

Polarizing Political Polls: How Visualization Design Choices Can Shape Public Opinion and Increase Political Polarization

Eli Holder and Cindy Xiong Bearfield

Abstract—While we typically focus on data visualization as a tool for facilitating cognitive tasks (e.g. learning facts, making decisions), we know relatively little about their second-order impacts on our opinions, attitudes, and values. For example, could design or framing choices interact with viewers' social cognitive biases in ways that promote political polarization? When reporting on U.S. attitudes toward public policies, it is popular to highlight the gap between Democrats and Republicans (e.g. with blue vs red connected dot plots). But these charts may encourage social-normative conformity, influencing viewers' attitudes to match the divided opinions shown in the visualization. We conducted three experiments examining visualization framing in the context of social conformity and polarization. Crowdforkers viewed charts showing simulated polling results for public policy proposals. We varied framing (aggregating data as non-partisan "All US Adults," or partisan "Democrat" / "Republican") and the visualized groups' support levels. Participants then reported their own support for each policy. We found that participants' attitudes biased significantly toward the group attitudes shown in the stimuli and this can increase inter-party attitude divergence. These results demonstrate that data visualizations can induce social conformity and accelerate political polarization. Choosing to visualize partisan divisions can divide us further.

Index Terms—Political Polarization, Public Opinion, Social Categorization, Survey Data, Social Influence, Attitude Change

1 INTRODUCTION

The United States is increasingly divided between Democrats and Republicans (or, more precisely, between anti-Republicans and anti-Democrats [37]). Americans describe their out-party peers as "close-minded" or "unintelligent" [66] and they may be right, as increased polarization is associated with many aspects of diminished judgment [71]. These partisan-induced impairments impact not just citizens, political polarization is also linked with impaired governance [39] which is a cause for concern given a widening ideology gap in Congress [18].

Could data visualization share some blame for this? Previous work from social and political psychology suggests that efforts to characterize polarization can exaggerate partisanship [57, 60, 73]. With the growing popularity of data visualizations reporting polarized public opinion, differences in attitudes and norms between social groups are increasingly visible and accessible to the general public. Perceived social norms can be influential [42, 52, 73], including through data visualization [4, 64]. This is consistent with a growing body of work showing that data visualization can influence viewers' attitudes and beliefs, beyond what is strictly entailed by the data [35, 36, 53, 70, 87]. Since polarization describes group-level attitude shifts [21], considered together, these dynamics suggest a plausible path wherein visualizing political polarization may actually make it worse.

This study explores pathways from visualization to polarization. We show that visualized attitudes can influence viewers' attitudes through social-normative conformity. Further, when visualizing polarized policy opinions, popular design conventions (i.e. red vs blue partisan dot plots) can influence readers toward more polarized attitudes. In some cases, this effect is stronger than reporting partisan divisions colloquially in text. We also explore alternative, less polarizing framings and find that, while they do not improve polarization, charts that focus on "common ground" at least avoid making it worse.

Contributions: We contribute three experiments demonstrating that data visualization can induce social conformity in readers and divergence in group-level attitudes. We tested conventional approaches to

visualizing public policy opinion and showed that visualizing partisan data can increase political policy polarization.

2 RELATED WORK

We discuss an interdisciplinary set of literature across social and political psychology and data visualization that inspired this work.

2.1 Social Conformity

We are inescapably influenced by the people around us [42]. If they look up, we look up [63]. If they imagine a blinking light, we imagine a blinking light [77]. If they steal ancient artifacts, we steal ancient artifacts [12]. Some argue that the very concepts that define our reality are inherently social [51, 79]. We are particularly influenceable by people who are like us, and indignant toward people who are not [40]. We like people who are like us [41]. We show favoritism toward people inside our social identity groups (our in-groups) and prejudice toward outsiders [82]. But, surprisingly, similarities do not need to be particularly meaningful to be impactful. People who share our birthdays, our hotels, letters in our first names, or our favorite books can have disproportionate influence over us [11, 17, 29].

We are also susceptible to conformity when making judgments under uncertainty [16]. When interpreting data visualizations of public opinion, readers might assume the visualized opinions are (at least) as informed as their own [19]; this implied similarity in knowledge might make these visualized opinions even more persuasive. Prior work also demonstrates that social information about visualized data can influence readers' trust in the data, such that if data disagrees with readers' and others' prior expectations, readers see the data as less credible; in this way social information may facilitate confirmation biases [46, 47]. This suggests that social information, such as visualizations of public policy opinion, has the potential to influence readers' own opinions towards the visualized policies.

2.2 Polarization and Partisanship

Polarization can be affective or ideological [37, 62]. It can imply different distributions of people and attitudes [21]: People can sort into opposing homogeneous groups, without any individual person changing their attitudes. Or whole groups can diverge away from each other toward more extreme attitudes. Or partisan groups can converge in on themselves; by attracting ambivalent moderates away from the center, and attracting radical partisans from the wings, attitude distributions may become more bimodal, but not necessarily more extreme. However it is defined, polarization is generally described as harmful. This may include stoking antipathy between fellow citizens [5, 25, 37, 38],

- Eli Holder is with 3iap (eli@3iap.com)
- Cindy Xiong Bearfield is with Georgia Tech

creating partisan gridlock [39], and even afflicting polarized individuals with cognitive maladies like overconfidence, intolerance, or motivated reasoning [71]. Polarization can arise from individual differences [24] or encouragement and accommodations from social actors such as policymakers or the media [1, 48, 56].

People are sensitive to how political information is framed. Data stories can polarize policy attitudes by pushing readers to conform with their political in-groups or diverge from political out-groups [14]. For example, information with "partisan cues" (e.g. "Democrats support and Republicans oppose allowing prescription drug imports from Canada") can make a policy with otherwise bipartisan support substantially more popular with Democrats and less popular with Republicans [13, 14, 61]. Priming people with partisan branding or concepts can bias beliefs and polarize political attitudes [32, 52]. Polling data can polarize attitudes toward candidates [60], drive support to (perceived) leading candidates [10, 73], and impact voter turnout [31].

Social judgment theory suggests that proposed attitudes are influential in proportion to their discrepancy with prior attitudes, up to a threshold of rejection [8, 30, 78]. The boundaries of this rejection-threshold depend on the information's source, such that ideas that might normally be rejected become acceptable when presented by a trusted party [30, 52]. This implies that seeing the attitudes of a trusted in-group may have a conforming effect, aligned and proportionate to the distance between a viewer's starting attitude and those of the in-group. (e.g. Seeing that your party's attitudes are more extreme than yours might push you toward more extreme attitudes.)

This makes identifying techniques for depolarizing public policy attitudes a critical topic for investigation. For example, reducing party salience in bipartisan communication networks can reduce partisan boomerang effects in interpreting climate change data [32]. Highlighting incidental similarity with peers can increase receptivity to opposing viewpoints [7]. Priming a common national identity can decrease affective polarization [58]. And correcting misperceptions about out-party members can reduce negative attributions and perceived social distance [3, 54]. As data visualizations begin to play a key role in influencing people's attitudes towards news and public policy [74], we are motivated to investigate visualization techniques to communicate public opinion while minimizing the harmful effects of polarization.

2.3 Visualization and Framing Effects

Data visualization can be influential beyond what is rationally supported by the data itself [22, 59]. Visualizations can be emotional [53, 65]. The same data can be framed in ways that give rhetorical advantage to some narratives over others [36, 53]. They can combine with our prior beliefs to encourage illusory correlation [44, 88], illusory causation [87], or motivated innumeracy [6, 32].

Visualization design can have profound impacts on readers' interpretations and takeaways. The spatial arrangement of bar charts can change which data values people compare [28, 86]. Color and annotation choices can enhance objective memory of visualized content [9]. Viewers' recollections of visualizations tend to align more closely with the title than the visualization itself, and title framing can lead to different perspectives [49, 50]. Chart choices for visualizing probability distributions (e.g. bar histograms vs strip plots) can impact estimation accuracy for central tendency or standard deviation [67]. However, design effects related to visualizing uncertainty and variance are not limited to comprehension [15, 34, 43, 85], they can also impact trust [69], perceived effect sizes [34], dichotomization fallacies [84], and even tendencies toward harmful social stereotypes [35].

Data visualization can play many roles in politics, such as guiding public deliberation or policy-making [26, 68]. However, attitudes are not always shaped deliberately, especially in politics where group dynamics can dominate [14]. When reading a narrative visualization [76], people can adjust their attitudes toward or away from the communicated message depending on how the visualized information aligns with their prior beliefs [33, 83, 88]. Visualized social norms can change behaviors like gym attendance or household energy consumption [4, 64], influence interpretations of climate change trends [32], and impact expectations, trust, and recall [46]. Even see-

ing raw numbers of candidate favorability can polarize attitudes toward those candidates [60]. Dichotomous, adversarial comparisons between Democrats and Republicans may further exacerbate polarization by exaggerating stereotypes about political out-parties [3, 35, 54] or diminishing voters' expected influence in elections [27]. Together, this prior work motivates our investigation into how visualizing political attitudes might channel social norms to shape public opinion.

3 STUDY MOTIVATION AND OVERVIEW

We conducted three experiments to understand how public opinion data visualizations can influence viewers' attitudes toward public policy, specifically investigating how visualizations can channel social-normative influences and how these combined psychological forces can contribute to political polarization. Stimuli, survey, data, and analysis scripts are available here: <https://osf.io/xqdw6/>.

3.1 Hypotheses

Building on the ideas above, we hypothesized:

H1: Influence Public opinion polling visualizations may have second-order impacts beyond informing viewers. They can actively influence viewers' attitudes regarding the topics being visualized.

H2: Social Conformity Polling visualizations may shift viewers' attitudes to match the visualized opinions of their relevant in-groups. Charts showing stronger (or weaker) group support for a given policy may bias viewers toward stronger (or weaker) support.

H3: Polarization Polling visualizations' social conformity effects may increase policy polarization (i.e. inter-party attitude divergence).

H4: Framing Effects Visualization framing may influence polarization. Partisan-framed visualizations (e.g. blue vs red dot-plots of party-aggregate opinions) may increase polarization, while emphasizing "common ground" may decrease it (e.g. dot-plots of national-level aggregations, or jitter-plots highlighting within-party variance).

3.2 Experiment Overview

Experiment 1: As an exploratory study, we tested whether realistic public polling visualizations could influence attitudes on gun policies in the United States, an already highly polarized topic (**H1**).

Experiment 2: Using a series of dynamically generated, single-issue charts, covering five different policy topics with randomly varied stimulus values representing group opinions, we investigated the influence of social conformity (**H2**), how this impacts polarization (**H3**), and partisan vs national framing effects (**H4**).

Experiment 3: We tried, and failed, to solve political polarization with chart design (**H4**), but instead replicated our previous results.

4 EXPLORATORY EXPERIMENT 1: GUN-RELATED POLICIES

In Experiment #1, we tested whether public opinion visualizations could influence viewers' opinions (**H1**) and explored a potential pathway between visualization-triggered social conformity and political polarization (as initial investigations of **H2** and **H3**). We used realistic stimuli to increase the external validity and real-world relevance of this initial exploration.

4.1 Experiment 1: Design

Through a between-subject design, we asked participants to report their attitudes toward gun policies after randomly viewing one of three stimuli charts (Fig 1). We did not ask for pre-treatment attitudes because our pilot studies and other prior work suggested this can diminish treatment effects for related (partisan cue) experiments, such that participants seek to keep their pre- and post-treatment responses consistent to avoid perceptions of partisan influence [13]. We looked for differences between **control** and treatment groups to estimate attitude change, using responses from politically aligned participants in the control group as a proxy for the prior attitudes of participants in the treatment groups.



partisan dot plot:**consensus dot plot:****control:**

Fig. 1. Experiment 1 conditions. Complete stimuli available on OSF.

4.1.1 Experiment 1: Stimuli Design

We tested three stimuli charts (Fig 1): The **consensus_dot_plot** and **partisan_dot_plot** both show stacks of 9 gun policy proposals, with the former showing total support for the policy amongst US adults and the latter additionally showing Democrats' and Republicans' support for each policy. The **control** shows public support for the not-yet-politicized topic of bake sale goods.

For **partisan_dot_plot**, the plots show plausible, but exaggerated support levels for each policy. For example, recent polling suggests 83% of democrats support "banning assault-style weapons," but the stimulus showed 87% [75]. The stimulus showed similarly exaggerated opposition, for example 37% of Republicans support assault-weapon bans, but the stimulus showed 25%. We chose more extreme values to increase the likely effect on people with more moderate attitudes toward each policy, based on social judgment theory's predictions that larger discrepancies between prior and proposed attitudes should lead to larger changes in attitudes [8, 30, 78]. We chose gun control because it is timely, divisive, and important, however this may have led to a more muted effect as prior work suggests that political cues are most influential for issues that are not currently politicized [14, 61]. For **consensus_dot_plot**, the values are based on the midpoint of the partisan values shown in **partisan_dot_plot**.

Each chart condition shows an "in-group" with whom participants can identify [52, 58]. For **consensus_dot_plot**, we expected participants to identify with the "total" markers representing US Adults. For **partisan_dot_plot**, we expected participants to identify with their political party. We expect participants' reported attitudes to bias toward their respective in-groups [8, 29, 30, 52, 78]. For example, for **consensus_dot_plot**, attitudes should converge toward the middle of the scale regardless of participants' politics. For **partisan_dot_plot**, attitudes should bias toward participants' respective parties: For moderates, their reported attitudes should diverge from the middle, toward their parties' (relatively) more extreme attitudes. For more extreme partisans, their reported attitudes should converge toward the middle, toward their parties' (relatively) more moderate attitudes.

4.2 Experiment 1: Procedure

The experiment starts with demographic questions (age, gender, ethnicity, political affiliation) and an attention check. On the next page, participants are randomly assigned to see one of the three stimuli: **consensus_dot_plot**, **partisan_dot_plot**, or **control**. We asked participants to answer six comprehension questions to encourage engagement with the stimuli, in addition to a simple attention check and a short free-response question where we asked for 1-2 sentences on their preferred policy or baked good, which we used as an additional attention check. On the next page, participants reported their attitudes toward six specific policy statements on a continuous support scale (0 = strongly oppose, 100 = strongly favor), as shown in the **consensus_dot_plot** and **partisan_dot_plot** stimuli in Fig. 1. They also reported their attitudes toward two more general gun policy statements (e.g. "Greater gun restriction laws are necessary to reduce violence"). On the last page, they reported political views as ideological alignment (0 = very liberal, 100 = very conservative), party affiliation, how strongly they identify with Democrats / Republicans (0 = not at all, 100 = very strongly), and approval of congressional Democrats / Republicans (0 = strongly disapprove, 100 = strongly approve).

4.3 Experiment 1: Participants and Exclusions

We recruited participants on Amazon's Mechanical Turk. We screened for people who did not participate in pilot experiments, reported locations in the United States, have an approval rating above 98%, correctly answered two simple attention checks, provided coherent responses to the short free-response question (e.g. excluding responses that only say "GOOD", or 500-word GPT outputs), and have reputable IP addresses (e.g. not VPNs, according to IPHub, per recommended best practices from social science research, [2, 45]). We further filtered for participants who correctly answered more than half of the chart comprehension checks and excluded independents, very liberal Republicans, or very conservative Democrats, to obtain a more comparable sample of Democrats and Republicans. Details on exclusions are available in the supplementary materials.

Mechanical Turk's worker population leans liberal and recruiting conservatives through their "Premium Qualifications" is slow and expensive. So we recruited toward a more balanced ratio, and ensured stimulus condition balance within party groups, but recognized that we would not achieve overall party parity. We account for this in the study design by choosing a balanced set of policies as our stimulus and measuring political alignment on a continuous scale. We also conducted a power analysis, based on pilot data comparing the effects of A and B ($Cohen's f = 0.02$), suggesting that a target sample of 274 participants per condition would yield sufficient power to detect an overall difference between chart conditions, assuming an alpha level of 0.05.

Given the above, we recruited 161 left-leaning participants and 145 right-leaning participants for **consensus_dot_plot**, 183 left and 132 right participants for **partisan_dot_plot**, and 173 left and 133 right participants for **control**, making up 927 participants total.

4.4 Experiment 1: Analysis Approach

We predicted that visualizing public opinion would influence public opinion (H1). To test this hypothesis, we compared reported attitudes between politically aligned members of the **control** and treatment groups. To analyze these between-group differences, we modeled reported attitudes toward each policy (our dependent variable) as a three-way interaction between the policy's partisan alignment, the participant's partisan alignment, and the stimulus chart condition. Specifically, we used a linear mixed effects model, including participants' IDs as a random effect to account for individual differences. Each policy's partisan alignment was included as a binary variable, either left- or right- favored. Participants' partisan alignment was included as a third-degree polynomial of their average response to the political alignment questions. The stimulus chart was one of: **control**, **consensus_dot_plot**, **partisan_dot_plot**. We also included an interaction between the policy's partisan alignment and factors for participants' backgrounds, to account for demographic differences in prior attitudes. The supplement includes the detailed specification.

We modeled participants' political alignment as a polynomial because we expected the direction of attitude change to alternate up to three times, based on the relative positions of participants' prior attitudes (their political alignment) and their in-group's reference attitude shown in the stimulus. A conformity effect suggests that participants will adjust their attitudes up or down to match their in-group [8,52,78]. So two participants, whose prior attitudes are on opposite sides of their in-group, should change their attitudes in opposite directions to conform with the group. For example, given a plot showing US adults' support for a given policy is 57, we would expect participants' responses to pivot around that reference point. Suppose a participant's prior attitude toward this policy is 19 (strong opposition), and they see that national support is 57 (moderate support). Their support should then increase, from 19 to 57 (a positive bias). Suppose a different participant's prior attitude toward this same policy is 83 (strong support) and they see the same plot. Their support should then decrease, from 83 to 57 (a negative bias). In this example, both participants' attitudes converge toward the reference point (57), but they change in opposite directions given their opposite starting positions. Thus the direction of attitude change pivots around the reference point. When we consider partisan charts (with both Democrat and Republican reference points) that are viewed by both liberal and conservative participants responding to their in-parties, we might expect up to three pivot points where the response curve flips directions: 1) the stimulus values shown for Democrats, 2) near $x=50$, which loosely delineates liberal vs conservative participants (and whether they identify with Democrat or Republican in-groups), and 3) the stimulus values shown for Republicans.

Cronbach's alpha for policy attitude questions showed consistency (liberal policies: $\alpha=0.9$, conservative: $\alpha=0.8$), with one policy more generally unpopular than the others (shorter waiting periods for gun purchases). Factor analysis showed that liberal and conservative policies loaded into separate factors. Political alignment measures also showed consistency ($\alpha=0.9$) and loaded into expected partisan factors.




Fig. 2. The partisan chart polarized moderate conservatives, increasing their support for conservative policies (bottom-right). Plots of reported attitude bias toward different sets of policies (y-axis: the difference between reported attitudes for the treatment, minus control, using estimated marginal means) as a function of their political alignment (x-axis: 0 = most liberal, 100 = most conservative). Uncertainty bands indicate 95% CIs. Positive bias indicates higher than expected support, negative indicates lower than expected. The left panels show responses to liberal policies (e.g. banning assault weapons), the right panels show responses to conservative policies (e.g. expanding concealed-carry). Cells with cross-hatching indicate polarization (i.e. increased support for in-party policies, or decreased support for out-party). Stars indicate significance at $p < 0.05$.

4.5 Experiment 1: Results

We found the expected, significant three-way interaction between chart condition, policy alignment, and participant's political alignment

($\chi^2(6) = 29.5, p < 0.0001$), supporting **H1**. Fig. 2 shows these relationships, as estimated marginal mean differences between the treatment conditions and control. The supplement contains the full statistics, including model fit. Fig. 2 shows three regions where participants' attitudes significantly differ from the control. For **consensus dot plot** participants, the top-left panel shows that moderates increased their support for nationally-popular liberal policies (at $x=[44,54]$, $p < 0.05$). For **partisan dot plot** participants, the bottom-right panel shows that moderate-conservatives increased their support for in-party policies (at $x=[46,64]$, $p < 0.05$). The bottom-left panel shows that extreme-conservatives increased their support for their out-party's policies (at $x[75,100]$, $p < 0.05$).

4.6 Experiment 1: Discussion.

In this exploratory study, we find that public opinion visualizations can influence public opinion (**H1**).

The observed attitude shifts are consistent with a social conformity effect (**H2**), suggesting that participants changed their attitudes toward their reference groups. The **consensus dot plot** stimulus chart showed that liberal gun policies are relatively popular with US Adults. These policies became more popular with politically moderate participants who saw this chart. The **partisan dot plot** stimulus chart showed that conservative gun policies are supported by 44–50% of US adults, and popular with conservatives (supported by 67–84% of Republicans). These policies became more popular with conservative-leaning moderates who saw this chart, indicating conservatives shifting toward the conservative reference point. The same **partisan dot plot** stimulus showed that liberal gun policies are very popular with Democrats (supported by 87–91%), relatively popular with US Adults, and unpopular with Republicans (but still supported by 25–33%). These liberal policies became more popular with extreme conservatives who saw this chart. While this shift indicates conservatives moving in a more liberal direction, it is still consistent with social conformity when considering their prior attitudes; for the more extreme conservatives in this group, whose prior attitudes are more extreme than the Republican reference point, conforming to their in-group implies adopting a more moderate attitude.

The results show mixed support for polarization (**H3**). Seeing the **consensus dot plot** chart depolarized moderates' attitudes, increasing support for consensus-popular policies. However, **partisan dot plot** had opposing effects on conservatives. It increased moderate conservatives' support for in-party policies (shifting away from consensus, consistent with polarization). It also increased extreme conservatives' support for out-party policies (shifting toward consensus, consistent with depolarization). The **partisan dot plot** shifts could be considered polarization based on within-party attitude consolidation [21], however, they do not clearly show inter-party attitude divergence. We explore this further in Experiment 2.

The shape of the curves in Fig. 2 may also be revealing: For moderates, the diverging biases between the treatments (top-row vs bottom-row) show different effects between the consensus and partisan framings (**H4**). The left-right asymmetry between liberals and conservatives may reflect differences in prior attitudes (e.g. liberals uniformly and strongly support gun-control policies) or asymmetry in support levels shown on the stimulus (the charts show, realistically, that liberal-favored gun policies are more generally popular).


While the results suggest a potential social-conformity effect, the story for polarization is less clear. We suspect this is related to design limitations in Experiment 1. First, gun policy is already deeply polarized in the United States. This implies ceiling effects, where participants have little room to become more extreme. It also implies high involvement and prior knowledge of either party's prior attitudes. These factors can make attitudes more resistant to change [14, 30]. Second, while **partisan dot plot** showed exaggerated attitudes for each party, they were based on realistic support levels. Realism, in this case, implies a relatively small gap between participants' prior attitudes and the group attitudes shown in the stimulus charts. This implies another possible ceiling effect, as we expect attitudes to change in proportion to the size of this gap [8, 30, 78]. We address these limitations next.

5 EXPERIMENT 2

We address several design limitations from Experiment 1 and further investigate the underlying processes that enable polling visualizations to lead to polarization. We first profiled attitude change as a function of social conformity (**H2**). Then we investigated how these normative influences lead to the group-level shifts in attitude distributions that underlie polarization (**H3**). Finally, we considered differences between partisan and consensus framing effects (**H4**).

We applied several lessons-learned from Experiment 1. Instead of gun policies, where partisan attitudes in the U.S. are already deeply entrenched and well-known [14, 30], we tested a mix of more obscure policies for which most participants would have little prior knowledge (or interest). We simplified the participant experience. Instead of reviewing a batch of 9 policies and responding on a separate page, we showed each policy individually and asked participants to respond in-line, similar to prior partisan cue studies [13]. To clarify discrepancy-driven social conformity effects [8, 30, 78], we switched to dynamically generated stimulus charts that randomized the values representing each reference group's support levels.

partisan range plot:



consensus dot plot:



partisan text:

Democrats tend to favor and Republicans tend to oppose allowing individuals to import prescription drugs from Canada

control:



Fig. 3. Experiment 2 conditions. Complete stimuli available on OSF.

5.1 Experiment 2: Design

Experiment 2 follows a similar between-subject design as Experiment 1. Based on our previous rationale, participants reported their attitudes toward various policies after randomly viewing one of four stimulus conditions, and we look at differences between control and treatment groups to estimate attitude change.

5.1.1 Experiment 2: Stimuli Design

We tested four stimulus chart conditions (Fig. 3). Two are manipulation conditions: **partisan_range**, which showed a range between blue and red markers, representing Democrat and Republican support, and **consensus_dot_plot**, which showed overall national support for “all US adults.” Two are control conditions, **partisan_text**, a pseudo-control which showed a verbal summary of partisan support levels (e.g. “Democrats tend to favor and Republicans tend to oppose allowing individuals to import prescription drugs from Canada”), and **control**, which showed a neutral image. We included **partisan_text** as a

control, to enable comparing our results with existing political science studies of partisan cues [13, 61], which use similar language to convey categorical partisan support. This condition also provides some limited capacity for exploring modality effects (text vs visualization), to the extent that participants’ respond similarly to categorical and numeric representations of group attitudes [72].

Policy Topics: We chose policy topics that were:

1. Non-polarized (at least compared to gun control, abortion, etc), so that participants’ prior attitudes were less entrenched [14, 61],
2. Less familiar, such that other groups’ support of the policy were not well known (letting us generate plausible synthetic data),
3. Balanced and heterogeneous in terms of plausible partisan support (e.g. estate taxes are unpopular with both parties, but it is plausible that conservatives dislike them more), and
4. Across a range of baseline popularity levels, for external validity and to ensure sufficient distance between participants’ prior attitudes and the simulated attitudes we showed in the charts.

Based on recent prior survey work [13, 55, 80], we chose two bipartisan popular policies (prescription drug imports; self-driving cars), two bipartisanly ambivalent policies with plausible leans (estate taxes; nuclear energy), and one wildcard that could plausibly be supported by either party (genetically modified foods).

Stimulus Values: The charts were generated dynamically, using randomly sampled values representing (fictional, but plausible) levels of support from Democrats and Republicans, based on the following:

1. Values should be continuous and follow a known distribution,
2. Most instances should show one party as clearly favoring the policy, and the other clearly opposing it (so we can observe potential polarization effects),
3. Values should seem realistic (e.g. showing Republicans favoring estate taxes might raise suspicions)
4. Topics and values should balance, such that participants see their party favors the policy half the time (and opposes the other half)

To accomplish this, we assigned four of the five topics to have fixed “favoring” and “opposing” parties (e.g. estate taxes and prescription drug imports typically showed support from Democrats and opposition from Republicans). GMOs were randomly assigned a favoring party. We then sampled two values to represent either party’s mean support for the policy: the “favoring” value was drawn from a beta distribution with mean 66, and the “opposing” value was from a distribution with mean 33. These partisan values were shown in **partisan_range** and **partisan_text**. We calculated a third value representing “All US Adults” as the mid-point of the two partisan values. This mid-point value was shown in **consensus_dot_plot**. We generated and tracked these stimulus values even when participants were assigned to stimuli that did not show the values directly (e.g. the value shown for **consensus_dot_plot** was based on the generated partisan values, but does not show the partisan values; **control** does not show any of the generated values). This lets us compare **control** to the treatment conditions by treating the stimulus values as latent variables representing hypothetical partisan support for each policy; the values existed independently of the chart, but the treatment conditions made them visually salient.

5.2 Experiment 2: Procedure and Participants

The experiment started with demographic questions, an attention check, and how closely one follows U.S. politics. Participants were randomly assigned to one of the four conditions and read the five policy proposals in that condition, one on each page. For each policy, participants reported their attitude on a sliding scale, from 0 (strongly oppose) to 100 (strongly support). For the treatment conditions that contained a visualization, participants answered two comprehension questions per chart to ensure engagement. They lastly reported their political views as ideological alignment (0 = very liberal, 100 = very conservative), party affiliation, and how strongly they identified with Democrats and Republicans (0 = not at all, 100 = very strongly). We




Fig. 4. Experiment 2 results: Visualized attitudes influenced reported attitudes. For all three treatment conditions, participants' attitudes biased toward their visualized in-groups' attitudes. These plots show participants' bias in their reported attitudes toward various policies (y-axis: the mean difference between reported attitudes for the treatment, minus **control**, using estimated marginal means) as a function of their in-group's visualized attitude (x-axis, 0 = in-group opposes, 100 = in-group supports). Positive bias values indicate higher than expected support, negative indicate lower than expected support. The uncertainty ranges indicate 95% confidence intervals. Stars indicate significant differences at $p < 0.05$.

also asked about their affinity toward peer Democrats and Republicans using a feelings thermometer (0 = very cold, 100 = very warm). The experiment ended with a debrief, explaining that the data was synthesized and should be disregarded. We followed the same recruitment criteria from Experiment 1, resulting in 707 left-leaning and 505 right-leaning participants from MTurk, for a total of 1212 participants. Participants were balanced across chart conditions, with a similar ratio of left-to-right leaning participants within each condition.

5.3 Experiment 2: Social Conformity

We predicted that participants' attitudes would conform with the values shown for their in-group (**H2**), where "in-group" could either be "US Adults" or participants' reported political party. We again compare reported attitudes between the control and treatment groups to show the influence of each treatment. To analyze these between-group differences, we modeled reported attitudes toward each policy (our dependent variable) as an interaction between the stimulus chart condition and the participant's in-group's value shown in the stimulus. Specifically, we used separate, but similarly defined linear mixed-effects models, with participants' IDs included as a random effect to account for individual differences. To account for participants' prior attitudes toward a given policy, we included an interaction between the policy topic and the participant's political alignment, as well as their demographic factors. Following our rationale from Experiment 1, we specified the in-group stimulus value and participants' political alignment as third-degree polynomials because we expected the direction of attitude change to pivot multiple times based on the relative positions of prior attitudes and the visualized reference points.

We modeled **partisan_range**, **partisan_text**, and **control** separately from **consensus_dot_plot** and **control** because they show different reference groups, with different simulated values. The **partisan_range** and **partisan_text** charts showed simulated partisan support for each party, with values ranging between 11 and 89. The **consensus_dot_plot** chart showed simulated support for "All US Adults", which we generated as the mid-point of the two partisan values, with values typically ranging between 30 and 70 (per 5.1.1).

To evaluate conformity toward political parties, we tested **partisan_range** and **partisan_text** against **control** and used the value shown for the participant's self-identified political party as the in-group stimulus value. We found the expected significant interaction between chart condition and in-group value ($\chi^2(6) = 33.8$, $p < 0.0001$), indicating that participants' attitudes were influenced by the visualized attitudes of their in-party (supporting **H2**). Fig. 4 (right two panels) shows this relationship. We also found significant main effects for in-group value ($\chi^2(3) = 28.7$, $p < 0.0001$), confirming that participants' attitudes aligned with simulated attitudes. The supplement reports the full statistics and model fit.

To evaluate conformity toward a national-identity group, we tested

consensus_dot_plot against **control**, using the value shown for "All US Adults" as the in-group stimulus value. We found the expected significant interaction between chart condition and in-group value ($\chi^2(2) = 15.0$, $p = 0.002$), indicating that participants' attitudes were influenced by the national in-group (supporting **H2**). Fig. 4 (first panel) shows this relationship. We also found significant main effects for in-group value ($\chi^2(3) = 14.9$, $p = 0.002$), confirming that participants' attitudes aligned with simulated attitudes. The supplement reports the full statistics and model fit.

As shown in Fig. 4, when participants in the **consensus_dot_plot** or **partisan_range** groups saw charts showing in-group opposition to a policy (where the simulated values were less than 50), they reported less support than expected. When the charts showed in-group support, they reported more support than expected. In both cases they conformed their attitudes toward the visualized group norms. A similar pattern holds for **partisan_text**, with an apparent tapering effect.

5.3.1 Aggregate conformity

The previous analysis shows attitudes change as a function of in-group stimulus values, but understanding polarization means evaluating group-level shifts in attitude distributions. For example, Fig. 4 shows that even when the **partisan_range** chart showed in-group support, it had opposing effects depending on the level of support (negative bias from $x=51-57$, positive bias from $x=57-89$). As an intermediate step toward demonstrating polarization, we examined whether these social-conformity effects aggregate to group-level shifts across the range of in-group stimulus values. While this does not necessarily imply polarization, polarization would be implausible without it. To test this, we repeat the previous analysis with one change: instead of modeling in-group stimulus values continuously, we convert them to a categorical factor, indicating whether the in-group showed support (>50) or opposition (≤ 50) to each policy. Experiment 2 was designed to support these categorical comparisons. Even though the in-group stimulus values were randomly generated, they were sampled from distributions centered around 33 (oppose) or 66 (support) to balance the overall instances where participants saw their in-groups supporting or opposing each policy. This lets us examine aggregate conformity across a range of visualized attitudes.

To examine aggregate conformity toward a political party, we tested **partisan_range** and **partisan_text** against **control**. We found the expected significant interaction between the (categorized) in-group value and the chart condition ($\chi^2(2) = 17.2$, $p = 0.0002$), suggesting collective conformity toward participants' in-parties (per **H2**). Dunnett-adjusted post-hoc comparisons (Fig. 5) reveal this is driven by significant differences between the **control** and **partisan_range** ($MD = -4.9$, $p = 0.001$, $d = -0.19$) and the **control** and **partisan_text** ($MD = -4.7$, $p = 0.002$, $d = -0.18$), when the generated charts showed in-group opposition. We also ran post-hoc comparisons between **partisan_range**




Fig. 5. Experiment 2 results: All charts showed aggregate social conformity effects. When participants saw in-group opposition, they reported lower than expected support. These plots show participants' reported attitude bias (x-axis: mean difference between reported attitudes for treatment minus control, using estimated marginal means) after viewing charts showing their in-party opposing (left column) or their in-party supporting a given policy (right column). Positive values indicate higher than expected support, negative indicate lower than expected. The uncertainty ranges indicate 95% CIs. Stars indicate significant differences (* $p < 0.05$, ** $p < 0.01$, *** $p = 0.001$).

and **partisan_text** and found no significant difference. The supplement reports the full statistics.

To examine aggregate conformity toward a national-identity group (US adults), we tested **consensus_dot_plot** against control. We found the expected significant interaction between the (categorized) in-group value and the chart condition ($\chi^2(1) = 10.3$, $p = 0.001$), suggesting collective conformity toward a national consensus (per H2). Dunnett-adjusted post-hoc comparisons (Fig. 5) reveal this is driven by significant differences between control and **consensus_dot_plot** (MD = -3.2, $p = 0.03$, $d = -0.13$) when the generated charts showed in-group opposition. We also note that an alternative analysis of **consensus_dot_plot**, modeling the in-group stimulus value as the hidden value generated for the participant's self-identified political party, might be a more analogous comparison between **consensus_dot_plot** and the two partisan charts. We offer this analysis in the supplement and note that, under these assumptions, the interaction between the categorized stimulus value and the stimulus chart condition is not significant. The supplement reports the full statistics for both approaches.

5.3.2 Social Conformity Discussion

These results show that viewers' attitudes conform toward the visualized attitudes of their in-groups and that data visualizations can channel social-normative influences (supporting H2). As Fig. 4 shows, for participants who saw **consensus_dot_plot**, their responses were biased toward "All US Adults." For those who saw **partisan_range**, their responses were biased toward their political in-party. Participants who saw **partisan_text** also biased toward their in-party, but this effect tapers or reverses at the extremes, possibly reflecting that **partisan_text** conveys categorical values (either "favors" or "opposes"), whereas the visualizations show numeric values.

Both Fig. 4 and Fig. 5 suggest a stronger effect for charts showing in-party opposition compared to in-party support. This is expected given the experiment design. Two of the five policies we tested are popular with both parties, but we showed them as much less popular than reality. For example, participants in the control condition reported strong bipartisan support (82-average) for requiring autonomous vehicles to have a human driver behind the wheel. However, in the treatment conditions, we showed this policy was opposed by Democrats (33-average stimulus values), supported by Republicans (66-average), and mixed nationally (50-average). Since the visualized reference points were typically lower than participants' prior attitudes, we would expect a social conformity effect to lean negatively for this policy.

The response curves in Fig. 4 are consistent a discrepancy mechanism for attitude change [8, 30, 78], where people change their attitudes in proportion to the discrepancy between their prior attitudes and an acceptable reference attitude. The x-axis in these charts can be interpreted as the discrepancy between participants' prior attitudes and the visualized reference attitude. Given a moderately supportive (62-average) overall response from control participants (a baseline estimate of overall prior attitudes), **consensus_dot_plot** and **partisan_range** show that participants' attitudes changed in proportion to the stimulus value's distance from this baseline prior attitude. The conformity effect for **consensus_dot_plot** hints at a viable intervention for reducing policy polarization, since by definition these charts would pull people toward a more moderate national consensus.

5.4 Experiment 2: Polarization and Framing

We predicted that social conformity effects can build to group-level polarization (H3), which we define as between-party attitude divergence. We also predicted that these effects would depend on how the results are visualized (H4). Specifically, we predicted that partisan charts would polarize attitudes and common-ground charts would depolarize. While our previous analysis suggests potential effects of conformity, this may not necessarily imply polarization, which is a relative change in attitude distributions between two groups.




Fig. 6. Experiment 2: **partisan_range** led to significantly more divergent polarization than the other three conditions. Horizontal bars show the mean inter-party attitude distance (gap) between left- and right-leaning participants across all policies. The symmetric distributions on the ends show the bootstrapped samples of how wide the bars could be. Bars are centered horizontally to avoid implying changes in absolute attitude positions for one particular party. Stars indicate significant differences-in-gaps from control based on non-overlapping CIs (* = 95%, *** = 99.9%).

Our dependent variable was the inter-party attitude gap, which is the average absolute difference in attitudes between left-leaning and right-leaning participants across all policies. We evaluate polarization for each treatment as the difference-in-gaps between the treatment and the control. A treatment is polarizing if it shows a wider gap than control, and depolarizing if the gap is smaller. We also evaluate the visualizations' unique impact on polarization (relative to categorical text descriptions used in prior political science work [13, 61]) by comparing **partisan_range** to **partisan_text**. While "polarization" has other colloquial or formal definitions, we focus on the inter-party attitude gap because it is more consistent with our experiment design. To clarify, this measure does not necessarily speak to radicalization or extremism (e.g. if two parties are both extreme and aligned, and an intervention moves one of them to more moderate attitudes, we consider this polarization as the two parties move away from each other). The measure also does not speak to DiMaggio's "consolidation" (intra-party kurtosis) definitions of polarization [21]), which we would not expect to see given the randomly generated stimuli.

We analyzed the inter-party attitude gaps and the difference-in-gaps via bootstrapping. We bootstrapped inter-party gaps for each condition by sampling the response data from the four conditions 10,000 times, with replacement, taking as many samples as responses. We calculated the inter-party gap as the absolute distance between each party's


mean attitude, for each topic, then averaged the gaps across topics. We bootstrapped differences-in-gaps between the treatment and **control** similarly by calculating inter-party gaps for the treatment condition and **control**, then subtracting the treatment's gap by the **control**'s. We repeated this procedure for **partisan_range** vs **partisan_text**.

Fig. 6 summarizes the inter-party gaps for each stimuli condition. We found that both **partisan_range** and **partisan_text** led to significantly wider gaps (more divergent polarization) than **control** (supporting **H3**). The inter-party gap for **partisan_range** was 19.8 points (95% CI = [17.0, 22.5]), or 8.1 points wider than **control** (MD = 8.1, 95% CI = [5.1, 12.3]). The inter-party gap for **partisan_text** was 15.1 points (95% CI = [11.9, 17.2]), or 3.4 points wider than **control** (MD = 3.4, 95% CI = [0.0, 7.1]). Differences between **consensus_dot_plot** and **control** were not significant. We also found that **partisan_range** led to significantly more polarization than **partisan_text**, with a 4.7 point wider gap (MD = 4.7, 95% CI = [1.5, 9.0]).

5.5 Experiment 2: Discussion

The results suggest that public opinion visualizations exert a social-normative influence on viewers' attitudes (supporting **H2**), leading to social conformity effects for both national and partisan identity framings. This conformity effect can build to inter-party polarization (supporting **H3**) when visualized with partisan-framed charts like **partisan_range**. Given that both partisan treatments increased polarization, while **consensus_dot_plot** did not, the results also imply framing effects for design choices (supporting **H4**). The wider inter-party gap for **partisan_range** vs **partisan_text** (19.8 vs 15.1 points) may suggest that visualizing polarization may be uniquely polarizing, at least relative to reporting similar gist information in text. However, this needs further research to differentiate between the effects of modality (text vs visualization) and specificity ("favor"/"oppose" vs numeric values). Even though **consensus_dot_plot** led to significant social-conformity effects, and by definition it shows a moderate reference point, it did not reduce between-party polarization.

partisan jitter plot:



consensus jitter plot:




Fig. 7. Experiment 3 conditions. Complete stimuli available on OSF.

6 EXPERIMENT 3

While we have shown that partisan framing can polarize, we have had limited success with designs intended to depolarize. Experiment 1 showed that **consensus_dot_plot** increased moderates' support for consensus gun-control policies. Experiment 2 showed that **consensus_dot_plot** also exerts a social-conformity force, but had no effect on divergent polarization. In this final experiment, we look to other theories for design ideas that might depolarize policy attitudes (per **H4**). In particular, we suspected that the dichotomization of partisan

opinions may contribute to further polarization. Viewing out-parties as monoliths can increase harmful stereotypes, potentially influencing policy attitudes through affective polarization [3, 23, 35, 54]. On the other hand, increasing perceived similarity [7] and correcting misconceptions of out-party composition and animus [3, 54] seem like viable interventions for depolarization. Through some combination of these mechanics, we suspected that highlighting the variance of intra-party attitudes may help depolarize viewers' attitudes.

6.1 Experiment 3: Participants, Design, and Procedure

We followed similar recruiting, exclusion protocols, experimental design, and procedures as the previous experiment, with a few changes. Instead of using MTurk's "Premium Qualifications" for partial partisan balance, we recruited until all cells had at least 125 participants, then randomly sampled 125 participants from each cell, for a fully-balanced design. This resulted in 750 participants total, 250 per condition, and 375 each of left- and right-leaning participants. The slightly lower participant count reduced our expected power from 90% to 87%.

Stimuli Design: We also tested two new designs (Fig. 7). **consensus_jitter_plot** used a common-ground framing (showing "U.S. Adults") and jitter dots representing the individual attitudes of several hundred (simulated) U.S. Adults. We suspected that showing a breadth of attitudes may additionally dispel narratives Americans' attitudes are irreconcilably divided. **partisan_jitter_plot** used a partisan-group framing, showing average support for either party, combined with jitter dots representing the individual attitudes of several hundred (simulated) partisans. We suspected that, if illusory dichotomization were an issue, explicitly showing the overlap between parties may help dispel exaggerated polarization beliefs and minimize inter-party stereotypes. We generated group-level stimulus values for the parties and U.S. adults following the procedure from Experiment 2. We then sampled an additional 250 values per group, from matching beta distributions, and used these values to place the jitter dots. To create a more fair test for depolarization, we also used only the two most partisan policy topics from Experiment 2: increasing estate taxes (liberal-favored) and expanding nuclear energy (conservative-favored).




Fig. 8. Experiment 3: Neither design reduced polarization. **partisan_jitter_plot** made it worse. See Fig. 6 for chart interpretation notes.

6.2 Experiment 3: Results

We predicted that charts showing within-group variance could depolarize participants' policy attitudes (per **H4**). Following our analysis from Experiment 2 to evaluate between-party attitude divergence, Fig. 8 shows results from bootstrapping inter-party attitude gaps. We found that the inter-party gap for **partisan_jitter_plot** was 31.8 points (95% CI [27.3, 36.5]), or 12.9 points wider than **control** (MD = 12.9, 95% CI [6.2, 19.7]). Differences-in-gaps between **consensus_jitter_plot** and **control** were not significant.

6.3 Experiment 3: Discussion

While visually dichotomizing groups contributes to harmful stereotypes, and out-party stereotyping contributes to polarization [3, 35, 54], our attempts to highlight intra-party variation in political attitudes had unexpectedly polarizing effects. Instead of improving polarization, highlighting intra-group attitude variance either did not help (**consensus_jitter_plot**) or made it much worse (**partisan_jitter_plot**), corroborating findings from Experiment 2.

For **partisan_jitter_plot**, there could have been a number of additional factors to consider, such as the relationship with affective polarization as a mediator, sensitivity to the range of values (i.e. showing individual dots implies showing people on both extremes, which may imply even more extreme social norms), or perhaps just that it is even more red and blue ink on the page to activate partisan identities. **consensus_jitter_plot** results were similar to **consensus_dot_plot**: The treatment did not help polarization, but it did not hurt.

While **partisan_jitter_plot** (Fig. 8) shows a wider gap compared to **partisan_range** from Experiment 2 (Fig. 6), this reflects differences between experiments, not differences between the charts. Experiment 2 included responses to three additional policies with more baseline bipartisan support, dampening polarization for all conditions.

7 EXPERIMENT SUMMARY

Our experiments show a path from social conformity to polarization.

In Experiment 1, we show that visualizing policy opinions can shape policy opinions, particularly for moderate partisans. We used a realistic visualization of gun policy attitudes to show that a popular chart type is not merely a passive source of political information; rather, it can actively shape our politics.

In Experiment 2, we showed that data visualization can induce social conformity, and as a consequence, visualizing polarization can increase polarization. We also found that visualization can have a stronger influence on partisan attitudes than the verbal gist summaries studied in political science.

In Experiment 3, our designs, intended to decrease polarization, either made it worse or had no effect, replicating our previous results.

8 DISCUSSION AND DESIGN IMPLICATIONS

Our results establish two effects between social psychology and data visualization: 1) social conformity in attitudes can be induced through data visualization and 2) visualizing political polarization can accelerate it. The effects further question visualizations' presumed neutrality. Other work questions rhetorical neutrality [36, 53] or suggests that causally-neutral visualizations lead to unsupported beliefs [20, 35]. These results question neutrality-as-passivity, showing that visualizations can actively shape the phenomenon that they portray.

In recent debates on newsroom neutrality, one dominating view encourages journalists to use intentionally narrow criteria when deciding what to publish ("Is it true? Is it important?") and be "profoundly skeptical" of subjective concerns about harmful outcomes [81]. These cultural norms can affect data journalists, who are aware of the persuasive influence of their work, but fear stigmatization if they "talk openly" about it [53]. Our results highlight the potentially harmful consequences for publishing partisan polling results, a popular news topic. They corroborate others' suggestions that stories of polarization may be self-perpetuating [14, 48, 57]. Together, this suggests not only a need for more open discussion but also provides grounds for reconsidering a default posture of consequential skepticism.

We extend prior political science work [13, 61], showing that partisan cue effects translate to data visualization. We also show that visualization can have a stronger effect than the verbal, categorical descriptions of partisan support used in these prior experiments, suggesting that visualization (or, at least, numerically represented attitudes) can add a unique dimension to established partisan-conformity effects.

Most urgently, the conformity effects may have consequences for public health communication. To promote health equity, institutions often publish charts highlighting adherence disparities between social groups (e.g. differences in vaccination rates). To the extent that viewers interpret these disparities as representing groups' normative attitudes (e.g. "vaccines are unpopular with people like me"), these charts may backfire and further reduce support within the (marginalized) communities they are intended to help. Further research on social conformity in public health visualizations is an urgent next step.

Design Implications: Designers visualizing social norms should understand that they influence their audiences, not just inform them.

Choosing to show that an attitude is popular (or not) within a certain social group can make it more (or less) popular for that group. This implies that, when deciding what to visualize, whatever benefits readers might gain from improved sense-making should be offset by the risks (or benefits) of influencing their attitudes (e.g. further polarization).

9 LIMITATIONS AND FUTURE DIRECTIONS

As an initial investigation of polarization through visualization, we discuss several limitations and suggest directions for future research.

Other visualizations. While we found that popular visualization approaches can increase polarization, we have yet to find designs that can reduce it. We suspect that other non-partisan aggregations (e.g. by geography, perceived expertise, other social affiliations), the salience of partisan branding (e.g. blue donkeys for Democrats), the visual range of markers across the scale, or other ways of conveying common ground between parties may still be viable interventions.

Other polarization dynamics. We focused on polarization as inter-party divergence, but other definitions may reveal other dynamics. We suspect Dimaggio's consolidation definitions might clarify a path toward depolarizing visualizations [21].

Other topics. Experiments 2 and 3 used relatively non-partisan topics as case studies, referencing examples used in prior work [13, 55, 80]. However, we found that even these showed baseline polarization. Exploring truly neutral topics may help clarify how opinion polling visualizations contribute to the early stages of a topic becoming polarized (e.g. masking, prior to Covid-19). There may also be asymmetries found in different framings of the same topic, for example where participants may support "increasing gun restrictions" more than they oppose "increasing gun access."

Individual differences. While the experiment was not intended to explore individual differences, our results showed that demographic factors could be significant (see the supplementary materials). Future work can explore the impacts of prior policy-related beliefs, affective polarization, actual political knowledge (vs self-reported familiarity), or needs for security, certainty, and individuation [24].

Positive vs negative partisanship. Negative partisanship (i.e. biasing away from an out-party) is often cited as more influential than positive partisanship (i.e. biasing toward an in-party). So while we might expect participants to be more repulsed by their out-party than attracted to their in-party, our experiments were not designed to differentiate between these two forces. Future work could explore even more obscure topics to isolate these effects.

Post-only design. Our experiments only gathered participants' attitudes and political alignment after they viewed the stimuli. To understand individual effects (e.g. how many people change their attitudes and by how much?), future studies could use repeated measures designs or explore other interesting attitude proxies, such as collecting predicted attitudes for *other* people.

United States focus. While political polarization exists around the world, this study was designed specifically for participants enculturated within the United States, including the topics chosen and the political parties featured in the visualizations. Further work is needed to see if these results generalize to other countries and cultures. The U.S. is also a predominantly two-party system; future work could also explore implications for multi-party politics.

ACKNOWLEDGMENTS

We are grateful for the participants who made this study possible and our reviewers for their insightful feedback. We also thank Jeremy Wilmer and Sarah Kerns for their advice and encouragement, as well as Amanda Holder for her support and infinite patience. This work was partially supported by NSF award IIS-2237585.

REFERENCES

- [1] A. I. Abramowitz. *The Disappearing Center: Engaged Citizens, Polarization, and American Democracy*. Yale University Press, 2010.

- [2] J. Agle, Y. Xiao, R. Nolan, and L. Golzarri-Arroyo. Quality control questions on amazon's mechanical turk (mturk): A randomized trial of impact on the usaudit, phq-9, and gad-7. *Behavior research methods*, 54(2):885–897, 2022.
- [3] D. J. Ahler and G. Sood. The parties in our heads: Misperceptions about party composition and their consequences. *The Journal of Politics*, 80(3):964–981, 2018.
- [4] H. Allcott and S. Mullainathan. Behavior and energy policy. *Science*, 327(5970):1204–1205, 2010.
- [5] S. Atske. Partisan antipathy: More intense, more personal, 2019. Pew Research Center.
- [6] M. Baekgaard, J. Christensen, C. M. Dahmann, A. Mathiasen, and N. B. G. Petersen. The role of evidence in politics: Motivated reasoning and persuasion among politicians. *British Journal of Political Science*, 49(3):1117–1140, 2019.
- [7] S. Baliotti, L. Getoor, D. G. Goldstein, and D. J. Watts. Reducing opinion polarization: Effects of exposure to similar people with differing political views. *Proceedings of the National Academy of Sciences*, 118(52):e2112552118, 2021.
- [8] S. Bochner and C. A. Insko. Communicator discrepancy, source credibility, and opinion change. *Journal of Personality and Social Psychology*, 4:614–621, 1966.
- [9] M. A. Borkin, Z. Bylinskii, N. W. Kim, C. M. Bainbridge, C. S. Yeh, D. Borkin, H. Pfister, and A. Oliva. Beyond memorability: Visualization recognition and recall. *IEEE transactions on visualization and computer graphics*, 22(1):519–528, 2015.
- [10] A. Boukouras, W. Jennings, L. Li, and Z. Maniadis. Can biased polls distort electoral results? evidence from the lab. *European Journal of Political Economy*, page 102383, 2023.
- [11] J. M. Burger, N. Messian, S. Patel, A. del Prado, and C. Anderson. What a coincidence! the effects of incidental similarity on compliance. *Personality and Social Psychology Bulletin*, 30(1):35–43, 2004.
- [12] R. B. Cialdini, L. J. Demaine, B. J. Sagarin, D. W. Barrett, K. Rhoads, and P. L. Winter. Managing social norms for persuasive impact. *Social Influence*, 1(1):3–15, 2006.
- [13] S. Clifford, G. Sheagley, and S. Piston. Increasing precision without altering treatment effects: Repeated measures designs in survey experiments. *American Political Science Review*, 115(3):1048–1065, 2021.
- [14] G. L. Cohen. Party over policy: The dominating impact of group influence on political beliefs. *Journal of Personality and Social Psychology*, 85(5):808–822, 2003.
- [15] M. Correll and M. Gleicher. Error bars considered harmful: Exploring alternate encodings for mean and error. *IEEE transactions on visualization and computer graphics*, 20(12):2142–2151, 2014.
- [16] L. Cosmides and J. Tooby. Are humans good intuitive statisticians after all? rethinking some conclusions from the literature on judgment under uncertainty. *cognition*, 58(1):1–73, 1996.
- [17] D. Cwir, P. B. Carr, G. M. Walton, and S. J. Spencer. Your heart makes my heart move: Cues of social connectedness cause shared emotions and physiological states among strangers. *Journal of Experimental Social Psychology*, 47(3):661–664, 2011.
- [18] D. Desilver. The polarization in today's congress has roots that go back decades, 2022. Pew Research Center.
- [19] M. Deutsch and H. B. Gerard. A study of normative and informational social influences upon individual judgment. *The journal of abnormal and social psychology*, 51(3):629, 1955.
- [20] C. D'ignazio and L. F. Klein. *Data feminism*. MIT press, 2020.
- [21] P. DiMaggio, J. Evans, and B. Bryson. Have americans' social attitudes become more polarized? *American Journal of Sociology*, 102:690–755, 1996.
- [22] E. Dimara, S. Franconeri, C. Plaisant, A. Bezerianos, and P. Dragicevic. A task-based taxonomy of cognitive biases for information visualization. *IEEE transactions on visualization and computer graphics*, 26(2):1413–1432, 2018.
- [23] J. N. Druckman, S. Klar, Y. Krupnikov, M. Levendusky, and J. B. Ryan. Affective polarization, local contexts and public opinion in america. *Nature Human Behaviour*, 5(1):28–38, 2021.
- [24] C. M. Federico. When do psychological differences predict political differences?: Engagement and the psychological bases of political polarization. In *The Psychology of Political Polarization*. Routledge, 2021.
- [25] E. J. Finkel, C. A. Bail, M. Cikara, P. H. Ditto, S. Iyengar, S. Klar, L. Mason, M. C. McGrath, B. Nyhan, D. G. Rand, L. J. Skitka, J. A. Tucker, J. J. Van Bavel, C. S. Wang, and J. N. Druckman. Political sectarianism in america. *Science*, 370(6516):533–536, 2020.
- [26] S. L. Franconeri, L. M. Padilla, P. Shah, J. M. Zacks, and J. Hullman. The science of visual data communication: What works. *Psychological Science in the Public Interest: A Journal of the American Psychological Society*, 22(3):110–161, 2021.
- [27] R. A. Furrer, K. Schloss, G. Lupyan, P. M. Niedenthal, and A. Wood. Red and blue states: dichotomized maps mislead and reduce perceived voting influence. *Cognitive Research: Principles and Implications*, 2023.
- [28] A. Gaba, V. Setlur, A. Srinivasan, J. Hoffswell, and C. Xiong. Comparison conundrum and the chamber of visualizations: An exploration of how language influences visual design. *IEEE Transactions on Visualization and Computer Graphics*, 29(1):1211–1221, 2022.
- [29] N. J. Goldstein, R. B. Cialdini, and V. Griskevicius. A room with a viewpoint: Using social norms to motivate environmental conservation in hotels. *Journal of Consumer Research*, 35:472–482, 2008.
- [30] E. A. Griffin. Chapter 15: Social judgment theory of muzafer sherif. In *A first look at communication theory*. New York : McGraw-Hill, 2012.
- [31] J. Großer and A. Schram. Public opinion polls, voter turnout, and welfare: An experimental study. *American Journal of Political Science*, 54(3):700–717, 2010.
- [32] D. Guilbeault, J. Becker, and D. Centola. Social learning and partisan bias in the interpretation of climate trends. *Proceedings of the National Academy of Sciences*, 115(39):9714–9719, 2018.
- [33] J. Heyer, N. K. Raveendranath, and K. Reda. Pushing the (visual) narrative: the effects of prior knowledge elicitation in provocative topics. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–14, 2020.
- [34] J. M. Hofman, D. G. Goldstein, and J. Hullman. How visualizing inferential uncertainty can mislead readers about treatment effects in scientific results. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, CHI '20, pages 1–12. Association for Computing Machinery, 2020.
- [35] E. Holder and C. Xiong. Dispersion vs disparity: Hiding variability can encourage stereotyping when visualizing social outcomes. *IEEE Transactions on Visualization and Computer Graphics*, 29(1):624–634, 2023.
- [36] J. Hullman and N. Diakopoulos. Visualization rhetoric: Framing effects in narrative visualization. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2231–2240, 2011.
- [37] S. Iyengar, G. Sood, and Y. Lelkes. Affect, not ideology: A social identity perspective on polarization. *Public Opinion Quarterly*, 76(3):405–431, 2012.
- [38] S. Iyengar and S. J. Westwood. Fear and loathing across party lines: New evidence on group polarization. *American Journal of Political Science*, 59(3):690–707, 2015.
- [39] G. C. Jacobson. Polarization, gridlock, and presidential campaign politics in 2016. *The ANNALS of the American Academy of Political and Social Science*, 667(1):226–246, 2016.
- [40] D. R. Jhangiani and D. H. Tarry. Ingroup favoritism and prejudice. 2022.
- [41] D. R. Jhangiani and D. H. Tarry. Initial attraction. 2022.
- [42] D. R. Jhangiani and D. H. Tarry. The many varieties of conformity. 2022.
- [43] A. Kale, M. Kay, and J. Hullman. Visual reasoning strategies for effect size judgments and decisions. *IEEE transactions on visualization and computer graphics*, 27(2):272–282, 2020.
- [44] A. Karduni, D. Markant, R. Wesslen, and W. Dou. A bayesian cognition approach for belief updating of correlation judgement through uncertainty visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):978–988, 2020.
- [45] R. Kennedy, S. Clifford, T. Burleigh, P. D. Waggoner, R. Jewell, and N. J. Winter. The shape of and solutions to the mturk quality crisis. *Political Science Research and Methods*, 8(4):614–629, 2020.
- [46] Y.-S. Kim, K. Reinecke, and J. Hullman. Data through others' eyes: The impact of visualizing others' expectations on visualization interpretation. *IEEE transactions on visualization and computer graphics*, 24(1):760–769, 2017.
- [47] Y.-S. Kim, K. Reinecke, and J. Hullman. Explaining the gap: Visualizing one's predictions improves recall and comprehension of data. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pages 1375–1386, 2017.
- [48] E. Klein. *Why We're Polarized*. Avid Reader Press / Simon & Schuster, 2020.
- [49] H.-K. Kong, Z. Liu, and K. Karahalios. Frames and slants in titles of visualizations on controversial topics. In *Proceedings of the 2018 CHI conference on human factors in computing systems*, pages 1–12, 2018.

- [50] H.-K. Kong, Z. Liu, and K. Karahalios. Trust and recall of information across varying degrees of title-visualization misalignment. In *Proceedings of the 2019 CHI conference on human factors in computing systems*, pages 1–13, 2019.
- [51] I. Y.-M. Lau, C.-y. Chiu, and S.-I. Lee. Communication and shared reality: Implications for the psychological foundations of culture. *Social Cognition*, 19(3):350–371, 2001.
- [52] A. Ledgerwood and S. Chaiken. Priming us and them: Automatic assimilation and contrast in group attitudes. *Journal of Personality and Social Psychology*, 93:940–956, 2007.
- [53] E. Lee-Robbins and E. Adar. Affective learning objectives for communicative visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 29(1):1–11, 2023.
- [54] J. Lees and M. Cikara. Inaccurate group meta-perceptions drive negative out-group attributions in competitive contexts. *Nature Human Behaviour*, 4(3):279–286, 2020.
- [55] R. Leppert. Americans continue to express mixed views about nuclear power, 2022. Pew Research Center.
- [56] M. Levendusky. *The Partisan Sort: How Liberals Became Democrats and Conservatives Became Republicans*. University of Chicago Press, illustrated edition edition, 2009.
- [57] M. Levendusky and N. Malhotra. Does media coverage of partisan polarization affect political attitudes? *Political Communication*, 33(2):283–301, 2016.
- [58] M. S. Levendusky. Americans, not partisans: Can priming american national identity reduce affective polarization? *The Journal of Politics*, 80(1):59–70, 2018.
- [59] Y. Luo and J. Zhao. Attentional and perceptual biases of climate change. *Current Opinion in Behavioral Sciences*, 42:22–26, 2021.
- [60] G. Madson and D. S. Hillygus. All the best polls agree with me: Bias in evaluation of political polling, 2018.
- [61] A. Malka and Y. Lelkes. More than ideology: Conservative-liberal identity and receptivity to political cues. *Social Justice Research*, 23(2):156–188, 2010.
- [62] L. Mason. “i disrespectfully agree”: The differential effects of partisan sorting on social and issue polarization. *American Journal of Political Science*, 59(1):128–145, 2015.
- [63] S. Milgram, L. Bickman, and L. Berkowitz. Note on the drawing power of crowds of different size. *Journal of Personality and Social Psychology*, 13:79–82, 1969.
- [64] K. L. Milkman, D. Gromet, H. Ho, J. S. Kay, T. W. Lee, P. Pandiloski, Y. Park, A. Rai, M. Bazerman, J. Beshears, L. Bonacorsi, C. Camerer, E. Chang, G. Chapman, R. Cialdini, H. Dai, L. Eskreis-Winkler, A. Fishbach, J. J. Gross, S. Horn, A. Hubbard, S. J. Jones, D. Karlan, T. Kautz, E. Kirgios, J. Klusowski, A. Kristal, R. Ladhania, G. Loewenstein, J. Ludwig, B. Mellers, S. Mullainathan, S. Saccardo, J. Spiess, G. Suri, J. H. Talloen, J. Taxer, Y. Trope, L. Ungar, K. G. Volpp, A. Whillans, J. Zinman, and A. L. Duckworth. Megastudies improve the impact of applied behavioural science. *Nature*, 600(7889):478–483, 2021.
- [65] F. F. Mueller, T. Dwyer, S. Goodwin, K. Marriott, J. Deng, H. D. Phan, J. Lin, K.-T. Chen, Y. Wang, and R. Ashok Khot. Data as delight: Eating data. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI ’21, pages 1–14. Association for Computing Machinery, 2021.
- [66] R. Nadeem. As partisan hostility grows, signs of frustration with the two-party system, 2022. Pew Research Center.
- [67] E. Newburger, M. Correll, and N. Elmqvist. Fitting bell curves to data distributions using visualization. *IEEE Transactions on Visualization and Computer Graphics*, 2022.
- [68] T. U. Nærlund. The political significance of data visualization: Four key perspectives. In M. Engebretsen and H. Kennedy, editors, *Data Visualization in Society*, pages 63–74. Amsterdam University Press, 2020.
- [69] L. Padilla, H. Hosseinpour, R. Fygenon, J. Howell, R. Chunara, and E. Bertini. Impact of COVID-19 forecast visualizations on pandemic risk perceptions. *Scientific Reports*, 12(1):2014, 2022.
- [70] A. V. Pandey, K. Rall, M. L. Satterthwaite, O. Nov, and E. Bertini. How deceptive are deceptive visualizations? an empirical analysis of common distortion techniques. In *Proceedings of the 33rd annual acm conference on human factors in computing systems*, pages 1469–1478, 2015.
- [71] J.-W. v. Prooijen. The psychology of political polarization: An introduction. In *The Psychology of Political Polarization*. Routledge, 2021.
- [72] V. Reyna. A theory of medical decision making and health: fuzzy trace theory. *Journal of the Society for Medical Decision Making*, 2008.
- [73] D. Rothschild and N. Malhotra. Are public opinion polls self-fulfilling prophecies? *Research & Politics*, 1(2):2053168014547667, 2014.
- [74] P. Sánchez-Holgado, C. Arcila-Calderón, and M. Frías-Vázquez. The effect of interest and attitude on public comprehension of news with data visualization. *Frontiers in Communication*, 8:1064184, 2023.
- [75] K. Schaeffer. Key facts about americans and guns, 2021. Pew Research Center.
- [76] E. Segel and J. Heer. Narrative visualization: Telling stories with data. *IEEE transactions on visualization and computer graphics*, 16(6):1139–1148, 2010.
- [77] M. Sherif. *The psychology of social norms*. The psychology of social norms. Harper, 1936.
- [78] M. Sherif and C. I. Hovland. *Social judgment: Assimilation and contrast effects in communication and attitude change*. Social judgment: Assimilation and contrast effects in communication and attitude change. Yale Univ. Press, 1961.
- [79] S. A. Sloman. How do we believe? *Topics in Cognitive Science*, 14(1):31–44, 2022.
- [80] A. Smith. Americans’ attitudes toward driverless vehicles, 2017. Pew Research Center.
- [81] A. Sulzberger. Journalism’s essential value. *Columbia Journalism Review*, 2023.
- [82] H. Tajfel, M. G. Billig, R. P. Bundy, and C. Flament. Social categorization and intergroup behaviour. *European Journal of Social Psychology*, 1(2):149–178, 1971.
- [83] R. Wesslen, S. Santhanam, A. Karduni, I. Cho, S. Shaikh, and W. Dou. Investigating effects of visual anchors on decision-making about misinformation. In *Computer Graphics Forum*, volume 38, pages 161–171. Wiley Online Library, 2019.
- [84] J. B. Wilmer and S. H. Kerns. What’s really wrong with bar graphs of mean values: variable and inaccurate communication of evidence on three key dimensions, 2022.
- [85] C. Xiong, A. Sarvghad, D. G. Goldstein, J. M. Hofman, and Ç. Demiralp. Investigating perceptual biases in icon arrays. In *CHI Conference on Human Factors in Computing Systems*, pages 1–12, 2022.
- [86] C. Xiong, V. Setlur, B. Bach, E. Koh, K. Lin, and S. Franconeri. Visual arrangements of bar charts influence comparisons in viewer takeaways. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):955–965, 2021.
- [87] C. Xiong, J. Shapiro, J. Hullman, and S. Franconeri. Illusion of causality in visualized data. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):853–862, 2020.
- [88] C. Xiong, C. Stokes, Y.-S. Kim, and S. Franconeri. Seeing what you believe or believing what you see? belief biases correlation estimation. *IEEE Transactions on Visualization and Computer Graphics*, 29(1):493–503, 2023.