Pre-Analysis Plan

Is online microtargeting increasing financial support of extreme candidates?

March 2026

Registered prior to data collection

Abstract

This pre-analysis plan (PAP) describes the design, hypotheses, estimands, and analysis strategy for a survey experiment examining whether online political microtargeting increases financial support for candidates. Respondents are randomly assigned to see either a fully microtargeted campaign ad (matching all four of their policy positions) or a fully mismatched ad (matching none), and report their willingness to donate a share of survey earnings to the candidate. The two primary research questions are whether microtargeting increases respondents' willingness to donate, and whether this effect is larger among respondents who hold extreme policy positions than among those who hold moderate ones. All hypotheses, estimands, estimators, and subgroup analyses specified here are pre-registered prior to data collection.

Contents

1 Research Question and Motivation **3**

2 Experimental Design **3**

2.1 Survey Platform and Sample . . . 3
2.2 Procedure . . . 3
2.3 Randomization Structure . . . 4
2.4 Ad Construction . . . 4
2.5 Issue Match Assignment Logic . . . 5
2.6 Mismatch Pool Restrictions . . . 5

3 Outcome Variable **5**

4 Hypotheses **6**

4.1 Primary Hypothesis 1: Any Microtargeting Effect . . . 6
4.2 Primary Hypothesis 2: Heterogeneous Effects by Position Extremity . . . 6
4.3 Secondary Hypothesis . . . 6
4.4 Null Hypothesis . . . 6

5 Estimands **7**

5.1 Primary Estimands . . . 7
5.2 Secondary Estimands . . . 7

6 Estimators and Statistical Analysis **7**

6.1 Primary Estimator: Difference in Means . . . 7
6.2 Subgroup Analyses (Pre-Specified, Exploratory) . . . 8

7 Power Analysis **8**

7.1 Priors from Related Literature . . . 8
7.2 Analytical Framework . . . 9
7.3 Sample Size and Cell Allocation . . . 9
7.4 Minimum Detectable Effects . . . 9
7.5 Power Curves . . . 9
7.6 Interpretation and Stopping Rule . . . 10

8 Notes on Validity **10**

1. Research Question and Motivation

This study examines two related questions about private campaign donations and algorithmic political advertising:

1. Do higher levels of issue-position microtargeting in campaign ads increase respondents' willingness to donate to a candidate?
2. Is issue-position microtargeting more effective at increasing donations among respondents holding extreme policy positions than among those holding moderate positions?

The broader theoretical motivation is whether algorithmic microtargeting—the ability to serve personalized political ads based on voter data—props up political extremism by disproportionately funding candidates who appeal to narrow ideological preferences. If microtargeting is differentially effective on voters who already hold extreme positions, it may generate outsized financial support for extreme candidates from the slice of the electorate most ideologically sympathetic to them. We isolate the targeting mechanism by exogenously varying whether a hypothetical campaign ad matches or does not match the respondent's own policy positions, and test whether that effect is heterogeneous by position extremity.

2. Experimental Design

2.1 Survey Platform and Sample

The study is conducted on the Qualtrics platform with sample recruited among American citizens through the Prolific platform. A pilot sample of approximately $N = 500$ respondents will be collected first to estimate outcome variance and calibrate power for the full study. The full study targets approximately $N = 8{,}000$ respondents, and may utilize a more complicated design that varies the level of microtargeting, depending on the effect sizes recovered from the pilot.

2.2 Procedure

Respondents complete the following sequence of tasks:

1. Informed consent
2. Baseline demographics (age, gender, income, ethnicity, education, region)
3. Two attention check questions
4. Party identification (7-point scale; lean question for Independents)

5. Issue positions on four topics, each administered in its own block: abortion, guns, energy, and immigration
6. Exposure to a single hypothetical campaign ad
7. Donation decision: respondent allocates a percentage of survey earnings to the candidate shown in the ad
8. Debrief: respondents are informed that all ads were hypothetical and that they retain all earnings regardless of their donation decision

2.3 Randomization Structure

The design is a single binary randomization. Each respondent is assigned with equal probability to one of two conditions:

Condition	**Description**	**Assignment Prob.**
Full match ($T_i = 1$)	Ad matches all of respondent's issue positions	50%
No match ($T_i = 0$)	Ad matches none of respondent's issue positions	50%

Table 1: Randomization structure: binary treatment.

In both conditions the candidate is always co-partisan (i.e., shares the respondent's party identification). The full-match vs. no-match contrast provides the sharpest possible test of the microtargeting effect (as opposed to e.g., partial match conditions, which may be tested in the future).

2.4 Ad Construction

A library of 600 unique hypothetical campaign ads was generated prior to data collection. Each ad corresponds to a unique combination of:

- Candidate party: Democrat or Republican (2 levels)
- Abortion position: codes 1–4 (4 levels)
- Guns position: codes 1–5 (5 levels)
- Energy position: codes 1–3 (3 levels)
- Immigration position: codes 1–5 (5 levels)

This yields $2 \times 4 \times 5 \times 3 \times 5 = 600$ unique ads. Every possible combination of respondent positions and treatment assignment maps to exactly one pre-written ad.

2.5 Issue Match Assignment Logic

Each respondent is assigned to either a full-match ($T_i = 1$) or no-match ($T_i = 0$) condition. Assignment is fully random and independent of all respondent characteristics, which licenses causal identification of the treatment effects described in Section 5.

Full-match condition. The ad displays the candidate's position on all four issues (abortion, guns, energy, immigration) as identical to the respondent's own reported position on each issue.

No-match condition. The ad displays a position on each of the four issues that differs directionally from the respondent's own position. Each mismatched position is drawn from the directional mismatch pool for that issue (see Section 2.6). No matched issue positions appear in the ad.

2.6 Mismatch Pool Restrictions

To preserve ecological validity, extreme position codes are never shown as mismatches. Extreme codes are defined as follows:

- **Abortion**: code 1 (legal in all cases; extreme pro-choice); code 4 (illegal in all cases; extreme pro-life)
- **Guns**: ad code 1 (ban all guns; extreme pro-regulation); ad code 5 (no restrictions; extreme anti-regulation)
- **Immigration**: ad code 1 (open borders; extreme pro-immigration); ad code 5 (closed borders; extreme anti-immigration)
- **Energy**: no extreme codes excluded; full directional mismatch pool is used

Mismatch pools consist of non-extreme codes whose directional label (pro/neutral/anti) differs from the respondent's own position. This ensures mismatches are substantively meaningful while avoiding implausible combinations.

3. Outcome Variable

The **primary outcome** (Y_i) is the respondent's hypothetical donation amount, measured as the percentage of their survey earnings they report being willing to contribute to the candidate shown in the ad. As a secondary outcome, we will analyze whether the respondent donates at all (binary indicator: $Y_i > 0$).

We also collected data on self-reported reasons for donating or not donating (open-ended text). Open-ended text will be summarized using LLMs, but is considered a tertiary and exploratory analysis to uncover mechanisms only.

4. Hypotheses

4.1 Primary Hypothesis 1: Any Microtargeting Effect

H1 (Microtargeting increases donations) Respondents assigned to the full-match condition ($T_i = 1$) donate more than those assigned to the no-match condition ($T_i = 0$):

$$\mathbb{E}[Y \mid T = 1] > \mathbb{E}[Y \mid T = 0].$$

4.2 Primary Hypothesis 2: Heterogeneous Effects by Position Extremity

H2 (Extremity moderates microtargeting) The effect of microtargeting on donations is larger for respondents holding extreme policy positions than for those holding moderate positions. Formally, denoting extremity as a binary indicator $E_i \in \{0, 1\}$ (where $E_i = 1$ if the respondent holds at least one extreme position):

$$\mathbb{E}[Y \mid T{=}1, E{=}1] - \mathbb{E}[Y \mid T{=}0, E{=}1] \;>\; \mathbb{E}[Y \mid T{=}1, E{=}0] - \mathbb{E}[Y \mid T{=}0, E{=}0].$$

A respondent is classified as holding an extreme position if their remapped ad code on any issue falls in the extreme category (abortion codes 1 or 4; guns ad codes 1 or 5; immigration ad codes 1 or 5). This hypothesis is tested via a $T_i \times E_i$ interaction term in the primary regression.

4.3 Secondary Hypothesis

H3 (Donation incidence) Microtargeting increases the probability of donating at all, not only the amount:

$$\Pr(Y > 0 \mid T = 1) > \Pr(Y > 0 \mid T = 0).$$

4.4 Null Hypothesis

The null hypothesis for all primary tests is that full microtargeting has no effect on donation amounts or probability relative to no microtargeting, and that this non-effect does not vary by respondent position extremity. In other words, that fully matched ads are equally persuasive

as fully mismatched ads, regardless of whether the respondent's own preferences are extreme or moderate.

5. Estimands

Let Y_i be the donation outcome for respondent i, and $T_i \in \{0, 1\}$ be the treatment indicator ($1 =$ full match, $0 =$ no match).[1]

5.1 Primary Estimands

Estimand	Definition	Hyp.
ATE	$\mathbb{E}[Y \mid T = 1] - \mathbb{E}[Y \mid T = 0]$	H1
$T \times E$ interaction	$\big(\mathbb{E}[Y \mid T{=}1, E{=}1] - \mathbb{E}[Y \mid T{=}0, E{=}1]\big)$ $-\big(\mathbb{E}[Y \mid T{=}1, E{=}0] - \mathbb{E}[Y \mid T{=}0, E{=}0]\big)$	H2

Table 2: Primary estimands. E_i denotes the respondent position extremity indicator.

5.2 Secondary Estimands

Estimand	Definition	Hypothesis
ATE on donation incidence	$\Pr(Y > 0 \mid T = 1) - \Pr(Y > 0 \mid T = 0)$	H3
ATE by party (Democrat)	$\mathbb{E}[Y \mid T = 1, \text{pid} = D] - \mathbb{E}[Y \mid T = 0, \text{pid} = D]$	Exploratory
ATE by party (Republican)	$\mathbb{E}[Y \mid T = 1, \text{pid} = R] - \mathbb{E}[Y \mid T = 0, \text{pid} = R]$	Exploratory

Table 3: Secondary and exploratory estimands.

6. Estimators and Statistical Analysis

6.1 Primary Estimator: Difference in Means

The estimator for the ATE is the simple difference in means, estimated using OLS with adjustment. The primary regression specification for H1 is:

$$Y_i = \alpha + \beta_1 T_i + \mathbf{X}_i^\top \boldsymbol{\gamma} + \varepsilon_i,$$

where $\mathbf{X}_i$ is a vector of pre-treatment covariates (age, gender, education, income, region, party identification), included for variance reduction only.

[1] This may be updated to $T_i \in \{0, 2, 4\}$ to accommodate multiple treatment groups with partial matches in the future.

To test the heterogeneous treatment effect by position extremity (H2), the following augmented specification is estimated:

$$Y_i = \alpha + \beta_1 T_i + \beta_2 E_i + \beta_3 (T_i \times E_i) + \mathbf{X}_i^\top \boldsymbol{\gamma} + \varepsilon_i,$$

where $E_i \in \{0, 1\}$ is a binary indicator equal to 1 if the respondent holds at least one extreme position (abortion codes 1 or 4; guns ad codes 1 or 5; immigration ad codes 1 or 5). The coefficient β_3 is the primary estimand for H2: a positive and significant $\hat{\beta}_3$ indicates that microtargeting is more effective among extreme respondents. We will also report the treatment effect estimated separately within the $E_i = 0$ and $E_i = 1$ subgroups for interpretability.

6.2 Subgroup Analyses (Pre-Specified, Exploratory)

The following subgroup analyses are pre-specified but exploratory, and are not subject to correction for multiple comparisons:

- **By `pid_binary` (Democrat vs. Republican)**: Does the microtargeting effect on donations differ by respondent party?
- **By issue salience** (if measured): Does matching on a respondent's most salient issue produce larger effects than matching on less salient issues?
- **By number of extreme positions**: Among respondents classified as extreme ($E_i = 1$), does holding extreme positions on more issues further amplify the microtargeting effect? This extends the binary extremity indicator in H2 to a count measure.

7. Power Analysis

7.1 Priors from Related Literature

Because the outcome variable—the percentage of survey earnings the respondent is willing to donate—has no direct prior in the literature, we anchor power calculations on the closest analogous evidence base: the distribution of effect sizes in survey experiments on political persuasion and donation-adjacent behavioral outcomes, which are typically estimated to be quite small to null.

Given these priors, we treat $d = 0.10$ as a *pessimistic* anchor (consistent with weak or null persuasion) and $d = 0.20$ as an *optimistic* anchor (consistent with the upper range of survey

persuasion effects, and with the possibility that issue matching is a stronger cue for donation than for vote choice). Our key MDE targets are set at $d = 0.10$ and $d = 0.15$.

7.2 Analytical Framework

For the interaction test (H2), power depends on the magnitude of the $T_i \times E_i$ coefficient and the proportion of respondents classified as extreme (π_E). The variance of the interaction estimator is inflated relative to the main effect by a factor of $1/[\pi_E(1 - \pi_E)]$, so power is maximized when $\pi_E = 0.5$ and decreases as π_E moves away from 0.5. We assume $\pi_E \approx 0.10$ as a conservative baseline.

7.3 Sample Size and Cell Allocation

Under the full study design ($N = 8{,}000$):

- **Full-match condition** ($T_i = 1$): $n = 4{,}000$ respondents.
- **No-match condition** ($T_i = 0$): $n = 4{,}000$ respondents.

7.4 Minimum Detectable Effects

Table 4 reports the minimum detectable effect (MDE, in standardized units) for each primary hypothesis at 80% and 90% power, at $\alpha = 0.05$ (two-tailed), without correction for multiple comparisons.

Hypothesis	**Test**	**MDE (d)**	**Power**	α
H1	ATE: full match vs. no match	0.044	80%	0.05
H1	ATE: full match vs. no match	0.056	90%	0.05
H2	$T \times E$ interaction ($\pi_E = 0.10$)	0.148	80%	0.05
H2	$T \times E$ interaction ($\pi_E = 0.10$)	0.186	90%	0.05
H3	ATE on donation incidence (binary Y)	0.044	80%	0.05

Table 4: Minimum detectable effects (MDE) under the full study design ($N = 8{,}000$, equal allocation), at $\alpha = 0.05$ two-tailed, without multiple comparisons correction. The interaction MDE (H2) assumes $\pi_E = 0.10$.

7.5 Power Curves

Table 5 reports power at three illustrative effect sizes for the primary ATE test (H1) under both the pilot and full study designs.

Design	Power at effect size d		
	$d = 0.05$	$d = 0.10$	$d = 0.20$
Pilot ($N = 500$, $n = 250$ per arm)	0.10	0.29	0.85
Full study ($N = 8{,}000$, $n = 4{,}000$ per arm)	0.40	0.97	>0.999

Table 5: Power as a function of effect size for the primary ATE test (H1), at $\alpha = 0.05$ two-tailed. The pilot is powered primarily for variance estimation; the full study achieves high power across all realistic effect sizes.

The binary design yields a substantially lower MDE than a multi-condition design in which the amount of microtargeting is varied (e.g., $d = 0.044$ vs. $d = 0.069$ at 80% power for a 5-condition design, because the entire sample is concentrated in two cells rather than five).

7.6 Interpretation and Stopping Rule

The pilot ($N \approx 500$) will be used for hypothesis testing in the event that the experiment remains unchanged and continues from the pilot to the full sample. The full study ($N \approx 8{,}000$) will be launched after power recalibration from the pilot is complete. Data collection will not be stopped early based on intermediate results.

8. Notes on Validity

Hypothetical donation outcome. While efforts have been made to increase the costliness of donations and respondents may reasonably believe they are making actual donations, the outcome is not actual donation behavior. Results should therefore be interpreted as effects on stated financial support preferences.

Ad content heterogeneity. Because mismatch codes are drawn randomly from the mismatch pool, the specific content of mismatched ads varies across respondents. This is intentional and by design (a conjoint-style approach that marginalizes over specific content), but it implies that the treatment effect averages over different types of mismatched content.

Neutral respondents. Respondents holding neutral positions (e.g., guns code 3 = "about right") face a mismatch pool that includes both pro- and anti-regulation codes. Their assigned mismatch may land on either side, which could attenuate treatment effects for this subgroup.

Extreme respondents. Respondents at the extremes (e.g., abortion codes 1 or 4) have

a narrower mismatch pool (only one non-extreme option). Their mismatches are less varied than those of moderate respondents.

External validity of the binary contrast. The full-match vs. no-match contrast is a maximal treatment, comparing the best-case microtargeting scenario against the worst-case. Effects in the real world, where targeting is partial, are likely smaller than those estimated here.