Pre-analysis plan for an experiment on gender group interests

Clara Vandeweerdt

September 12, 2018

1 Introduction

Theories of political opinion, both old (Campbell et al., 1960; Kinder, 1998; Conover, 1984) and new (Achen and Bartels, 2016) have emphasized social groups as a relatively easy reference point for ordinary people trying to make sense of the political world. Compared to concepts such as ideology or values, group memberships are concrete, visible, and they play are large role in daily life. They are also connected to emotion in a way that makes them highly salient (Green et al., 2004; Mason, 2018). As a result, people's group memberships, people's group memberships, and their attitudes toward their in- and out-groups, have the potential to be strong drivers of political opinions. In particular, we expect people to have policy opinions that benefit their in-group (i.e., a group that they are a member of), even or especially at the cost of an out-group.

In this pre-analysis plan, I introduce the first study in a series of experiments that ask whether in-group interests structure people's opinions. In this study, I investigate whether people change their opinions about societal problems when they learn that those problems affects their gender group especially. Later iterations will investigate the same question for racial groups and age groups.

The main hypothesis in this survey experiment, is that people change their attitudes about an issue when they are told that the issue happens more frequently to people of their own gender. In particular, I am interested in three kinds of attitudes: concern, salience, and support for government spending. In addition, I investigate whether the key mechanism is in-group favoritism. I do this, first, by checking whether in-group identification moderates the effect of the new information. Second, I try to rule out an alternative mechanism: that people use the group information as a heuristic for their own self-interest. In other words, I see whether the effect is due to respondents changing their beliefs about whether the issue will affect them personally. Finally, if we find a significant treatment effect, I ask whether the effect is due to priming, and I use instrumental variables analysis to focus on the effect of actual (rather than intended) changes in gender disparity beliefs.

2 Design

This project involves a survey experiment in two phases. In phase 1 of the study, I measure people's group identities, their pre-treament attitudes about several

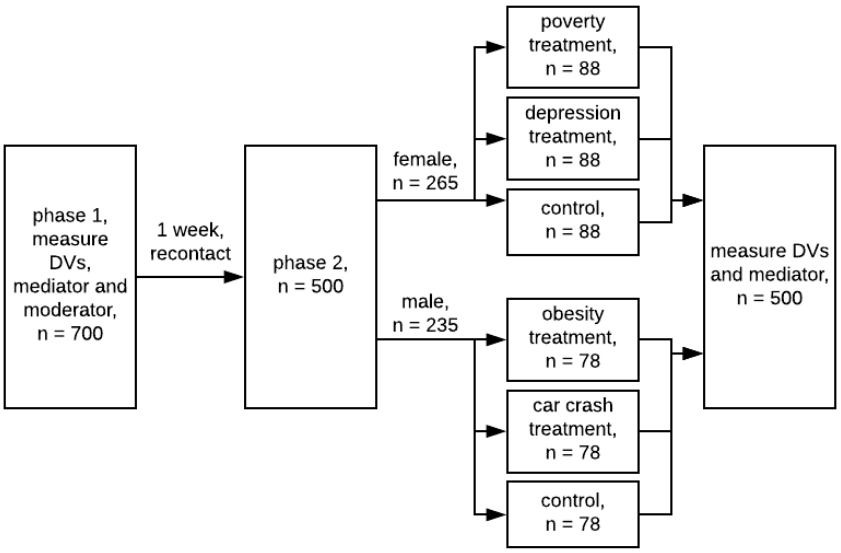

Figure 1: Flow diagram of the experimental procedure, including sample sizes.

social issues, their prior beliefs about who is affected by those issues, and their perceptions about whether those issues are a threat to themselves. One week later, respondents are recontacted for phase 2, which includes the experimental manipulation (in-group interest information) and post-treatment measures of issue attitudes, beliefs and self-interest.

In the following sections, I introduce the sample and give an overview of the experimental procedure. Next, I dive deeper into descriptions of the treatment and of the dependent variable measurement. Figure 1 summarizes the procedure, including sample sizes at each point.

2.1 Sample

The sample for this study will consist of about 500 American respondents (18 years and older) recruited in September 2018 through the Amazon Mechanical Turk platform. In the Mechanical Turk task (or *HIT*) description, respondents are told that the survey take about three minutes and is about "social groups and the issues that people face in their lives."

Some participants drop out between phases–I expect a retention rate of about 65-75%. Drop-out is not a problem for internal validity here, since dropping out is independent of the treatment. As a result, in order to obtain a final sample size of 500, I plan to recruit about 667 respondents initially. Depending on the retention rate, I will increase the sample for phase 1 until the sample size for phase 2 is approximately 500. The expected gender ratio, based on earlier MTurk surveys, is 53% female, 47% male.

Two of the three dependent variables are measured on a four-point Likert scale. With 500 participants, and a panel design, the minimum detectable effect

size (with a power of .80) for the effect of in-group interests would be between .19 and .31, depending on the amount of noise added between waves. The third dependent variable, measured as an eight-place ranking, would have a minimum detectable effect size between .36 and .49.

2.2 Experimental procedure

The procedure for phase 1 of the study is as follows:

- Ask respondent about their gender, and level of identification with their gender group.
- Dependent variables, pre-treatment: for each of the issues, ask respondent about:
  - whether they see the issue as a (serious) problem
  - how important the issue is compared to others
  - support for government spending on issue
- Mediator, pre-treatment: Ask respondent how likely it is that the problem will be a threat to them in the future.
- In-group beliefs, pre-treatment: ask respondent which group is disproportionately affected by each of the issues, and how certain respondent is about his or her answer.
- Four-item version of Just World Belief scale (Lerner, 1980).

To measure membership in a gender group, I ask people to choose which gender group(s) they belong to. Next, to measure group identification, I use the Centrality subscale developed by Leach et al. (2008). The scale consists of three items (with seven-point agree-disagree scales as answer options), which we average into an overall identification 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 question about beliefs consists of two parts for each issue. First, we ask whether the issue happens more to men, more to women, or whether it is about the same. Next, we probe whether the respondent has a lot of confidence, a moderate amount, or only a little confidence in their answer. This results in a seven-point scale. Respondents who answer "about the same" are placed at the midpoint. Respondents who answer 'more to men" or "more to women" are place on either side of the midpoint, with the most confident respondents sitting at the ends. Finally, the scale is coded in such a way that 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 the second phase, I implement the following procedure:

- Randomly assign respondent to an issue matching their gender (poverty and depression for women; obesity and car accidents for men) or to the control group. Treated respondents receive information that their in-group is disproportionately affected by the issue. Control respondents receive no information.
- Dependent variables, post-treatment: for each of the issues, ask respondent about:
  - whether they see the issue as a (serious) problem
  - how important the issue is compared to others
  - support for (a) policy solution(s)
- Mediator, post-treatment: Ask respondent how likely it is that the problem will be a threat to them in the future.
- Beliefs, post-treatment: ask respondent which group is disproportionately affected by each of the issues, and how certain respondent is about his or her answer.

2.3 Treatment: gender groups and their issues

This study covers four issues: poverty, depression, obesity, and car accidents. Each of these issues affects one gender disproportionately: poverty and depression for women; obesity and car accidents for men. In the second phase of the experiment, I randomly assign participants to one of two issue treatment groups, or to the control group. Treated participants read a statistic about how one of the two issues affects their gender group especially. The treatment statements are:

- 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.
- 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.

2.4 Dependent variables: concern, salience, spending support

I measure three outcome variables: concern about the issue (i.e. whether the issue is seen as a problem, cf. Wlezien 2005), salience of the issue (i.e. its relative importance compared to other issues, cf. Conover 1984), and support for government spending to help tackle the issue. I ask each respondents about all three dependent variables for each of the four issues involved.

Problem perception is measured on a four-point likert scale. from "Not at all serious/not a problem" to "Very serious". In the question about salience, people

rank the four studied issue along with four other ones: smoking, unemployment, air pollution and climate change. The question about government spending is once again a four-point Likert scale ranging from "Do not favor" to "Favor very much".

3 Analyses

In this section, I describe the analyses I plan to do on the data from this study. In addition, I plan to pool these data with results from two follow-up experiments on race and age, to be reported in the same paper. The pooled analyses will be identical to the analyses outlined below, but combining data from all three studies. the goal is to understand the effect of learning group interests in general, as well as by social group type.

3.1 Dependent variables and units

The dependent variables in this survey are three types of attitudes about social issues: concern, perceived issue importance, and support for spending. We measure each of the variables for four issues that have a gender bias: poverty, depression, obesity, and car accidents. Since I measure attitudes in both phases of the experiment, the outcome variable in the analyses is the difference between a participant's attitudes before treatment, and those after treatment.

I conduct three main analyses, one for each dependent variable (attitude type). For treated participants, I leave non-treated issues out of the analysis, in case the treatment spills over into other issue opinions (more on this in section A.1). For control participants, I include their responses on both issues related to their gender, i.e. both issues on which they textitcould have been treated. Thus, the unit of analysis here is the person-issue, and n will be about 667 person-issues: one issue coming from each of the roughly 333 treated participants, and two issues coming from each of the roughly 167 control participants. All analyses use standard errors clustered at the person level.

3.2 Effects of in-group interests

The main research question is whether the change in issue attitudes between phases is larger for treated person-issues than for non-treated person-issues. For each of the dependent variables, we start out with a basic model including only the following indicators:

- Treated: indicates whether the person received an informational treatment connecting their group to the issue
- Issue: indicates which issue the attitude question was about

The coefficient on Treated gives us an overall estimate of whether respondents change their attitudes (concern, salience, spending support) more when they learn that their group is particularly affected by an issue. The effect is averaged across genders and issues. All models, including this one, are linear models.

Next, we may want to know whether the effect of learning is particularly large for one gender group or issue. To answer that question, we can estimate a separate model for each of the four issues, including only the observations (from both the treatment and control groups) on that issue. For clarity, I will show the effect size estimates (and their uncertainty) in a graph rather than reporting a regression table for each of the 12 (3 dependent variables $\times$ 4 issues) analyses.

This issue-by-issue analysis is relatively low-powered: the minimum detectable effect size is around .45 on a 4-point scale, if we want a power of .80. For that reason, I plan to combine it with the issue-by-issue analyses of the two follow-up experiments mentioned above. By presenting all the issue-specific results at once, we are able to explore possible patterns in the kinds of issues that seem to show most group-centric opinion formation.

3.3 Group identification as a moderator

Once we have established whether there is a main effect of in-group interests, we want to know whether the effect is larger for people who identify more with the in-group. To address that hypothesis, I add the following covariates to the basic model:

- Group identification: constructed from four items measuring the connection between self and group (see above)
- interaction between Treated and Group identification

The interaction effect indicates whether or not the in-group interest treatment is most effective for those respondents who are strong identifiers. To interpret its size, we may calculate the estimated effect of the treatment for respondents who are one standard deviation above or below the mean identification level, respectively.

3.4 Self-interest as a mediator

Finally, if there is an effect of in-group interests on any of the three issue attitudes, we may be interested to see whether this effect is mediated by people's perception of their self-interest. To that end, I construct another first-differenced variable: the difference between phase 2 and phase 1 in the respondent's perceived likelihood that the problem will be a threat to them in the future.

The indirect effect of the information treatment on attitudes–that is, the effect by way of self-interest–is identified under rather stringent conditions. The key assumption here is that the value of the mediator has to be as-if random given the value of the treatment (Imai et al., 2010). Here, we get a leg up from the fact that both the mediator (self-interest) and the outcome (issue attitude) are first-differenced between phase 1 and phase 2. As a result, the only confounders that can interfere with as-if randomness are ones that change between phases, within person-issues. That is why we can feel fairly confident estimating the effect using two simple linear regressions: one of the mediator on the treatment indicator, and one of the outcome on treatment indicator and mediator (Baron and Kenny, 1986). We can calculate standard errors in the way suggested by Imai et al. (2010).

First-differencing protects us against any confounders that are stable within person-issue combinations. For example, it is not a problem if a person's general pessimism predicts both how seriously they take a societal problem, and how much they anticipate being harmed by that problem themselves. But we can still imagine that some time-varying confounders exist. For example, if a respondent loses their job in the week between the two phases, he or she might come to see poverty as more of a threat to him- or herself, and as more of a societal problem. As this example illustrates, most confounders will affect mediator and outcome in the same direction, causing an upward bias in our indirect effect estimates. For that reason, the effect estimates should be read as an upper bound.

Appendices

A Priming and change in beliefs

In the following analyses, I first investigate whether the treatment is effective because it primes gender as a political division. Next, I conduct an instrumental variables analyses to see how effective it is to change people's beliefs about the gender imbalance in a particular issue. These analyses will be included in an appendix to the paper; only their conclusions will be referenced in the main text. Since both analyses will be underpowered if the treatment effect is small, I will only conduct the analyses if we find a statistically significant main effect of the treatment.

A.1 Priming

The experimental treatment in this study involves telling participants about how an issue affects people of their gender especially. The effect of the treatment can be due to one of two things: the impact of the information itself, or the fact that the treatment primes gender as a lens that respondents can use to view the world. We cannot ask people directly whether the treatment primed them–first, because the effect of priming may be unconscious, and second, because asking such a question would itself prime all participants, including ones in the control group.

Still, we can indirectly address priming by looking at opinion change in issues that people were not treated on. Specifically, if the treatment highlights gender as a social divide, then people should also change their opinion on other, non-treated issues that they know to be related to gender. To check this, we can run an analysis on only the issues that respondents were not treated on (meaning all issues for control respondents, and the three non-treated issues for treated respondents). I will limit the analysis to those issues that respondent believed, in phase 1, to be more frequent for people of their own gender. Finally, we can only do this for the outcome variables of concern and spending support. On the salience measure, non-treated issues are by definition impacted by any treatment effect, since they must go down in the ranking if the treated issue is to go up.

I regress opinion change on the following covariates:

- Treatment Group: indicates whether the person was in a treated group as opposed to the control group
- Certainty: indicates pre-treatment level of certainty about the association between in-group and issue
- the interaction between Treatment Group and Certainty.

plus an indicator for Issue. If the treatment works as a prime, we would expect there to be a main effect of Treatment Group in this analysis. The reason is that for people in the treated groups, the gender-priming effect of the treatment should change their opinion about non-treated issues, as long as they believe those issues to be associated with their own gender group. Moreover, we may expect Treatment Group to interact with Certainty, as this will work best if a respondent is quite certain that an issue is skewed toward his or her own gender group.

A.2 Change in beliefs

In the main analysis, I measure the average effect of receiving in-group information. To be precise, this is an intention-to-treat effect rather than a treatment effect. For various reasons, the treatment may not actually have affected some respondents' belief in the connection between their group and the issue at hand. For example, some respondents are no doubt aware of the connection beforehand; for those people, the treatment cannot have an effect on their attitudes. Fortunately, the design includes a measure of belief change: it is the first-difference between pre- and post-treatment belief in the association between in-group and issue (measured on a seven-point scale, see above).

In order to estimate a true average treatment effect, we need to conduct an instrumental variables analysis, using the treatment as an instrument for belief change. The effect we are estimating is now a conditional one: it is the effect of the informational treatment on compliers, people who would actually change their beliefs about the issue when faced with an informational treatment like ours. In practice, this means running a two-stage least-squares regression. First, I regress the first-difference in respondents' beliefs on the treatment. Then, I regress the outcome on the predicted values from this first-stage regression.

The identification of a treatment effect hinges on the exclusion restriction: the assumption that the instrument (information about the gender imbalance in an issue) can only affect the outcome (issue attitudes) through the treatment of interest (beliefs about the gender imbalance in an issue). This restriction is violated if the analyses above suggest an effect of priming. For that reason, if we have strong suspicions that there is a priming effect (meaning either a statistically significant effect of Treatment in the analysis from section A.1, or a statistically significant interaction between it and Certainty), we will not conduct these instrumental variables analyses.

References

C. H. Achen and L. M. Bartels. *Democracy for realists: Why elections do not produce responsive government.* Princeton University Press, 2016.

R. M. Baron and D. A. Kenny. The moderator–mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations. *Journal of personality and social psychology*, 51(6):1173, 1986.

A. Campbell, P. E. Converse, W. E. Miller, and D. E. Stokes. *The American voter*. New York: John Wiley, 1960.

P. J. Conover. The influence of group identifications on political perception and evaluation. *The Journal of Politics*, 46(3):760–785, 1984.

D. P. Green, B. Palmquist, and E. Schickler. *Partisan hearts and minds: Political parties and the social identities of voters*. Yale University Press, 2004.

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.

C. W. Leach, M. Van Zomeren, S. Zebel, 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.

L. Mason. *Uncivil agreement: How politics became our identity*. University of Chicago Press, 2018.

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