Parallel Ads to Mint Partisans

PRE-ANALYSIS PLAN

November 26, 2018

Introduction

In late October and early November 2018, a two-wave survey experiment was conducted to evaluate the effects of eight original 30-second video advertisements on political party identification. The experiment was conducted over Amazon's Mechanical Turk service.

Building on past experiments, this experiment was designed to evaluate whether video messages can increase party identification. The experiment tests eight ads, four of which were meant to increase identification with the Democratic Party and four meant to increase identification with the Republican Party.

Each of the ads tests one of four theories about partisan identification. For each motivating theory, two ads were produced, one for Republicans and the other for Democrats. The scripts are almost entirely the same, while the party being promoted varies. Some ads test *performance,* or whether claims to superior performance when a party is in office can increase identification with that party. Other ads test *persona,* or whether the charisma of individual party leaders can increase identification with that party. Still others test *social identity,* or whether conveying information about the social groups that make up a party's coalition can increase identification with that party. Finally, other ads test *issues,* or whether messages which emphasize policies a party supports can increase identification with that party.

Hypotheses

The principal hypothesis is that exposure to the ads will increase identification with the party being advertised. Our outcome measure will be the 7-point Party ID scale. A one-sided test will be used to assess this hypothesis at alpha .05.

We will also evaluate how the ads affect Party ID when they are grouped by motivating theory and compared to placebo. We do not have strong expectations about which theories, if any, will prove better than others at affecting Party ID.

In addition, we will also measure how the ads affect partisan intensity---that is, how strongly people feel about their attachment to a political party. Details appear below.

Sample

Subjects were recruited over Amazon's Mechanical Turk platform. Demographic information, including Party ID, was collected in the first wave. Randomized exposure to videos occurred in the second wave, in accordance with a block randomization scheme described below. 1,082 subjects completed the second wave. This pre-analysis plan is prospective, insofar as we have not yet accessed the outcome data.

Random Assignment of Treatment

The accompanying randomization script shows how the block-random assignment was conducted. Blocks were constructed using BlockTools on R, with blocks based on pre-treatment Party ID, age, responses to 0-100 feeling thermometers for the Republican and Democratic Party, education on a 5-point scale, race and responses to a question about how closely the respondent follows politics as measured on a 4-point scale.

Intervention

The eight treatment videos used in this study may be found in the attached supplementary materials. The scripts were written by the researchers. The videos were produced by an undergraduate research assistant without the aid of any party committee or political candidates. A same-length commercial for Bounty paper towels was used as a placebo.

Data and Outcome Measures

Our primary outcome measure will consist of subjects' post-treatment self-reported Party ID as measured on the standard 7-point scale. To measure partisan intensity, we will "fold" the 7-point scale so that 4 becomes 0, 3 and 5 becomes 1, 4 and 6 become 2, and 1 and 7 become 3.

Method for Estimating Average Treatment Effects

One regression model will measure post-treatment Party ID on a seven-point scale, with indicator variables for the eight treatment conditions and the placebo condition omitted. As a covariate, we will include subjects' pre-treatment Party ID.

Another regression will be constructed identically but use the partisan intensity scale described above as the dependent variable. Partisan intensity measures overall partisan attitudes as opposed to a focus solely on party identification.

For all regression estimates, we will report 95% confidence intervals for the average treatment effect, constructed from HC2 robust standard errors multiplied by the appropriate critical value from the *t*-distribution. Given our large sample size, we will obtain two-sided p-values by comparing observed values to a null distribution approximated by the appropriate *t*-distribution. We will call estimates with p-values less than 0.05 "significant."

We will also report regressions with only treatment indicators, omitting pre-treatment Party ID. We expect these results to be similar but less precisely estimated. When interpreting the results, we will rely primarily on the covariate-adjusted estimates.

Heterogeneous effects analysis

We might reasonably suspect that the effects of the videos could be different for subjects with different partisan attachments. To explore this possibility, we will conduct the above analysis separately for the seven partisan subgroups defined by the seven-point party id scale. Because any particular comparison between the estimated effects for two subgroups will be quite noisy, we will conduct a joint test of the hypothesis that the effects are the same for all seven subgroups, via an F-test (using the HC2 robust variance-covariance matrix).

Default Procedures for Decisions Not Explicitly Specified

For any decisions not explicitly specified in this pre-analysis plan, we plan to follow the "standard operating procedure" document of Donald P. Green's research group (version 1.05, June 7, 2016), which can be found on GitHub.