, individuals follow their ingroup significantly more than their outgroup. Our study also extends the scope of their conclusions by focusing on the entire cycle of belief formation across prior beliefs, information demand, and information processing.

In another recent study, Dimant et al. (2023) theoretically and empirically analyze how identity affects investment decisions. In their theoretical model, they posit that identity distorts individual beliefs about uncertain outcomes and drives preferences by imposing psychic costs on identity-incongruent actions. Their finding that soccer fans have overoptimistic (underconfident) beliefs about identity-congruent (incongruent) outcomes in comparison to neutral outcomes is consistent with our wave 1 results.

Finally, Liu and Zhang (2023) use an online experiment with the issue prompt of genetically modified mosquitoes to study whether subsequent information acquisition can mitigate the impact of biased narratives (interpretations of objective facts or events) on belief formation. Their design is similar to the information demand and processing part of our wave 1 design. We compare their results with ours in Section 3.

Heterogeneity in intergroup preferences. Evolutionary psychologists argue that historically humans have lived in tribes, in which individuals cooperate to ensure survival and compete with rival tribes. The term "groupish" is used to describe the tendency of individuals to prioritize their ingroup over outgroups (Haidt 2012).[1] However, a precise metric for groupishness is lacking.

In this paper, we use the bystander allocation game of the minimal group paradigm (Tajfel et al. 1971) but with real incentives (Chen and Li 2009) to measure a participant's groupishness. An individual is groupish if they allocate more than 50% of monetary endowment to their ingroup member between members of randomly assigned minimal groups. We show that (minimal) groupishness is (1) highly correlated with political groupishness (bystander allocation game with partisan groups); (2) stable across two waves; and (3) explains heterogeneous treatment effects in unlabeling information sources as well as equalizing beliefs about the instrumental value of information sources.

Closely related to our measurement of groupishness is Kranton and co-authors' definition of a "groupy" type as the intersection of minimal and political groupishness (Kranton and Sanders 2017, Kranton et al. 2020). Compared to the "groupiness" measure, we argue that (minimal) groupishness has two main advantages: (1) groups are randomly assigned, and (2) the measurement is portable across contexts.

The complement of groupishness is universalism. Using bystander allocation games across different domains of natural identities, Enke, Rodriguez-Padilla and Zimmermann (2022) define a behavioral type called *universalist*, which corresponds to those who are not groupish in our setting. Using nationally representative surveys in 60 countries, Cappelen, Enke and Tungodden (2025) show that universalism is strongly linked to a broader radius of trust across countries and are predictive of many left-wing political views.

[1] See Chapter 9 in Haidt (2012) for details, pages 219-255.

Second, we contribute to the literature on **information preferences**. While classic theories on the value of information have focused on its instrumental value (Stigler 1961), there is a growing body of literature exploring non-instrumental information preferences. The latter might come from anticipatory utility (Caplin and Leahy 2001), motivated beliefs (Bénabou and Tirole 2016), or valence (Golman and Loewenstein 2018). Experimental studies include Eliaz and Schotter (2007, 2010), who evaluate agents' demand for non-instrumental information. Other experiments examine the willingness of subjects to pay for non-instrumental information in the context of social learning (Kübler and Weizsäcker 2004, Goeree and Yariv 2015) and matching (Chen and He 2021). We extend the information gap theory of Golman and Loewenstein (2018) by introducing the concept of "source utility," or a preference for information based solely on the group association of the source which is independent of the information's content or instrumental value, and evaluate its role in belief formation.

Lastly, we contribute to the broader social science literature surrounding the sources, forms, and consequences of **political polarization**. Recent research documents the increasing prevalence of *affective* polarization – negative feelings towards people identifying with political parties other than their own (Iyengar et al. 2012). Using social identity theory, several studies posit that affective polarization stems from partisans increasingly viewing each other as disliked outgroups (Iyengar et al. 2012, 2019). Our research complements this literature by validating the bystander allocation game as an incentivized measure of affective polarization, and by demonstrating that groupish and universalist types respond to different policy interventions designed to decrease political polarization.

3 Wave 1: Belief Formation in a Political Context

Overview. To investigate the role of group identity in belief formation, we conducted three waves of online experiments between October 2020 and November 2021. In wave 1, we recruited a representative sample of the US adult population and investigated individuals' belief formation in the context of the 2020 US presidential election. Wave 2 was implemented three months later with participants from the same sample. Finally, we implemented wave 3 between 23 October and 3 November 2021 after all the policy predictions from the previous waves were realized. Each of the three waves of our study was pre-registered at the AEA RCT Registry (Bauer et al. 2020, 2021a, b). We used the platform *Prolific.co* to implement our experiment as a computerized online study. The experiment was coded using Python libraries including oTree (Chen, Schonger and Wickens 2016). Experimental instructions, screen shots, surveys, and responses are included in Appendix A.

Due to space limitation, we focus on the first two waves in this paper, and refer the reader to our SSRN working paper for the complete instructions and results in wave 3 (Bauer et al. 2023).

3.1 Wave 1: Experimental Design

For wave 1, we deployed an online experiment with 1,005 participants the week before the 2020 US presidential election. Our sample is nationally representative of the US population in terms of gender, urban versus rural location, race and ethnicity (except for the "Other" category), but is slightly younger and better educated than the US population average (Appendix Table B.1). Furthermore, our sample contains more Democrats (0.48), fewer Republicans (0.26) and Independents (0.22) compared to the population average (Appendix Table B.1). These participants form the basis for our panel whom we re-contact in wave 2.

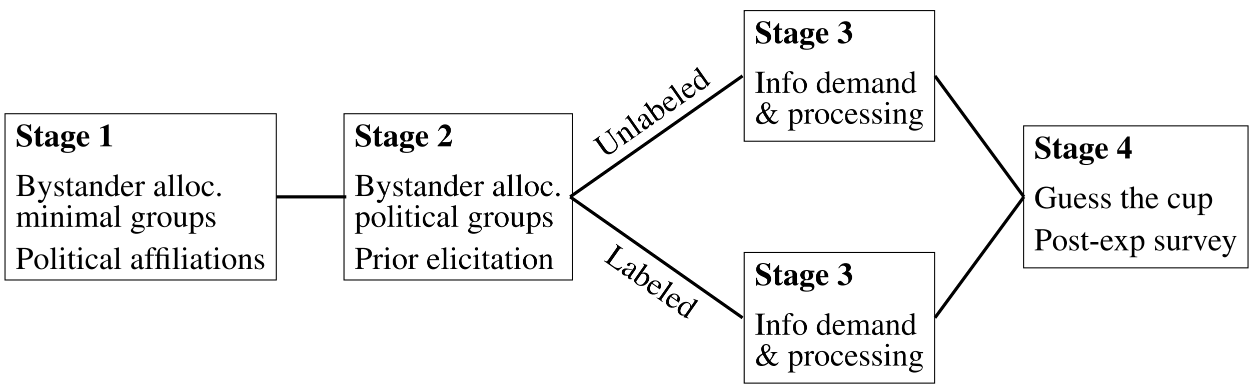

Figure 2: Wave 1 experiment (26-29 October 2020): Alloc is short for allocation game.

Figure 2 illustrates the four stages of our experiment in wave 1. As our goal is to explore the relationship between group identity and belief formation, we start by measuring participants' sensitivity to group identity. To do so, we draw on classic studies on social identity theory (Tajfel et al. 1971) and use incentivized bystander allocation games with randomly assigned minimal groups (Chen and Li 2009) and political groups (Kranton et al. 2020). We use the results of these games to classify participants into groupish and universalist types based on their propensity to exhibit ingroup favoritism.

Stage 1. We begin with minimal groups, where we randomly assign participants to a "circle" or "triangle" group. Each participant receives $6 to allocate between two randomly selected participants, one from each group or both from the same group. The minimum amount that participants can allocate to each individual is \$1. We use their choices in this game to identify a participant's tendency to differentiate between ingroup and outgroup members, favoring the former over the latter, and refer to this tendency as the participant's level of *groupishness*. Participants are classified as *minimally groupish* if they allocate at least one dollar more to an ingroup member, i.e., exhibit ingroup favoritism, in the context of minimal groups. This stage will be repeated in wave 2 to assess the stability of minimal groupishness over time.

Stage 2. After observing participant behavior in this minimal group setting, we ask participants to answer survey questions designed to measure their political party affiliation and degree of affiliation.[2] We use their answers to classify them as Democrat (Democrats and Independents leaning Democratic) or Republican (Republicans and Independents leaning Republican). Using these political identifications, we again implement our bystander allocation game, this time requiring participants to allocate money between a randomly selected member of the Democrat group (henceforth Democrat) and a randomly selected member of the Republican group (henceforth Republican), or between two participants from the same group. Participants are classified as *politically groupish* if they allocate at least one dollar more to those in their party, i.e., exhibit ingroup favoritism in the context of partisan groups. This part will also be repeated in wave 2 to assess the stability of political groupishness.

In this stage, we also measure participants' prior beliefs in the context of the 2020 US Presidential Election. Specifically, we ask participants to predict the outcome of the election, as well as the trajectory of the unemployment rate and an official public health system ranking (based on the annual rankings from *US News and World Report*) in September 2021, conditional on which candidate becomes president, i.e., they make predictions under both possible election outcomes. We incentivize participants by paying them \$3 per correct forecast in February (for their forecast of the election outcomes) and \$10 per correct forecast in November 2021 (for their forecasts of the unemployment rate and the public health system ranking), respectively. Since unemployment rates and public health system rankings are tied to specific candidate platforms, conditional beliefs about these statistics reflect participants' prior beliefs about the efficacy of each candidate's policies. These incentivized forecasting tasks create a demand for relevant information about each candidate.

Stage 3. To examine participants' information demand and processing behavior, we next implement two respective substages where participants can update their prior predictions about the

[2] Validating our classification through a set of factual statements which prior research has shown to demonstrate partisan gaps, we find that we replicate the partisan gap (Peterson and Iyengar 2021). See Appendix B.2.

trajectory of public health and unemployment statistics based on curated news articles. We randomize the order of the two substages at the individual level.

In the *information demand* substage, participants are presented with two titles of articles that discuss the potential policy consequences for either the public health system or unemployment under different election outcomes. Participants can choose to read one of the two articles and subsequently update their predictions on the respective topic. As a default, one of the two articles is selected with 50% probability and displayed on the participant's computer screen. Participants use a slider to indicate how much of a \$3 endowment they would spend to adjust the probability of receiving one article versus the other, where each 10% probability change costs \$1. Participants are allowed to keep any remaining portion of their endowment.[3] In this stage, participants make two consecutive slider decisions, one of which we randomly select for implementation. In both cases, one article is from a nonpartisan news source while the second is from either a left- or right-leaning source. We curate articles from well-known news media outlets with different political leanings and then synthesize these articles to ensure that they are similar in length, format, and facts (see Appendix A.1.9).[4] After reading the article provided to them, participants answer incentivized review questions about the article, indicate its perceived leaning and quality, and decide whether to update their predictions on the corresponding topic. The share of their endowment they choose to spend to increase the probability of receiving information from the (relatively) more party-favorable source reflects their intergroup preferences in information demand; e.g., the share a Democrat spends to receive a left-leaning (nonpartisan) article versus a nonpartisan (right-leaning) one.

In the *information processing* substage, each participant exogenously receives in random order two news articles that both discuss the potential policy consequences for the public health system or unemployment under different election outcomes. Regardless of the topic, one article is always curated from left-leaning sources and another from right-leaning news sources. The topic (public health or unemployment) of the articles in this substage differs from the topic a participant receives in the information demand substage. Again, the articles cover the same facts but with a different slant. After reading both articles, participants answer incentivized review questions about the

[3] Note that they can increase the probability of reading an article to at most 80%, if they spend their entire endowment.

[4] Our left-leaning news outlets include the *New York Times*, the *Washington Post*, and NBC or MSNBC; our nonpartisan outlets include The Bureau of Labor Statistics (for the economy), the *Economist*, *Nature* (for health), and Reuters; our right-leaning outlets include the *Wall Street Journal*, the *Washington Examiner*, and Fox News. Our classification of the political leanings of the news media outlets is consistent with that based on a combination of machine-learning and crowdsourcing techniques (Budak, Goel and Rao 2016).

articles and indicate their perceived leaning and quality. Importantly, participants have the option to update their predictions on the topic chosen for this stage.

Intervention – Unlabeling of News Sources. At the beginning of stage 3, half of the participants are randomized to see the labels of the news sources ("Labeled" - control condition) in the demand and processing stage, whereas the other half of the subjects see identical articles without the labels of the news sources ("Unlabeled"- treatment condition). Since labels provide a meaningful cue regarding the group association of information, the unlabeling of this information should reduce the salience of group identity associated with the information source and hence intergroup preference in selecting and processing information. Figure 3 illustrates how participants see articles with (left panel) and without labels (right panel) in the information demand stage, while Appendix A.1.8 presents the corresponding screen shots in the information processing stage.

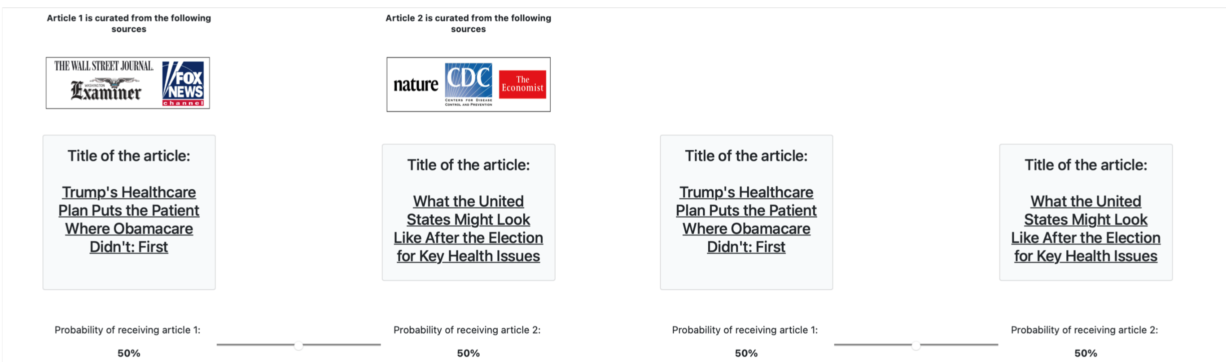

Figure 3: Information demand screenshots in wave 1: The left (right) panel depicts the information demand task with (without) source labels. The slider is positioned at 50% as a default.

Stage 4. In the fourth stage, participants play an incentivized, neutral "guess-the-cup" game of the following structure. We first show participants a Green and a Yellow cup and explain that the Green (Yellow) cup contains two green (yellow) marbles and one yellow (green) marble. We then randomly choose one of the two cups with 50% probability without informing the participant about its color. For each participant, the computer randomly draws a marble from the chosen cup and reveals its color to the participant. Upon seeing the color of the marble, a participant guesses the color of the cup, receiving $1 for a correct guess and zero otherwise. This game is designed to examine the participants' Bayesian inference in a neutral context. As such, it provides a control variable for participant updating behavior ("Bayesian rationality" in subsequent regression tables) and establishes the nature of the wave 2 task for participants.

Lastly, participants complete a questionnaire containing items on their socio-demographics and their news consumption behavior (see Appendix A.1.6 for questionnaire and survey items).

On average, it took 26 minutes for our participants to finish the experiment. The average earnings in this wave are $13.04.[5]

3.2 Wave 1 Results

In this section, we first report our results regarding participants' types – groupish or universalist. We then provide the results of their prior beliefs, information demand and processing experiments.

We present a theoretical framework in Appendix C which derives hypotheses for our experiment, to which we refer when discussing the results.[6]

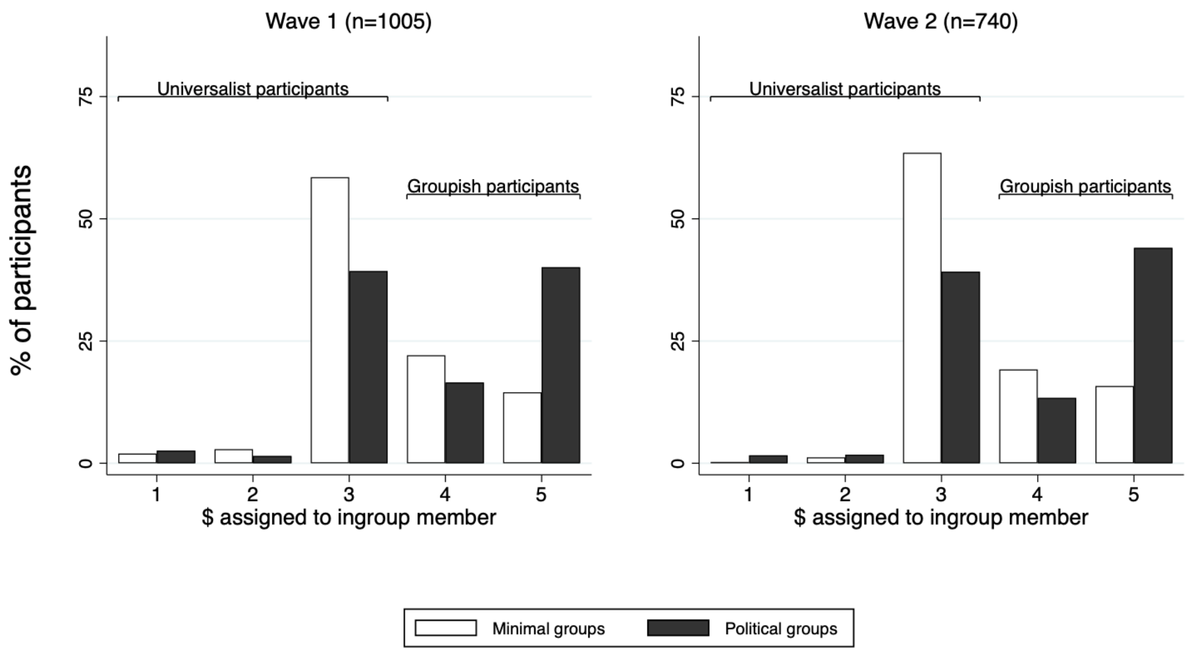

Figure 4: Distribution of allocation decisions in the minimal (white bars), and political (black bars) bystander allocation games in waves 1 (left panel) and 2 (right panel): The horizontal axis shows the amount (in dollars) allocated to an ingroup member out of a total budget of $6. The maximum amount a participant can allocate to another participant is $5.

Groupishness. Figure 4 illustrates the participants' decisions in our wave-1 and wave-2 bystander allocation game with minimal groups (white bars) and political groups (black bars). We

[5]Note that this amount only includes payments participants received at the end of wave 1. The payments associated with the incentivized forecasts that the participants made were received at the end of waves 2 and 3.

[6]When presenting wave 1 results, we compare our subjects' behavior with the forecasts by 34 academic experts from the Social Science Prediction Platform (shortened as SSPP henceforth) (DellaVigna and Pope 2017). See Appendix B in our SSRN working paper (Bauer et al. 2023) for our forecast survey and expert responses.

focus on minimal groups as the key measure of participants' behavioral sensitivity to group contexts in allocation choices, as the minimal group membership is randomly assigned, stable across waves, and highly correlated with political groupishness in wave 1 ($\rho = 0.361$, $p < 0.01$) and wave 2 ($\rho = 0.466$, $p < 0.01$).

Table 1: Stability of minimal groupishness across waves: Linear panel specifications.

DV: Being Groupish	In Wave 1		In Wave 2				
	Politically Groupish		Minimally Groupish		Politically Groupish		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Minimally Groupish (wave 1)	0.256***	0.260***	0.441***	0.421***	0.256***	0.260***	0.068*
	(0.036)	(0.036)	(0.035)	(0.036)	(0.036)	(0.036)	(0.037)
Minimally Groupish (wave 2)							0.457***
							(0.034)
Constant	0.487***	0.325***	0.199***	0.226**	0.487***	0.325***	0.222***
	(0.023)	(0.094)	(0.018)	(0.093)	(0.023)	(0.094)	(0.082)
Controls	No	Yes	No	Yes	No	Yes	Yes
Observations	1005	1005	740	740	740	740	740
Adj. R-squared	0.059	0.130	0.191	0.201	0.059	0.130	0.284

[a.] Robust standard errors are reported in parentheses.

[b.] As controls, we include participants' political leaning, party identification, factual bias, frequency of consuming media from opposing viewpoints, gender, education level, ethnicity, area of living, age, and Bayesian rationality.

[c.] We denote significance levels by * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

From the left panel in Figure 4, we see that in the minimal (political) setting, 59% (39%) of our wave-1 participants split their endowment equally, 5% (4%) strictly favor outgroup participants, and 37% (57%) strictly favor ingroup participants. We refer to this latter group as *minimally groupish* (politically groupish) individuals, and the first two groups as *universalists*. We do not find any difference in the distributions of the minimal group allocation choices between Democrats and Republicans ($p = 0.93$, Kolmogorov-Smirnov test). Note that these distributions are comparable to those found in previous studies (Chen and Li 2009, Kranton, Pease, Sanders and Huettel 2020).[7] Similar to Kranton et al. (2020), we find that Independents are significantly less likely to be minimally groupish compared to Democrats and Republicans (-9.8 pp and -13.6 pp, respectively; $p < 0.02, \chi^2$-test for both) and less likely to be politically groupish compared to Democrats and Republicans (-19.9 pp and -10.6 pp, respectively; $p < 0.02, \chi^2$-test for both).

[7] We do observe a statistically significant difference in the distribution of political group allocation choices between Democrats and Republicans ($p < 0.01$, Kolmogorov-Smirnov test), with Democrats being more likely to be politically groupish (64.3% v. 54.9% of participants).

The right panel of Figure 4 shows that the shares of groupish and universalist types in wave 2 are remarkably similar to those in wave 1.[8]

With repeated measurements across different waves of our experiment, our panel design enables us to investigate both the consistency and stability of the distribution of groupishness: in wave 1 (2), respectively, we find that 37% (35%) of our participants are minimally groupish. For individuals who participate in both waves ($n = 740$), we consistently classify 74.6% of them as groupish (21.9%) or universalist (52.7%). Linear panel regression analysis in Table 1 reveal a strong and significant predictive power of wave-1 groupishness for wave-2 groupishness.

Result 1 (The Groupish Type)**.** *In wave 1 (2), 37% (35%) of our participants are groupish in randomly assigned minimal groups. Being minimally groupish is stable across waves and highly predictive of political groupishness.*

On the methodology front, we also find a strong positive correlation between affective polarization and political groupishness ($\rho = 0.48, p < 0.01$, Spearman). The former is measured as the gap between individuals' feeling thermometer ratings of their supported and opposing political parties (Iyengar et al. 2019). Therefore, the political bystander allocation game can be interpreted as a simple behavioral measure of affective polarization. We also find a positive correlation between affective polarization and minimal groupishness, albeit to a lesser extent ($\rho = 0.14, p < 0.01$, Spearman), which could be related the absence of a political context in minimal groups.

In subsequent analyses, we will use minimal groupishness for our heterogeneity analysis, investigating its moderating role on information demand and processing.

Information Demand. Our theoretical framework (Appendix C.1) is an extension of recent research on information preferences (Golman and Loewenstein 2018). Proposition 1 decomposes an individual's demand for an information source into three parts: (1) its instrumental value, (2) its content valence (the goodness or badness of having a subjective belief), and (3) its source utility (a preference for or an aversion to a specific information source *per se*, i.e., independent of its content or instrumental value). We derive the following pre-registered hypotheses from this framework.

First, we expect that participants will be willing to pay more for information from ingroup versus outgroup sources (Hypothesis 1). The intuition here is that ingroup sources may have both

[8] Also analogues to wave 1, we do not observe a statistically significant difference in the distribution of minimal group allocation choices between Democrats and Republicans ($p = 0.37$, Kolmogorov-Smirnov test) but we do observe a significant difference for political group allocation choices ($p < 0.02$, Kolmogorov-Smirnov test) with Democrats being more likely to be politically groupish (66.8% v. 52.4%).

higher content valence and source utility. In addition, they may perceive the information from ingroup sources to be more accurate, i.e., to be of higher instrumental value.

Second, we expect that any treatment that reduces source utility should reduce intergroup preferences in the demand for information. In particular, the *unlabeling* treatment in wave 1 will cause a decrease in the demand for a favored information source (Hypothesis 2).

Third, we expect that the treatment effect of unlabeling will be stronger for groupish participants (Hypothesis 3). We now proceed to test these hypotheses.

Recall participants can spend up to $3 to increase the probability of reading an article from a more favored information source. Figure 5 illustrates participants' demand for ingroup versus outgroup information relative to receiving information from a nonpartisan information source. The white (black) bar indicates their willingness to pay for a nonpartisan versus an outgroup (ingroup) information source. We show separate results for the baseline group (with labels) and the treatment group (without labels).

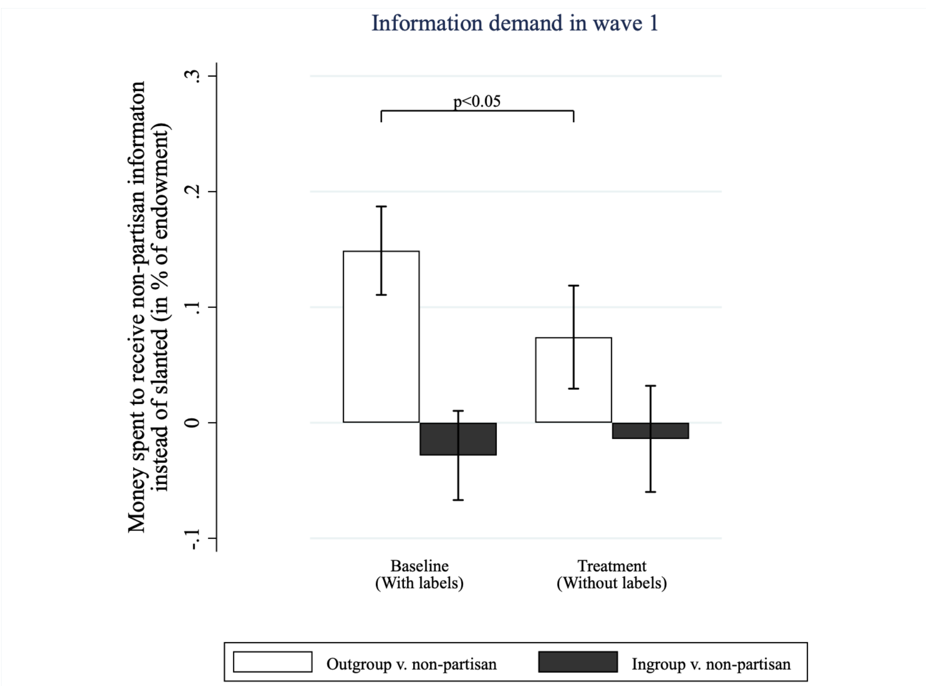

Figure 5: Information demand: The difference in participants' willingness to pay to obtain information from a non-partisan versus an outgroup or an ingroup information source in wave 1. Error bars represent 95% confidence intervals. Estimates are based on OLS regressions in Table 2.

Table 2 presents the results of six OLS specifications with robust standard errors reported in parentheses. The dependent variable measures the share of their endowment that participants spend to increase the likelihood of reading an article from a more-favored information source. Independent variables include the treatment dummy, the groupish dummy, and their interactions.

From Figure 5 and Table 2 , we see that our baseline participants spend, on average, 14.9% of their endowment to avoid information from the outgroup source in favor of the nonpartisan source ($p < 0.01$, F-test; see Table 2 column 1). On the other hand, they spend only 2.8% of their endowment to receive information from ingroup versus nonpartisan sources, suggesting near indifference between these sources ($p = 0.15$, F-test, see Table 2 column 4).

Figure 5 also shows that our treatment intervention reduces the avoidance of outgroup information by approximately 50% (from 14.9% in the baseline to 7.4% in the treatment, $p < 0.05$, F-test; Table 2 column 1), suggesting that reducing the salience of information sources is effective in reducing outgroup information avoidance. This effect is robust to the inclusion of demographic controls (Table 2 column 2). Adding the Groupish dummy and its interaction with the Treatment dummy in column (3) suggests that the treatment effect is driven by groupish participants (-0.11, $p < 0.05$). At the same time, we see no significant treatment effect on participants' demand for ingroup versus nonpartisan sources (Table 2 columns 4 and 6). Results are summarized below.

Table 2: Participants' information demand in wave 1: OLS.

DV: Share of endowment	Outgroup v. nonpartisan			Ingroup v. nonpartisan		
spent for nonpartisan source	(1)	(2)	(3)	(4)	(5)	(6)
Unlabeling Treatment (β_1)	-0.075**	-0.066**	-0.041	0.014	-0.003	0.038
	(0.030)	(0.029)	(0.035)	(0.031)	(0.030)	(0.036)
Minimally Groupish (β_2)		0.074**	0.102**		-0.038	0.006
		(0.032)	(0.042)		(0.033)	(0.044)
Treatment $\times$ Groupish (β_3)			-0.069			-0.108
			(0.063)			(0.066)
Constant	0.149***	-0.020	-0.038	-0.028	0.077	0.049
	(0.019)	(0.083)	(0.083)	(0.020)	(0.089)	(0.090)
F-test ($\beta_1 + \beta_3$)			-0.110**			-0.071
Controls	No	Yes	Yes	No	Yes	Yes
Observations	1,005	1,005	1,005	1,005	1,005	1,005
Adj. R-squared	0.005	0.055	0.055	-0.001	0.037	0.039

[a.] Controls include those in Table 1, and the topic a participant encountered in the information demand stage.

[b.] We denote significance levels by * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Result 2 (Information Avoidance and Treatment Effects of Unlabeling)**.** *(a) In the baseline condition when information sources are labeled, participants spend approximately 15% of their endow-*

ment to obtain information from a nonpartisan versus outgroup source. (b) Unlabeling information sources decreases outgroup information avoidance by 7.5 pp, or 50% relative to baseline. (c) The treatment effect is driven by groupish participants.

By Result 2(a), (b) and (c), we reject the null in favor of Hypotheses 1, 2 and 3, respectively.[9]

So far, our results indicate that outgroup information avoidance is an important factor in explaining the observed divergence in political opinions in the US. Our unlabeling treatment reduces outgroup information avoidance by 50%, demonstrating the efficacy of interventions aimed at decreasing the salience of information sources in disrupting information bubbles. We next examine whether people also differ in the way they process information, that is, the weight they accord to prior beliefs versus new information in their posterior beliefs, as well as to information coming from ingroup versus outgroup sources.

Information Processing. Our hypotheses for information processing when signals are open to interpretation, such as news articles, are based on Fryer et al. (2019). The key feature of their belief updating model is that agents first *interpret* the signal according to their prior, and then update their beliefs using the interpreted signal. We briefly summarize their model and adapt it to our experimental setting to derive hypotheses for our experiment in Appendix C.2.

This model has two implications in our context. First, it implies that priors are sticky in that they hold greater weight in an agent's posterior than do signals (Hypothesis 6). This overweighting of priors prevents agents with heterogeneous priors from converging, even if they observe the same information. Second, it implies that there is an order effect, as an early signal, which receives less interpretation, is weighted more than a later one (Hypothesis 7).[10] We now proceed to test these hypotheses using our wave 1 data.

Information Processing – Prior Beliefs. In wave 1, participants were asked to predict the unemployment rate and the public health system ranking in September 2021, conditional on the outcome of the presidential election.

Figure 6 presents participants' average prior (white bars) and posterior beliefs (black bars) elicited before and after they read two respective news articles from ingroup and outgroup sources. The beliefs are conditional on whether an ingroup or outgroup candidate wins the election. We

[9] Compared to the actual 50% reduction in outgroup information avoidance, SSPP experts predict a 100% reduction, more extreme than the actual observation ($p < 0.001$). 85% of the SSPP experts correctly predicted that groupish participants would respond more strongly to a reduction in source utility.

[10] Note that Hypotheses 6 and 7 are not pre-registered.

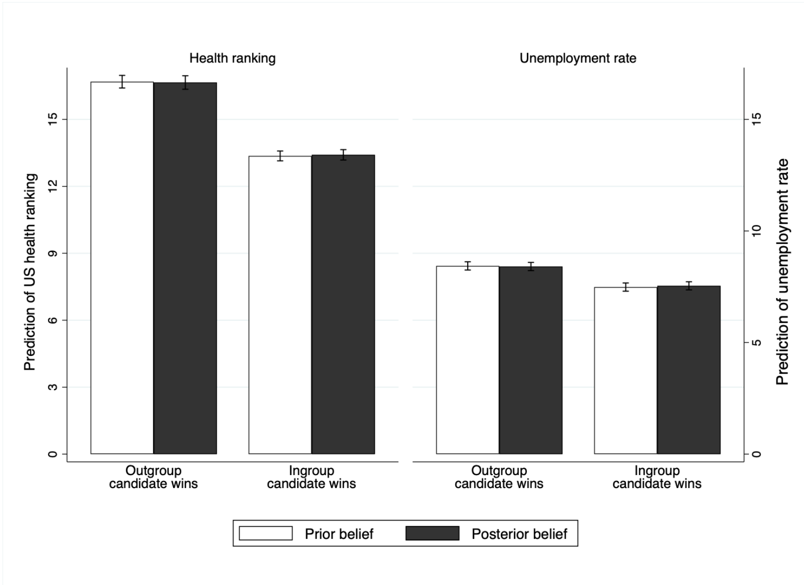

Figure 6: Wave 1 baseline prior and posterior beliefs about the US public health system ranking (left panel) and unemployment rate (right panel) 10 months after the 2020 election. Error bars represent 95% confidence intervals.

observe notable partisan favorable gaps in prior beliefs. 63.3% of the participants predict a greater improvement (lower decline) for both statistics should their ingroup candidate win.[11]

Table 3: Groupishness in prior beliefs in both waves: OLS.

DV: Strict partisan favorable gap in priors	Wave 1 (1)	Wave 1 (2)	Wave 2 (3)	Wave 2 (4)
Minimally groupish	0.059*	0.063**	0.186***	0.183***
	(0.032)	(0.032)	(0.037)	(0.038)
Constant	0.400***	0.170**	0.289***	0.268***
	(0.019)	(0.086)	(0.021)	(0.097)
Controls	No	Yes	No	Yes
Observations	1,005	1,005	740	740
Adj. R-squared	0.002	0.063	0.033	0.067

a. The controls are identical to those included in Table 1.

b. Significant at * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

c. Results are robust in logit specifications.

To investigate whether groupish participants differ from their universalist counterparts in their

[11] The mean SSPP expert forecast is 62.6% (stdev 0.256), which is remarkably accurate ($p = 0.23$). As an additional benchmark, of the 31 experts on the US Economic Experts Panel who were asked in October 2020 to predict the December 2020 unemployment rate, 10% predict the range correctly (IGM COVID-19 Economic Outlook Survey Series). In comparison, 12% of our participants correctly predicted the September 2021 unemployment rate back in October 2020, whereas 7.5% correctly predicted the US public health system ranking.

prior beliefs, Table 3 presents four OLS specifications with robust standard errors in parentheses in waves 1 (columns 1 and 2) and 2 (columns 3 and 4). The dependent variable measures participants' likelihood to exhibit a strict partisan gap in their priors. We find that groupish participants are 6.3 pp, or 37% more likely to exhibit a partisan gap in their prior beliefs than their universalist counterparts in wave 1 ($p < 0.05$, column 2 in Table 3).[12]

Information Processing – Posterior Beliefs. Participants were presented with summaries of left-leaning and right-leaning news articles in random order. After reading each summary, they answered comprehension questions, rated the perceived quality of the information, and updated their predictions on the unemployment rate and the public health system ranking 10 months later.[13]

In Figure 6, we see that our initial partisan gap in predictions persists, even after receiving information from diverse sources. Specifically, 68.5% (64.1%) of participants adhere to their initial beliefs if the ingroup (outgroup) candidate wins ($p < 0.05, F$-test). Meanwhile, approximately equal percentages of participants update their beliefs in an optimistic (15.3% for ingroup win, 17.6% for outgroup win; $p = 0.25, F$-test) or pessimistic (16.2% for ingroup win, 18.2% for outgroup win; $p = 0.35, F$-test) direction. This heterogeneity in updating enables us to test Hypotheses 6 and 7.

To test the two hypotheses, we directly estimate a linear model based on Eq. (6) in Appendix C.2, suggested by Fryer et al. (2019), to decompose participants' posterior beliefs into their prior beliefs and signals as follows:

$$\text{(Posterior Belief)}_i = \alpha + \beta_0 \times \text{(Prior Belief)}_i + \beta_1 \times \text{(Ingroup Article First)}_i + \varepsilon_i, \quad (1)$$

where i indexes individuals, β_0 corresponds to the weight on the prior, α (β_1) corresponds to the weight on signals when the first signal comes from an outgroup (ingroup) source.

Table 4 reports OLS estimations of Eq. (1) in the unemployment domain (columns 1 and 2) and public health domain (columns 3-5). Results suggest that participants' prior beliefs significantly influence their posteriors across the domains ($p < 0.01$ in columns 1-5). Furthermore, participants' prior beliefs account for at least 62% and up to 80% of their posterior beliefs. This result supports Hypothesis 6.

[12] SSPP experts correctly predict the direction of the groupish participants' prior beliefs.

[13] On average, participants answered 1.5 out of the two review questions correctly, with 64.5% answering both questions correctly. Notably, participants are 3.2 pp more likely to answer the review questions correctly for the article from the ingroup source (77.4% v. 74.2%, $p < 0.04$, χ^2-test).

Table 4: Information processing in wave 1: Estimating posterior from Eq. (1)

DV: Posterior	Unemployment Rate		Public Health System Ranking		
	Outgr. wins	Ingr. wins	Outgr. wins	Ingr. wins	Ingr. wins
	(1)	(2)	(3)	(4)	(5)
Prior	0.618***	0.649***	0.801***	0.689***	0.599***
	(0.059)	(0.061)	(0.043)	(0.052)	(0.059)
Ingroup article first	-0.013	0.117	-0.076	-0.313**	-0.278*
	(0.174)	(0.168)	(0.170)	(0.152)	(0.148)
Constant	3.229***	2.613***	3.253***	4.380***	5.376***
	(0.497)	(0.452)	(0.712)	(0.724)	(0.945)
Controls	No	No	No	No	Yes
Observations	594	594	606	606	606
Adj. R-squared	0.386	0.426	0.559	0.422	0.453

[a] In addition to the controls included in Table 1, we control for the topic encountered in the information processing stage.

[b] We report robust standard errors in parentheses. We denote significance levels by * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

The regression results in Table 4 further suggests that the influence of a signal on a participant's posterior beliefs, albeit small, depends on the order in which it is presented. However, this effect only holds in the US public health system ranking domain (-0.313, $p < 0.05$, column 4), and becomes marginally significant when we control for demographics (-0.278, $p < 0.10$, column 5). These findings offer indicative evidence supporting Hypothesis 7.

To investigate factors driving the overweighting of priors, we use alternative OLS specifications presented in Table 5. Specifically, columns (1) and (2) regress the likelihood of exhibiting a partisan-favorable gap in posterior beliefs on the presence of a similar gap in prior beliefs, the treatment dummy, the groupishness dummy, and their interactions. Columns (2) and (3) additionally incorporate the perceived quality of the news articles. The results confirm the significant stickiness in participants' priors, with a majority maintaining partisan gaps even after reading articles from opposing sources. The inclination to exhibit a partisan gap in posterior beliefs is positively influenced by the perceived quality of ingroup information (+0.026, $p < 0.01$, F-test) and negatively by the perceived quality of outgroup information (-0.028, $p < 0.01$, F-test). Importantly, participants' perceptions of information quality themselves are shaped by their prior beliefs. Those with partisan gaps in prior beliefs rate ingroup information as more reliable (0.562, $p < 0.01$, F-test, column 4) and outgroup information as less reliable (-0.465, $p < 0.01$, F-test, column 5).

Table 5: Participants' information processing in wave 1: OLS.

Dependent Variable:	Posterior partisan gap (binary)			Perceived article quality	
				Ingroup	Outgroup
	(1)	(2)	(3)	(4)	(5)
Prior partisan gap (binary)	0.623***	0.597***	0.596***	0.562***	-0.465***
	(0.028)	(0.030)	(0.030)	(0.095)	(0.094)
Unlabeling Treatment		0.028	-0.005	-0.074	0.095
		(0.024)	(0.030)	(0.089)	(0.088)
Ingroup article quality (1-7)		0.026***	0.027***		
		(0.008)	(0.008)		
Outgroup article quality (1-7)		-0.027***	-0.027***		
		(0.009)	(0.009)		
Minimal groupish			-0.045		
			(0.031)		
Treatment × Minimal groupish			0.090*		
			(0.050)		
Constant	0.134**	0.140*	0.167**	2.867***	3.526***
	(0.067)	(0.076)	(0.078)	(0.272)	(0.258)
Controls	Yes	Yes	Yes	Yes	Yes
Observations	1,005	1,005	1,005	1,005	1,005
Adj. R-squared	0.440	0.449	0.450	0.143	0.288

[a] Controls include those in Table 1, and the topic a participant encountered in the information demand stage.

[b] We denote significance levels by * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Using either the continuous outcome variables in Table 4 or the binary partisan gap variables in Table 5 yields the robust finding that participants' prior beliefs significantly influence their posteriors across the domains. The latter further corroborates Fryer et al's theoretical model by showing that priors also indirectly influence belief updating by shaping perceived signal quality. Results are summarized below.

Result 3 (Biased Information Processing: Sticky Priors)**.** *Participants' prior beliefs significantly influence their posteriors across the domains. The prior beliefs influence information processing by shaping the perception of signal quality.*

Result 3 provides support for Hypothesis 6. Consistent with Result 3, Liu and Zhang (2023) also find that prior beliefs are sticky and that the opportunity to read additional arguments does not prompt participants to adjust their attitudes shaped by the initial narrative. Relatedly, Levy (2021) reports that exposure to counter-attitudinal news decreases negative attitudes toward the opposing political party, yet it does not affect political opinions.

4 Wave 2: Underlying Mechanisms

While the timing of wave 1 creates a realistic demand for policy-relevant information, each news article summary varies not only in its source utility but potentially also in its content valence and perceived instrumental value. To uncover the underlying mechanisms in a clean way, we use a simpler design in wave 2. Specifically, we use variants of the guess-the-cup game, which removes content valence while retaining source utility. We further implement a debiasing treatment that equalizes the instrumental value of information sources, which enables us to identify the underlying mechanisms more precisely.

4.1 Wave 2 Experimental Design

We implemented wave 2 of our online experiments at the end of January 2021, after Joe Biden's inauguration. We conduct this wave to assess intergroup preferences in belief formation in a simple context where we control prior beliefs. We also pay participants for their accuracy in election prediction at this time.[14] Of the 1,005 participants of the first wave, 740 participated in this second wave, which yields a retention rate of 74%. Despite attrition, wave 2 participants continue to be representative of the US population in terms of gender, age, race and ethnicity (except the "Other" category), and urban/rural distribution (Appendix Table B.1). Similar to wave 1, Democrats (0.52) are over-represented in wave 2 while Republicans (0.23) and Independents (0.25) are under-represented. Figure 7 presents a schematic diagram of wave 2.

Stage 1. As in wave 1, we again randomly assign participants to one of two minimal groups (triangle or circle) and repeat the bystander allocation game with the minimal groups. Participants then answer the same survey questions designed to measure their degree of political party affiliation.

Stage 2. Given that Joe Biden was inaugurated as the President of the United States in January 2021, we next give participants the opportunity to update their predictions about the US unemploy-

[14] Of our wave 1 participants, 63% predicted the election outcome correctly.

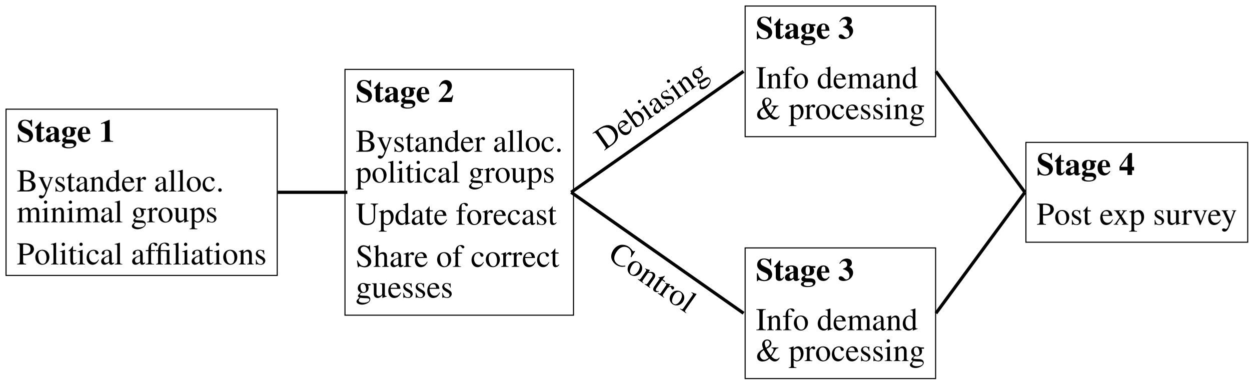

Figure 7: Wave 2 of the experiment (29 January - 8 February 2021): alloc. is short for allocation game.

ment rate and public health system ranking in September 2021 for the Biden-win case. To reduce noise originating from imperfect memory, we remind participants of their updated policy predictions conditional on Biden winning at the end of wave 1.[15] Subsequently, participants again play the bystander allocation game with participants of the same or different political party affiliations or leanings. This second measurement of participants' allocation decisions enables us to analyze the stability of their previously-revealed level of intergroup preferences.

At the end of this stage, we remind participants of the "guess-the-cup" game they played in wave 1 and inform them that, across all participants, 67% guessed their cup correctly. After anchoring their belief on the guess accuracy, we ask participants to estimate the share of correct guesses among Democrats and Republicans, respectively.

Intervention – Debiasing. As in wave 1, we include a between-subject treatment variation in wave 2. Specifically, at the end of stage 2, participants in the treatment condition learn that the share of correct guesses is 67% for both Democrats and Republicans, whereas those in the control condition do not receive any information on the share of correct guesses. Put differently, treated participants learn that the observed guesses of Democrats and Republicans are equally accurate. Thus, this intervention equalizes the perceived instrumental value of ingroup and outgroup sources.

Stage 3. In this stage, we explore participants' information demand and processing behavior in various "guess-the-cup" games. The order of the information demand and processing substages is randomized at the individual level.

The structure of the *information demand* substage in the second wave of our experiment closely mirrors that of its counterpart in wave 1. Specifically, it comprises two independent rounds of a

[15]70% of participants in wave 2 did not change their prediction for the health system ranking, whereas 59.3% did not change their prediction for the unemployment rate. 44.7% did not change either prediction. On average, participants' predictions became slightly more optimistic.

"guess-the-cup" game, i.e., we randomly draw a new cup for each round with the composition of marbles in each cup being identical to that in wave 1. In both rounds, participants observe either a randomly-drawn marble or the wave-1 guess of another participant whose marble had been drawn from the same cup (with replacement), each with a 50% probability. Using a slider on a computer screen, participants have the option to change their marble versus guess observation probability, with each 10 percentage points in either direction costing 10 of their 40-cent endowment. Participants are shown both the guess and political party affiliation of the wave-1 guess individual. In a random order, the other participant is either a Democrat or a Republican. Based on the chosen probabilities, we randomly determine whether participants observed a marble or guess. Subjects earn $1 for a correct guess, and zero otherwise. At the end of the experiment, we randomly select one of the two rounds for the final payment.

We next have participants observe the independent wave-1 guesses of two other participants, each of whom had observed a randomly-drawn marble from the same colored cup (with replacement). Each of three rounds includes two groups comprised of two other participants. Each participant observes the guesses from one of these two groups. The default probability of observing a given group's guesses is 50%. Again, participants observe both the guesses and the guesser political party affiliations and are able to change their observation probability, with each 10 percentage point change costing 10 cents of their 40-cent endowment. The composition of the two groups varies across the three rounds: (i) two Democrats v. two Republicans; (ii) two Democrats v. one Democrat and one Republican, and (iii) two Republicans v. one Democrat and one Republican. We randomize the order of the different scenarios. Based on the chosen probabilities, we randomly draw one of the two groups and show the corresponding guesses to the given participant. Participants earn $1 for a correct guess, and zero otherwise. At the end of the experiment, we randomly select one of the three rounds in this stage for payment.

In all of these rounds, we use the endowment share a participant spends to increase the probability of seeing the guesses of participants with the same political affiliation to indicate that participant's intergroup preferences in information demand.

The structure of the *information processing* substage in the second wave of our experiment similarly mirrors that of its counterpart in wave 1 and is related to the design by Agranov, Lopez-Moctezuma, Strack and Tamuz (2022). For each participant in each of six independent rounds, we randomly draw a new cup, and they observe a randomly drawn marble from that cup along with the independent wave-1 guesses of two other participants whose marble had been drawn from the same cup (with replacement). In three of the six rounds, both of the observed guesses contradict their private signal (the observed marble). In different rounds, the contradicting guesses come from two Democrats, two Republicans, or two participants from unknown party affiliations, in random order. In the other three rounds, only one of the observed guesses contradicts their private signal. In different rounds, the contradicting guess comes from a Democrat, a Republican, or a participant with unknown party affiliation. This structure allows us to identify when participants are willing to abandon their private signal conditional on the respective party affiliations of the two individuals whose guesses they observe. One of the six rounds is randomly selected for payment at the end of the experiment. Participants earn $1 for a correct guess and zero otherwise.

Stage 4. In the final stage of the experiment, participants complete a questionnaire designed to elicit their attitude toward the then-ongoing debates about Covid-19 vaccinations, beliefs about the legitimacy of the 2020 presidential election, and perceptions about the social status of Democrats and Republicans (see Appendix A.2.6).

On average, it took 15 minutes for our participants to finish wave 2. The average earning in this wave is $16.81.

4.2 Wave 2 Results

Prior Beliefs Before our intervention, we inform participants that, across all participants, 67% correctly guessed their cup in wave 1. We then ask participants to separately predict the proportion of wave-1 Democrats and Republicans who guessed the cup correctly. On average, participants believe that ingroup members' guesses are 5.5 percentage points (pp) more likely to be correct than outgroup members' guesses.

As in wave 1, groupish participants are more biased in their prior beliefs. More specifically, in Table 3 in Section 3.2, columns (3) and (4) present two OLS specifications with robust standard errors reported in parentheses. The dependent variable measures the likelihood that participants exhibit a strictly partisan favorable gap in their estimates of others' guess accuracy, and hence the perceived difference in the instrumental value of different information sources. Groupish participants are 18 pp more likely to exhibit such partisan-favorable bias as they overestimate (underestimate) the guess accuracy of their ingroup (outgroup) members by a wider margin (7.31 pp v. 4.69 pp) compared to their universalist counterparts. This indicates that participants consider their ingroup members' guesses to be more accurate in this neutral belief updating task, where Democrats and Republicans are actually equally accurate (67%).

Information Demand. In the following, we present the results from the two rounds where participants may choose between drawing a marble (a private signal) versus observing another participant's guess knowing the political party affiliation of the guesser.[16]

Figure 8 presents the proportion of their endowment participants are willing to spend to observe a marble instead of another participant's guess under the baseline and the debiasing treatment condition. From Figure 8, we see that participants generally prefer to observe a signal themselves instead of observing another participant's guess. On average, they spend 11.5% of their endowment to increase their likelihood of drawing a marble.[17] More importantly, and in line with our findings for wave 1, this preference differs according to the group membership of the guesser: if the guess comes from an outgroup (ingroup) member participants spend, on average, 14.9% (8.2%) of their endowment to see a marble. The difference is significant ($p < 0.05$, F-test; see Table 6 columns 1 and 4). These findings provide support for the persistence of outgroup information avoidance in a domain absent of content valence.

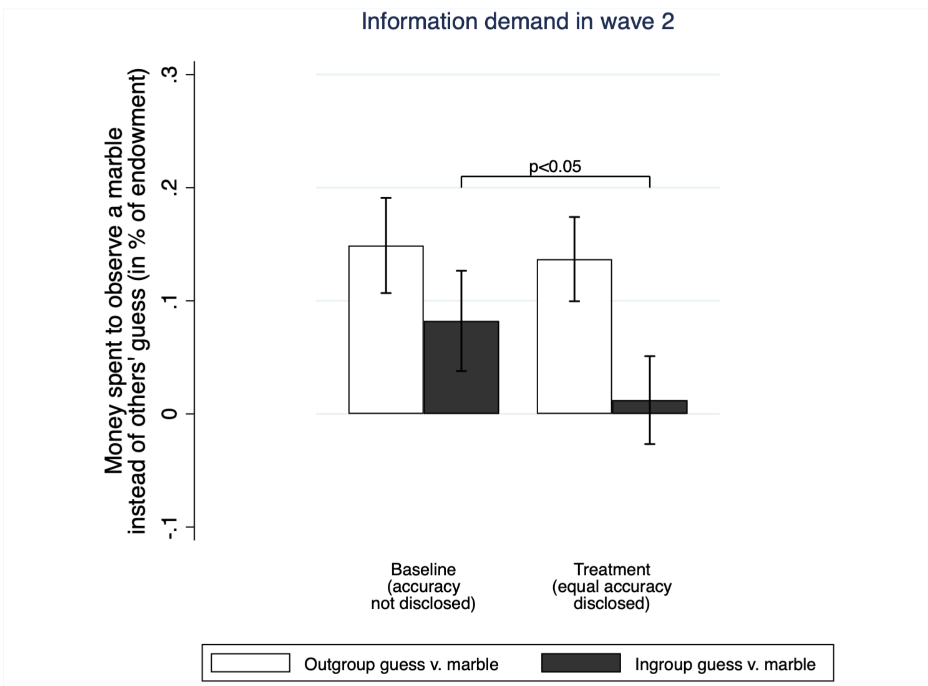

Figure 8: Information demand: Participants' willingness to pay to observe a marble versus an outgroup or an ingroup member's guess in wave 2. Error bars represent 95% confidence intervals. Estimates are based on OLS regressions in Table 6.

[16]Results from the rounds, in which participants do not themselves observe a marble but may choose between guesses from different groups, are provided in our SSRN working paper (Bauer et al. 2023). The results of these tasks mirror the results for information demand reported in the main text, providing additional evidence of avoidance of outgroup information.

[17]This is consistent with findings by Conlon et al. (2022) in whose study participants observe signals of others, however, not guesses.

Next, we consider whether our observed outgroup information avoidance reflects the instrumental value of information.

In our debiasing treatment, we truthfully communicated to half of the participants that the accuracy of Democrats and Republicans is identical (67%), and equal to the diagnosticity of a marble (2/3), thus removing any instrumental value for demanding information from a particular source. The results in Table 6 show that informing participants of the identical accuracy has no effect on their outgroup source avoidance. Compared to baseline participants who spend 14.9% of their endowment, their counterparts in the treatment condition spend on average 13.7% of their endowment to avoid information from the outgroup ($p = 0.67$, F-test, column 1). However, the treatment significantly reduces participants' willingness to pay for observing a marble instead of an ingroup guess by 7 percentage points (8.2% v. 1.2%, $p < 0.05$, F-test, column 4).

Table 6: Participants' information demand in wave 2: OLS.

DV: Share of endowment	Outgroup guess v. a marble			Ingroup guess v. a marble		
spent to observe a marble	(1)	(2)	(3)	(4)	(5)	(6)
Debiasing Treatment (β_1)	-0.012	-0.023	-0.022	-0.070**	-0.093**	-0.091***
	(0.029)	(0.034)	(0.034)	(0.030)	(0.034)	(0.035)
Minimally Groupish (β_2)		0.097**	0.104**		-0.038	-0.025
		(0.031)	(0.04)		(0.05)	(0.051)
Treatment × Groupish (β_3)		0.026	0.024		0.066	0.077
		(0.061)	(0.061)		(0.067)	(0.067)
Constant	0.149***	0.116 ***	0.087	0.082***	0.095***	-0.032
	(0.021)	(0.026)	(0.085)	(0.023)	(0.026)	(0.087)
$\beta_1 + \beta_3$		0.003	0.002		-0.027	-0.014
p-value, F-test		0.95	0.96		0.64	0.80
Controls	No	No	Yes	No	No	Yes
Observations	740	740	740	740	740	740
Adj. R-squared	0.005	0.055	0.055	-0.001	0.037	0.039

[a.] Controls include those in Table 1, and the topic a participant encountered in the information demand stage.

[b.] We denote significance levels by * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

We now examine the heterogeneity in information demand. Results in columns (2), (3), (5) and (6) in Table 6 show treatment heterogeneities conditional on groupishness in information demand. We see that universalist participants in the baseline spend 11.6% of their endowment to avoid the guess of an outgroup participant, compared to 21.3% for groupish participants ($p < 0.05$, column 2) indicating that outgroup information avoidance is consistently greater for groupish participants in the baseline condition of both waves. Turning to our treatment results, we find that the only effect of our debiasing treatment is a reduction in the amount spent on drawing a marble oneself versus seeing an ingroup guess. The results in column (5) show that our observed reduction is significant only for universalist participants (-0.093, $p < 0.05$, a decrease from 9.5% to 0.2% compared to a decrease from 5.7% to 3% for the groupish participants).

Result 4 (Information Avoidance and Treatment Effects of Debiasing)**.** *(a) In the baseline condition, when deciding between observing a marble directly or the guess of another participant, participants spend 6.7 pp or 87% more of their endowment to observe a marble when the guess comes from an outgroup vs. an ingroup participant. (b) The debiasing treatment decreases the demand for observing a marble rather than an ingroup guess by 7 pp, or 85% over the baseline, while it has no effect on outgroup information avoidance. (c) The treatment effect is driven by universalists.*

Removing differences in beliefs about the instrumental value of information sources (debiasing treatment) only reduces universalists' preference for drawing a marble instead of seeing an ingroup guess (-9.3 pp). Groupish types do not respond significantly.[18] Result 4 thus suggests that the role of source utility in outgroup information avoidance is persistent and independent of the instrumental value or content valence of the information sources. By Result 4(b) and 4(c), we reject the null in favor of Hypotheses 4 and 5, respectively.

Synthesizing our results on information demand from waves 1 and 2, we find consistent evidence that participants' information demand is influenced by the mere group affiliation of information sources, a phenomenon that is independent of content valence, or the instrumental value of information sources. Interestingly, groupish participants respond strongly to the reduction of salience of information sources (the unlabeling treatment in wave 1), whereas universalists respond to the equalization of instrumental values (the debiasing treatment in wave 2).

Information Processing. Unlike the news article summaries in wave 1, the signals in wave 2 are not open to interpretation, nor is there any content valence. This simple setting thus enables us to estimate the effects of group identity on participants' belief updating.

[18] Bursztyn and Yang (2022) show in a meta-study that experimental treatments intended to debias beliefs about others are generally effective and can lead to behavioral change. Our result shows that it only works for universalists in this setting.

To derive our hypotheses related to how participants process information, i.e., update their subjective beliefs after observing signals, we add group identity to a Bayesian updating model in which signals are discrete and unambiguous (Appendix C.2).

Let β_S, β_I and β_O be the weight an agent attaches to a signal from their own observation (self), an ingroup, and an outgroup source, respectively. Multiple signals from the same source are pooled by the agent (Observation 1 in Appendix C.2). Based on our extension of the classic empirical Bayesian updating model (Grether 1980), we state the following hypotheses regarding information processing. First, we hypothesize that the information from the ingroup will resonate more with an individual, and thus the individual accords this information more weight than the corresponding outgroup information (Malmendier and Veldkamp 2022), in other words, $\beta_I > \beta_O$ (Hypothesis 8).

Second, we expect that our experimental treatment that targets perceived source accuracy (*debiasing* treatment) will reduce the level of observed ingroup bias in information processing. The reason for this expectation is that the treatment decreases the participants' differential beliefs related to the instrumental value of the signals, in other words, $\beta_I = \beta_O$ (Hypothesis 9).

In our context, when each participant is presented with a private signal and two guesses, a Bayesian decision maker should follow the majority rule (Observation 1 in Appendix C.2).[19] However, we find that only 63.3% of the decisions in our baseline condition follow this majority rule. To estimate the effects of group identity on belief updating, we employ the following estimation equation based on Eq. (9) in Appendix C.2:

$$\text{Posterior}_{igj} = \alpha + \beta_S \times \text{Self}_{ij} + \beta_I \times N^I_{ij} + \beta_O \times N^O_{ij} + \gamma X_{igj} + \varepsilon_{igj}, \tag{2}$$

where Posterior_{igj} represents individual i in group g guessing cup j, Self_{ij} indicates individual i's own signal is of color j, N^I_{ij} (N^O_{ij}) is the number of i's ingroup (outgroup) members guessing cup j, X_{igj} is the set of demographic controls, and ε_{igj} is an error term.

Table 7 presents the results from six OLS specifications estimating Eq. (2), with robust standard errors clustered at the individual level. The dependent variable is whether a participant guesses the green cup. Independent variables include whether a participant observes a green marble, the number of green cup guesses from ingroup members, the number of green cup guesses from outgroup members, a debiasing treatment dummy, and interaction terms. From Table 7, we see that participants put significantly more weight on the guesses of ingroup members than those of outgroup

[19] Under the assumption that everyone follows Bayes' Rule, seeing a guess is as good as seeing another independently drawn marble.

members ($\beta_I > \beta_O$, $p < 0.01$; F-test, columns 1 and 2), and more weight on their own signal than the corresponding guesses of ingroup members ($\beta_S > \beta_I$, $p < 0.05$; F-test, columns 1 and 2), regardless of the treatment condition ($p > 0.10$ for the Treatment dummy and its interactions, column 3).

Table 7: Information processing (Wave 2): OLS estimates of Eq. (2).

DV: Guessing green	Baseline	Treatment	Pooled
	(1)	(2)	(3)
Self green marble (β_S)	0.282***	0.250***	0.283***
	(0.025)	(0.027)	(0.025)
Ingroup # of green guesses (β_I)	0.201***	0.180***	0.202***
	(0.018)	(0.018)	(0.018)
Outgroup # of green guesses (β_O)	0.151***	0.110***	0.151***
	(0.019)	(0.018)	(0.018)
Debiasing Treatment			0.044
			(0.031)
Treatment×Self Green marble			-0.033
			(0.037)
Treatment×Ingroup # of green guesses			-0.022
			(0.025)
Treatment×Outgroup # of green guesses			-0.041
			(0.026)
$\beta_I - \beta_O > 0$	0.050***	0.070***	0.051***
p-value, F-test	0.006	0.005	0.000
$\beta_S - \beta_I > 0$	0.081***	0.070**	0.081***
p-value, F-test	0.004	0.02	0.000
Controls	Yes	Yes	Yes
Observations	2,226	2,214	4,440
Adj. R-squared	0.082	0.063	0.073

[a.] Controls include those in Table 1, and the topic a participant encountered in the information demand stage.

[b.] Robust standard errors in parentheses are clustered at the individual level. We denote significance levels by * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Result 5 (Biased Information Processing: Guesses)**.** *In the baseline condition, participants place 5 pp, or 33% more weight on ingroup than outgroup guesses. They also place 8 pp, or 40% more weight on their own signal than ingroup guesses. The debiasing treatment has no significant effect, indicating that these differential weights are not associated with any perceived instrumental value of the information, pointing to the role of source utility in information processing.*

By Result 5, we reject the null in favor of Hypothesis 8 ($\beta_I > \beta_O$). We further reject Hypothesis

9 that the debiasing treatment would equalize the weights on ingroup and outgroup guesses ($\beta_I = \beta_O$). While groupish participants have higher β_S, β_I and β_O than their universalist counterparts, none of the differences is statistically significant. Furthermore, we do not find any heterogeneous treatment effect between groupish and universalist participants in their response to the debiasing treatment in wave 2 (see Appendix Table B.3).

Synthesizing the results from waves 1 and 2, we find evidence that participants' information processing is influenced by the group affiliation of information sources, albeit in a more nuanced way than observed for information demand. In wave 1, participants' evaluation of signal quality is shaped by group affiliations, which in turn influences how information is processed - a pattern consistent with the theoretical framework proposed by Fryer et al. (2019) that participants overweigh their priors and an early signal carries more weight than a later one. In wave 2, conducted in a simple setting that allows for precise estimates of the weight participants place on signals, we observe that outgroup signals are significantly discounted, even in the debiasing treatment, driven solely by their group association. Notably, neither treatment intervention alters participants' information processing patterns. Furthermore, we find no substantial treatment heterogeneity with respect to groupishness, suggesting that the influence of group identity on information processing is pervasive across both groupish and universalist participants.

5 Discussion

In two waves of online experiments, we examine how group identity influences the full cycle of belief formation, from prior beliefs to the demand for and processing of new information, when participants are incentivized to predict policy outcomes (wave 1) and the color of an urn (wave 2). We document three sets of results.

Group identity influences the entire cycle of the belief formation process. First, our analyses show that participants hold partisan-favoring prior beliefs, are willing to pay to avoid outgroup information in favor of nonpartisan information sources, and that they give less weight to outgroup-sourced information when forming their posterior beliefs. Interestingly, we also find that participants are nearly indifferent between ingroup and nonpartisan information sources.

Source utility serves as an important underlying mechanism. Second, through two interventions, we identify *source utility* as an additional underlying mechanism driving group effects in belief formation. In wave 1, the unlabeling intervention reduces outgroup information avoidance by 50%, although it cannot entirely eliminate it. Meanwhile, information processing is not affected at all by the intervention. In wave 2, outgroup information avoidance and the underweighting of outgroup information during information processing remain even when differences in the *content valence* of different information sources as well as the *instrumental value* of different information sources are eliminated through the debiasing intervention.

Groupish and universalist participants exhibit systematic differences across their belief formation cycle. Investigating individual heterogeneity, we find systematic differences between groupish and universalist participants across their belief formation process. Specifically, we find that groupish individuals: (i) hold more pronounced partisan-favoring prior beliefs, (ii) exhibit greater outgroup information avoidance, and (iii) react strongly to a reduction of source utility (unlabeling treatment), but not to a removal of the differential instrumental values of information sources (debiasing treatment) in information demand. By contrast, universalists respond to changes in instrumental values and drive the debiasing treatment effect, suggesting that universalists' outgroup information avoidance is an artifact of inaccurate beliefs about information quality, whereas groupish types' information avoidance is a more affective inclination that providing correct facts cannot resolve. This heterogeneity suggests that policies designed to reduce polarization need to consider the composition of types in its target audience.

Our study lends insight to foundational research on group identity by illustrating (1) the extent to which group identity affects each component of the belief formation process, from prior beliefs, to the demand for information, and the processing of new information; and (2) the robust heterogeneity in individual responses to interventions designed to alleviate political polarization. On the methodology front, we demonstrate that typing individuals to groupish versus universalist based on their behavior in the bystander allocation game in the minimal group context predicts characteristics of their belief formation as well as their political polarization. These findings suggest that the tendency to understand the world through ingroup-outgroup distinctions might be a deeply-rooted aspect of personality and constitutes a fundamental individual trait (Kranton and Sanders 2017, Hett, Mechtel and Kröll 2020). Our methodology could inspire future work exploring the nature of this trait, both empirically and theoretically.

On a broader level, our results provide guidance for designing policies to alleviate political polarization documented across the world, by suggesting that such policies take into account the group-identity roots of belief formation. In light of our wave 1 results, policymakers could reduce the salience of group and partisan identity associated with a policy to decrease outgroup information avoidance and increase policy uptake. For example, Bursztyn, Kolstad, Rao, Tebaldi and Yuchtman (2024) show that the effectiveness of the Affordable Care Act (ACA), aka Obama Care, is diminished by "political adverse selection" due to its political association, as Republicans enroll at lower rates than Democrats and Independents. In a related field experiment, Lerman, Sadin and Tratchman (2017) show that emphasizing the market versus government-based aspect of the ACA substantially increased the insurance uptake by Republicans, an intervention effectively reducing the source utility of the ACA. Policymakers might also reduce political group salience by providing opportunities for individuals to humanize those outside of their group through social interaction (Bruneau, Cikara and Saxe 2015). Furthermore, intergroup preferences in belief formation might also apply directly to politicians and decision-makers themselves: Using a large-scale RCT in Spain, Garcia-Hombrados et al. (2024) demonstrate that policymakers are significantly more likely to adopt a policy when the information comes from ideologically aligned sources. In general, the heterogeneous patterns we find for groupish and universalist participants indicate that policies aiming at breaking information bubbles should take into account the composition of the target audience in their groupishness.

Lastly, an unexpected finding of our experiment is that participants do not favor ingroup over nonpartisan information sources. People from both ends of the political spectrum treat nonpartisan information sources, such as *Nature* and *The Economist*, the same as their ingroup information sources.

Appendix

This section contains three appendices. Appendix A contains experimental instructions and screen shots. Appendix B contains summary statistics and additional analyses. Appendix C contains our theoretical framework and hypotheses.

A Experimental instructions and screen shots

Note that texts in square brackets are included for organization purposes. They do not appear on the participants' interface.

A.1 Instructions Wave 1

A.1.1 Preamble/Consent Form

Consent to Participate in Research Study

Title of the Project: Decision-making experiment
Principal Investigators: Kevin Bauer, Yan Chen, Florian Hett and Michael Kosfeld

Invitation to Participate in a Research Study: Researchers from Goethe University Frankfurt, Johannes Gutenberg University Mainz, Leibniz Research Institute SAFE, and the University of Michigan invite you to be part of an online research study to better understand how different types of information affects our judgment and decisions. The study is funded by the three universities and SAFE.

Description of Your Involvement: If you agree to be part of the research study, you will be prompted to participate in a short demographic survey and a sequence of research games, and respond to a short questionnaire. The total time taken today will be about 30 minutes. In addition, we will send you a follow-up survey in about three months time if you agree.

Benefits of Participation: During the experiment, you will have the opportunity to earn an income that will be paid to you after the experiment. The amount of income you earn depends on your decisions and the decisions of other participants in the experiment. In addition to you directly benefit from being in this study, others may benefit because the results from the study may inform public policy.

Risks and Discomforts of Participation: Some of the survey questions may touch on sensitive topics and cause you discomfort. However, we stress that your participation is entirely voluntary. You may choose at any time to abandon the study or to skip a particular question.

Confidentiality: The results of this study will be published. We will not include any information that would identify you. Your privacy will be protected and your research records will be confidential. It is possible that other people may need to see the information you give us as part of the study, such as organizations responsible for making sure the research is done safely and properly like the University of Michigan.

Storage and Future Use of Data: We will store your answers for possible use in future research studies, for a period of up to ten years. Your study answers will be secured and stored at the University of Michigan School of Information. Only the researchers involved in this study will have access to your research files and data. Research data may be shared with other investigators but will never contain any information that could identify you.

Voluntary Nature of the Study: Participating in this study is completely voluntary. Even if you decide to participate now, you may change your mind and stop at any time. You can also skip any question you do not want to answer. Your data will not be used if you abandon the survey before reaching the end.

Contact Information for the Study Team: If you have any questions about this study, click this link to be taken to a question form. A member of the research team will see your question and reply within two days. With the answer, you will receive a new link allowing you to participate in the study if you are interested. The University of Michigan Health Sciences and Behavioral Sciences Institutional Review Board (IRB) has determined that this research is exempt from IRB oversight.

Consent: By checking the box, I agree to participate in the study. I understand that if I complete it, I will be re-contacted for a follow-up in about three months. I also understand that my responses will be saved after the expiration of the study, for a period of up to ten years.

- (radio button) I agree to participate.
- (radio button) I agree to be re-contacted in January 2021.
- (radio button) I agree to be re-contacted in October 2021.

A.1.2 Stage 1. [Group Assignment and Bystander Allocation Games]

Thank you for participating in this experiment. The objective of this experiment is to study how people make decisions. There is no deception in this experiment - and we want you to understand everything about the procedures. The amount of money you earn will depend upon the decisions you make and on the decisions other people make. This experiment has three parts. Your total earnings will be the sum of your payoffs in each part.

1. [A bystander allocation game based on minimal groups]

Choice task in part 1. Before the experiment starts, every participant is randomly assigned to one of two groups, Triangle or Circle. Half of the participants are assigned to the Triangle group, while the other half to the Circle group.

You are a member of the Circle group. [Display of a green circle.]

Instructions part 1. In Part 1 of the experiment, you will be asked to make decisions in three scenarios. For each scenario, you will have $6. You will be asked to allocate these $6 between two other participants under three scenarios:

1. if both are from your own group (Circle group)
2. if both are from the other group (Triangle group)
3. if one is from your own group (Circle group), and one is from the other group (Triangle group).

For each scenario, you must allocate all dollars between the two participants. Allocations have to be integers. You can not allocate any dollars to yourself. Your answers will be used to determine other participants' payoffs. Similarly, your payoff will be determined by others' allocations.

[Decision Screen: Each participant makes three decisions, each involving splitting $6 between two other participants in the amount of ($1, $5), . . ., ($5, $1).]

2. A Political Identity Survey:

Part 1. Please answer the following question:

1. Do you consider yourself a(n):

   (a) Democrat (b) Republican (c) Independent (d) None of the above

2. (Conditional on choosing option a or b) Are you a strong or moderate Democrat/Republican?

   (a) Strong (b) Moderate

3. (Conditional on choosing option c or d) Do you consider yourself closer to the:

   (a) Democratic party (b) Republican party

3. A [Political] Quiz:

Part 1. Please check true or false for each statement. [Statements appear in randomized order.]

1. The vast majority (over 90%) of climate scientists believe that global warming is an established fact and that it is most likely caused by man-made emissions.

2. Michael Cohen, Donald Trump's personal lawyer, pleaded guilty to fraud and illegal campaign finance charges in August 2018.
3. 40% of firearm sales in the US occur without a background check.
4. Illegal immigrants commit violent crime at a significantly higher rate than legal American citizens.
5. Former President Obama ordered wire taps on Donald Trump's phones.
6. Millions of illegal votes were cast in the 2016 presidential election.

4. [Prior elicitation.]

Part 1. Information Now we will ask you to make a number of predictions regarding the 2020 Presidential Election and some of its consequences 12 months from now.

Please answer the following questions: We will check the official results of the election at the end of January 2021 and inform you whether your predictions were correct. For every correct prediction we will pay you $3 in January 2021.

- By January 20, 2021, who will win the 2020 presidential election? [Candidate names appear in randomized order.]
  - (radio button) Donald Trump – (radio button) Joe Biden
- By January 20, 2021, who will win the majority of the popular votes in the 2020 presidential election? [Candidate names appear in randomized order.]
  - (radio button) Joe Biden – (radio button) Donald Trump

Next, we ask you questions about the consequences of the election outcomes. In 12 months we will check the official statistics and inform you whether your predictions were correct. For each correct prediction, we will pay you $10.

Note: in subsequent parts of the experiment, you will be given opportunities to update your initial predictions on the consequences of the election, after we provide you information that might help you make a correct prediction.

- **According to the Bureau of Labor Statistics, the unemployment rate in September 2020 was 7.9%.**

  What will the unemployment rate be in September 2021 if Joe Biden (Donald Trump) wins the election in November 2020?
  - Strong increase (10 % or higher)
  - Moderate increase (Higher than or equal to 8.5 %, but less than 10 %.)
  - Stable (Higher than or equal to 7.5 %, but less than 8.5 %.)

  - Moderate decrease (Higher than or equal to 6 %, but less than 7.5%.)
  - Strong decrease (6 % or lower)

- **According to US News and World Report, Canada ranks 1st among countries with the most developed public health care systems in 2020, while the United States ranks 15th.**

  What will be the ranking of the United States in September 2021 if Joe Biden (Donald Trump) wins the election in November 2020?

  - Strong improvement (Rank 12 or better)
  - Moderate improvement (Rank 13 or 14)
  - No change (Rank 15)
  - Moderate decline (Rank 16 or 17)
  - Strong decline (Rank 18 or worse)

5. [A bystander allocation game based on political groups.] In this part of the experiment, you will be asked to make decisions in three scenarios. For each scenario, you will have $6. You will again be asked to allocate these $6 between two other participants under three scenarios.

- if both are Democrats;
- if both are Republicans;
- if one is a Democrat, and the other a Republican.

Whether someone is labeled as Democrat or Republican depends on her/his responses in the previous questionnaire:

Democrats and those closer to the Democratic party are labeled Democrats. Similarly, Republicans and those closer to the Republican party are labeled Republicans.

[Decision Screen: Each participant makes three decisions, each involving splitting $6 between two other participants in the amount of ($1, $5), . . ., ($5, $1).]

A.1.3 [Information Processing Stage]

In part 2 (3) of this experiment, you will be given two articles, each of which summarizes current events from three different news sources. You will [will not] see the names of the news sources.

The two articles contain information that might help you improve your predictions about the United States unemployment rates 12 months from now. You will be given the chance to change your initial predictions after you have read both articles.

Additionally, you will be asked to answer two multiple-choice questions, one for each article. You will earn 50 cents for each correct answer. You will also be asked to evaluate the political leaning and reliability of the articles.

- Display screen for Article A;

- Display screen for Article B;
- Updating screen:

  Now you have the option to change your initial predictions about the unemployment rates in the United States 12 months from now. Your initial predictions will be overwritten. You will not have the chance to change your predictions about the unemployment rates again. In 12 month we will check the official statistics and inform you whether your predictions were correct. For a correct prediction, we will pay you $10.

  Your initial prediction for the unemployment rate in September 2021 if Joe Biden wins the 2020 presidential election: __

  Your initial prediction for the unemployment rate in September 2021 if Donald Trump wins the 2020 presidential election: __

  Please answer the following questions: [Display the same unemployment questions again.]
- Please answer the following questions about the articles you have just read. [Evaluation screen:]
  - On a scale of -3 to +3, with negative numbers representing left leaning or liberal skew, positive numbers representing right leaning or conservative skew, and 0 representing true nonpartisan, how would you rate article A [Title]?
  - On a scale of 1 to 7, 1 being not reliable at all and 7 being very reliable, how would you rate the information in article A [Title]?
  - On a scale of -3 to +3, with negative numbers representing left leaning or liberal skew, positive numbers representing right leaning or conservative skew, and 0 representing true nonpartisan, how would you rate article B [Title]?
  - On a scale of 1 to 7, 1 being not reliable at all and 7 being very reliable, how would you rate the information in article B [Title]?
- Please answer the following MC questions about the articles you have just read. There is one question for each article. For a correct answer, you receive $0.5 paid out at the end of the experiment.. [See Appendix A8 for the articles and their corresponding multiple-choice questions.]

A.1.4 [Information Demand]

You will see the titles of two pairs of articles. Each article summarizes current events from three different news sources.

Under each pair of articles, a slider indicates the likelihood that you will be randomly given one of the articles to read. The default probability is 0.5. That is, each article is equally likely to be selected. You have a $3 budget which you can use to change the probabilities. Moving the slider to the left increases the likelihood that the article on the left is chosen, and vice verse. It costs $1 to change the probability by 10 percentage points. The amount of budget you do not spend will be paid out to you at the end of the experiment.

One of the two decisions you make will be randomly drawn and implemented and an article will be drawn with your chosen probabilities. You will then be able to read that article.

After you finish reading the article, you will be given the opportunity to change your initial predictions.

Additionally, you will be asked to answer a multiple-choice question. You will earn $0.5 for each correct answer. You will also be asked to evaluate the political leaning and reliability of the articles.

- Decision screen: Display titles of article pairs in random order: left and nonpartisan; right and nonpartisan; display sources if in treatment; slider
- Reading screen
- Updating screen (if the chosen article is about the economy):

  Now you have the option to change your initial predictions about the unemployment rates in the United States 12 months from now. Your initial predictions will be overwritten. You will not have the chance to change your predictions about the unemployment rates again. In 12 month we will check the official statistics and inform you whether your predictions were correct. For a correct prediction, we will pay you $10.

  Your initial prediction for the unemployment rate in September 2021 if Joe Biden wins the 2020 presidential election: __

  Your initial prediction for the unemployment rate in September 2021 if Donald Trump wins the 2020 presidential election: __

  Please answer the following questions: [Display the same unemployment questions again.]

- Please answer the following questions about the articles you have just read. [Evaluation screen:]
  - On a scale of -3 to +3, with negative numbers representing left leaning or liberal skew, positive numbers representing right leaning or conservative skew, and 0 representing true nonpartisan, how would you rate this article?
  - On a scale of 1 to 7, 1 being not reliable at all and 7 being very reliable, how would you rate the information in this article?
- Now we ask you to answer a question about the article you have just read. For a correct answer, you receive $0.5 paid out at the end of the experiment. [See Appendix A8 for the articles and their corresponding multiple-choice questions.]

A.1.5 The last game

At the beginning of this round, we will use the computer to simulate the draw of a marble from a "cup." There are two cups, with different mixes of colored marbles, and you will be asked to guess the cup that is being used.

First, we draw a computer-generated random number which will be either 1, 2, ... 6. Think of this as the throw of a die with 6 sides, with each side being equally likely.

- If the roll of the die yields 1 - 3, then the draw will be from the Green cup, which contains 2 green marbles and 1 yellow marble.

- If the roll of the die yields 4 - 6, then the draw will be from the Yellow cup, which contains 2 yellow marbles and 1 green marble.

You will not be told in advance the result of the die throw, so you will not know which cup is being used. Once the computerized die throw determines the cup to be used, you will be shown a randomly selected marble from that cup.

You will get a chance to indicate the cup that you think is being used. Your money payoff will depend on whether your prediction turns out to be correct. You will earn $0.5 for a correct prediction, and zero for an incorrect prediction.

A.1.6 Closing survey

Thank you! You have completed all the tasks in the experiment. Before we show you your generated income, we kindly ask you to answer the following questions about yourself.

[Demographics:]

- How old are you?

- What is your gender?

- What is your highest academic degree achieved/in progress?

- What is your ethnic background?

[News consumption:]

- Which do you think are your main sources of news?

  - ABC, NBC, or CBS
  - CNN
  - Fox News
  - Local TV or radio
  - MSNBC
  - NPR (National Public Radio) or PBS
  - Newspapers, magazines, online or in paper
  - Facebook
  - Twitter
  - Other

- (Conditional on 'Newspapers, magazines, online or in paper' being mentioned in prior question) Which of the following newspapers and magazines do you read on a regular basis (at least once a week)?

– Economist
– New York Times
– Wall Street Journal
– Washington Examiner
– Washington Post
– Nature
– Science
– Other:

A.1.7 [Payoff Screen]

On this screen, subjects will see his or her payoffs from each part of the experiment tabulated and summarized.

In the next couple of days, you may receive an additional bonus payment, based on the decision of other players in the first stage of the experiment. We will use the computer to randomly divide all participants in two halves.

You also receive the minimum of $1 from the allocation decision in stage 1. Within the one week (once the study is finished), you may receive an additional bonus payment, based on the allocation decision of another participant, i.e., if someone allocated more than $1 to you. We will use the computer to randomly divide all participants in two halves. For one half of participants, the allocation based on their group membership in the experiment (Triangle group or Circle group) will be payoff relevant. For the other half, the allocation based on political leaning (Democrat or Republican) will be payoff relevant. In both cases, payoffs will be determined as follows: the computer will generate a random sequence of the ID numbers. The first number in the sequence will be the ID number of the person who allocates to the second and third IDs. The second ID drawn will allocate to the third and fourth IDs, ... and so on. The last ID will allocate to the first and second IDs. Therefore, your payoff will be the sum of dollars allocated to you by the two participants preceding you.

We will contact you again in late January 2021 to pay you for your predictions of the presidential election, and in late October 2021 to pay you for your predictions on unemployment and health ranking. Thank you very much for participating in our experiment!

A.1.8 Screenshots in Wave 1

A.1.9 Articles used in experiment

In this section, we display the articles used in the Information Demand and Processing stages. We display summaries of news articles in the "Labels" treatment, where the names of the news sources are labelled. In the treatment with "No Labels," the names of the news sources are replaced by "A news source" the first time we mention it, and "the news sources" the last time we mention it.

The Economy - Left-leaning news sources

Title: 'Staggeringly High': U.S. Jobless Claims Remained Elevated Last Week

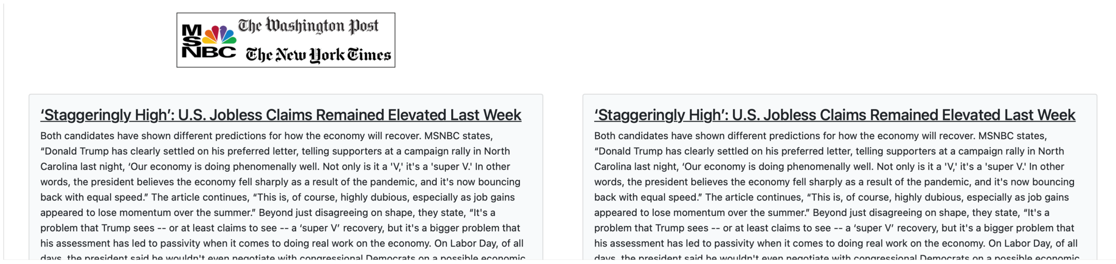

Figure A.1: Interface processing stage (wave 1).

Notes: We depict the interfaces baseline and treatment participants in wave 1 encountered in the processing stage; left panel with source labels, right panel without source labels.

Both candidates have shown different predictions for how the economy will recover. MSNBC states, "Donald Trump has clearly settled on his preferred letter, telling supporters at a campaign rally in North Carolina last night, 'Our economy is doing phenomenally well. Not only is it a 'V,' it's a 'super V.' In other words, the president believes the economy fell sharply as a result of the pandemic, and it's now bouncing back with equal speed." The article continues, "This is, of course, highly dubious, especially as job gains appeared to lose momentum over the summer." Beyond just disagreeing on shape, they state, "It's a problem that Trump sees – or at least claims to see – a 'super V' recovery, but it's a bigger problem that his assessment has led to passivity when it comes to doing real work on the economy. On Labor Day, of all days, the president said he wouldn't even negotiate with congressional Democrats on a possible economic aid package. In other words, despite high unemployment and widespread economic distress, many on Team Trump are prepared to do very little to improve the status quo – and that's even more alarming than the White House's misguided choice in letters used to represent the state of the 'recovery."' (MSNBC, 9/9/2020)

Trump recently stopped talks for a stimulus package before the election. This action came as a surprise to many, and according to the Washington Post, there is fear the President's action will "stall the U.S. economic recovery – or even trigger a backslide." The article continues, "Many economists and business leaders were quick to dub the move disheartening and irresponsible. . . 'Corporations were holding off on laying off employees in the hopes of further stimulus. With this afternoon's news, I expect that we will see businesses capitulate and begin to announce large scale layoffs,' said Peter Atwater, an adjunct lecturer in economics at the College of William & Mary," hinting that more unemployment is coming. With his decision to shut down stimulus talks, "Trump risks the nation backsliding economically, putting more jobs and business in danger of going away. He wanted a V-shaped recovery, but a W is looking more likely." (Washington Post, 10/6/2020).

The job market is at the forefront of everyone's mind. According to the New York Times, "Monthly jobs data released last week showed that job growth slowed sharply in September, and that last spring's temporary furloughs are increasingly turning into permanent job losses. Major corporations like Disney and Allstate have announced thousands of new job cuts. And with winter approaching, restaurants and other businesses that were able to shift operations outdoors during warmer weather could be forced to pull back anew." This is worrisome for future prospects and the report states, "The net result is that potentially millions of workers could see their benefits expire this winter. Epidemiologists warn that cases of the coronavirus are likely to rise as temperatures drop, and winter weather could reduce job opportunities." (New York Times, 10/8/2020).

The Economy - Right-leaning news sources

Title: New Jobs, Unemployment Numbers Point to 'V'-Shaped Recovery – Only a Biden Lockdown Could End It

Both candidates have shown different predictions for how the economy will recover. According to Fox News, "We increasingly appear to be experiencing a "V"-shaped recovery, with four straight months of the most robust job creation ever recorded – 10.6 million new jobs in just one-third of a calendar year and nearly half of the jobs lost in the shutdown." Continuing, they predict "a full and lasting recovery, a return to the prosperity and historically strong job market we enjoyed throughout the vast majority of the president's first term ... While the market has its ups and downs, it has basically recovered its pre-pandemic valuations and was recently again setting all-time records." However, they warn, "The only thing that could mess things up now is if – heaven forbid – politicians plunge the country into another lockdown for political reasons. Unfortunately, that's exactly what one of America's great parties is willing to do. 'I would shut it down. I would listen to the scientists,' Democratic presidential nominee Joe Biden said [on] whether he would subject the American economy to another devastating trip through the 'total lockdown' ringer that's still fresh in our minds." (Fox News, 9/4/2020).

Trump recently stopped talks for a stimulus package before the election. According to the Wall Street Journal, "President Trump halted talks with Democrats about a Covid-19 stimulus package earlier this week, instead pushing for individual relief bills, including aid for airlines and another round of direct checks to many U.S. households." Trump is still thinking of a stimulus package: "'I shut down talks two days ago because they weren't working out. Now they're starting to work out...We're talking about airlines, and we're talking about a bigger deal than airlines,' he said, mentioning the $1,200 stimulus checks to taxpayers that both parties have said they support." Further, there were other reasons for stopping talks for a stimulus package, including "Disagreements over state and local aid in the package . . .. Republicans have criticized the Democratic approach for being too sprawling and expensive, while Democrats say that the GOP plan is insufficient for the scale of the crisis." (Wall Street Journal, 10/8/2020).

The job market is at the forefront of everyone's mind. According to the Washington Examiner, Biden's job projection is questionable, stating "Here's the reality check. Despite Biden predicting 4% economic growth under Barack Obama, the economy barely averaged 2% – rather pathetic for a 'recovery.' The people who made these preposterous bullish predictions are the ones who now say the Biden economic plan will gain millions of jobs." They back this with public opinion data, stating "Throughout nearly all of the Biden-Obama presidency, roughly 1 out of 3 people in the United States rated the economy as good or great...That number surged to about 65% rating the economy as good or great within a year of Trump's presidency." Further, the path is clear: "Now, the question is which game plan gets the economy and employment back to normal as quickly as possible. Biden promises a $4 trillion tax hike almost all on U.S. businesses and investors. That's roughly 5% of everything we produce that gets snatched away in higher taxes. If you believe that this will get America back on the fast track, you probably believe Obama got us 4% growth." (Washington Examiner, 10/8/2020).

The Economy - nonpartisan news sources

Title: The Good and Bad of Trumponomics, and an Overview of the Labor Market

Both candidates have shown different predictions for how the economy will recover. Trump has been adamant in the greatness of Trumponomics. According to the Economist, "President Donald Trump says Americans should re-elect him because of his record on the economy. Before Covid-19, America enjoyed its lowest unemployment rate in 50 years, fast annual wage growth of almost 5% among the lowest-paid workers and a buoyant stock market. Mr Trump attributes all this to his three-pronged strategy of tax cuts, deregulation and confrontational trade policy, and says more of the same will revive the economy after the pandemic." The article states this is one of his strong suits: "Many voters agree. The economy is one issue where Mr Trump does not face a big deficit in the polls." (Economist, 10/17/2020).

Speaking to the overall picture of Trumponomics, the Economist states, "his administration's economic record from before the pandemic is mixed. It got one thing right: when Mr Trump took office the economy was still in need of stimulus, which tax cuts and more spending helped provide. But that success has also helped conceal the damage done by his protectionism." Evaluating his trade war with China and other countries, they state, "Recent research suggests that Mr Trump's tariffs destroyed more American manufacturing jobs than they created, by making imported parts more expensive and prompting other countries to retaliate by targeting American goods. Manufacturing employment barely grew in 2019. At the same time tariffs are pushing up consumer prices by perhaps 0.5%, enough to reduce average real household income by nearly $1,300." In conclusion, there are both positives and negatives: "In 2019 Mr Trump presided over the best labour-market conditions America had seen in several decades. He deserves some of the credit. Despite that, he is overselling Trumponomics. It was both a help and a hindrance." (Economist, 10/17/2020)

The job market is at the forefront of everyone's mind, given that the U.S. economy contracted at "31.4% annualized rate in the second quarter, the deepest drop in output since the government started keeping records in 1947" (Reuters, 9/30/2020). According to the U.S. Bureau of Labor Statistics, employment is going up, but more slowly than before. Breaking down the 661,000 new jobs in September 2020, they state, "In September 2020, almost two-thirds of the gain in leisure and hospitality employment occurred in food services and drinking places (+200,000). Despite job growth totaling 3.8 million over the last 5 months, employment in food services and drinking places is down by 2.3 million since February. Amusements, gambling, and recreation (+69,000) and accommodation (+51,000) also added jobs in September." Further, layoffs have been hardest on smaller establishments: "The layoffs and discharges rate in private nonfarm establishments reached a historical high of 8.8 percent in March 2020 as the Covid-19 pandemic began. The rate remained high in April at 6.9 percent. Layoffs and discharges rates in May, June, July, and August were similar to those before the pandemic. Establishments with 10 to 49 employees had the highest layoffs and discharges rates in March and April, followed by establishments with 50 to 249 employees." (U.S. Bureau of Labor Statistics, 10/6/2020 and 10/14/2020).

Health - Left-leaning news sources

Title: Bidencare Would Be a Big Deal

Both candidates have shown very different approaches to healthcare. The New York Times states, "[Trump] is definitely lying when he claims to have a plan that's better and cheaper than Obamacare. No such plan exists, and he has to know that." In contrast, "Independent estimates suggest that under Biden's plan, 15 million to 20 million Americans would gain health insurance. And premiums would fall sharply, especially for middle-class families." Further, they state, "The plan would also provide significant aid for long-term care, rural health, and mental health. None of this amounts to revolutionary change – in contrast to Trump's efforts to kill Obamacare, which would drastically change American health care, for the worse." (New York Times, 10/5/2020).

With the upcoming case against the Affordable Care Act (ACA), NBC news reports that "Democrats have warned that Barrett's record shows that she would be just as conservative as her mentor, Justice Antonin Scalia. They have argued Barrett could vote to dismantle the Affordable Care Act, with the Supreme Court set to hear oral arguments in a case challenging the health care law Nov. 10. Democratic lawmakers also say they fear her confirmation could lead to a reversal of the landmark 1973 Roe v. Wade decision that protects a woman's right to abortion." Further, they say "one of the main reasons President Donald Trump and the Republicans are trying to ram Barrett's nomination through the Senate ahead of the election is because Trump wants her installed on the bench in case there's a dispute over the election results that rises to the Supreme Court, as it did in the 2000 Bush v. Gore case." Nonetheless, "despite Democrats' fierce opposition to her nomination, Senate Republicans are poised to confirm Barrett, filling the vacancy left by the late Justice Ruth Bader Ginsburg, as Democrats don't have the votes to block her nomination." (NBC news, 10/15/2020).

The Covid-19 vaccine has also been a hot topic, with many questioning its effectiveness and safety. The Washington Post notes, "To receive authorization, a vaccine or drug should seem effective, but its efficacy doesn't definitively need to be proved. For example, a vaccine could generate an immune response, but it might not prevent infection or serious illness. This is a much lower bar. Rather than going through the usual process of consulting panels of experts, the FDA chief alone is able to make this determination – and Commissioner Stephen Hahn has already signaled that he is willing to do so as soon as Oct. 22." Adding to the uncertainty, "Despite studies showing that hydroxychloroquine does not work, Trump has continued to insist that it is safe and effective, citing his own preventive use as evidence." Further, there are concerns about deadlines. The Washington Post continues, "medical research does not operate well on politically imposed deadlines. If this effort is going to be successful, it will need to be done in a way that builds on public trust in science and government, rather than falling into existing partisan divisions. " (Washington Post, 9/4/2020).

Health - Right-leaning news sources

Title: Trump's Healthcare Plan Puts the Patient Where Obamacare Didn't: First

Both candidates have shown very different approaches to healthcare. According to the Washington Examiner, "President Trump unveiled his own health system overhaul, the America First Healthcare Plan. Trump's healthcare plan is exceedingly specific in its diagnosis of all the damage Obamacare has wreaked upon the country and the solutions to reverse it." Further, they state, "It's the patient whom Trump had in mind, not the insurance companies or federal regulators, when he outlined the plan. It focuses on reforms that specifically address what patients want: lower costs, more options, and better care." while contrasting Biden, who "instead promises to double down on the failures of Obamacare and make them worse with his own Medicare-for-all scheme thrown in." (Washington Examiner, 10/1/2020).

With the upcoming case against the Affordable Care Act, "Former Vice President Biden has repeatedly and falsely alleged that President Trump plans to 'destroy the Affordable Care Act, and with it the protections for preexisting conditions."' However, according to Fox News, "President Trump signed an executive order that stated, in part, 'It has been and will continue to be the policy of the United States to give Americans seeking healthcare more choice, lower costs, and better care and to ensure that Americans with pre-existing conditions can obtain the insurance of their choice at affordable rates."' Although "The Trump administration has argued in court that the Affordable Care Act is no longer legal and should thus be struck down," however, "Trump and congressional Republicans have already promised to pass new legislation to protect people with preexisting conditions should the court strike the entire law down." (Fox News, 9/29/2020).

The Covid-19 vaccine has also been a hot topic, with many questioning its effectiveness and safety. As a result, according to the Wall Street Journal, "Several drug makers developing Covid-19 vaccines plan to issue a public pledge not to seek government approval until the shots have proven to be safe and effective, an unusual joint move among rivals that comes as they work to address concerns over a rush to mass vaccination." Further, "The statement would join a growing number of public assurances by industry executives that they aren't cutting corners in their rapid testing and manufacturing of the vaccines." The draft says, "We believe this pledge will help ensure public confidence in the Covid-19 vaccines that may ultimately be approved and adherence to the rigorous scientific and regulatory process by which they are evaluated." However, there have been events adding to the uncertainty. They state, "The emergency-use authorization of hydroxychloroquine, an antimalarial touted by Mr. Trump, which was rescinded over concerns of safety and efficacy, also drew criticism." Despite that, they continue, "Pfizer's CEO said the company would never submit any vaccine for authorization or approval before 'we feel it is safe and effective.' He also said Pfizer hasn't felt any political pressure to rush a vaccine out. 'We will not cut corners,' Mr. Bourla said." (Wall Street Journal, 9/4/2020).

Health - nonpartisan news sources

Title: What the United States Might Look Like After the Election for Key Health Issues

Both candidates have shown very different approaches to public health. Nature states, "Despite public-health agencies counting more than 200,000 Covid-19 deaths in the country, some Trump supporters feel that the impact of the virus has been exaggerated in an effort to control the populace." Further, "the Biden campaign has stated that his administration would direct the CDC to issue transparent, evidence-based guidance around the public-health risks of reopening restaurants, schools and public spaces. This could also go a long way towards restoring morale within the CDC and the FDA." Further, Biden will support the World Health Organization (WHO), "providing badly needed funds to the WHO to fight the coronavirus, polio and other diseases globally[;] reinstating the United States' commitment to the organization would pave the way for joining its international COVAX facility, which aims to accelerate the search for and manufacture of coronavirus vaccines." Beyond commitment to the WHO, "Biden has committed to continue supporting coronavirus-vaccine research, and has pledged that an eventual vaccine will be priced fairly by the federal government." (Nature, 10/1/2020).

With the upcoming case against the Affordable Care Act, according to the Economist, "each Democrat probed Ms Barrett on whether she would vote to scrap it–and strip coverage from some

23m Americans–days after taking Ms Ginsburg's seat. The interrogation was accompanied by stories and photos of sick constituents with pre-existing conditions who could be left without affordable coverage should the high court toss the law." The article follows, "The line of attack is not without footing. In 2017, Ms Barrett criticised NFIB v Sebelius, the 2012 Supreme Court decision upholding the constitutionality of the law's requirement that most Americans buy health insurance. When Chief Justice John Roberts anchored a 5-4 majority interpreting the mandate as a tax within Congress's revenue-raising power, she wrote, he 'pushed the Affordable Care Act beyond its plausible meaning to save the statute'. Juxtaposing Chief Justice Roberts with 'staunch textualists' such as her mentor, Antonin Scalia, Ms Barrett then used a footnote to detail several other cases in which the chief 'depart[ed] from ostensibly clear text' in order to achieve his 'preferable result'. She also favourably quoted Mr Scalia's condemnation of Chief Justice Roberts in Sebelius and in King v Burwell, as having turned the ACA into 'scotuscare'." (The Economist, 10/17/2020).

The Covid-19 vaccine has also been a hot topic, with many questioning its effectiveness and safety. As stated on CDC's website, "In the United States, there is currently no authorized or approved vaccine to prevent coronavirus disease 2019 (Covid-19). Operation Warp Speed has been working since the pandemic started to make a Covid-19 vaccine(s) available as soon as possible. CDC is focused on vaccine planning, working closely with health departments and partners to get ready for when a vaccine is available. CDC does not have a role in developing Covid-19 vaccines." On the website, CDC states 8 things to know about Covid-19 vaccine plans since they may become available before the end of the year, including that "The U.S. vaccine safety system ensures that all vaccines are as safe as possible." (Center for Disease Control, 10/14/2020).

A.2 Instructions in Wave 2

Note that instructions for repeated measurements are omitted as they are identical to the corresponding instructions in wave 1.

A.2.1 Preamble/Consent Form

[Omit: Similar to that of wave 1.]

A.2.2 Stage 1. Bystander allocation games, prediction updating

1. [A bystander allocation game based on minimal groups] [Omit]

2. [A political identity survey.] [Omit]

Please answer the following questions: . . . . . . [Questions from wave 1.]

- Since your participation in our experiment in early November, your political party affiliation, leaning, or strength of affiliation:
  - remains the same – has changed

3. Prediction updating

Instructions. In the first part of this study in October 2020, we asked you to make a number of predictions regarding the consequences of the 2020 presidential election in October 2021. Now that Joe Biden has become the President of the United States, we will give you the opportunity to change your predictions. Remember that for each correct prediction, we will pay you \$10 in October 2021, after the official statistics come out.

We will show you your previous predictions. If you do not want to change your answers, please select the option 'I do not want to change my previous prediction.'

Please answer the following questions: You can now change your previous predictions about the consequences of the election outcomes. In October 2021, we will check the official statistics and inform you whether your predictions were correct. For each correct prediction, we will pay you \$10.

Note: if you decide to change your answer, your previous guess will be overwritten and we will compare your new prediction to official statistics to determine your payoffs.

If you do not want to change your answers, please select the option 'I do not want to change my previous prediction'.

- **According to the Bureau of Labor Statistics, the unemployment rate in September 2020 was 7.9%.**

  What will the unemployment rate be in September 2021, now that **Joe Biden** is the president of the United States?

  Your previous prediction: [Copy previous prediction here.]

  – Strong increase (10 % or higher)
  – Moderate increase (Higher than or equal to 8.5 %, but less than 10 %.)
  – Stable (Higher than or equal to 7.5 %, but less than 8.5 %.)
  – Moderate decrease (Higher than or equal to 6 %, but less than 7.5%.)
  – Strong decrease (6 % or lower)
  – I do not want to change my previous prediction

- **According to US News and World Report, Canada ranks 1st among countries with the most developed public health care systems in 2020, while the United States ranks 15th.**

  What will be the ranking of the United States in September 2021, now that **Joe Biden** is the president of the United States?

  Your previous prediction: [Copy previous prediction here.]

  – Strong improvement (Rank 12 or better)
  – Moderate improvement (Rank 13 or 14)
  – No change (Rank 15)
  – Moderate decline (Rank 16 or 17)
  – Strong decline (Rank 18 or worse)
  – I do not want to change my previous prediction

4. [A bystander allocation game based on political groups] [Omit]

5. Belief elicitation, anchoring and treatment.

Instructions. Recall that, in late October, everyone participated in a "Guessing which cup is used" task. The task is described below to help refresh your memory.

In that task, we used the computer to simulate the draw of a marble from a "cup". There were two cups, with different mixes of colored marbles, and you were asked to guess the cup that was being used.

First, for every participant individually, we drew a computer-generated random number which would be either 1, 2, ... 6. Think of this as the throw of a die with 6 sides, with each side being equally likely.

- If the roll of the die yielded 1 - 3, then the draw would be from the Green cup, which contained 2 green marbles and 1 yellow marble.
- If the roll of the die yielded 4 - 6, then the draw would be from the Yellow cup, which contained 2 yellow marbles and 1 green marble.

We did not reveal the result of the die throw in advance. Once the computerized die throw determined the cup to be used, each participant was shown a randomly selected marble from his or her personal cup. Each participant was then asked to guess which cup was used.

You are now asked to estimate the proportion of Republicans and Democrats who guessed their cup correctly. You will earn 50 cents for each correct estimate.

Note: across all participants in October 2020, 67% guessed their cup correctly.

Please answer the following questions.

Among the Democrats, what is the share of participants who correctly guessed the cup?

- More than 80%
- More than 70% but less than or equal to 80%
- Between 60% and 70%, inclusive
- Less than 60%, but more than or equal to 50%
- Less than 50%

Among the Republicans, what is the share of participants who correctly guessed the cup?

- More than 80%
- More than 70% but less than or equal to 80%
- Between 60% and 70%, inclusive
- Less than 60%, but more than or equal to 50%
- Less than 50%

Results. [*Note this screen is shown to participants in the treatment group only.*]

The answers to the previous questions about the share of Democrats and Republicans who guessed the cup correctly are as follows:

Democrats: 67% **Republicans: 67%**

Put differently: There is no difference in the share of Democrats and Republicans who guessed the cup correctly.

A.2.3 Stage 2. Information Demand

Instructions. This stage consists of three rounds of the task "Guessing which cup is used" that resembles the task from October 2020.

At the beginning of each round, we draw a computer-generated random number which will be either 1, 2, ... 6. Think of this as the throw of a die with 6 sides, with each side being equally likely.

- If the roll of the die yields 1 - 3, then the Green cup will be used, which contains 2 green marbles and 1 yellow marble.
- If the roll of the die yields 4 - 6, then the Yellow cup will be used, which contains 2 yellow marbles and 1 green marble.

You will not be told in advance the result of the die throw, so you will not know which cup is being used. In every round your task is to guess which cup is used. **A new die throw occurs in each round, which determines the cup used for that round.**

Details: Before you make your guess, **you will see the guesses of a group comprised of two other participants from October 2020 who were both shown a randomly selected marble from the cup being used.**

Please note that **(i) after showing a randomly selected marble to a participant, the marble is put back into the cup, before the computer randomly selects and shows a marble to another participant, and (ii) the other two participants did not see each other's guesses.** You yourself will not be shown a marble.

In each round, there are two groups that comprise two other participants. You will see the guesses of one of the two groups. Under each group, a slider indicates the likelihood that you will see their guesses. The default probability is 0.5. That is, each group is equally likely to be selected. You have a 40-cent budget which you can use to change the probabilities. Moving the slider to the left increases the probability that the group on the left is chosen, and vice versa. It costs 10 cents to change the probability by 10 percentage points. The remaining budget will be paid out to you at the end of the experiment.

Based on your choice of probabilities, the computer will randomly draw a group of two participants and show you their guesses. Based on the information, you will get a chance to indicate the cup that you think is being used. Your money payoff will depend on whether your guess turns out to be correct. You will earn $1 for a correct guess, and zero for an incorrect guess. At the end of the experiment, one of the rounds in this stage will be randomly selected for payment.

Please make a decision [Decision screen: remaining budget, group composition with partisan labels, slider,]

Guesses of other people whose marble has been drawn from your cup

Person 1, party affiliation and party image [donkey or elephant]

Person 1 guesses that the cup's color is [green or yellow].

Person 2, party affiliation and party image [donkey or elephant]

Person 2 guesses that the cup's color is [green or yellow].

Considering the guesses of the other people, which cup do you think was used in the current round?

- Yellow cup
- Green cup

A.2.4 Stage 3: Information Processing

Instructions. This stage consists of six rounds of the task "Guessing which cup is used" that resembles the task in stage 2.

At the beginning of each round, we will again use the computer to simulate the random selection of a "cup". As before, there are two cups, with different mixes of colored marbles, and you will be asked to guess the cup that is being used. You will not be told in advance which cup is being used. A new die throw occurs in each round, which determines the cup used for that round.

Details: Once the computerized die throw determines the cup to be used, you will be shown a randomly selected marble from that cup.

Furthermore, you will be shown the guesses of two other participants from October 2020 whose marble has been drawn from the same cup as yours. Note that (i) after showing a randomly selected marble to a participant, the marble is put back into the cup, before the computer randomly selects and shows a marble to the next person, i.e., you and the other two participants may have seen the same or different marble(s), and (ii) the other two participants did not see each other's guesses.

You will then get a chance to indicate the cup that you think is being used. Your monetary payoff will depend on whether your guess turns out to be correct. You will earn $1 for a correct guess, and zero otherwise. At the end of the experiment, one of the six rounds in this stage will be randomly selected for payment.

Round 1 of 4

Guesses of other people whose marble has been drawn from the same cup as yours

Person 1, party affiliation and party image [donkey or elephant], guesses that the cup's color is [green or yellow].

Person 2, party affiliation and party image [donkey or elephant], guesses that the cup's color is [green or yellow].

Your own marble The color of the marble that has been randomly drawn for you is [green or yellow].

Considering the guesses of the other participants and the marble shown to you, which cup do you think was used in the current round?

- Yellow cup
- Green cup

Please answer the following question: How have you come up with your decisions in the previous 4 rounds?

Next we ask you to make 2 more decisions in a similar fashion as before.

Round 1 of 2

Guesses of other people whose marble has been drawn from the same cup as yours

Person 1 guesses that the cup's color is [green or yellow].

Person 2 guesses that the cup's color is [green or yellow].

Your own marble The color of the marble that has been randomly drawn for you is [green or yellow].

Considering the guesses of the other participants and the marble shown to you, which cup do you think was used in the current round?

- Yellow cup
- Green cup

A.2.5 Stage 4: Information Demand - Signal vs. Guess

Instructions This stage consists of two more rounds of the task "Guessing which cup is used" that resembles previous tasks.

At the beginning of each round, we will again use the computer to simulate the random selection of a "cup". As before, there are two cups, with different mixes of colored marbles, and you will be asked to guess the cup that is being used. You will not be told in advance which cup is being used. A new die throw occurs in each round, which determines the cup used for that round.

Details: At the beginning of each round, you will get to choose the likelihood that you see a randomly drawn marble from that cup, or another participant's guess who saw a randomly selected marble from the same cup as yours.

A slider indicates the likelihood that you will see a marble from your cup or another person's guess whose marble has been drawn from the same cup as yours. The default probability is 0.5. That is, each is equally likely to be selected. You have a 40-cent budget which you can use to change the probabilities. Moving the slider to the left increases the probability that you see a marble drawn from the cup, whereas moving the slider to the right increases the likelihood that you see another person's guess. It costs 10 cents to change the probability by 10 percentage points. The remaining budget will be paid out to you at the end of the experiment.

Based on the probability you choose, the computerized die throw determines whether you will be shown a randomly selected marble from that cup or another person's guess.

You will then get a chance to indicate the cup that you think is being used. Your monetary payoff will depend on whether your guess turns out to be correct. You will earn $1 for a correct guess, and zero otherwise. At the end of the experiment, one of the two rounds in this stage will be randomly selected for payment.

A.2.6 A final questionnaire

The final part of the experiment consists of a questionnaire. Please read each question carefully and answer it truthfully.

Once you have answered all questions, please press the "Next" button on your screen.

Please answer the following questions.

- Which state do you live in? (drop-down menu)
- Who did you vote for in the 2020 presidential election?
  - Joe Biden
  - Donald Trump
  - Other
- Do you think that the 2020 presidential election was rigged? (yes/no)
- Think of a ladder as representing where people from different groups stand in our society. At the top of the ladder are the people who have the highest standing in society. At the bottom are the people who have the lowest standing in society.
  - Where would you place Democrats on this ladder? (value from 1-10, where 10 represents the top of the ladder and 1 the bottom)
  - Where would you place Republicans on this ladder? (value from 1-10, where 10 represents the top of the ladder and 1 the bottom)
- On January 26, 2021, President Biden said that his administration was nearing a deal with two manufacturers that would enable 300 million Americans to have their shots by the end of the summer.
- Have you received a vaccine shot? (yes/no)
- If not, do you intend to get vaccinated when the opportunity arrives? (yes/no)
- Do you think close to 300 million Americans will have their shots by September 1, 2021? (yes/no)

A.2.7 Screenshots in Wave 2

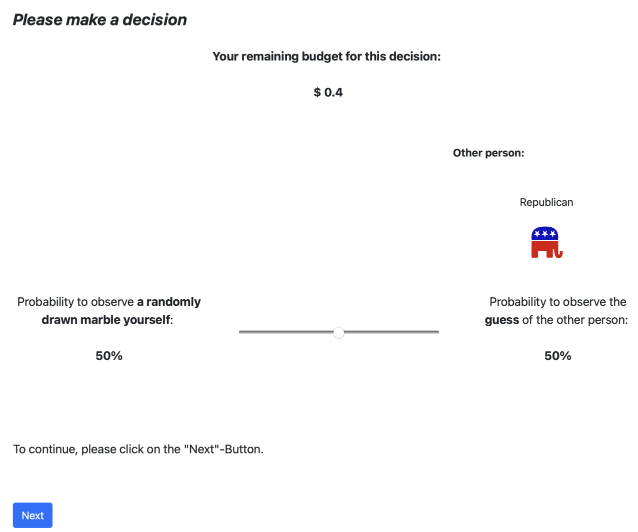

Figure A.2: Interface demand stage; marble v. guess (wave 2).

Notes: We depict the interfaces baseline and treatment participants encountered in wave 2 in the demand stage, when they had to choose between observing a marble themselves or the guess of another participant.

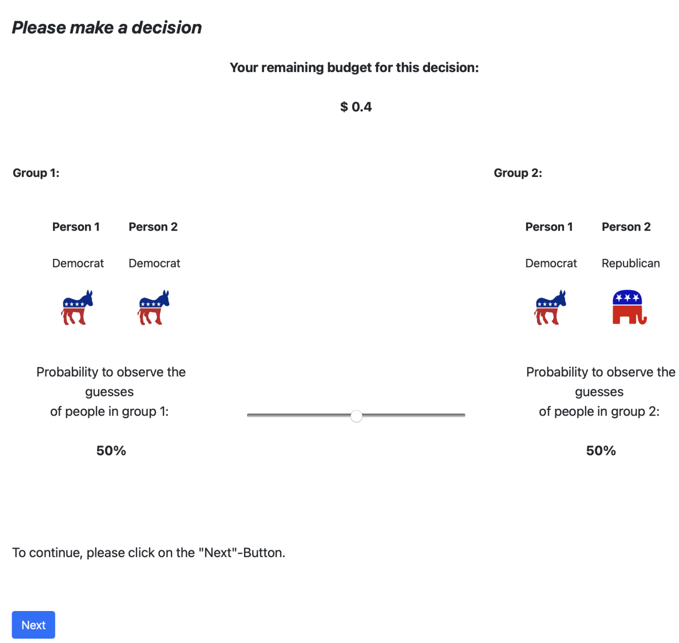

Figure A.3: Interface demand stage; guess v. guess (wave 2).

Notes: We depict the interfaces baseline and treatment participants encountered in wave 2 in the demand stage, when they had to choose between observing guesses from different participants.

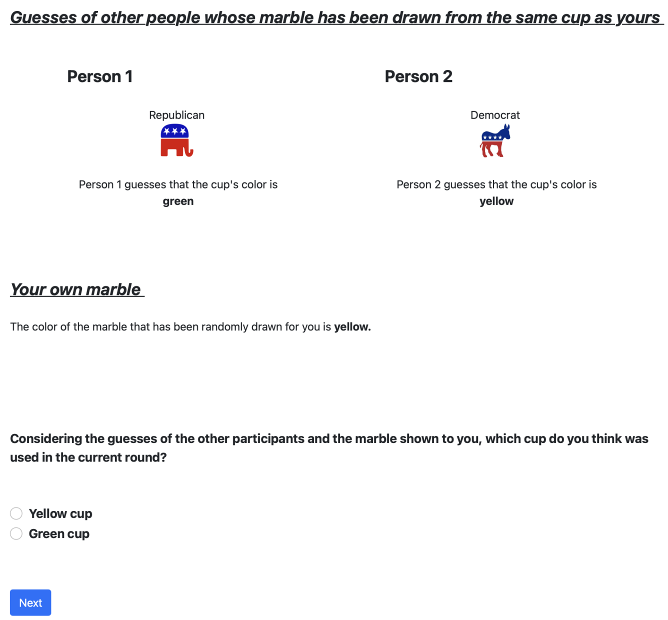

Figure A.4: Interface processing stage (wave 2).

Notes: We depict the interfaces baseline and treatment participants encountered in wave 2 in the processing stage.

A.3 Instructions Wave 3

As we do not use any data from wave 3, instructions are omitted. See our SSRN working paper (Bauer et al. 2023).

B Summary statistics and additional analyses

Table B.1 presents information about the demographic representation of our participants across gender, age, education level, residential areas, and ethnicity. When we compare our sample with the U.S. population data derived from the 2020 US Census and the 2019 American Community Survey, we observe a close resemblance in terms of gender, residential areas, race and ethnicities. However, our participants tend to be slightly younger and better educated. A notable difference emerges in the participants' self-reported political leanings. Notably, our sample has a higher (lower) proportion of Democrats (Republicans and Independents) compared to their proportion in the population.

B.1 Summary statistics on participants' socio-demographic characteristics across waves

Table B.1: Summary statistics on participants' socio-demographic characteristics.

Variable	US Average	Wave 1 (N=1005)		Wave 2 (N=740)		Wave 3 (N=530)	
		Mean	P-value	Mean	P-value	Mean	P-value
Gender							
Male	0.49	0.48	0.57	0.47	0.20	0.47	0.40
Female	0.51	0.51	0.95	0.53	0.35	0.52	0.56
Non-binary	-	0.01	-				
Age	47.8	45.7	0.00	48.35	0.35	50.96	0.00
Area of living							
Rural	0.19	0.17	0.07	0.19	0.88	0.2	0.42
Urban and suburban	0.81	0.83	0.95	0.81	0.88	0.8	0.42
Ethnicity							
Caucasian	0.72	0.74	0.19	0.75	0.09	0.77	0.01
African American	0.13	0.13	0.61	0.11	0.20	0.10	0.04
Asian American	0.06	0.07	0.26	0.07	0.18	0.06	0.85
Other	0.10	0.06	0.00	0.06	0.00	0.06	0.00
Academic Degree							
No degree	0.10	0.01	0.00	0.12	0.00	0.01	0.00
High School	0.55	0.30	0.00	0.33	0.00	0.34	0.00
Bachelor	0.22	0.44	0.00	0.46	0.00	0.45	0.00
Master	0.11	0.22	0.00	0.17	0.00	0.17	0.0003
PhD	0.02	0.03	0.07	0.03	0.12	0.03	0.12
Democrats							
Overall	0.29	0.48	0.00	0.52	0.00	0.5	0.00
Moderate		0.26		0.24		0.21	0.00
Strong		0.22		0.28		0.29	0.00
Republicans							
Overall	0.30	0.26	0.00	0.23	0.00	0.24	0.00
Moderate		0.16		0.14		0.16	
Strong		0.10		0.09		0.08	
Independents							
Overall	0.38	0.22	0.00	0.25	0.00	0.26	0.00
Leaning Democrat		0.14		0.15		0.17	
Leaning Republican		0.08		0.1		0.09	
groupishness							
Min. groupish		0.37		0.35		0.30	
Pol. groupish		0.57		0.57		0.56	

[a.] US statistics are based on the 2019 American Community Survey for gender, age, and ethnicity, the 2020 Census data for Area of living, and the Current Population Survey, 2020 Annual Social and Economic supplement, for academic degree. P-values are from proportion of t-tests.

B.2 Wave 1: Factual questions with partisan valence (Peterson and Iyengar 2021)

Table B.2: Responses to factual questions

Factual question	Correct Response	Party Valence	Share of corrects answers (Democrats)	Share of corrects answers (Republicans)	Partisan Divide
Illegal immigrants commit violent crime at a significantly higher rate than legal American citizens.	False	Rep	89.5	57.3	32.2
Millions of illegal votes were cast in the 2016 presidential election.	False	Rep	80.9	71.3	9.6
Former President Obama ordered wire taps on Donald Trump's phones.	False	Rep	92.3	51.7	40.6
40% of firearm sales in the US occur without a background check.	False	Dem	23.4	43.3	19.9
The vast majority (over 90%) of climate scientists believe that global warming is an established fact and that it is most likely caused by man-made emissions.	True	Rep	96.5	75.6	20.9
Michael Cohen, Donald Trump's personal lawyer pleaded guilty to fraud and illegal campaign finance charges in August 2018.	True	Rep	91.4	76.7	14.4

[a] Summary statistics on participants' partisan divide in correctly answering factual questions. Shares of participants, conditional on their partisanship, who provide a correct answer to factual questions with a given valence. We pool observations for both label and no-label conditions.

Table B.2 presents participants' responses to factual questions sourced from (Peterson and Iyengar 2021). We detail the percentages of Democrats and Republicans who correctly answered factual questions that present a particular party in a less favorable light. Take, for instance, the statement: "Millions of illegal votes were cast in the 2016 presidential election." The accurate response to this is 'false'. This response casts the Republicans in a negative light, given their repeated assertions to the contrary, making the statement lean Republican in its valence. The table indicates the percentage of respondents who answered each question correctly. We combine observations from both our label and no-labels conditions. Additionally, the table measures the partisan divide for each question, defined as the absolute difference in the percentages of Democrats and Republicans answering correctly. The answers to these factual questions closely align with findings by (Peterson and Iyengar 2021). On average, there's a partisan disparity of 22.9 percentage points in the propensity to correctly answer a factual question. This suggests that participants are significantly less likely to answer correctly if the truthful response casts their affiliated party in an unfavorable light.

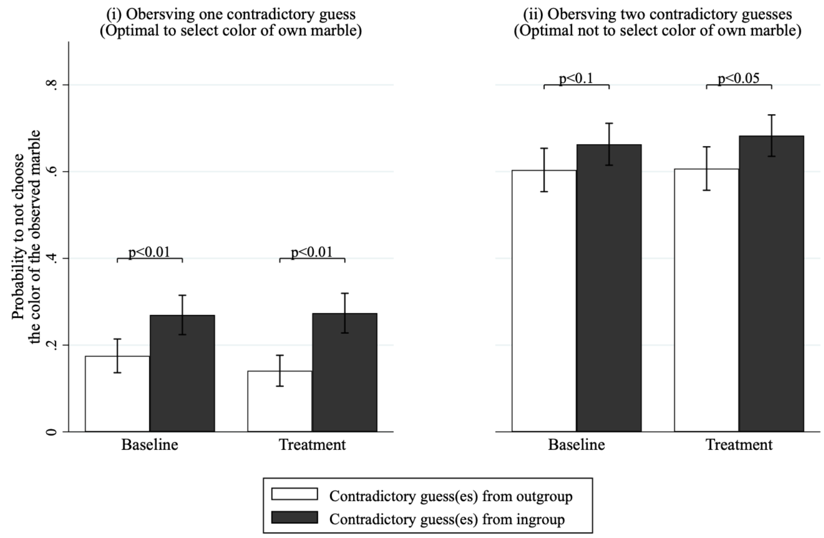

Figure B.5: Probability of abandoning a private signal in favor of contradictory guess(es) when participants observe one (panel i) or two contradictory guesses (panel ii) in wave 2.

B.3 Wave 2: Information processing

Table B.3: Information processing (Wave 2): OLS estimates of Eq. (2).

DV: Guessing green	All participants			Universalist	Groupish
	Baseline	Treatment	Baseline & Treatment		
	(1)	(2)	(3)	(4)	(5)
Self green marble (β_S)	0.282***	0.250***	0.283***	0.271***	0.306***
	(0.025)	(0.027)	(0.025)	(0.033)	(0.039)
Ingroup # of green guesses (β_I)	0.201***	0.180***	0.202***	0.188***	0.228***
	(0.018)	(0.018)	(0.018)	(0.022)	(0.030)
Outgroup # of green guesses (β_O)	0.151***	0.110***	0.151***	0.142***	0.169***
	(0.019)	(0.018)	(0.018)	(0.022)	(0.033)
Debiasing Treatment			0.044	0.003	0.123**
			(0.031)	(0.037)	(0.055)
Treatment×Self Green marble			-0.033	-0.015	-0.062
			(0.037)	(0.045)	(0.061)
Treatment×Ingroup # of green guesses			-0.022	-0.004	-0.058
			(0.025)	(0.031)	(0.043)
Treatment×Outgroup # of green guesses			-0.041	0.002	-0.123***
			(0.026)	(0.031)	(0.046)
$\beta_I - \beta_O > 0$	0.050***	0.070***	0.051***	0.045**	0.059*
p-value, F-test	0.006	0.005	0.000	0.000	0.04
$\beta_S - \beta_I > 0$	0.081***	0.070**	0.081***	0.084**	0.078*
p-value, F-test	0.004	0.02	0.000	0.000	0.04
Controls	Yes	Yes	Yes	Yes	Yes
Observations	2,226	2,214	4,440	2,886	1,554
Adj. R-squared	0.082	0.063	0.073	0.072	0.082

a. In addition to the controls included in Table 1, we control for the topic encountered in the information processing stage.

b. Robust standard errors in parentheses are clustered at the individual level. We denote significance levels by * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

C Theoretical Framework and Hypotheses

We begin by outlining a theoretical framework to incorporate the role of group identity in belief formation. The framework is to provide guidance for us to identify potential channels and develop our main hypotheses. The proposed framework builds on and extends existing models of information demand (subsection C.1) and information processing (subsection C.2).

C.1 Information Demand

Recent studies on information preferences provide evidence that people value information beyond its instrumental value (Thornton 2008, Chen and He 2021, Möbius, Niederle, Niehaus and Rosenblat 2022). Contemporary models typically specify a utility function with both an instrumental and a non-instrumental component (Caplin and Leahy 2001, Bénabou and Tirole 2016, Golman and Loewenstein 2018). Here we adapt the information gap theory of Golman and Loewenstein (2018) to our context.

Let $\mu \in \Omega$ be a possible state of the world and $m(\mu) \geq 0$ be the instrumental (monetary) value associated with state μ. In our experiment, an agent will receive a monetary prize $m(\mu)$, if they correctly guess the state of the world, and will receive zero otherwise. Let $\pi(\mu)$ be an agent's subjective belief that state μ is the true state and μ^0 denote their prior belief. Signals come from various sources. An information source, $z \in \{f, n\}$, might be favored or non-favored.[20] An agent's expected utility function has the following components:[21]

(IN) Expected instrumental value, $\sum_\mu \pi(\mu)m(\mu)$, which motivates the formation of accurate beliefs;

(VA) Expected valence, $\sum_\mu \pi(\mu)v_z(\mu)$, where $v_z(\mu)$ is the goodness (or badness) of having a subjective belief; and

(SU) Source utility, $g_z \in \mathbb{R}$, which captures a preference for or aversion to a specific information source *per se*, i.e., independent of its content or instrumental value. We expect that groupish participants have source utility, i.e., $|g_z| \neq 0$, whereas universalist ones do not, $g_z = 0$.

Golman and Loewenstein (2018) posit that information *content* has valence. This framework has been used in the health care context to illustrate the case of why people might avoid learning their HIV status when testing is free and treatments are highly effective (Thornton 2008). Content valence likely matters as well in the political sphere, where people may make choices that help them maintain the view that their preferred policy options or candidates are effective.

In our experiment, agents have the possibility to learn by acquiring signals from one of two possible sources. We set the default probability of learning from either source at 0.5. An agent can spend resources, $x \in [0, \bar{x}]$, to change the likelihood of receiving a signal from one source versus the other. Let $c(x)$ be the cost of changing the likelihood of observing a signal from a specific

[20] In wave 1, an information source might come from an ingroup, nonpartisan or outgroup source. Therefore, we use favored versus non-favored to compare them.

[21] The original model in Golman and Loewenstein (2018) contains a third term, which is the entropy of the probability distribution over the states of the world. Since our design cannot separately identify the information entropy from its instrumental value, we drop this term from our model. We note that all our theoretical results hold when we include the entropy term.

source, where $c'(x) > 0$ and $c''(x) \geq 0$. An agent incurs a cost of $c(x)$ to increase the probability of receiving a signal from a favored source by $\Delta p(x)$. For simplicity, we assume linear pricing, i.e., $\Delta p(x) = \Delta p \cdot x$. If an agent invests x to increase the probability of observing a signal from a favored source, the agent's expected utility equals:

$$u(\pi, x) = (0.5 + \Delta p(x))u_f + (0.5 - \Delta p(x))u_n - c(x), \tag{3}$$

where

$$u_z = \sum_{\mu} \pi(\mu|s_z)[m(\mu) + v_z(\mu)] + g_z, \tag{4}$$

for $z \in \{f, n\}$. After an agent receives a signal from an information source, z, they update their posterior to $\pi(\mu|s_z) = (1 - \beta_z)\mu^0 + \beta_z s_z$. An agent maximizing (3) with respect to x obtains $x^* = \Delta p \cdot c'^{-1}[\Delta Em(\mu) + \Delta Ev(\mu) + \Delta g]$, where $c'^{-1}(\cdot)$ is the inverse marginal cost function, $\Delta Em(\mu)$, $\Delta Ev(\mu)$, and Δg respectively represents the difference in expected instrumental value, content valence and source utility between a favored and a non-favored information source.[22]

Proposition 1 (Information Demand)**.** *Ceteris paribus, an agent's willingness to pay to change the information source increases in the information's instrumental value, content valence, and source utility.*

From Proposition 1, we derive the following pre-registered hypotheses. First, we expect participants will be willing to pay more for information from ingroup versus outgroup sources. The intuition here is that ingroup sources may have both higher content valence and source utility. In addition, they may perceive the information from ingroup sources to be more accurate, i.e., to be of higher instrumental value.

Hypothesis 1. *Participants will be willing to pay more for ingroup information sources than outgroup ones.*

Second, we expect any treatment targeting either the source utility or the perceived instrumental value should reduce intergroup preferences in the demand for information. In particular, when source utility is reduced (as in our wave 1 *unlabeling* treatment), we expect a decrease in the demand for a favored information source.

Hypothesis 2. *When information sources are unlabelled, outgroup information avoidance is reduced compared to the control condition.*

To the extent that valence and source utility are rooted in group identity, we expect that groupish participants will attach higher utility to ingroup versus outgroup sources. We expect that the treatment effect of unlabeling will be stronger for groupish participants.

Hypothesis 3. *The treatment effect of unlabeling is stronger for groupish participants.*

In wave 2, if participants learn that both information sources are equally accurate (as in our wave 2 *debiasing* treatment), we again expect a decrease in the willingness to pay for ingroup information.

[22]More precisely, $\Delta Em(\mu) \equiv \sum_{\mu} \pi(\mu|s_f)[m(\mu)] - \sum_{\mu} \pi(\mu|s_n)[m(\mu)]$, $\Delta Ev(\mu) \equiv \sum_{\mu} \pi(\mu|s_f)[v(\mu)] - \sum_{\mu} \pi(\mu|s_n)[v(\mu)]$, and $\Delta g \equiv g_f - g_n$.

Hypothesis 4 (Treatment effect of debiasing)**.** *The ingroup bias in information demand is smaller in the debiasing treatment than in the control group.*

Similarly, since the debiasing treatment removes the differential instrumental values of information sources, we expect a different treatment effect from groupish versus universalist participants.

Hypothesis 5. *The treatment effect of debiasing is different for groupish and universalist participants.*

C.2 Information Processing

To derive our hypotheses related to how participants process information, i.e., update their subjective beliefs after observing signals, we first consider a model of information processing, developed by Fryer et al. (2019), where signals are continuous and open to interpretation, and adapt it to our setting in wave 1. We then add group identity to a simple Bayesian updating model in which signals are discrete and unambiguous (as in wave 2 of the experiment).

In wave 1, signals, e.g., newspaper articles, are continuous and open to interpretation. To address this possibility, Fryer et al. (2019) extend the Bayesian updating framework and show that differences in priors can prevent agents who receive exactly the same signals from converging in their posteriors. The key feature of their belief-updating model is that agents first *interpret* the signal according to their prior, and then update their beliefs using the interpreted signal. We adapt their model to our experimental setting to derive additional hypotheses for our experiment. For completeness, we summarize their model specifications below and refer the reader to the original paper for more details.

In Fryer et al. (2019), the true state is denoted by $\mu \in \mathbb{R}$. An agent's prior μ_0 is the expectation of the true mean based on a normal distribution over potential means with variance σ_0^2. There is a sequence of signals, s_t, arriving in period t that are i.i.d. according to a normal distribution centered around the true mean μ and with variance $\sigma_s^2 \sim N(\mu, \sigma_s^2)$. Let $\hat{\mu}_t$ denote the posterior of an agent after observing t signals.

This model comprises a two-stage updating process. The first stage is the interpretation stage. Let $\hat{s}_t$ be the interpreted signal:

$$\hat{s}_t = \frac{\hat{\mu}_{t-1} + x_t s_t}{1 + x_t}, \tag{5}$$

where $x_t = \sigma_{t-1}^2 / \sigma_s^2$ is the precision-adjusted weight an agent applies when interpreting the signal.

The second stage is a standard Bayesian updating stage using interpreted signals. Proposition 5 in Fryer et al. (2019) characterizes the agent's posterior as follows:

$$\hat{\mu}_t = \left[\mu_0 \left(\frac{\sigma_s^2}{\sigma_s^2 + \sigma_0^2}\right) + \sum_{\tau=1}^{t} s_\tau \left(\frac{\sigma_0^2}{\sigma_s^2 + \tau\sigma_0^2}\right)^2 \left(\frac{\sigma_s^2 + \tau\sigma_0^2}{\sigma_s^2 + (\tau+1)\sigma_0^2}\right)\right] \left[1 + \frac{\sigma_0^2}{\sigma_s^2 + t\sigma_0^2}\right]. \tag{6}$$

Eq. (6) implies that the estimating equations should take the form of posteriors as a linear function of priors and signals.

Adapting the model to our experimental setting in wave 1, we set $t = 2$, as each participant reads two summaries of news articles (signals). For simplicity, we assume that the signals and prior are equally noisy, $\sigma_s^2 = \sigma_0^2$. After an agent receives the first signal, they interpret the article using their own prior. Applying Eq. (5) for the interpreted signal, $\hat{s}_1$, and Eq. (6) for their posterior, $\hat{\mu}_1$, we obtain:

$$\hat{s}_1 = \frac{1}{2}\mu_0 + \frac{1}{2}s_1, \text{ and } \hat{\mu}_1 = \frac{3}{4}\mu_0 + \frac{1}{4}s_1.$$

Thus, the interpreted signal is closer to the prior belief μ_0. In the updating stage, the agent uses the interpreted signal, $\hat{s}_1$, which effectively weights the prior belief twice. Similarly, after receiving the second signal and interpreting it, the agent has the following posterior:

$$\hat{\mu}_2 = \frac{2}{3}\mu_0 + \frac{2}{9}s_1 + \frac{1}{9}s_2. \tag{7}$$

By contrast, a standard Bayesian agent who does not interpret any signal weights the prior and all subsequent signals equally, i.e., $\hat{\mu}_2 = \frac{1}{3}\mu_0 + \frac{1}{3}s_1 + \frac{1}{3}s_2$.

This model has two implications in our context. First, it implies that priors are sticky in that they hold greater weight in an agent's posterior than do signals. This overweighting of priors prevents agents with heterogeneous priors from converging, even if they observe the same information.

Hypothesis 6. *An agent's prior holds more weight in their posterior than do the signals.*

Second, it implies that there is an order effect, as an early signal, which receives less interpretation, is weighted more than a later one.[23]

Hypothesis 7. *An earlier signal receives more weight than a later one.*

We now turn to a simple environment as in wave 2 of our experiment. Suppose there are two possible states of the world, $\mu \in \{A, B\}$. Let $P(\mu)$ denote an agent's prior belief. Let S be a set of signals and $P(S|\mu)$ be the likelihood of observing S in state μ. When a Bayesian decision maker observes multiple independent signals, they should pool the signals. More precisely, let N_a and N_b be the numbers of a and b signals, respectively. A Bayesian decision maker thus treats $N_a - N_b$ as the sufficient statistic for making an inference, when both states are ex ante equally likely and signals are informative. We formulate this known result as an observation (Kahneman and Tversky 1972).

Observation 1. *Given a uniform prior, $P(A) = P(B)$, a Bayesian decision maker infers the state of the world to be: (1) A iff $N_a > N_b$, (2) A with probability 1/2 iff $N_a = N_b$, and (3) B otherwise.*

To allow for biases in information processing, we use $\pi(\cdot)$ to denote an agent's subjective (posterior) belief. Following Grether (1980) and subsequent research on belief updating, we use the following reduced-form representation:

$$\frac{\pi(A|S)}{\pi(B|S)} = \left[\frac{P(S|A)}{P(S|B)}\right]^{\beta} \left[\frac{P(A)}{P(B)}\right]^{\gamma},$$

[23]Note that Hypotheses 6 and 7 are not pre-registered. Also note that these two results do not rely on the assumption that signals and prior are equally noisy, or $\sigma_s^2 = \sigma_0^2$.

where the parameter $\beta \geq 0$ measures the biased use of the signals and $\gamma \geq 0$ measures the biased use of the prior. Bayes' Rule emerges when $\beta = \gamma = 1$. In the case of a uniform prior (as in wave 2 of our experiment), $P(A) = P(B)$, the prior ratio drops out.

In our setting, an agent may receive multiple independent signals, $s_1, \cdots, s_n$, from different sources that may be treated differently. We represent this scenario as follows:

$$\frac{\pi(A|s_1, s_2, \cdots, s_n)}{\pi(B|s_1, s_2, \cdots, s_n)} = \left[\frac{P(s_1|A)}{P(s_1|B)}\right]^{\beta_1} \left[\frac{P(s_2|A)}{P(s_2|B)}\right]^{\beta_2} \cdots \left[\frac{P(s_n|A)}{P(s_n|B)}\right]^{\beta_n}. \tag{8}$$

We can then use Eq. (8) to distinguish among private signals, as well as signals from ingroup and outgroup sources. Let β_S, β_I and β_O be the weight an agent attaches to a signal from their own observation (self), an ingroup, and an outgroup source, respectively. Multiple signals from the same source are pooled by the agent (Observation 1). This leads to the following equation that will form the basis for our estimation strategy:

$$\ln\left[\frac{\pi(A|S)}{\pi(B|S)}\right] = \beta_S(N_a^s - N_b^s) + \beta_I(N_a^i - N_b^i) + \beta_O(N_a^o - N_b^o), \tag{9}$$

where $N_a^k - N_b^k$ is the difference between the number of a and b signals that come from an agent's own observation, an ingroup source, or an outgroup source, respectively, with $k \in \{s, i, o\}$. Eq. (9) suggest that an agent's posterior belief about the true state of the world increases in the number of supporting signals that come from their own observation, ingroup sources, and outgroup sources, respectively.

Based on the above analysis, we state the following hypotheses regarding information processing when signals are discrete and unambiguous. First, we hypothesize that the information from the ingroup will resonate more with an individual, and thus the individual accords this information more weight than the corresponding outgroup information (Malmendier and Veldkamp 2022).

Hypothesis 8 (Group-contingent Information Processing)**.** *Participants will put more weight on the guesses of ingroup members than those of outgroup members when updating in the "guessing the cup game." In other words,* $\beta_I > \beta_O$.

Second, we expect that our experimental treatment that targets perceived source accuracy (*debiasing* treatment in wave 2) will reduce the level of observed ingroup bias in information processing. The reason for this expectation is that the treatment decreases the participants' differential beliefs related to the instrumental value of the signals.

Hypothesis 9 (Treatment Effects of Debiasing)**.** *Participants in the debiasing treatment will put equal weights on the guesses of ingroup and outgroup members when updating in the "guessing the cup" game. In other words,* $\beta_I = \beta_O$.

References

Agranov, Marina, Gabriel Lopez-Moctezuma, Philipp Strack, and Omer Tamuz, "Learning through imitation: an experiment," Technical Report, National Bureau of Economic Research 2022.

Akerlof, George A. and Rachel E. Kranton, "Economics and Identity," *The Quarterly Journal of Economics*, August 2000, *115* (3), 715–753.

Alesina, Alberto, Armando Miano, and Stefanie Stantcheva, "The polarization of reality," in "AEA Papers and Proceedings," Vol. 110 American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203 2020, pp. 324–328.

Allcott, Hunt, Levi Boxell, Jacob Conway, Matthew Gentzkow, Michael Thaler, and David Yang, "Polarization and public health: Partisan differences in social distancing during the coronavirus pandemic," *Journal of public economics*, 2020, *191*, 104254.

Andre, Peter, Teodora Boneva, Felix Chopra, and Armin Falk, "Misperceived social norms and willingness to act against climate change," *Review of Economics and Statistics*, 2024, pp. 1–46.

Bauer, Kevin, Yan Chen, Florian Hett, and Michael Kosfeld, "Political Identities, Information Demand and Processing," October 2020. AEA RCT Registry trial number AEARCTR-0006670.

__ , __ , __ , and __ , "Political Identities, Information Demand and Processing (Wave 2)," February 2021. AEA RCT Registry trial number AEARCTR-0007106.

__ , __ , __ , and __ , "Political Identities, Information Demand and Processing (Wave 3)," November 2021. AEA RCT Registry trial number AEARCTR-0008584.

__ , __ , __ , and __ , "Group Identity and Belief Formation: A Decomposition of Political Polarization," 2023. SAFE Working Paper No. 409, http://dx.doi.org/10.2139/ssrn.4670473.

Bénabou, Roland and Jean Tirole, "Mindful economics: The production, consumption, and value of beliefs," *Journal of Economic Perspectives*, 2016, *30* (3), 141–164.

Bruneau, Emile G, Mina Cikara, and Rebecca Saxe, "Minding the Gap: Narrative Descriptions about Mental States Attenuate Parochial Empathy," *PloS One*, 2015, *10* (10), e0140838.

Budak, Ceren, Sharad Goel, and Justin M. Rao, "Fair and Balanced? Quantifying Media Bias through Crowdsourced Content Analysis," *Public Opinion Quarterly*, 04 2016, *80* (S1), 250–271.

Bursztyn, Leonardo and David Y Yang, "Misperceptions about others," *Annual Review of Economics*, 2022, *14* (1), 425–452.

__ , Jonathan Kolstad, Aakaash Rao, Pietro Tebaldi, and Noam Yuchtman, "Polarization and Public Policy: Political Adverse Selection under Obamacare," 2024. Working Paper.

Caplin, Andrew and John Leahy, "Psychological Expected Utility Theory and Anticipatory Feelings*," *The Quarterly Journal of Economics*, 02 2001, *116* (1), 55–79.

Cappelen, Alexander W., Benjamin Enke, and Bertil Tungodden, "Universalism: Global

Evidence," *American Economic Review*, January 2025, *115* (1), 43–76.

Charness, Gary and Yan Chen, "Social Identity, Group Behavior and Teams," *Annual Review of Economics*, 2020.

__ **, Luca Rigotti, and Aldo Rustichini**, "Individual behavior and group membership," *The American Economic Review*, September 2007, *97*, 1340 – 1352.

Chen, Daniel L., Martin Schonger, and Chris Wickens, "oTree—An open-source platform for laboratory, online, and field experiments," *Journal of Behavioral and Experimental Finance*, 2016, *9*, 88–97.

Chen, Yan and Sherry Xin Li, "Group Identity and Social Preferences," *The American Economic Review*, March 2009, *99* (1), 431–457.

__ **and YingHua He**, "Information acquisition and provision in school choice: An experimental study," *Journal of Economic Theory*, 2021, *197*, 105345.

Chowdhury, Subhasish M., "The Economics of Identity and Conflict," 05 2021.

Conlon, John J, Malavika Mani, Gautam Rao, Matthew W Ridley, and Frank Schilbach, "Not Learning from Others," Technical Report, National Bureau of Economic Research 2022.

Dekel, Inbal and Moses Shayo, "Follow the Crowd: But Who Follows, Who Counteracts, and Which Crowd?," 2023. Available at SSRN.

DellaVigna, Stefano and Devin Pope, "What motivates effort? Evidence and expert forecasts," *Review of Economic Studies*, 2017, *85* (2), 1029–1069.

Dimant, Eugen, Kwabena Donkor, Lorenz Goette, Michael Kurschilgen, and Maximilian Mueller, "Identity and Economic Incentives," 2023. Working Paper.

Druckman, James N and Jeremy Levy, "Affective polarization in the American public," in "Handbook on politics and public opinion," Cheltenham, United Kingdom: Edward Elgar Publishing, 2022, pp. 257–270.

Eckel, Catherine C. and Philip J. Grossman, "Managing Diversity by Creating Team Identity," *Journal of Economic Behavior & Organization*, November 2005, *58* (3), 371–392.

Eliaz, Kfir and Andrew Schotter, "Experimental Testing of Intrinsic Preferences for NonInstrumental Information," *The American Economic Review*, 05 2007, *97* (2), 166–169.

__ **and** __ , "Paying for confidence: An experimental study of the demand for non-instrumental information," *Games and Economic Behavior*, 2010, *70* (2), 304 – 324.

Enke, Benjamin, Ricardo Rodriguez-Padilla, and Florian Zimmermann, "Moral universalism: Measurement and economic relevance," *Management Science*, 2022, *68* (5), 3590–3603.

Fryer, Roland, Philipp Harms, and Matthew O Jackson, "Updating Beliefs when Evidence is Open to Interpretation: Implications for Bias and Polarization," *Journal of the European Economic Association*, 08 2019, *17* (5), 1470–1501.

Garcia-Hombrados, Jorge, Marcel Jansen, Ángel Martínez, Berkay Özcan, Pedro Rey-Biel, and Antonio Roldán-Monés, "Ideological Alignment and Evidence-Based Policy Adoption," 2024.

Goeree, Jacob K. and Leeat Yariv, "Conformity in the lab," *Journal of the Economic Science Association*, July 2015, *1* (1), 15 – 28.

Goette, Lorenz, David Huffman, and Stephan Meier, "The Impact of Group Membership on Cooperation and Norm Enforcement: Evidence Using Random Assignment to Real Social Groups," *The American Economic Review*, May 2006, *96* (2), 212–216.

Golman, Russell and George Loewenstein, "Information gaps: A theory of preferences regarding the presence and absence of information," *Decision*, 2018, *5*, 143–164.

Grether, David M., "Bayes Rule as a Descriptive Model: The Representativeness Heuristic*," *The Quarterly Journal of Economics*, 11 1980, *95* (3), 537–557.

Haidt, Jonathan, *The Righteous Mind: Why Good People Are Divided by Politics and Religion*, 1st ed., New York City, NY: Pantheon Books, 2012.

Hett, Florian, Mario Mechtel, and Markus Kröll, "The structure and behavioral effects of revealed social identity preferences," *The Economic Journal*, 2020, *130* (632), 2569–2595.

Iyengar, Shanto, Gaurav Sood, and Yphtach Lelkes, "Affect, not ideologya social identity perspective on polarization," *Public Opinion Quarterly*, 2012, *76* (3), 405–431.

__ **, Yphtach Lelkes, Matthew Levendusky, Neil Malhotra, and Sean J Westwood**, "The origins and consequences of affective polarization in the United States," *Annual Review of Political Science*, 2019, *22*, 129–146.

Kahneman, Daniel and Amos Tversky, "Subjective probability: A judgment of representativeness," *Cognitive Psychology*, 1972, *3* (3), 430–454.

Kranton, Rachel E and Seth G Sanders, "Groupy versus non-groupy social preferences: Personality, region, and political party," *The American Economic Review*, 2017, *107* (5), 65–69.

Kranton, Rachel, Matthew Pease, Seth Sanders, and Scott Huettel, "Deconstructing bias in social preferences reveals groupy and not-groupy behavior," *Proceedings of the National Academy of Sciences*, 2020, *117* (35), 21185–21193.

Kübler, Dorothea and Georg Weizsäcker, "Limited Depth of Reasoning and Failure of Cascade Formation in the Laboratory," *The Review of Economic Studies*, 2004, *71* (2), 425–441.

Lerman, Amy E., Meredith L. Sadin, and Samuel Tratchman, "Policy Uptake as Political Behavior: Evidence from the Affordable Care Act," *American Political Science Review*, 2017, *111* (4), 755–770.

Levy, Ro'ee, "Social Media, News Consumption, and Polarization: Evidence from a Field Experiment," *American Economic Review*, March 2021, *111* (3), 831–70.

Li, Sherry Xin, "Group identity, ingroup favoritism, and discrimination," *Handbook of labor, human resources and population economics*, 2020, pp. 1–28.

Liu, Manwei and Sili Zhang, "The Persistent Effect of Narratives: Evidence from an Online Experiment," 2023. Working paper.

Malmendier, Ulrike and Laura Veldkamp, "Information resonance," Technical Report, Working Paper 2022.

Möbius, Markus M., Muriel Niederle, Paul Niehaus, and Tanya S. Rosenblat, "Managing

Self-Confidence: Theory and Experimental Evidence," *Management Science*, 2022, *68* (11), 7793–7817.

Peterson, Erik and Shanto Iyengar, "Partisan Gaps in Political Information and Information-Seeking Behavior: Motivated Reasoning or Cheerleading?," *American Journal of Political Science*, 2021, *65* (1), 133–147.

Shayo, Moses, "Social Identity and Economic Policy," *Annual Review of Economics*, 2020, *12* (1), 355–389.

Stigler, George J, "The economics of information," *Journal of political economy*, 1961, *69* (3), 213–225.

Tajfel, Henri and John Turner, "An Integrative Theory of Intergroup Conflict," in Stephen Worchel and William Austin, eds., *The Social Psychology of Intergroup Relations*, Monterey, CA: Brooks/Cole, 1979, pp. 94–109.

__ **, Michael Billig, R. Bundy, and Claude L. Flament**, "Social Categorization and Inter-Group Behavior," *European Journal of Social Psychology*, 1971, *1*, 149–177.

Thornton, Rebecca L., "The Demand for, and Impact of, Learning HIV Status," *American Economic Review*, December 2008, *98* (5), 1829–63.