

Appendix to In-group interest cues do not change issue attitudes

Clara Vandeweerd

Massachusetts Institute of Technology
30 Wadsworth Street, Cambridge, MA 02142
ORCID: 0000-0002-3024-1441

March 24, 2021

Contents

1	Design details and justifications	4
1.1	Procedure	4
1.2	Samples and recruitment	5
1.2.1	Experiment 1: Gender	7
1.2.2	Experiment 2: Race/ethnicity	7
1.2.3	Experiment 3: LGBT	8
1.3	Dependent variables: concern, importance, spending support	9
1.4	Group identification	10
1.5	Group–issue connection beliefs	11
1.6	Self-interest and linked fate	13
1.7	Model specifications	14
2	Effects and confidence intervals	22
3	Robustness checks	24
3.1	Pooling all experiments	25
3.2	Looking at single issues	27
3.3	Change in beliefs	28
3.3.1	Treatment as an instrument for belief change	28
3.3.2	Moderation by prior beliefs	30
3.3.3	Group identity centrality and beliefs	30
3.4	Ceiling effects and anchoring	33
3.5	Alternative explanations	34
4	Questionnaire	38

4.1	Informed consent	38
4.2	Group membership	39
4.3	Group identity	40
4.3.1	Centrality	40
4.4	Issue attitudes	40
4.4.1	Concern	40
4.4.2	Importance	41
4.4.3	Spending support	42
4.5	Self-interest	44
4.6	Prior/Posterior Beliefs	45
4.7	Just world beliefs	47
4.8	Treatment	47
4.8.1	Gender	48
4.8.2	Race	48
4.8.3	LGBT	49

1 Design details and justifications

1.1 Procedure

Figure 1 shows how participants moved through the different steps in Experiment 1 (on gender) and 3 (on LGBT issues), including which questions were asked before and after treatment. Experiment 1 occurs in two phases: a pre-treatment and (one week later) a treatment phase. The design for Experiment 2 (on race) is analogous to Experiment 1, with different groups and issues. In Experiment 3, a pre-post design was not a good option, because the pool of potential LGBT participants was small—even a moderate drop-out rate would have led to an undersized final sample. After measuring covariates, participants were immediately treated. Then they answered post-treatment attitude questions.

I preferred a pre-post design where possible because, first, the measurement of group identities is separated from the treatment. This avoids making identities salient for all participants (including the control group). Second, being able to control for pre-treatment measures of the dependent variables (in this case by first-differencing the dependent variables) makes the treatment effect estimates more precise. Third, prior beliefs about the group–issue connection can be measured before and after treatment, without the treatment coming across as a correction. First-differencing beliefs has statistical benefits as well (see section 3.3 below). Fourth, first-differencing improves causal identification when I investigate whether the effect of group interests is mediated by self-interest. Briefly put, the mediator must be as-if random controlling for the treatment (Imai et al., 2010). Because both the mediator

(self-interest) and the outcome (issue attitude) are first-differenced, the only confounders that can interfere with as-if randomness are ones that change between phases, within person-issue combinations.

In all experiments, participants can be assigned to learn about one of two group-related issues (or to be in the control group). I include two issues per group to diminish (though of course not eliminate) the possibility that any inter-group differences in effect sizes can be ascribed to differences in the issues they were linked to. Covering a total of ten issues in all experiments combined also helps me find any patterns (or a lack of them) in the kinds of issues that are most sensitive to treatment. Section 3.5 below investigates a few of these patterns.

As the design flowcharts show, the experiment include two measures that are not featured in the paper’s main analyses: prior/posterior beliefs, and just world beliefs. These measures are used in sections 3.3 and 3.5 of this Appendix.

1.2 Samples and recruitment

Below, I describe the sample characteristics and recruitment strategies for each experiment. All of the respondents are adults from the United States. They were recruited on Amazon’s Mechanical Turk platform. Mechanical Turk workers are younger, lower-income, and more likely to be unemployed than the average American (Levay et al., 2016). However, there is more and more evidence that this recruitment strategy does not affect experimental results (Berinsky et al., 2012; Coppock, 2018; Mullinix et al., 2015). We

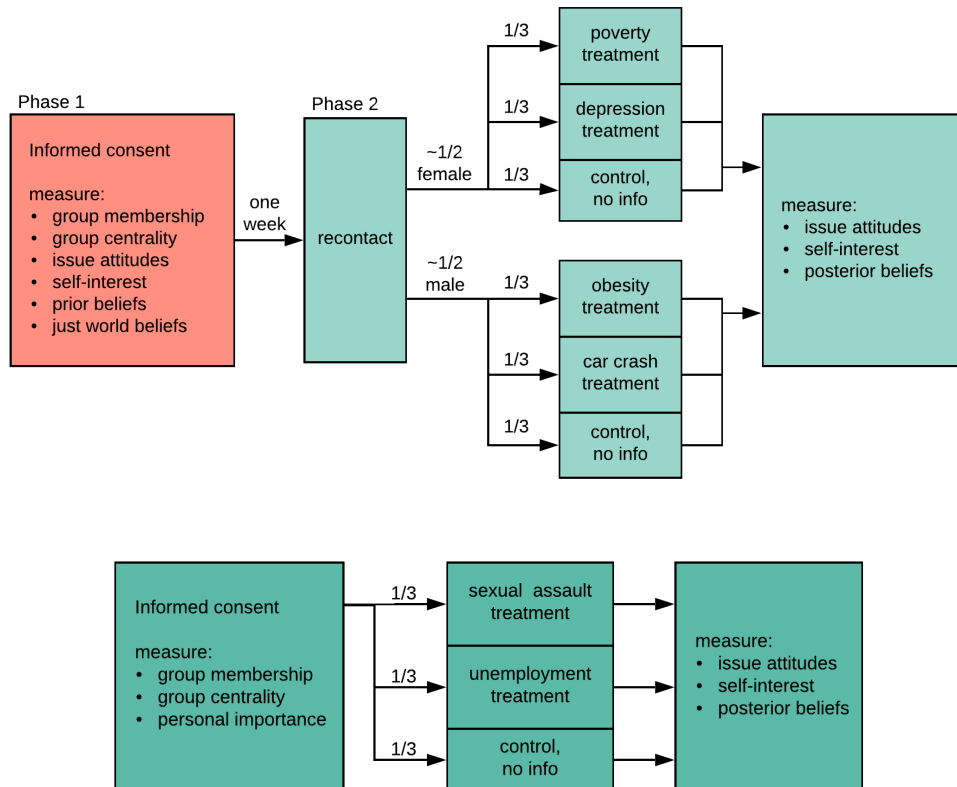


Figure 1: Designs for Experiment 1 (gender) and Experiment 3 (LGBT). Participants were treated (or not) with information about gendered issues or LGBT issues. Then, they answered post-treatment measures. In Experiment 1, participants also filled out pre-treatment measures of the DVs one week earlier. The design structure of Experiment 2 is identical to that of Experiment 1.

might worry that workers on crowd-sourcing platform pay less attention to textual treatments. However, section 1.5 shows that respondents' beliefs are indeed systematically shifted by the treatment.

1.2.1 Experiment 1: Gender

The sample for this study consists of 484 respondents recruited in September 2018. 262 are male, and 222 are female. Some participants dropped out between phases—the retention rate from the original sample was 65%.

As reported in the PAP (pre-analysis plan), this experiment was originally planned to have approximately 500 participants. Given this sample and the panel (pre-post) design, the minimum detectable effect size (with a power of .80) for the effect of in-group interests on the four-point outcome scales (for concern and spending) would have been between .19 and .31. The exact number depends on the amount of noise added between phases, in other words, the autocorrelation between attitudes in phase 1 and 2. The eight-place importance ranking would have had a minimum detectable effect size between .36 and .49.

1.2.2 Experiment 2: Race/ethnicity

The sample for this study consists of 451 White and 267 Black or Latino participants, recruited in September–October 2018. Most respondents were newly recruited, as in Experiment 1. In addition, I engaged Black and Latino participants from a pool of workers who had completed another survey task on Mechanical Turk at least one month, and up to 18 months, prior to being recruited for this survey. Respondents who identified as neither White, Black

nor Latino were filtered out of the sample.

This sample size (718) was larger than the anticipated sample size (610). The power calculation reported in the PAP was based on this slightly smaller sample size. The minimum detectable size (with a power of .80) for the effect of in-group interests on four-point-scale outcomes was calculated to be between 0.17 and 0.28, depending on the amount of noise added between waves. The eight-place issue ranking would have had a minimum detectable effect size between 0.32 and 0.43.

I grouped into the Latino category all participants who identified as Hispanic/Latino, including those who also identified as White or Black. Black and Latino participants were pooled into one group and received the same treatments, though the identity centrality measure applied to their group only (e.g., “I often think about the fact that I am Black”).

1.2.3 Experiment 3: LGBT

The 198 participants for this study (126 female, 66 male, 6 other or unknown gender) were recruited in November 2018. These workers had completed another survey task one month before being recruited for this survey. All participants had previously indicated that they identified as LGBT and were willing to be recontacted for another survey. In addition, I filtered participants by including an LGBT identification question at the beginning of this survey. I found no detectable difference between LGBT women and LGBT men in the strength of their LGBT identity.

This experiment’s PAP anticipated a slightly larger sample (around 215). This PAP drew on the results of Experiment 1 and 2 to estimate the variance

of each dependent variable. For concern about the issue, measured on a four-point Likert scale, the minimum detectable effect size (with power .80) was calculated to be .25. Issue importance, measured as the issue's place in an eight-place ranking, would have a minimum detectable effect size of .69. Finally, support for government spending, measured once again on a four-point scale, would have a minimum detectable effect size of .33.

1.3 Dependent variables: concern, importance, spending support

I took pre- and post-treatment measurements of three outcome variables: concern about the issue (i.e. whether the issue is seen as a problem), importance of the issue (when ranked with other issues), and support for government spending to help tackle the issue. Figures 2–4 show the distribution of the dependent variables for each issue in Experiment 1–3.

These three issue attitudes are somewhat dependent on each other (Wlezien, 2005), but they should be progressively harder to change. Self-reported concern is the most “costless” attitude in this study—people could change their response only to show their allegiance with the in-group (cf. Bullock et al. 2015; Prior et al. 2015). In contrast, because the importance measure is a ranking, upgrading one issue comes at the price of downgrading another. Finally, even if a respondent thinks of an issue as concerning and important, they may still not be in favor of increased spending on it. Indeed, in a review on self-interest in politics, Kinder (1998) finds that interests change perceived issue importance more often than they change positions on policy

(p. 802). Changes in support for government spending should be the most demanding test.

In the importance ranking question, the treatment effect size may depend somewhat on the alternative issues among which subjects needed to rank the issue of interest. Section 4.4.2 below shows, for each experiment, which other issues were present in the ranking. In section 3, I discuss the consequences of these choices along with concerns about ceiling effects and anchoring of the dependent variables by measuring them pre- and post-treatment.

1.4 Group identification

There is a vast literature on the measurement of group identification, and nothing resembling a consensus about which scales are best suited for which applications. I used the Centrality subscale developed by Leach et al. (2008). It has three items, which I average into an identity centrality score:

- I often think about the fact that I am a [man/woman/...]
- The fact that I am a [man/woman/...] is an important part of my identity
- Being a [man/woman/...] is an important part of how I see myself.

The Centrality scale measures the overlap between a person's self-concept and their concept of the group. This matches Conover's (1984) definition of group identification as the group schema becoming a self-schema. Greenwald et al. (2002) also speak about group identity as an association between the self and a social category. Similarly, Deaux (1996) writes that a person

identifies with a group if he or she accepts that membership of the group defines him or her in some way. Figure 5 shows the distribution of identity centrality scores by group.

Of course, there are other in-group attitudes that could moderate the connection between group membership and political opinions. One of them is group consciousness—a set of political beliefs including the idea that collective action is needed to improve the group’s social standing (McClain et al., 2009). Group consciousness is more demanding than group identity: identifying with a group would seem to be a necessary, but not a sufficient condition for group consciousness (cf. McClain et al. 2009; Miller et al. 1981). In fact, group-conscious members should act on in-group interests almost by definition. Complicating the situation, however, is the fact that three key components of group consciousness—closeness to the group, perceived discrimination and collective action—are empirically quite distinct (Sanchez and Vargas, 2016). Future research might investigate whether only group members with group consciousness are likely to translate in-group interests into political opinions.

1.5 Group–issue connection beliefs

All experiments included a two-part question on respondents’ beliefs about group disparities in each issue. First, respondents indicated whether they believed the issue happens more to their in-group, more to their our-group, or whether it is about the same. Next, they specified whether they had a lot of confidence, a moderate amount, or only a little confidence in their

answer. I used these questions to create a seven-point scale. Respondents who answered “don’t know” or “about the same” are at the midpoint. The other respondents are on either side of the midpoint, with the most confident respondents sitting at the ends. The top half of the scale represents a correct answer (e.g., poverty happens more to women), and the bottom half represents an incorrect one. In Experiment 1 and 2, I took both pre- and post-treatment measurements of belief. In Experiment 3, I only took a single, post-treatment measure of belief.

Descriptive data show that without the information treatment, relatively few respondents know (with high certainty) about the connections between their in-group and the issues in this study. In Experiment 1 (gender), before treatment, just 14% of respondents were highly confident about the correct answer (and 5% were confident, but gave the wrong answer). In Experiment 2 (race), this was 11% (with 2% confident and wrong). In Experiment 3, looking at untreated issues only, 24% of respondents were confident and correct (almost none were confident and wrong). Section 1.5 below shows that the lack of treatment effects cannot be ascribed to these respondents.

Figures 6–8 illustrate the strong effects of treatment on beliefs about each issue. In Experiment 1 and 2, untreated respondents rarely move their beliefs between phases, whereas treated respondents move towards the correct, high-confidence end of the scale. In Experiment 3, treated respondents are far more likely to give high-confidence correct answers than untreated respondents.

1.6 Self-interest and linked fate

The theory of linked fate implies that people use the in-group impact of an issue to estimate how that issue might affect them individually. If there is an effect of group interest cues on issue attitudes, we might wonder whether linked fate is the mechanism. However, there are some concerns with the classical linked fate survey item: “Do you think that what happens to [group members] in this country will have something to do with what happens in your life?”. First, on its face, this item does not exactly capture whether or not people are using the groups’ interest as a heuristic for their own. Second, studies show that the item actually seems to capture something closer to a general personality trait. Individuals tend to have a sense of linked fate either with all of their social groups or with none of them, and linked fate is generally not a good predictor of political opinions or behavior Gay et al. (2014).

For this reason, I opted to test linked fate theory directly, by asking participants about their perceived self-interest before and after they received an in-group interest cue. Specifically, I asked them whether they believed the issue would affect them in the future.¹ Follow-up studies could shed light on whether linked fate, as measured with the classic survey item, heightens the effect of in-group interest cues—perhaps because it captures empathy with fellow group members.

¹Because one of the answer options was “is already affecting me”, privacy concerns caused me to modify the question slightly for White and LGBT participants. Since these participants were being asked about sensitive issues (suicide, opioid addiction, unemployment and sexual assault), I added the phrase “or someone close to you” to the question wording, so respondents would not have to reveal that they had personally dealt with one of these problems.

1.7 Model specifications

In Experiment 1 and 2, the basic model for each of the issue attitudes is:

$$Y_{ij,t=2} - Y_{ij,t=1} = \alpha_j + T_{ij} + \epsilon_{ij}$$

where Y_{ijt} is respondent i 's attitude on issue j at time t , T_{ij} indicates whether the respondent got treated on this issue, and α_j is an issue-specific intercept. The error terms ϵ_{ij} are clustered at the level of the respondent, i .

In Experiment 3, since pre-treatment dependent variables were not available, they could not be used as controls in the analyses. To make up for this, before treatment, I measured the personal importance of the issue for the respondent on a four-point scale, as defined by Krosnick (1990). The baseline estimating equation for each of the issue attitude types becomes:

$$Y_{ij} = \alpha_j + T_{ij} + P_{ij} + \epsilon_{ij}$$

where Y_{ij} is respondent i 's attitude on issue j , and P_{ij} is the personal importance of issue j to respondent i .

For treated participants, I leave non-treated issues out of the analyses, in case the treatment spills over into other issue attitudes. For control participants, I include their responses on both of the issues related to their group.² The unit of analysis is the person-issue. All analyses are OLS regressions with issue fixed effects. Standard errors are clustered at the person level in all cases.

²This means that there are equally many control and treatment person-issues in the final sample: there are twice as many treated as control respondents, but control respondents contribute two issues each.

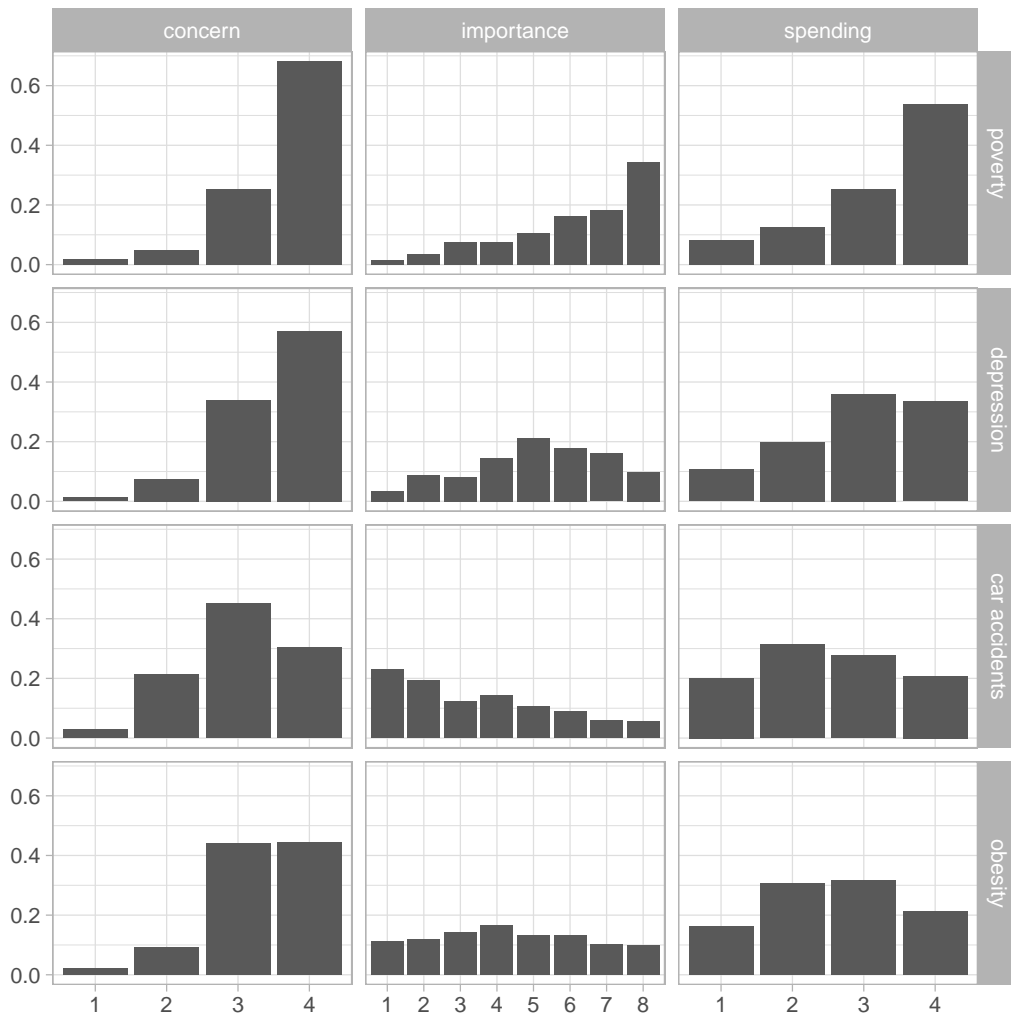


Figure 2: Distribution of respondents across concern, importance and spending support for each issue in Experiment 1 (gender), pre-treatment.

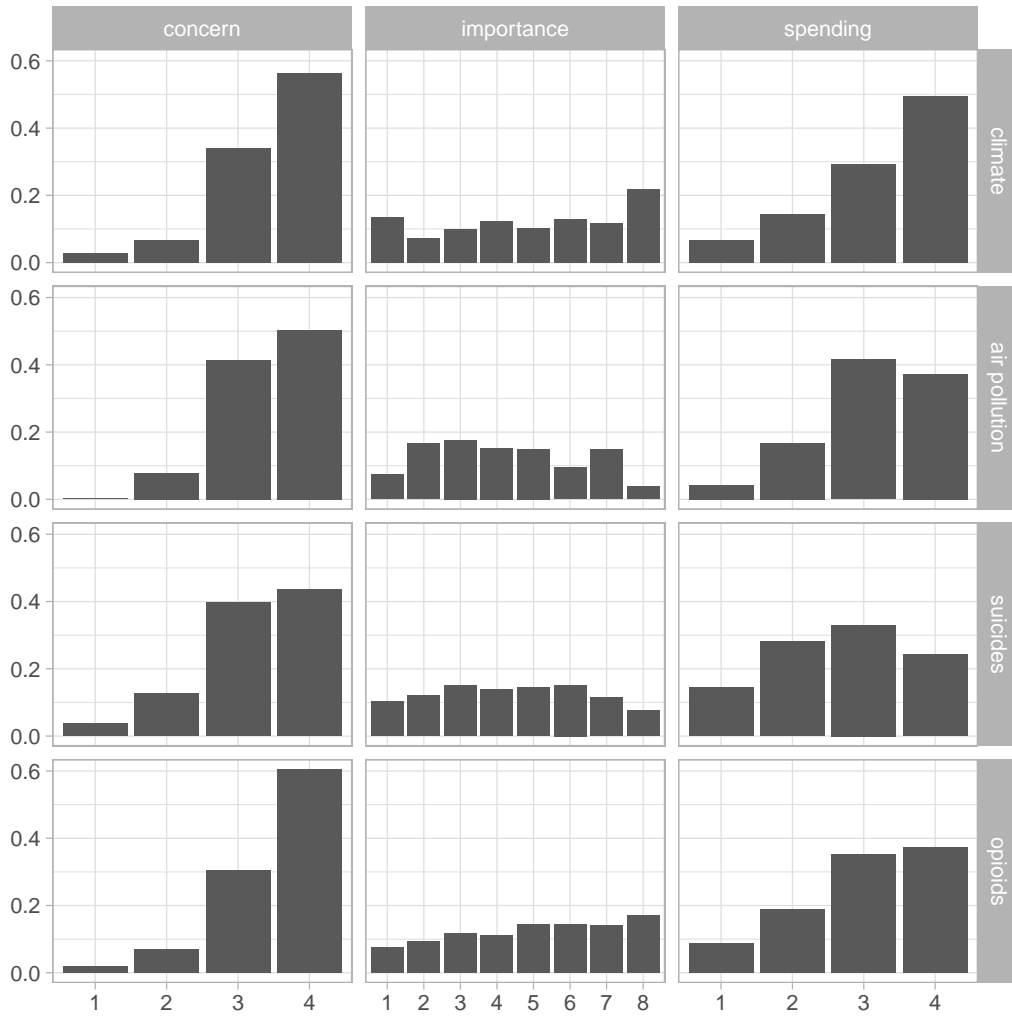


Figure 3: Distribution of respondents across concern, importance and spending support for each issue in Experiment 2 (race), pre-treatment.




Figure 4: Distribution of respondents across concern, importance and spending support for each issue in Experiment 3, control group.

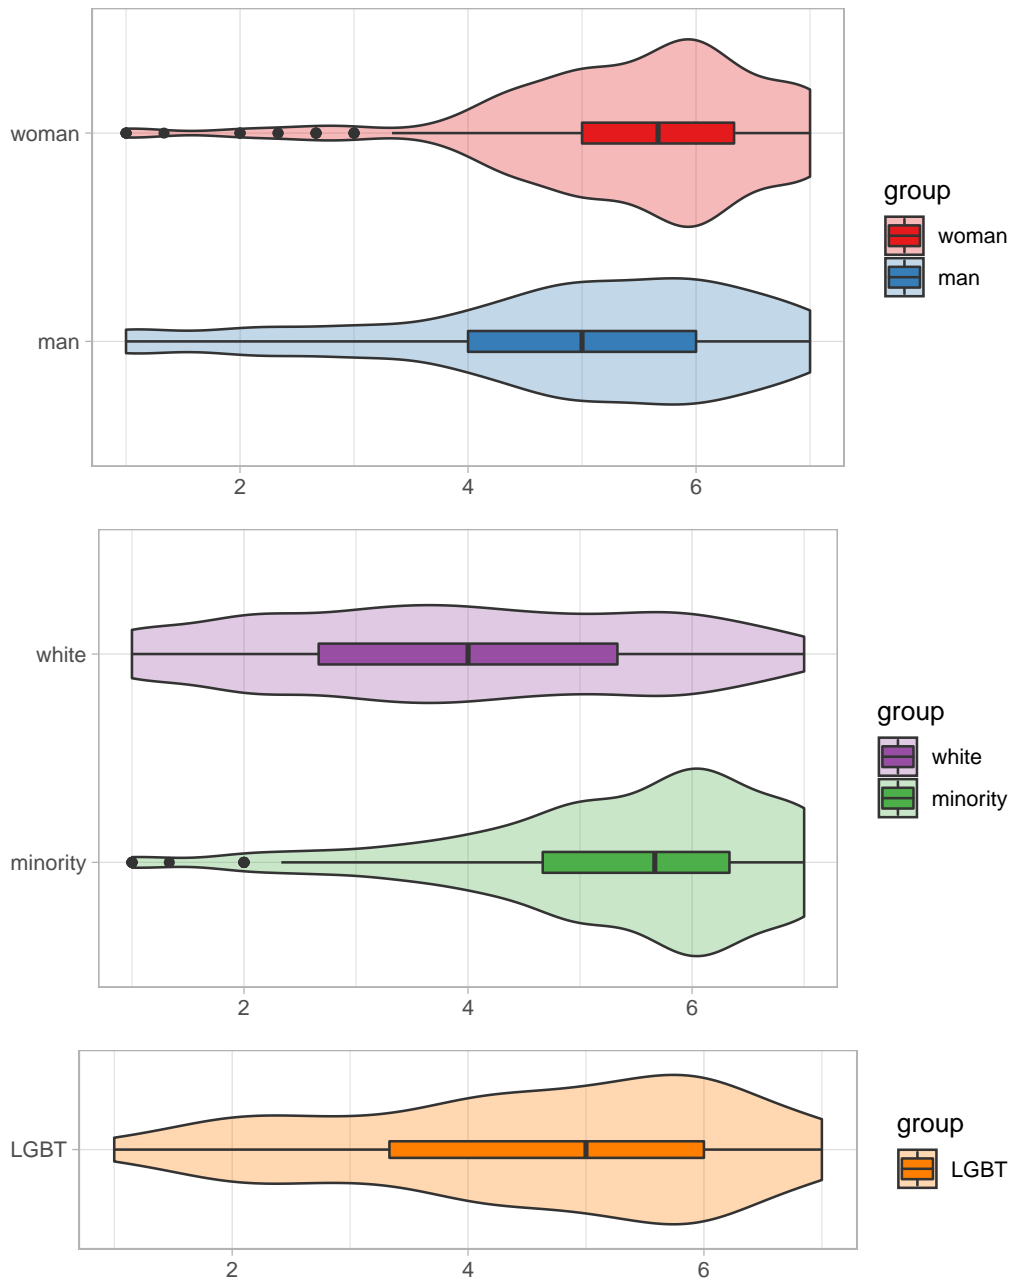


Figure 5: Violin plots (smoothed distributions plus boxplots) of scores on identity centrality scale (average of three 7-point items), by group.

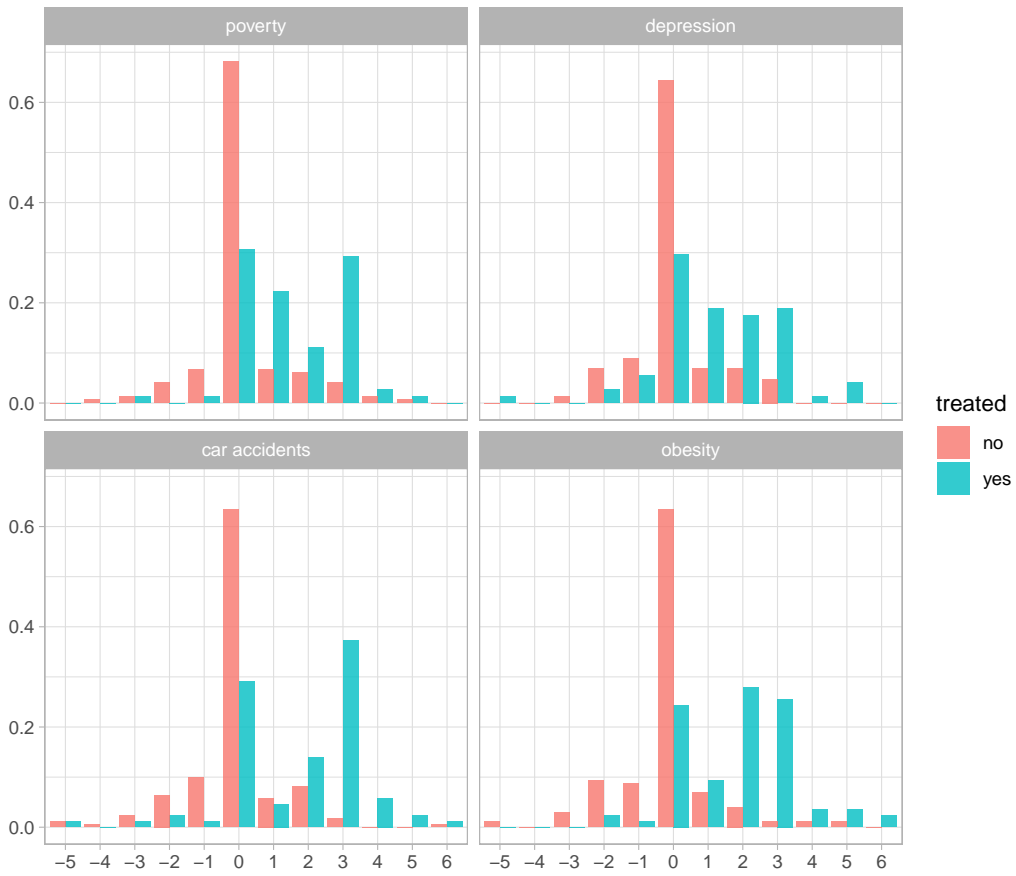


Figure 6: Distribution across respondents of *change* in beliefs about gender-issue connections between phases in Experiment 1 (gender), for treated and untreated respondents. Belief is a 7-point scale.

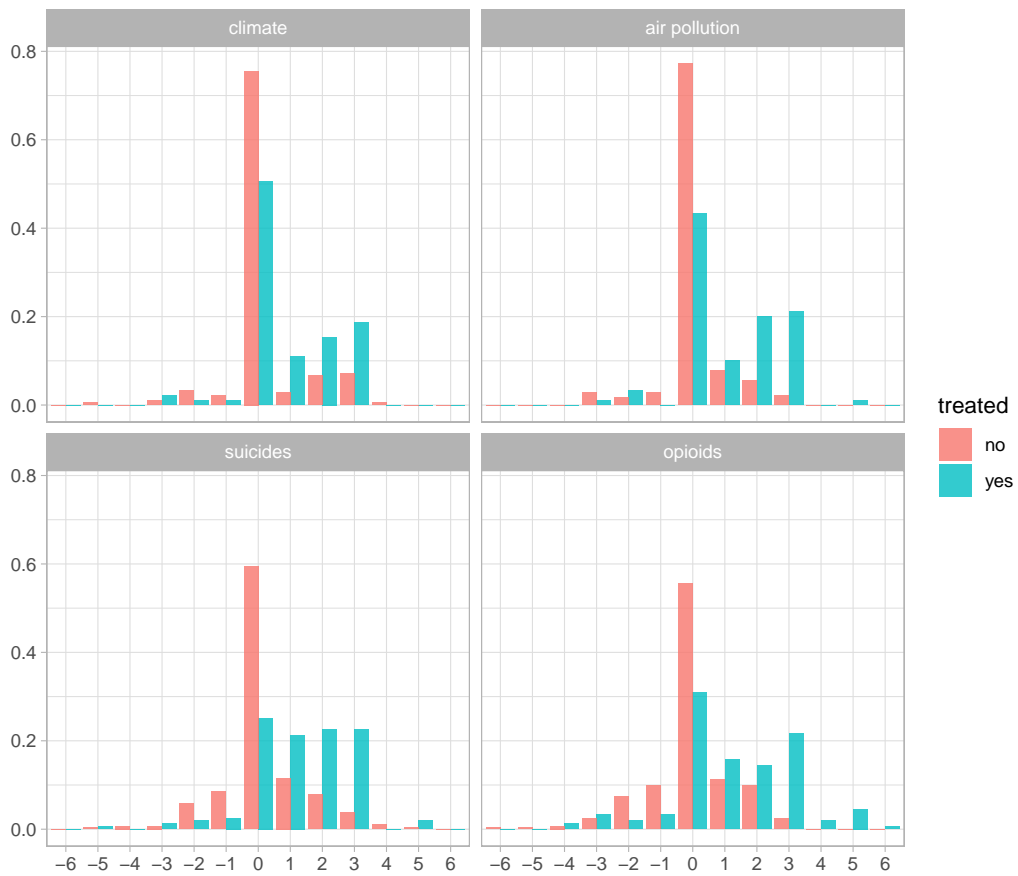


Figure 7: Distribution across respondents of change in beliefs about race-issue connections between phases in Experiment 2, for treated and untreated respondents.



Figure 8: Distribution across respondents of belief about LGBT-issue connections in Experiment 3, for treated and untreated respondents.

2 Effects and confidence intervals

Figure 9 illustrates the effect sizes of the interest cue treatment on each issue attitude in Experiment 1. Concern and spending are 4-point scales; importance is an 8-issue ranking. The figure shows two models: the basic model with the treatment only, and a model that includes an interaction with centrality (identity strength). 90% confidence intervals illustrate the precision of the estimates. The upper limits of such intervals have an 95% probability of being larger than the true average effect a priori (i.e., before data collection, Rainey 2014). Figures 10 and 11 show the same for Experiment 2 and 3.

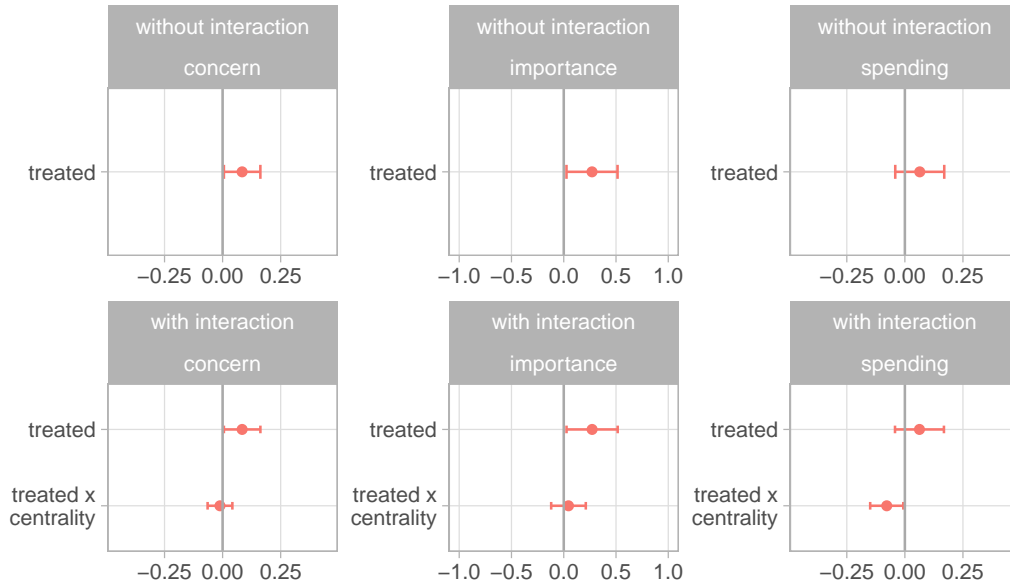


Figure 9: Effect of learning gender group interests on pre-post difference in three issue attitudes.

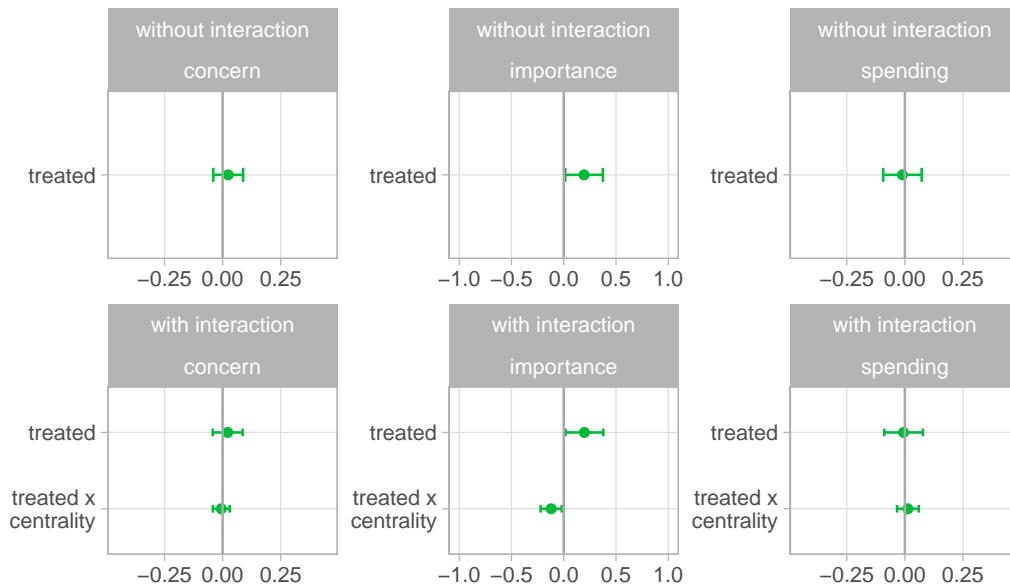


Figure 10: Effect of learning racial/ethnic group interests on pre-post difference in three issue attitudes.

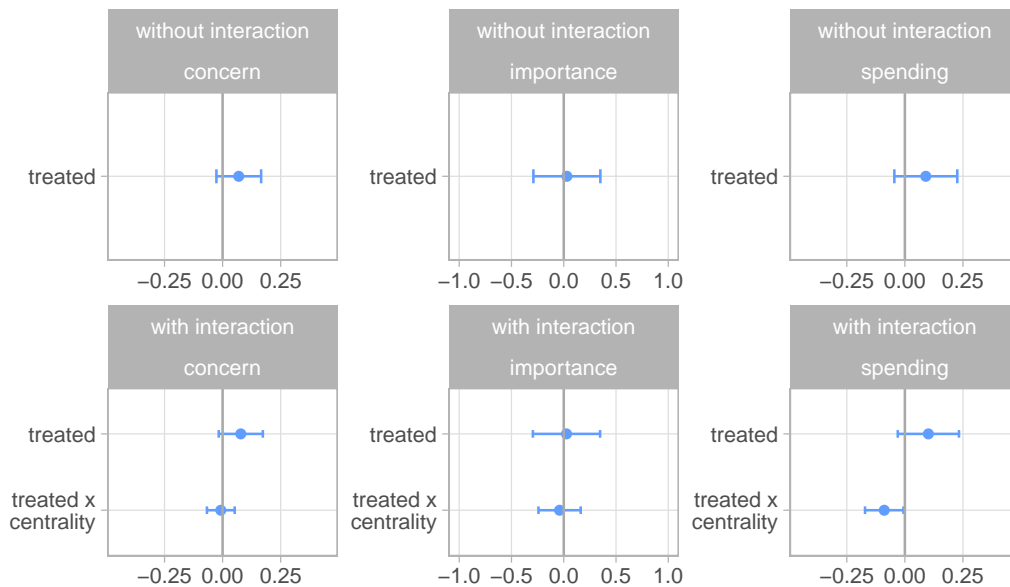


Figure 11: Effect of learning LGBT group interests on three issue attitudes.

3 Robustness checks

Throughout Experiment 1–3, I find that group interest cues have little to no effect on issue attitudes. In this section, I solidify this finding by showing that it holds even when I pool data across experiments; when I examine the group–issue combinations one by one; and when I control for the fact that the treatment may not have affected every respondent’s beliefs about interests (e.g., because the information was not new). Finally, I examine the possibility of ceiling and anchoring effects, and discuss substantive explanations for the null results.

3.1 Pooling all experiments

Since all the experiments have similar designs, I can pool their data to get a more general estimate of the effect of in-group interest cues.³ The LGBT experiment has no pre-treatment measurements, so I combine the data in two ways: leaving out the LGBT data ($n = 1179$); and using all data, but ignoring pre-treatment measurements for the gender and race studies ($n = 1376$). Figure 12 reports treatment effect estimates for both pools, on each of the three issue attitudes. The only noticeable effect is the effect on issue importance. Respondents move treated issues up by about .25 of a place on average in their rankings. The estimated effect on concern is .06 or less; the effect on spending support is .07 or less. Both pools allow us to reject effect sizes greater than .11 (concern), .40 (importance) and .15 (spending).

Interactions between the treatment and identity centrality, when added to the model, are never significant. To further investigate this, I split up identity centrality into within-group centrality (how strongly the person identifies with their group, compared to other members of that group) and between-group centrality (the group's average identification level).⁴ Pooling data from all experiments, within-group centrality has very small negative interaction effects with the treatment (concern: $-.01$; importance: $-.01$; spending: $-.05$, all $p > .1$). Between-group centrality has very small positive interaction effects (concern: $.01$; importance: $.12$; spending: $.05$; all $p > .1$)

³Note that here, the issue fixed effects absorb any differences between experiments in the average levels of the dependent variables without treatment.

⁴This analysis was not pre-registered.

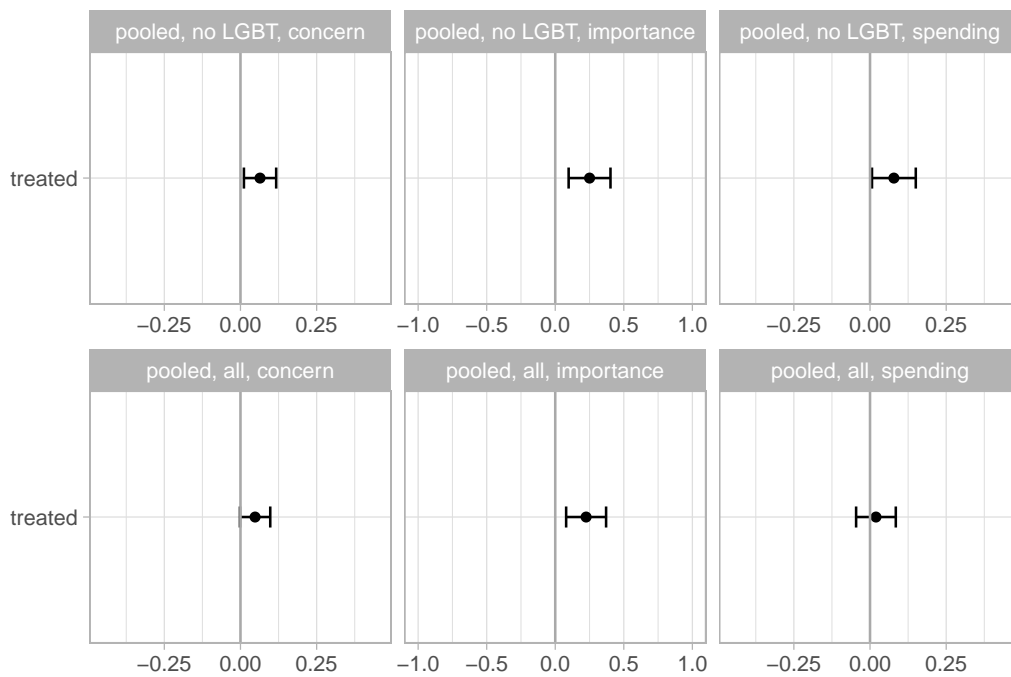


Figure 12: Effect of connecting an issue to in-group interests on three attitudes, with 90 % confidence intervals. Data pooled either across gender and racial groups, without LGBT ($n = 1179$); or across gender, racial and LGBT groups ($n = 1376$).




Figure 13: Effect of connecting an issue to in-group interests, by issue–group combination, on three attitudes, with 90 % confidence intervals.

3.2 Looking at single issues

We might be interested to know whether any issue/group combinations were more effective than others at moving attitudes. Figure 13 shows the issue-by-issue effect sizes from Experiment 1–3, obtained by adding issue-treatment interactions to the models specified in section 1.7. There are no obvious patterns across all three dependent variables. Three issue/group/dependent variable combinations narrowly cross the threshold of marginal significance (car accidents/men/concern, air pollution/minorities/importance, and unemployment/LGBT/concern, $p < .10$). However, with 30 (10 issues x 3 attitudes) comparisons, this is precisely what we would expect if all of the true effects are zero.

3.3 Change in beliefs

3.3.1 Treatment as an instrument for belief change

We may be concerned that the treatment did not actually change respondents' belief in the connection between their group and the issue at hand. Some respondents could have been aware of the connection beforehand; their beliefs would not have been moved by the treatment, because they were already aligned with the treatment to begin with. Other respondents perhaps did not believe the information provided in the treatment. While I cited a source for each treatment, perhaps not all respondents trusted these sources (or my interpretation of them). Finally, some participants may not have paid attention to the treatment. Here, I conduct instrumental variables analyses to show that the treatment is ineffective even when it actually changes people's beliefs.⁵

As we could see from Figures 6–8, the treatment is a strong instrument for belief (change) in all cases (gender: $F(1, 641) = 188.7, p < .001$, race: $F(1, 938) = 148.1, p < .001$, LGBT: $F(1, 393) = 21.9, p < .001$). Figure 14 shows the results of two-stage least squares regressions for each experiment (using first-differenced beliefs and issue attitudes for Experiment 1 and 2). The only effects that reach marginal significance are effects on issue importance, for gender and racial/ethnic groups. An average respondent who went up one point on the belief scale would upgrade an issue by less than .2 places in his or her importance ranking. In other words, the treatment

⁵The pre-analysis plans for Experiment 1 and 2 specified that I would only conduct this instrumental variables analysis if I found a statistically significant main effect of the treatment. However, I now believe that the instrumental variables analysis is also, and even especially, useful for making sure I am interpreting the null findings correctly.

generally does not change attitudes even for those respondents whose beliefs were moved. This rules out the explanation that the treatment effect is being suppressed by, for example, respondents who already knew their group is more affected by the issue.

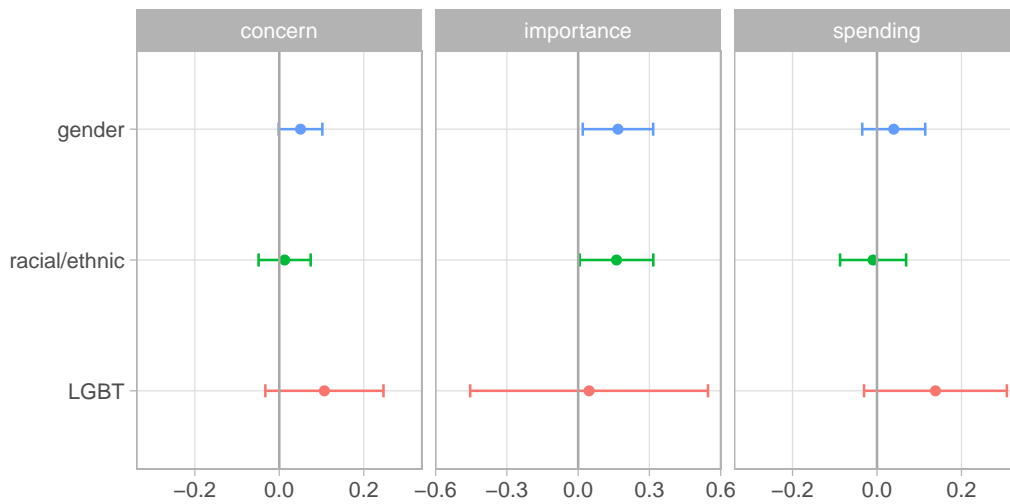


Figure 14: Effect of changing beliefs about in-group interests, as instrumented by an information treatment, on three issue attitudes, with 90% confidence intervals. Belief is on a 7-point scale. Concern and spending are 4-point scales, importance is an 8-issue ranking.

The instrumental variables approach leads to consistent estimates only if the so-called exclusion restriction holds. That is, the treatment can only affect issue attitudes through people’s beliefs about their in-group’s interests. It is possible that the information in the treatment has non-informational effects—for example, being treated might increase the salience of the social group and its interests, even for respondents who already knew about the issue-group connection. However, this and other plausible violations of the exclusion restriction would cause an *upward* bias in the belief change effect estimates. If anything, these estimates should therefore be read as upper

bound estimates.

3.3.2 Moderation by prior beliefs

One of the problems that the instrumental variables analysis above addresses, is that an informational treatment should not have any effect on respondents who already know the information. This section presents results for a different (not pre-registered) approach, where prior beliefs are used as a moderator.

Here, I re-analyze the data from Experiments 1 and 2, using respondents' pre-treatment beliefs as a moderator for the treatment effect. Pre-treatment beliefs are measured on a seven-point scale, from high confidence in the wrong answer (meaning there is a lot of room for the treatment to change beliefs) to high confidence in the right answer (meaning the treatment will likely not change beliefs).

Tables 1 and 2 show that pre-treatment belief is never a substantively or statistically significant moderator of the treatment effect. Including this moderator does not change conclusions about the (usually negligible) effect of the treatment. Like the instrumental variables analyses, this suggests that small treatment effects are not due to respondents already having the information beforehand.

3.3.3 Group identity centrality and beliefs

Finally, beliefs could help explain the null or negative effects of group identity centrality. For example, strong identifiers could be reluctant to connect their own group to a societal problem—as most of the issues in this study have

Table 1: Effect of learning about gender in-group interests on three issue attitudes, interacted with prior beliefs.

	<i>Dependent variable:</i>		
	concern	importance	spending
treated	0.08* (0.05)	0.26* (0.15)	0.06 (0.07)
belief	0.02 (0.02)	0.09 (0.07)	-0.002 (0.03)
treated:belief	0.01 (0.03)	0.06 (0.11)	0.02 (0.05)
Observations	641	638	637
R ²	0.02	0.02	0.01
<i>Note:</i>	*p<0.1; **p<0.05; ***p<0.01		

Table 2: Effect of learning about racial/ethnic in-group interests on three issue attitudes, interacted with prior beliefs.

	<i>Dependent variable:</i>		
	concern	importance	spending
treated	0.03 (0.04)	0.22* (0.12)	-0.03 (0.06)
belief	0.01 (0.02)	-0.07 (0.06)	-0.02 (0.03)
treated:belief	-0.01 (0.03)	-0.07 (0.09)	0.03 (0.04)
Observations	962	948	950
R ²	0.01	0.03	0.01
<i>Note:</i>	*p<0.1; **p<0.05; ***p<0.01		

a negative connotation. Alternatively, people for whom the group is central to their identity might know more about the group’s interests ahead of time, and therefore be less affected by the treatment. The analyses below (not pre-registered) investigate these possibilities.

For Experiment 1 and 2, I regress change in beliefs on identity centrality, treatment status, and their interaction (with issue fixed effects). In the gender case, I find a very small and non-significant interaction ($-.06, p > .1$). For racial groups, the interaction is small but statistically significant, meaning that strong identifiers are slightly less moved in their beliefs by the treatment ($-.11, p < .05$). For LGBT people, running the same regression using just post-treatment beliefs, I once again find a very small, non-significant interaction ($.05, p > .1$). Thus, the beliefs strong group identifiers are slightly less affected by information about in-group interests, but only for racial groups.

Further analyses show that strong identifiers do not have particularly correct prior beliefs. Without treatment, on a belief scale from -3 (high confidence in the wrong answer) to 3 (high confidence in the right answer), even a person with the highest possible identity centrality is predicted to score close to zero: 0.5 (gender), 0.4 (race) and 1.1 (LGBT). In fact, in the only case where centrality is significantly connected to smaller belief changes (race), respondents with strong identity centrality are actually slightly less confident that the issue negatively affects their in-group ($r = -.1$). Perhaps strong identifiers prefer not to associate their racial or ethnic group with societal problems—before the treatment and after.

3.4 Ceiling effects and anchoring

Figures 2–4 above illustrate the distribution of concern, importance and spending scores of the treated issues among untreated respondents. They show that we might worry about ceiling effects, where respondents would have picked the top option even without treatment—especially for the concern variable, but also for importance and spending on issues such as poverty, climate change and sexual assault.

For that reason, I repeat the analyses pooling Experiment 1 and 2. For each dependent variable, I exclude those observations (in the treatment and control groups) where a respondent already chose the top option for that issue on that variable in the pre-treatment phase. Estimated effects increase slightly (concern: .10, $p < .05$; importance: .25, $p < .05$; spending: .03, $p > .1$). However, they remain substantively small.

A more subtle variant of a ceiling effect could affect the issue importance ranking. Even if a respondent places the issue of interest in the middle of the scale, the issues ranked above it could be so important that it could never displace them. While I cannot identify respondents for whom this might be the case, it is reassuring that in Experiment 1 and 2, control group respondents commonly switch the issues of interest around in their rankings between phases (the issue with the smallest average movement in importance rank is poverty, which moves one place on average). Moreover, in all experiments, untreated respondents are quite spread out in their importance rankings of the treated issues (see Figures 2–4; the issue with the lowest spread is again poverty, with a standard deviation of 1.8 places). At least some respondents

found it reasonable to rank each issue in each of the eight available places.

A final concern related to dependent variable changeability, is that the pre-treatment measures may have anchored people's attitudes. The experimental designs partly address this worry. There is at least one week between pre-treatment and post-treatment measures; anchoring is less likely for the eight-place issue importance ranking (in fact, the numbers above show that even untreated issues are commonly moved around); and the effects in the LGBT experiment are still null, even though there were no pre-treatment measurements.

3.5 Alternative explanations

Finally, I use robustness checks to examine possible substantive explanations for the null results. Some of these are easily ruled out; for others, the counterevidence is more tentative.⁶

First, in-group bias is likely to be less socially acceptable when the in-group is not traditionally seen as disadvantaged. Members of those groups may be reluctant to openly favor their in-group (cf. White 2007). To test this, I pool respondents from groups seen as less (men, White people) or more (women, minorities, LGBT) disadvantaged. Treatment effects do not differ significantly from each other, and the effect sizes for disadvantaged groups are almost identical to the effects in the overall pool (concern: .06, $p < .1$; importance: .27, $p < .05$; spending: .13, $p < .05$).

A second explanation is that the treatment might only make an impression on respondents if the group's relative risk is high enough. For example,

⁶The additional analyses below are exploratory and not pre-registered.

Black and Latino people only have a 15% higher risk of dying from heat-related diseases, but White people are three times more likely to commit suicide. For that reason, pooling all data, I do an analysis where I interact the treatment with the group's relative risk for the issue being asked about (e.g. 1.15 or 3). Interaction effects are small and negative (concern: $-.01$; importance: $-.05$; spending: $-.11$; all $p > .1$).⁷ This suggests that, if anything, cues about larger relative risks have slightly smaller effects.

Similarly, information about the relative risk of an in-group member being affected by an issue might not be sufficient. Instead, issue attitudes may only move if the in-group is especially affected, *and* the absolute risk is significant. For example, white respondents may have disregarded the information about opioid overdose deaths, on the basis that opioid overdoses only represent a tiny fraction (about 2%) of deaths in the United States. This explanation is somewhat weakened by the fact that issue-by-issue analyses do not show particularly large effects for issues that are clearly large-scale, such as obesity (which affects around 40% of Americans) and depression (around 7.1%). The same is true for issues that were ranked as more important in society, such as poverty (median rank: 2 out of 8 pre-treatment) and sexual assault (median rank in the control group: 3 out of 8). Still, future iterations might use a treatment that communicates both relative and absolute risk to the group, for instance, by showing the number of group members that are affected by the problem.

Another alternative interpretation is that some (in particular conserva-

⁷From a model without an interaction with identity centrality; adding this interaction does not change results noticeably.

tive) respondents believe that people are, and should be, rewarded based on their individual efforts rather than their group memberships (cf. Carney and Enos 2017). To investigate this, I included Lerner’s (1980) just world belief scale (four items on seven-point scales) at the end of the gender and race experiments. I found no significant negative interaction between just world belief (averaged over items) and the treatment in either study, or when pooling the data from both studies. Interaction effects were tiny and positive (pooled estimate for concern: .00; importance: .09; spending: .01; all $p > .1$).

A related explanation would be that respondents do not identify with fellow group members who are affected by the issues, for one of two reasons. First, Marques et al. (1988) suggest that in-group members who deviate from a norm are judged more harshly than out-group members who do the same (the so-called “black sheep effect”). If respondents believe that being affected by a problem is a result of a personal choice or failure of the in-group member (for example, that being in a car accident is due to being a careless driver), then they may not be inclined to help. This idea is related to respectability politics, where members of a minority believe that their in-group will benefit if its members behave according to dominant out-group norms (Higginbotham, 1993). However, the issue-by-issue analysis (section 3.2 above) contains hints against this explanation. Issues that seem less influenced by personal choices, like climate change, air pollution or sexual assault, are no more affected than issues that can more easily be connected to individual actions, like car accidents or suicide.

A second reason is that respondents may think of the affected members

as part of a different subgroup. For example, older men may assume that car accidents happen more to *young* men, a group that they do not belong to. While difficult to falsify (there will almost always be a subgroup that is more affected by a problem), this explanation would still have important consequences for group thinking in politics. Many group theories refer to large social categories such as gender, race or even American citizens. If only subgroup interests are relevant, that would be an important constraint on the power of group-based political opinions. It would also support theories that emphasize cross-cutting identities and disadvantaged subgroups (e.g. Cohen 1999; Harris 2014; Strolovitch 2006).

Finally, it is possible that the dependent variables did not pick up any effect, because they do not make any explicit reference to the social group at hand. For example, women are not asked for their concern about poverty among women, or about their approval of spending to combat female poverty. It is possible that group interest cues only affect attitudes about the issue as it applies to the in-group, and not about the issue as a whole. This explanation, too, is quite difficult to test, because referring to groups as part of the dependent variable measurement would make those groups salient to the control group as well. On the other hand, accepting this explanation would again place strong limits on the relevance of group thinking in politics. Essentially, it implies that in-group members favor policies that benefit their in-group, but not if those policies also benefit some out-group members. This makes group interests a less useful explanation of political attitudes in general.

4 Questionnaire

This section details all of the questions that were asked of participants in all three experiments, as well as the wording of the information treatment. Figure 1 above details the order in which these questions were asked (as some questions were asked both pre- and post-treatment in Experiment 1 and 2).

4.1 Informed consent

I agree to participate in a research study conducted by Clara Vandeweerd at the Massachusetts Institute of Technology (MIT). In order to analyze responses to the questionnaire, my answers will be recorded. Researchers will have no access to any personal information about me, except for my MTurker ID, the time at which I took the survey and the answers I filled out. No identifying information about me will be made public and any views I express will be kept completely confidential.

Findings from this study will be reported in scholarly journals, at academic seminars, and at research association meetings. The data will be stored at a secured location and retained indefinitely. My participation is voluntary. I am free to withdraw from the study at any time.

By participating in this survey, I confirm that I am 18 or older. I also give the researchers permission to invite me for a (paid) follow-up study.

[in Experiment 3, LGBT:] Please note that this survey touches on the topic of sexual assault, and that you are free to skip questions on that topic if you need to.

Please select one of the following options. If you choose not to participate,

the survey will end immediately and no data will be recorded. Should you have questions, please send an e-mail to claravdw@mit.edu .

- I agree to participate
- I do not agree to participate

4.2 Group membership

First, we would like to ask a little more about you.

- | | | |
|--|---|--|
| Experiment 1 and 3: What is your gender? | Experiment 2: Which of these group(s) would you say you belong to? | Experiment 3: Do you identify as LGBT (Lesbian, Gay, Bisexual or Transgender)? |
| <ul style="list-style-type: none">• Male• Female• Another gender• Prefer not to say | <ul style="list-style-type: none">• White/Caucasian• Black/African American• Hispanic/Latino• Asian• Native American• Pacific Islander• Other | <ul style="list-style-type: none">• Yes• No |

4.3 Group identity

4.3.1 Centrality

[Experiment 1–3] How much would you say you agree or disagree with the following statements?

- The fact that I am [a man / a woman / White/ Black / Latino / LGBT] is an important part of my identity.
- I often think about the fact that I am [a man / a woman / White / Black / Latino / LGBT].
- Being [a man / a woman / White / Black / Latino / LGBT] is an important part of how I see myself.

answer options: Strongly disagree – Disagree – Somewhat disagree – Neither agree nor disagree – Somewhat agree – Agree – Strongly agree

4.4 Issue attitudes

[issue presentation order always randomized]

4.4.1 Concern

Next, we would like to ask your opinion about a few social issues. For each of the issues below, please tell us how serious of a problem you think this issue is for our society.

Experiment 1: • Poverty • Depression

- | | | |
|---|---|--|
| <ul style="list-style-type: none"> • Obesity (being seriously overweight) • Car accidents | <p>Experiment 2:</p> <ul style="list-style-type: none"> • Climate change • Air pollution • Suicides • Addiction to opioids (strong painkillers) | <p>Experiment 3:</p> <ul style="list-style-type: none"> • Unemployment • Sexual assault • Poverty • Climate change |
|---|---|--|

answer options: Not at all serious/not a problem – Not very serious – Somewhat serious – Very serious

4.4.2 Importance

Please rank the issues below by how important you think they are as problems in our society. You can drag and drop issues to change their order.

Experiment 1:

- Poverty
- Depression
- Obesity (being seriously overweight)
- Car accidents
- Smoking
- Unemployment
- Air pollution
- Climate change

Experiment 2:

- Poverty
- Suicides
- Addiction to opioids (strong painkillers)
- Car accidents
- Smoking
- Unemployment
- Air pollution
- Climate change

Experiment 3:

- Poverty
- Suicides
- Addiction to opioids (strong painkillers)
- Car accidents
- Sexual assault
- Unemployment
- Air pollution
- Climate change

4.4.3 Spending support

For each of the issues below, how much would you favor extra government spending to tackle them?

Experiment 1:

- Poverty
- Depression
- Obesity (being seriously overweight)
- Car accidents

Experiment 2:

- Climate change
- Air pollution
- Suicides
- Addiction to opioids (strong painkillers)

Experiment 3:

- Unemployment
- Sexual assault
- Poverty
- Air pollution

answer options: Do not favor – Favor a little – Favor moderately – Favor very much

4.5 Self-interest

Experiment 1, gender: Now, we would like you to think about whether some issues could happen to you personally. For each of the issues below, do you think this is something that will happen to you in the future?	Experiment 2, race: Now, we would like you to think about whether some issues could affect you personally. For each of the issues below, do you think this is something that could affect you in the future?	Experiment 3, LGBT: For each of the issues below, do you think this is something that could happen to you or someone close to you in the future? <ul style="list-style-type: none">• Unemployment• Sexual assault
<ul style="list-style-type: none">• Poverty• Depression• Obesity (being seriously overweight)• Car accidents	<ul style="list-style-type: none">• Climate change• Air pollution <p>And for each of the issues below, do you think this is something that could happen to you or someone close to you in the future?</p> <ul style="list-style-type: none">• A suicide attempt• Addiction to opioids (strong painkillers)	

answer options:

- Will probably not [happen/affect me]
- May or may not [happen/affect me]
- Will probably [happen/affect me]
- [Has/is] already [happened/affecting me]
- Prefer not to answer

4.6 Prior/Posterior Beliefs

Next, we would like to ask you about a few issues in the United States, and whether they happen:

- more often to [men / White people / LGBT people],
- more often to [women / Black and Latino people / heterosexual (non-LGBT) people], or
- about as often to [men / White people / LGBT people] as to [women / Black and Latino people / heterosexual (non-LGBT) people].

If you are unsure about an answer, please don't look up more information—instead, just give us your best guess.

Experiment 1:

In your opinion, [does poverty/does depression/does obesity (being seriously overweight)/do car accidents] happen more to men, more to women, or is it about the same?

Experiment 2:

In your opinion, [do suicides/does addiction to opioids (strong painkillers)] happen more to White people, more to Black and Latino people, or is it about the same?

Experiment 3:

In your opinion, does [unemployment/sexual assault] happen more to LGBT people, more to heterosexual people, or is it about the same?

In your opinion, [does climate change/does air pollution] affect White people, does it affect Black and Latino people more, or is it about the same?

answer options:

- More to [men / White people / LGBT people]
- More to [women / Black and Latino people / heterosexual (non-LGBT) people]
- About the same
- Don't know

[after every belief question, unless respondent said "Don't know"]: How much

confidence do you have in your answer?

- A lot
- A moderate amount
- A little

4.7 Just world beliefs

[Experiment 1–3] Finally, we would like to know a bit more about how you view the world. Please tell us how much you agree or disagree with each of the statements below.

- I feel that people get what they are entitled to have
- I feel that a person's efforts are noticed and rewarded
- I feel that people who meet with misfortune have brought it on themselves
- I basically feel that the world is a fair place

answer options: Strongly disagree – Disagree – Somewhat disagree – Neither agree nor disagree – Somewhat agree – Agree – Strongly agree

4.8 Treatment

We would like to share with you a piece of information about a social issue in the United States. Please take a moment to read it.

4.8.1 Gender

women:

- In the US, poverty happens more to women than to men. Women are 30% more likely to be living in poverty than men.
- In the US, depression happens more to women than to men. Women are twice as likely as men to be have depression.

men:

- In the US, car accidents happen more to men than to women. Men are twice as likely as women to die in a car crash.
- In the US, obesity happens more to men than to women. Men are 20% more likely than women to be seriously overweight.

This info comes from [the US Census Bureau, Gallup, the Centers for Disease Control and Prevention, the Kaiser Foundation].

4.8.2 Race

minorities:

- In the US, climate change affects minorities more than white people. Black and Latino people are already 15% times more likely than white people to die from causes related to very hot weather.
- In the US, air pollution affects minorities more than white people. Black and Latino people live in places where the air has 40% more of the harmful chemical NO₂ compared to white people.

white people:

- In the US, suicide affects white people more than minorities. White people are three times more likely to commit suicide than Black or Latino people.
- In the US, addiction to opioids (strong painkillers) affects white people more than minorities. White people are twice as likely to die from an opioid overdose than Black or Latino people.

This info comes from [the American Journal of Epidemiology, Environmental Health Perspectives, the Suicide Prevention Resource Center, the Kaiser Foundation].

4.8.3 LGBT

- In the US, unemployment affects LGBT people more than straight people. LGBT Americans are 50% more likely to be jobless than straight Americans.
- In the US, sexual assault affects LGBT people more than straight people. LGBT Americans are three times more likely to have been sexually assaulted than straight Americans.

This info comes from [the Williams Institute, the American Journal of Public Health].

References

- A. J. Berinsky, G. A. Huber, and G. S. Lenz. Evaluating online labor markets for experimental research: Amazon. com’s mechanical turk. *Political analysis*, 20(3):351–368, 2012.
- J. G. Bullock, A. S. Gerber, S. J. Hill, and G. A. Huber. Partisan bias in factual beliefs about politics. *Quarterly Journal of Political Science*, 10: 519–578, 2015.
- R. K. Carney and R. D. Enos. Conservatism and fairness in contemporary politics: Unpacking the psychological underpinnings of modern racism. In *Presented at the 2017 NYU CESS Experiments Conference*, 2017.
- C. J. Cohen. *The boundaries of blackness: AIDS and the breakdown of black politics*. University of Chicago Press, 1999.
- P. J. Conover. The influence of group identifications on political perception and evaluation. *The Journal of Politics*, 46(3):760–785, 1984.
- A. Coppock. Generalizing from survey experiments conducted on mechanical turk: A replication approach. *Political Science Research and Methods*, pages 1–16, 2018.
- K. Deaux. Social identification. In E. T. Higgins and A. W. Kruglanski, editors, *Social psychology: Handbook of basic principles*, pages 777–798. New York: Guilford Press, 1996.
- C. Gay, J. L. Hochschild, and A. White. Americans’ belief in linked fate: A wide reach but limited impact. APSA 2010 Annual Meeting Paper, 2014.

- A. G. Greenwald, M. R. Banaji, L. A. Rudman, S. D. Farnham, B. A. Nosek, and D. S. Mellott. A unified theory of implicit attitudes, stereotypes, self-esteem, and self-concept. *Psychological review*, 109(1):3, 2002.
- F. C. Harris. The rise of respectability politics. *Dissent*, 61(1):33–37, 2014.
- E. B. Higginbotham. *Righteous discontent: The women’s movement in the Black Baptist church, 1880-1920*. American Mathematical Soc., 1993.
- K. Imai, L. Keele, and T. Yamamoto. Identification, inference and sensitivity analysis for causal mediation effects. *Statistical science*, pages 51–71, 2010.
- D. R. Kinder. Opinion and action in the realm of politics. In S. T. F. D. T. Gilbert and G. Lindzey, editors, *The Handbook of Social Psychology*, pages 778–867. New York, NY: McGraw-Hill, 1998.
- J. A. Krosnick. Government policy and citizen passion: A study of issue publics in contemporary america. *Political behavior*, 12(1):59–92, 1990.
- C. W. Leach, M. Van Zomeren, S. Zebl, M. L. Vliek, S. F. Pennekamp, B. Doosje, J. W. Ouwerkerk, and R. Spears. Group-level self-definition and self-investment: a hierarchical (multicomponent) model of in-group identification. *Journal of personality and social psychology*, 95(1):144, 2008.
- M. J. Lerner. *The belief in a just world*. Springer, 1980.
- K. E. Levay, J. Freese, and J. N. Druckman. The demographic and political composition of mechanical turk samples. *Sage Open*, 6(1):2158244016636433, 2016.

- J. M. Marques, V. Y. Yzerbyt, and J.-P. Leyens. The “black sheep effect”: Extremity of judgments towards ingroup members as a function of group identification. *European Journal of Social Psychology*, 18(1):1–16, 1988.
- P. D. McClain, J. D. Johnson Carew, E. Walton Jr, and C. S. Watts. Group membership, group identity, and group consciousness: Measures of racial identity in American politics? *Annual Review of Political Science*, 12: 471–485, 2009.
- A. H. Miller, P. Gurin, G. Gurin, and O. Malanchuk. Group consciousness and political participation. *American journal of political science*, pages 494–511, 1981.
- K. J. Mullinix, T. J. Leeper, J. N. Druckman, and J. Freese. The generalizability of survey experiments. *Journal of Experimental Political Science*, 2(2):109–138, 2015.
- M. Prior, G. Sood, K. Khanna, et al. You cannot be serious: The impact of accuracy incentives on partisan bias in reports of economic perceptions. *Quarterly Journal of Political Science*, 10(4):489–518, 2015.
- C. Rainey. Arguing for a negligible effect. *American Journal of Political Science*, 58(4):1083–1091, 2014.
- G. R. Sanchez and E. D. Vargas. Taking a closer look at group identity: The link between theory and measurement of group consciousness and linked fate. *Political research quarterly*, 69(1):160–174, 2016.
- D. Z. Strolovitch. Do interest groups represent the disadvantaged? advocacy

at the intersections of race, class, and gender. *The Journal of Politics*, 68 (4):894–910, 2006.

I. K. White. When race matters and when it doesn't: Racial group differences in response to racial cues. *American Political Science Review*, 101(2):339–354, 2007.

C. Wlezien. On the salience of political issues: The problem with 'most important problem'. *Electoral Studies*, 24(4):555–579, 2005.