

# Minting Partisans: Experimental Interventions to Test Theories of Partisan Change

Alexander Coppock, Donald P. Green, and Ethan Porter\*

August 21, 2018

## Abstract

For more than a half-century, scholars have characterized party attachments and evaluations as causes of vote choice. Because partisan attitudes have rarely been altered via experimental manipulation, this interpretation is the subject of extensive and ongoing debate, with skeptics charging that the putative causal effect of partisanship is properly attributed to policy views, retrospective evaluations, or feelings about party leaders. We attempt to break this deadlock by conducting a series of experiments designed to exogenously shift subjects' partisan affinities through the use of short videos that dramatize performance evaluations, party platforms, political personae, and social group cues. Our experiments have thus far produced only short-lived shifts in partisan attitudes, in keeping with the vast literature emphasizing the stability of party attachments. We conclude by laying out our agenda for further research.

---

\*Alexander Coppock is Assistant Professor of Political Science, Yale University. Donald P. Green is Burgess Professor of Political Science, Columbia University. Ethan Porter is Assistant Professor of Media and Public Affairs, George Washington University. Paper prepared for presentation at the Annual Meeting of the American Political Science Association, Boston, MA, August 30 - September 2, 2018. This research was reviewed and approved by the Institutional Review Board of Columbia University (IRB-AAAN3800). No university funds were used in the production of the treatments or the collection of survey responses. We thank Josh Handelman, Jennifer Mueller and Emily Mellencamp-Smith for their assistance in securing funding, and the Democratic Legislative Campaign Committee and the American Center for facilitating support.

Few literatures in political science are more extensive than the accumulated body of scholarship on the subject of partisanship. Spanning several decades and drawing data from a host of different countries, the study of partisan attitudes encompasses a broad array of theoretical perspectives and research methods. In large part, the vitality of the literature on partisan attitudes – a capacious term that includes the affective sense of identification with a party as well as the cognitive assessments of parties’ attributes and competencies – grows out of their strong correlation with vote preference. People who hold the Labour Party in high esteem are much more likely to vote for its candidates than those who evaluate it negatively or hold a more favorable view of another party. The question is what, if any, causal interpretation should be attached to this correlation? Do party preferences form, develop into feelings of attachment and identification, and express themselves in vote choice? Or are vote choice and partisan attitudes correlated because they are each correlated with the same confounding variables, such as ideological proximity or assessments of party competence?

Over the past several decades, scholars have staked out a variety of positions in debates about partisan attitudes and their causal influence. Campbell et al. (1960, 1966) launched this literature with a forceful characterization of party identification as a longstanding attachment to partisan groups and symbols. This attachment colors voters’ perceptions and anchors their preferences in ways that dampen their reaction to short-term forces, such as political events or the nomination of a particular set of candidates. This view contrasts with theoretical accounts that regard partisan attitudes as a by-product of other evaluations. Voters are said to choose candidates based on which is the most ideologically proximal (Downs, 1957) or, indeed, to select their preferred party based on their perceived issue positions (Jackson, 1975; Franklin and Jackson, 1983). The idea that party attachments evolve over time in response to short-term forces gained increasing prominence during the 1980s and 1990s. Fiorina (1981) argued that partisan attachments represented a running tally of evaluations of past performance; recessions and scandals could tarnish a party’s reputation and put off voters who might otherwise identify with it. Fiorina’s work anticipated the time-series analysis of the partisan balance in the electorate (MacKuen, Erikson and Stimson, 1989) that showed how party attachments change according to the incumbent president’s popularity and the strength of the economy over which he presides.

Each of these arguments has generated a vast empirical literature, and debates continue to rage about almost every important claim, especially those related to the malleability of partisan attitudes in the wake of changing evaluations of the incumbent or the parties’ issue stances (Green and Palmquist, 1990; Schickler and Green, 1997; Green, Palmquist and Schickler, 1998). Although the sheer quantity of data has increased dramatically due to the advent of Internet panel surveys and extensive time-series data, controversies about the causal role of partisan attitudes increasingly hinge on methodological challenges that cannot be resolved with existing data sources. Non-experimental data cannot settle a fundamental inference problem: without knowing how people

came to hold their partisan attitudes, we can only speculate about whether partisanship affects vote choice. In order to break this deadlock, it is necessary to devise randomized interventions that have the potential to shape partisan attitudes. This line of experimental research remains underdeveloped, with just a handful of lab (Rahn, 1993; Cowden and McDermott, 2000), survey (Tomz and Sniderman, 2005), and field studies (Gerber, Huber and Washington, 2010).

Experimental investigation in this domain involves three interrelated research programs. The first is the systematic study of the conditions under which partisan change occurs. The literature has to date suggested an assortment of potential strategies: provide information about the parties' issue stances (Franklin and Jackson, 1983; Tomz and Sniderman, 2005), their record of performance in office (Fiorina, 1981; Brody and Rothenberg, 1988), or a profile of their supporters' ethnic, regional, or class profile (Green, Palmquist and Schickler, 2002). This program of study is potentially quite far-reaching, as it involves experimental variation of messages, messengers, the mode of communication, and a host of other factors. The second research program assesses which facets of partisanship change in the wake of exposure to different messages. Here the question is whether partisan evaluations and attachments change simultaneously or whether one moves first, perhaps gradually bringing about change in the other. The third research program builds on the first two. Once a reliable method for changing partisan attitudes has been discovered, the next task is to track its downstream consequences. Suppose, for example, one were to change party attachments by explaining the issue stances of the parties; would these newly minted partisans become more likely to vote for candidates of their party? Would this pattern hold regardless of whether the candidates in question took positions on the issues that induced partisan change in the first place?

This paper represents an attempt to scout what is effectively uncharted terrain.<sup>1</sup> We begin by laying out a set of theoretical propositions to guide the development of experimental interventions. Next, we discuss some of the methodological challenges that arise when attempting to reliably gauge the effects of interventions on partisan attitudes. We consider in addition the special complications that arise when attempting to assess the downstream consequences of partisan change on variables such as vote choice. We then present a pair of experiments in which Mechanical Turk participants are shown different persuasive videos in the course of an online survey, with follow-up measurements after several days. We conclude by describing how we intend to develop a more elaborate set of experiments, eventually culminating in unobtrusive interventions conducted in real-world settings.

---

<sup>1</sup>As the present-draft is very much a work in progress, we only present analysis of treatments designed to increase support for one party—the Democrats. To enhance verisimilitude, these ads have been made by party professionals. Yet they limit our ability to speak about party identification in general terms. In short order, we plan to test a series of ads, theoretically similar to those we test below, that make parallel cases for both parties.

## Four Theories of Partisan Change

In this section, we review four leading theories of partisan attitude change and their testable implications. The first and arguably most prominent theory emphasizes the role of *spatial proximity*. The voter, whose ideological location may be situated on one or more ideological dimensions, prefers the party that is closest according to some loss function that weights each of the dimensions. For example, suppose the issue dimensions of concern to the voter were social welfare liberalism and international interventionism. If this voter were staunchly opposed to social welfare spending and supportive of increased US military presence abroad, the Republican Party would currently be more spatially proximal than the Democratic Party. In the absence of further information about the issue stances of specific candidates, this voter might make a “standing decision” in favor of Republican candidates due to their presumed like-mindedness. Affiliation with a political party, from this theoretical vantage point, may be viewed as a decision shortcut, with party labels serving as shorthand for ideological location.

Several testable implications flow from the spatial proximity hypothesis. First, information about party platforms may lead voters to update their perceptions of their spatial location vis-à-vis the parties. For example, information about Democrats’ hawkishness and hostility toward social welfare spending could diminish the intensity with which this voter prefers the Republican Party over the Democratic Party. This type of information-based persuasive strategy could be applied to either party or to both parties. As discussed below, the success of this approach depends on whether the information is credible, which means that even nonpartisan messages could be influential. This nonpartisan persuasive strategy seems to underlie websites such as Pew Research’s Political Party Quiz ([www.people-press.org/political-party-quiz/](http://www.people-press.org/political-party-quiz/)) and [www.isidewith.com](http://www.isidewith.com), which attempt to inform voters which political parties most closely match their issue preferences.

Another testable implication concerns the relative importance of different issue dimensions. For example, messages that emphasize the importance of issues on which the Democratic Party is more proximal to the voter than the Republican Party (e.g., whether to permit stem cell research) may change the weights that the voter places on different issue dimensions, which may in turn change the voter’s sense of which party is closer in multi-dimensional space. The extensive literature on “issue ownership” (Petrocik, 1996; Holian, 2004), the idea that parties stake out distinct positions that appeal to voters in specific issue domains, emphasizes the strategic imperative of declaring an election to “be about” an issue dimension along which a candidate or party is favored. A somewhat different persuasive strategy is to question the credibility of a party’s declared position on a given issue. For example, Republicans charge that the Democratic Party’s declared sympathy for middle class tax cuts is disingenuous given its alliance with interest groups that oppose tax cuts for fear of jeopardizing the support they receive from public spending. Voters who care not simply about which party sympathizes with their position but rather focus on how party positions will translate into policy actions may be persuaded by arguments of this kind.

An enduring concern for issue-based theories of partisanship is the possibility that individuals' party attachment cause their issue positions rather than the other way around. Indeed, some experiments have discovered evidence that partisans "follow the leader" (Lenz, 2013) or adopt legislators' positions even when no evidence in favor of them is offered (Broockman and Butler, 2017). Of course, any of four possibilities might hold for an individual voter: positions cause partisanship, partisanship causes positions, both cause one another in a dynamic process, or the coincidence of positions and partisanship is the result of unobserved third factors.

A second class of theories stresses the role of performance evaluations, especially evaluations of the chief executive. As noted above, Fiorina (1981) pioneered research in this area, demonstrating that individuals in panel surveys changed their party attachments as their evaluations of the president changed. Although this specific demonstration has drawn criticism on methodological lines (Green, 1991), it seems clear that performance evaluations shape beliefs about party competence even if they do not markedly alter respondents' party attachments (Green, Palmquist and Schickler, 2002, chp. 5). In aggregate data, evaluations of the party in power deteriorate during economic recessions and improve during expansions, and even the balance of party identification shifts gradually during economic swings.

A subtle but important distinction may be drawn between two types of performance evaluations. The form stressed by Fiorina (1981) is retrospective; voters look at outcomes under the party in power, form an overall assessment, and reward or punish accordingly. An alternative formulation offered by (Achen, 1992): voters consider past performance when updating their assessment of how the parties are likely to perform in the future. From the standpoint of developing persuasive interventions, a retrospective message would stress blame for past failures or credit for past successes, while a prospective message would emphasize which party is better suited to govern in the future. A prospective message might still refer to past outcomes, but it strives to convince the voter that, looking to the future, one party represents a better bet.

A third theory of partisan change focuses on the outsized role that specific political figures play in party politics (Harris, 1954). Presidents such as Eisenhower or Reagan are said to have the capacity to win converts to their party in part because of their personal popularity but also because a large segment of the electorate understands politics at a fairly low "level of conceptualization" (Converse, 1964) that revolves around individuals rather than policies. Consistent with this hypothesis is evidence suggesting that the balance of partisanship shifts in response to changes in presidential popularity (MacKuen, Erikson and Stimson, 1989), a pattern that seems to hold outside the US as well (Clarke, Stewart and Whiteley, 1997). Among the testable implications of this theory is the hypothesis that partisan attitudes can be changed through words and images that dramatize the personal magnetism of partisan figures such as Reagan or, conversely, the failings of figures such as Nixon. This view echoes the way that popular media often portray presidential campaigns, with candidates rising and falling based on their supposed charm and charisma.

A final theory of partisan attitude change stresses the role of group affinities and stereotypes. This theory is premised on the idea that people have a sense of which groups they like or dislike (Brady and Sniderman, 1985; Kinder and Kalmoe, 2017), as well as a mapping of which groups are associated with which parties. For example, a person may dislike corporate executives and associate them with the Republican Party; all else being equal, this perception leads to more negative assessments of the Republican Party. Like the spatial model of issue proximity, a model of group affinities implies that partisan attitudes represent a weighted average of feelings toward a range of different groups that are, to varying degrees, associated with one party or another.

One important implication of this model is that partisan change occurs when salient social groups become associated with a political party. A classic example of this phenomenon occurred after the 1965 Voting Rights Act, which brought large numbers of blacks into the Democratic Party and, in turn, put off Southern whites who were hostile to them. A corollary point is party attitudes may change in the wake of events that call attention to existing ties between a political party and a social group. Achen and Bartels. (2016) argue that the southern realignment occurred because of the importance of “southern identity” for whites, with policy preferences “probably of secondary importance.”<sup>2</sup>

Group imagery naturally lends itself to persuasive political messaging. Ethnic, gender, age, and social class cues are often prominent in political mailings and television commercials. This imagery may affect recipients of this information through either of two channels. First, it may clarify or alter a recipient’s understanding of which social groups support each party; second, it may raise or lower the salience of particular social groups, which may in turn change a recipient’s global evaluation of the parties. For example, a message that links gay people with the Democratic Party may change global evaluation of the Democratic Party if recipients were either unaware of the connection or had accorded this connection little weight in their overall evaluation of the party. Depending on the recipient’s attitudes toward gay people, this information could raise or lower the assessment of the party.

These four theories have somewhat different implications for how, in practice, attitude change might be brought about. First, the persuasive role of the messenger varies across theories. In theories that stress issue positions or performance evaluations, source credibility may be important, as persuasive effects depend on whether the recipient believes the messenger’s assessment of where the parties stand or their past and future competence. In that sense, ostensibly nonpartisan messengers may be highly persuasive. In theories that stress candidates or groups, a credible messenger is certainly a plus, but the messenger also personifies the social supporters of candidates and parties. For example, a black messenger who extols the virtues of the Democratic Party implicitly underscores the connection between the party and racial minorities.

Our goal is emphatically not to adjudicate among these four theories to discover which one is

---

<sup>2</sup>See (Fowler, 2018) for a dissenting view on the causal logic of this argument.

correct and which three are incorrect. Doubtless, all four are true to some extent, at least under some conditions. The challenge is that no one has to date conducted a sustained research program of searching for effective instantiations of these theories. Accordingly, our goal is to learn which of these theories—if any—are useful for developing treatments that cause partisan change.

As the reader will soon see, we have preliminary evidence that treatments based on each of these theories *can* affect some measures of partisan identification in the intended direction. But none can do so for very long. The effects we observed immediately after treatment were no longer detectable one week later. Worth noting, however, is that in all of the studies discussed here, we have only presented subjects with one treatment, and assessed that treatments ability to affect attitudes. Future experiments will show subjects different numbers of treatments, to determine what level of dosage can effectively and durably change partisan identification, if indeed it can be done.

## Methodological Challenges

In this section, we discuss four threats to the validity that arise in the context of experiments that present subjects with persuasive messages concerning political parties. These threats concern (1) the measurement of outcomes, (2) control over the content of the message and the context in which it is received, (3) violations of the excludability assumption when messages are used to set in motion downstream outcomes, such as vote choice, and (4) pre-treatment measurement of individual differences that are thought to interact with the treatment.

### Outcome measurement

Party identification is a specific type of evaluative orientation. It is a self-conception: To what extent does one think of oneself as a belonging to a social group (e.g., Democrat)? The canonical measure of partisan self-image is the traditional American National Election Study branched question, which asks: “Generally speaking, do you usually think of yourself as a Republican, a Democrat, an Independent, or what?” The introductory phrase, “generally speaking” is designed to focus respondents’ attention on their broad orientation rather than their recent voting behavior. The follow-up questions to this stem question<sup>3</sup> are less obviously valid as measures of self-categorization (Green, Palmquist and Schickler, 2002), but we use them in order to generate the traditional seven-point party identification scale. The distribution of this scale for the control group in Study 2 may be found in Figure 1 and shows the overrepresentation of Democrats typical of Mechanical Turk samples (Huff and Tingley, 2015).

In addition to the commonly used seven-point scale, we also deploy other measures of party identification that have been proposed and validated in previous research. Green and Schickler

---

<sup>3</sup>Democrats and Republicans are asked, Would you call yourself a strong [Democrat/Republican] or a not very strong [Democrat/Republican]? while Independent are asked, Do you think of yourself as closer to the Republican Party or to the Democratic party?

(1993) show that a self-labeling exercise produces a valid and reliable indicator of net identification with Democrats versus Republicans. Respondents were asked, “On a scale from 1 to 10, where 10 represents a description that is perfect for you, and 1 a description that is totally wrong for you, how well do each of the following describe you? . . . A Republican . . . A Democrat.”<sup>4</sup> The resulting measure ranges from -9 to 9. Among subjects in the control group of Study 2, for example, the correlation between this measure and the seven-point party identification scale is 0.86.

Another measure of partisan social identity, drawing on the work of Greene (2002) and Huddy, Mason and Aarøe (2015), assesses the social distance that partisans of one party feel toward partisans of the other party. For brevity, we do not employ the entire scale; instead, we focus on the most widely discussed component of this measure, which focuses on marriage: “How would you react if an immediate family member were to marry a Republican? . . . a Democrat” The response options are “I would be unhappy,” “It doesn’t matter,” or “I would be happy.” The net response for the two questions creates a scale ranging from -2 to 2. Due to the attenuated nature of the social distance scale, it is not surprising that the correlation with the seven-point party identification scale is just 0.54. In part, the problem is that the scale is unable to detect gradations of partisanship, and the majority of respondents have a net score of zero. Still, it remains to be seen whether advertisements that stir up partisan sentiments have detectable effects on feeling of social distance.

Our final two measures gauge partisan sentiment rather than party identification per se. One measure is simply the “feeling thermometer” ratings of the Democratic party and Republican party commonly used in the American National Election Studies. Another measure is similar in spirit but perhaps more straightforward in its wording: On a scale from 1 to 7, with 1 being the least favorable and 7 being the most favorable, how would you rate the [Democratic party/Republican party]? As expected, these two measures are highly correlated with one another (0.94) and highly correlated with party identification (0.85 and 0.86, respectively), and in this sample the balance of evaluations favors the Democratic party.

Our core hypothesis is that exposure to Democratic ads should increase net identification with Democrats and net favorability toward the Democratic party. But making a convincing case that attitude change has occurred means showing that the persuasive effects of communication persist over time. This requirement places important practical demands on any experiment, as the subjects must be exposed to the treatment at one point in time and outcomes measured at some point in the future. Although most lab and survey experiments present stimuli and measure outcomes immediately thereafter (perhaps after a few minutes of intervening questions), some notable studies have tracked participants over time. For example, Chong and Druckman (2010) traced the enduring effects of issue frames over the course of 10 days. Internet panels, in particular, facilitate this type of follow-up measurement (Clinton and Lapinski, 2004).

---

<sup>4</sup>Respondents were also asked about other labels, in random order: an evangelical Christian, a feminist, a Southerner, and an environmentalist.



Separating in time the treatment from the measurement of outcomes serves other methodological purposes as well. Presenting subjects who know they are participating in a research study with a video about one or both political parties and shortly thereafter asking questions about the parties makes apparent the connection between the stimulus and the response. If participants sense that the intelligent or desired response is one that praises the party being extolled in the video, its apparent effects may be spurious. To some extent, this concern may be addressed by providing subjects with a cover story (e.g., they are asked to watch the video as though they were evaluating its cinematography) or by burying the treatment video in a series of other irrelevant videos, but the fact that subjects are being paid to watch a party propaganda video in itself is sufficiently unusual and obtrusive to threaten the symmetry of outcome measurement in treatment and control groups. Conducting the treatment and outcome measurement at separate points in time does not fully eliminate this threat to validity, but it does strengthen the credibility of the results, especially if new tasks are introduced in the follow-up session that have no apparent connection to political parties.

### **Control over messages and context**

In principle, messaging experiments could be designed so as to manipulate each aspect of the script, images, and messengers, as well as overall length and the context in which messages are presented. In practice, the number of possible variations becomes unmanageable. As is the case for any experimental agenda, some manipulations must be given priority over others. Because the persuasive effects of video messages have seldom, if ever, been explored systematically, a basic question for pilot testing is whether varying subtle aspects of presentation seems to generate effects. If so, the resulting research program necessarily becomes more extensive.

Another constraint is that certain kinds of combinations are inherently odd or confusing. For example, a narrator may extol the virtues of small government as the camera pans to an image of Ronald Reagan, but a script about the struggle for civil rights could not plausibly use Reagan as a visual image. Scripts and images must often be manipulated as packages, and researchers must devise ways to assess whether one package (small government and an image of Reagan) is comparable to another (civil rights and an image of Martin Luther King). For example, do supporters of each party find each scene equally appealing or persuasive, after standardizing their ratings according to how they each rate politically neutral material?

The testing of messages in field settings raises a more fundamental trade-off between competing experimental aims. An approach that sacrifices unobtrusiveness for a high level of compliance with the assigned treatment is to have canvassers display videos at the doorstep using tablet computers, as in the experiments conducted in the lab and field by Valentino, Traugott and Hutchings (2002). Following Green and Vavreck (2008), one could randomly assign small cable TV stations to treatment and control conditions and conduct follow-up surveys of residents. This approach

has the advantage of presenting the treatment in a naturalistic and unobtrusive manner, with no apparent connection to the outcomes measures, the questions on partisan attitudes. The downside of this approach is the prospect of failing to treat the intended treatment group; some of the people living in the cable TV markets slated for treatment may never watch the ads. To some extent, this problem may be addressed by advertising during a specific TV show, in which case one could screen for people who viewed the show during the advertising period. If viewership of this show is relatively small, however, one may be forced to discard a large number of follow-up interviews with non-viewers. An alternative is to saturate the treatment markets with ads on a variety of shows, but this approach is expensive, especially if considerations of statistical power force the researcher to deploy ads in a large number of target cable TV zones. With the advent of individually-targeted ads, this logistical problem may eventually be solved, but only if outcome measures can be obtained from those individuals.

### **Excludability as a design consideration**

A further set of constraints on the design and deployment of experimental ads stems from the theoretical aspiration of assessing the downstream consequences of changing partisan attitudes on vote choice. The core assumption in any downstream analysis (Green and Gerber, 2002) is that the random intervention that affects the immediate outcome (partisan attitudes) does not directly affect the downstream output (vote choice). The researcher must therefore take care not to deploy persuasive messages that might directly affect vote preference. For example, a treatment video that extols the virtues of a current presidential candidate in an effort to burnish the image of the candidate's party might have a direct impact on voter preference for the candidate mentioned in the ad. To discover that the ads caused the treatment group to later support the candidate is subject to two competing interpretations: the ad changed partisan attitudes, which in turn affected vote choice, or the ad affected vote choice directly by changing perceptions of the candidate. To the extent that candidate-centered ads are used to change partisan attitudes, a convincing test of downstream effects must focus on preferences for candidates in races that are unconnected to the ads. For example, do ads that focus on a Republican presidential candidate have downstream effects on preferences for Republican gubernatorial candidates?

### **Investigation of individual differences**

A final design consideration is how best to facilitate the study of heterogeneous treatment effects. Who changes most in the wake of the experimental messages? Is change, for example, greatest among those with weaker initial partisan attitudes? Assessing individual differences requires additional data collection prior to treatment. If prior partisan attitudes are measured immediately before treatment, respondents may be tipped off to the purpose of the study, or they may answer

follow-up questions in ways that maintain consistency with their pre-treatment answers<sup>5</sup> For this reason, pre-treatment measurement is best conducted in a prior wave of a panel survey, so that time can wash out priming or anchoring effects.

## Design

In this section we describe three studies that share many design features. The first two studies are three-wave panel survey experiments conducted on Mechanical Turk (MTurk). Some scholars take a negative view of studies conducted on MTurk because the population of “workers” is quite different from the national population of American adults in the sense that the distribution of demographic characteristics on MTurk is different from the national population (that American MTurk subjects are just as American as non-MTurker Americans is not in dispute). We chose to use MTurk because conducting multi-wave experiments on the platform is – by an order of magnitude – less expensive than on putatively “nationally representative” samples (Gross and Wood, 2018). Further, recent investigations have demonstrated that both Average Treatment Effect estimates (Mullinix et al., 2015; Coppock, forthcoming) and Conditional Average Treatment Effect estimates (Coppock, Leeper and Mullinix, 2018) obtained on probability and MTurk samples correspond quite tightly. On this basis, our guess is that the results we present below will generalize well to other populations of Americans, though confirmation of this guess will have to await replication on new, more representative, samples.

The third study is a single-shot experiment conducted on Lucid, an alternative source of online convenience samples. Coppock and McClellan (2018) show that treatment effect estimates obtained on Lucid match those obtained on both MTurk and probability samples quite well. For other examples of studies conducted on Lucid, see Flores and Coppock (2018) and Graham (2018). Our Lucid sample consists exclusively of registered voters living in Michigan because we appended this experiment to the survey questionnaire for an unrelated study of Michiganders. As with any convenience sample, our inferences are limited to the set of people who participate. However, all of the theories of partisan change enumerated above apply equally well to people living in Michigan as to the national population.

## Recruitment

In studies 1 and 2, Subjects were recruited to take part in a Wave 1 survey in which demographic characteristics (age, gender, education, race and ethnicity) political variables (ideology, interest in politics, political knowledge) and baseline measures of our outcome measures (feeling thermometers towards various leaders and groups, spending attitudes, and 7-point party ID). In Wave 2,

---

<sup>5</sup>Consistency bias is a concern because it potentially breaks the symmetry between treatment and control. The desire to maintain consistency with prior answers may be especially pressing for members of the control group, who are not subsequently exposed to ads with partisan content.

conducted approximately 10 days later, subjects were randomly assigned to watch either no video or a treatment video before responding to an outcome survey. In Wave 3, conducted approximately 15 days later, we conducted a second outcome survey. Most of the survey questions were identical across the two outcome surveys (and across the two studies), though we did add and subtract a few questions across implementations. See the online appendix for the wordings of all survey items.

Because Study 3 was single-shot, we were unable to obtain a pre-treatment measure of partisanship. We asked all of the demographic questions except partisanship before treatments were allocated, but we did not ask directly about partisanship for fear of contaminating responses to the post-treatment partisanship question. Because liberal-conservative ideology correlates quite well with partisanship, we will reap most of the precision gains we would have gotten from partisanship from ideology.

Table 1 summarizes the dates and sample sizes of all studies at each wave.

Table 1: Dates of implementation

	Wave 1: Collect demographics and baseline attitudes	Wave 2: Allocate treatments and collect post-treatment attitudes	Wave 3: Collect post-treatment attitudes 15 days after treatment
Study 1 (MTurk)	April 30, 2018 (N = 1,548)	May 7 - 10, 2018 (N = 1,1329)	May 23 - 29, 2018 (N = 1,154)
Study 2 (MTurk)	June 6, 2018 (1,599)	June 15 - 21, 2018 (N = 1,359)	July 2 - 9, 2018 (N = 1,138)
Study 3 (Lucid)		July 7 - August 5, 2018 (N = 4,634)	

## Measurement

We measure five outcome variables. Partisanship is measured using the standard branching format. We reverse code this variable so that higher values indicate stronger attachments to the Democrats. For the four dependent variables that are “Net” we take the difference between the Democratic response and the Republican response. Figure 1 shows the distributions and bivariate correlations of these five outcome measures. Table 2 gives the exact wording of all five outcomes and indicates in which studies and waves they were asked. Because we added and subtracted some survey items in each version of the study, the precise set of outcome measure is not identical in each instance.

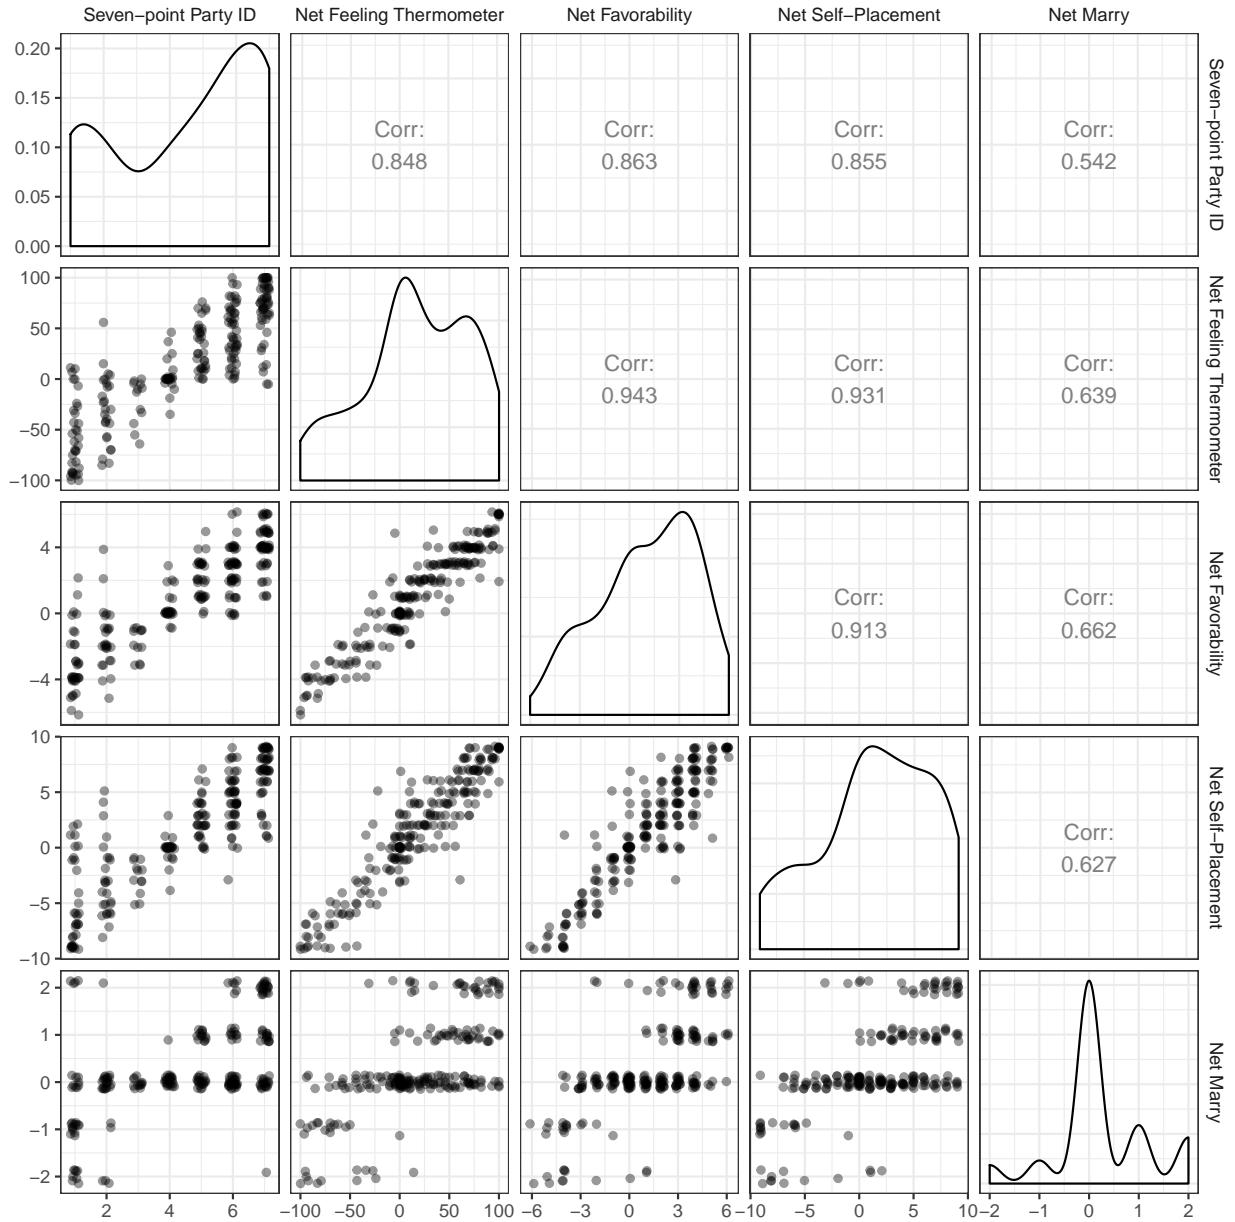
## Treatments

With one exception (A video produced by the DNC that combines aspects of all four theories), we commissioned our own treatment videos from media firms. The main reason we had to do this is that the vast majority of political advertisements are for or against a particular candidate or issue, rather than for or against a particular party. In order study how to change party attachments, we had to create videos de novo. We were also aware that as political scientists, we lack the requisite skills to create compelling advertisements on our own, so we relied on the creativity of advertising professionals. We collaborated with two unnamed advertising agencies to develop treatments that

Table 2: Outcome Measures, by Study and Wave

	Study 1	Study 2	Study 3
<i>Seven-point Party ID</i> Generally speaking, do you usually think of yourself as a Republican, a Democrat, an Independent, or what?  Would you call yourself a strong Republican or a not very strong Republican?  Would you call yourself a strong Democrat or a not very strong Democrat?  Do you think of yourself as closer to the Republican Party or to the Democratic party?	Waves 2 & 3	Waves 2 & 3	Wave 2
<i>Net Feeling Thermometer</i> Please use the sliders to indicate how warm or favorably you feel towards these political groups. Ratings between 50 degrees and 100 degrees mean that you feel favorable and warm toward the person. Ratings between 0 degrees and 50 degrees mean that you don't feel favorable toward the person and that you don't care too much for that person. You would rate the person at the 50 degree mark if you don't feel particularly warm or cold toward the person. [The Democratic Party / The Republican Party]	Waves 2 & 3	Waves 2 & 3	Wave 2
<i>Net Favorability</i> On a scale from 1 to 7, with 1 being the least favorable and 7 being the most favorable, how would you rate the [Republican/Democrat] party?	Waves 2 & 3	Waves 2 & 3	Wave 2
<i>Net Self-Placement</i> On a scale from 1 to 10, where '10' represents a description that is perfect for you, and '1' a description that is totally wrong for you, how well do each of the following describe you? [A Democrat / A Republican]	Not asked	Waves 2 & 3	Wave 2
<i>Net Marry</i> How would you react if an immediate family member were to marry a [Democrat / Republican]? [I would be happy, I would be unhappy, It doesn't matter]	Waves 2 & 3	Wave 2	Wave 2

Figure 1: Outcome Measures



Source: Control group subjects in Study 2 (N = 227). A very small amount of jitter has been added to each point to visually separate each observation.

correspond to each of the four theoretical paths to partisanship outlined above. In the appendix, we provide scripts and links to all the videos. Table 3 shows the number of subjects assigned to each treatment condition. The second study omitted “low performing” videos from the first study and added one video. Firm B produced two versions of an issues ad that discussed Democratic positions on climate change and marriage equality. Version 1 also included material about gun control and version 2 included material about drug decriminalization. The firm produced a third version that did not include any additional issue content over and above climate change and marriage equality, against which we could gauge what, if any, special impact the gun control or drug decriminalization versions may have exerted. For Study 3, Firm B also produced a “Negative Partisanship” ad that portrays the Republican party in a negative light.

Table 3: Assignment of Treatments to Subjects

	Study 1	Study 2	Study 3
No video	129	227	532
Issues (Firm A)	145	0	0
Performance (Firm A)	100	221	580
Charisma (Firm A)	111	0	0
Social Identity (Firm A)	131	0	0
Generic (DNC)	118	231	0
Performance (Firm B)	102	0	554
Charisma (Firm B)	137	0	554
Social Identity (Firm B)	124	0	500
Issues (Firm B, version 1)	123	240	521
Issues (Firm B, version 2)	109	218	547
Issues (Firm B, version 3)	0	222	534
Negative Partisanship	0	0	533

Cell entries are numbers of subjects assigned to each condition.

## Analysis

We will conduct all analyses among subjects who responded in all three waves in order to hold sample composition constant for all analyses. We make the (unverifiable) assumption that there are only two types of subjects, Always-Reporters and Never-Reporters (Gerber and Green, 2012, chp. 7) Always-Reporters respond in wave 3 regardless of treatment condition and Never-reporters do not respond in wave 3, regardless of treatment condition. The “always-reporters” assumption is violated if there are subjects who do or do not respond in wave 3 depending on what video they were assigned to in wave 2. In the online appendix, we report the results of a missingness analysis which shows that treatment assignment does not appear to be related to reporting (Study 1:  $p = 0.69$ , Study 2:  $p = 0.80$ ). Our main estimand is therefore the Average Treatment Effect (ATE) among Always-

Reporters. Just as we do not expect treatment effects to differ dramatically between probability and convenience samples, we do not have strong reasons to suspect that treatment response is different between Always- and Never-reporters, though of course such a claim is fundamentally unverifiable.

We will estimate ATEs by regressing the outcome on a vector of treatment indicators. We will exclusively report models that adjust for seven-point party ID measured pre-treatment, using the adjustment procedure described in Lin (2013). We employ this control exclusively to increase precision and not to adjust for unobserved confounders (Gerber and Green, 2012, chp. 4). Because of random assignment, all such unobserved confounders are balanced in expectation and thus do not induce bias in our estimates of the ATEs.

Our main outcomes of interest are seven-point party ID, the Democrat - Republican difference in 101-point feeling thermometers, and the Democrat - Republican difference in seven-point favorability rating. We report ATE estimates for the other dependent variables in the online appendix.

## Results

Figures 2, 3, and 4 display the estimated average treatment effect of each video (versus control) for seven-point party ID, the Democrat - Republican difference in 101-point feeling thermometers, and the Democrat - Republican difference in seven-point favorability rating, respectively. The red points refer to estimates from Study 1 and the blue triangles to estimates from study 2. The Wave 2 estimates derive from outcome measures taken immediately following the videos and the Wave 3 estimates from measurements taken approximately two weeks after treatment.

## Issues

In study 1, the best performing ad was “Issues (Firm B, version 1).” Version 1 included the additional content about gun control. It exerted a 0.43 point average effect (SE: 0.15) on party ID and a 0.71 point effect (SE: 0.21) on net favorability. These effects appear to have evaporated in the two weeks between waves 1 and 2. In study 2, we see similar effects of the Issues ad, though it was no longer ranked highest of the five videos in study 2.

## Performance

The performance ad commissioned by Firm A performed reasonably well in both studies 1 and 2. As Figure 2 shows, that ad was able to move Party ID toward the Democrats in the second study in both the second and third waves. Yet it was unable to do so in the third study. The performance ad produced by Firm B had no identifiable effect on party ID in any study. As the reader will note, and as was the case here, ads can increase positive feelings about the party, as shown in Figures 3 and 4 without having any concomitant effects on party identification.



## Charisma

While most ads that touted the charisma of Democratic Party leaders proved unable to move subjects in the intended direction, the charisma ad produced by Firm B improved favorability ratings and net feeling thermometer ratings of the party in study 3, again as depicted in Figures 3 and 4. Yet once again, party ID was resistant to this treatment, not only in study 3 but in those that preceded it.

## Social Identity

Regardless of the firm or the outcome measurement, ads which focused on social identity (specifically the Democratic party's inclusiveness of non-whites and non-heterosexuals) did not leave an effect on subjects. The version produced by Firm B did manage to move net feeling thermometer in the third study, as shown in Figure 4, albeit by a very small amount.

The relatively large amount of sampling variability attending to each estimates makes ranking the videos difficult. However, the overall pattern emerges that indeed, the videos can shift partisanship and attitudes about the party in the short term by a modest amount but that any such effects appear to be short-lived.

Figure 2: Effect of Treatments on Seven-point Party ID




Figure 3: Effect of Treatments on Dem - Rep Seven-point favorability

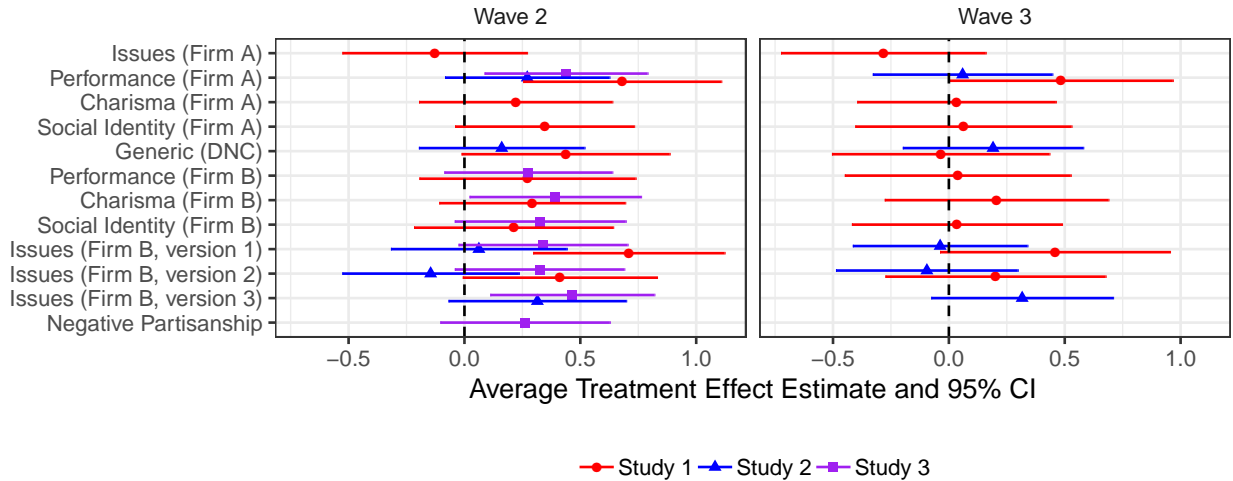



Figure 4: Effect of Treatments on Dem - Rep 101-point Thermometer



## Mapping Out a Research Program

The systematic study of persuasive political messages focused on partisan identification is a vast undertaking, and our approach is to build up a stock of robust results inductively, through testing and replication of multi-arm trials. In the first phase of our investigation, we have presented respondents one random message, produced by professional political ad-makers, designed to increase the extent to which people identify as Democrats. While we have found that some messages are able to push people toward identifying with Democrats, the observed effects have been fleeting, no longer discernible even a week later. More broadly, many messages did not work; that is, they had no detectable short-term effect on favorability or feeling thermometer ratings.

Yet as noted earlier, none of the tests described and analyzed here have tried to measure the effects of multiple messages. In an experiment now in the planning stages, we intend to randomly vary how many messages a subject sees, with some seeing as many as three treatment videos. (We will show only the best-performing videos.) Subjects in other conditions will see fewer treatment videos, and more placebo videos. We anticipate that identification with the Democratic Party will increase as dosage increases. In addition, the studies in this paper have only examined whether effects are still discernible a week later. It is surely possible that, while effects may dissipate over the course of a week, they may in fact be evident a few days after treatment. For this reason, the dosage study will also vary when we re-contact subjects after treatment. Resolving the question of effect decay is especially important for this project; as discussed above, we must be on the watch for subjects adjusting their responses to fit their perceptions of our objectives. Finding that some messages, or some combination of messages, are able to produce lasting effects on our primary outcome of interest would alleviate, though not eliminate, this concern.

Future research will also dive deeper into the plethora of outcome variables. As of now, it appears that measures of *affect* toward the party, as captured by feeling thermometers and approval ratings, are more responsive to messages than measures of party identification. If this continues to be the case, it will suggest that partisan identity is rooted deeper than one message, or multiple messages, can easily penetrate. Finally, we have not yet examined the effects of these messages on arguably the most important outcome measure: Vote choice. If resources permit, we intend to deploy the best-performing ads in the field, either on television or over the Internet, prior to an election.

We expect that the reader may be wondering if our effects would be different had we tested ads made for Republicans. In another future study, we are working to test version of parallel ads that are manipulated to appeal to both parties. For example, in the Republican version of the social identity ad, mostly white families are featured, while the Democratic social identity ad features a more diverse group of putative partisans. The script remains the same for both. These ads will not be made by professional ad-makers (who we have found are not very interested in learning what attracts people to the opposition party). Despite their comparatively limited production and creative value, these ads will allow us to more sharply distinguish between the four theories of

partisan identification that originally motivated this project.

## References

- Achen, Christopher H. 1992. "Social Psychology, Demographic Variables, and Linear Regression: Breaking the Iron Triangle in Voting Research." *Political Behavior* 14(3):195–211.
- Achen, Christopher and Larry Bartels. 2016. *Democracy for Realists*. Princeton, New Jersey: Princeton University Press.
- Brady, Henry E. and Paul M. Sniderman. 1985. "Attitude Attribution: A Group Basis for Political Reasoning." *American Political Science Review* 79(4):1061–1078.
- Brody, Richard A. and Lawrence S. Rothenberg. 1988. "The Instability of Partisanship: An Analysis of the 1980 Presidential Election." *British Journal of Political Science* 18(4):445–465.
- Broockman, David E. and Daniel M. Butler. 2017. "The Causal Effects of Elite Position-Taking on Voter Attitudes: Field Experiments with Elite Communication." *American Journal of Political Science* 61(1):208–221.
- Campbell, Angus, Philip E. Converse, Warren E. Miller and Donald E. Stokes. 1960. *The American Voter*. New York: John Wiley & Sons.
- Campbell, Angus, Philip E. Converse, Warren E. Miller and Donald E. Stokes. 1966. *Elections and the Political order*. New York: John Wiley & Sons.
- Chong, Dennis and James N. Druckman. 2010. "Dynamic Public Opinion: Communication Effects over Time." *American Political Science Review* 104(04):663–680.
- Clarke, Harold D., Marianne C. Stewart and Paul Whiteley. 1997. "Tory Trends: Party Identification and the Dynamics of Conservative Support Since 1992." *British Journal of Political Science* 27(2):299–319.
- Clinton, Joshua D. and John S. Lapinski. 2004. "'Targeted' Advertising and Voter Turnout: An Experimental Study of the 2000 Presidential Election." *Journal of Politics* 66(1):69–96.
- Converse, Philip E. 1964. The Nature of Belief Systems in Mass Publics. In *Ideology and Discontent*, ed. David E. Apter. Vol. 18 Glencoe, Illinois: Free Press.
- Coppock, Alexander. forthcoming. "Generalizing from Survey Experiments Conducted on Mechanical Turk: A Replication Approach." *Political Science Research and Methods* .
- Coppock, Alexander and Oliver A. McClellan. 2018. "Validating the Demographic, Political, Psychological, and Experimental Results Obtained from a New Source of Online Survey Respondents." *Unpublished manuscript* .
- Coppock, Alexander, Thomas J. Leeper and Kevin J. Mullinix. 2018. "The Generalizability of Heterogeneous Treatment Effect Estimates Across Samples." *Unpublished manuscript*.
- Cowden, Jonathan A. and Rose M. McDermott. 2000. "Short-term Forces and Partisanship." *Political Behavior* 22(3):197–222.

- Downs, Anthony. 1957. *An Economic Theory of Democracy*. New York: Harper.
- Fiorina, Morris. 1981. *Retrospective Voting in American National Elections*. New Haven: Yale University Press.
- Flores, Alejandro and Alexander Coppock. 2018. “Do Bilinguals Respond More Favorably to Candidate Advertisements in English or in Spanish?” *Political Communication* .
- Fowler, Anthony. 2018. “Partisan Intoxication or Policy Voting?” Unpublished manuscript.
- Franklin, Charles H. and John E. Jackson. 1983. “The Dynamics of Party Identification.” *American Political Science Review* 77(4):957–973.
- Gerber, Alan S. and Donald P. Green. 2012. *Field Experiments: Design, Analysis, and Interpretation*. New York: W.W. Norton.
- Gerber, Alan S., Gregory A. Huber and Ebonya Washington. 2010. “Party Affiliation, Partisanship, and Political Beliefs: A Field Experiment.” *American Political Science Review* 104(4):720–744.
- Graham, Matthew H. 2018. “Self-Awareness of Political Knowledge.” *Political Behavior* . Forthcoming.
- Green, Donald, Bradley Palmquist and Eric Schickler. 1998. “Macropartisanship: A Replication and Critique.” *American Political Science Review* 92(4):883–899.
- Green, Donald P. 1991. “The Effects of Measurement Error on Two-Stage, Least-Squares Estimates.” *Political Analysis* 2(1):57–74.
- Green, Donald P. and Alan S. Gerber. 2002. “The Downstream Benefits of Experimentation.” *Political Analysis* 10(4):394–402.
- Green, Donald P. and Bradley Palmquist. 1990. “Of Artifacts and Partisan Instability.” *American Journal of Political Science* 34(3):872–902.
- Green, Donald P., Bradley Palmquist and Eric Schickler. 2002. *Partisan Hearts and Minds*. New Haven: Yale University Press.
- Green, Donald P and Lynn Vavreck. 2008. “Analysis of Cluster-Randomized Experiments: A Comparison of Alternative Estimation Approaches.” *Political Analysis* 16(2):138–152.
- Green, Donald Philip and Eric Schickler. 1993. “Multiple-measure Assessment of Party Identification.” *Public Opinion Quarterly* 57(4):503–535.
- Greene, Steven. 2002. “The Social-Psychological Measurement of Partisanship.” *Political Behavior* pp. 171–196.
- Gross, Kimberly, Ethan Porter and Thomas J. Wood. 2018. “Identifying Media Effects Through Low-Cost, Multi-Wave Field Experiments.” *Political Communication* . Forthcoming.
- Harris, Louis. 1954. *Is There a Republican Majority?: Political Trends, 1952-1956*. New York: Harper.

- Holian, David B. 2004. "He's Stealing My Issues! Clinton's Crime Rhetoric and the Dynamics of Issue Ownership." *Political Behavior* 26(2):95–124.
- Huddy, Leonie, Lilliana Mason and Lene Aarøe. 2015. "Expressive partisanship: Campaign involvement, political emotion, and partisan identity." *American Political Science Review* 109(1):1–17.
- Huff, Connor and Dustin Tingley. 2015. "“Who are these people?” Evaluating the demographic characteristics and political preferences of MTurk survey respondents." *Research & Politics* 2(3):2053168015604648.
- Jackson, John E. 1975. "Issues, Party Choices, and Presidential Votes." *American Journal of Political Science* 19(2):161–185.
- Kinder, Donald R and Nathan P Kalmoe. 2017. *Neither Liberal nor Conservative: Ideological Innocence in the American Public*. University of Chicago Press.
- Lenz, Gabriel S. 2013. *Follow the Leader?: How Voters Respond to Politicians' Policies and Performance*. University of Chicago Press.
- Lin, Winston. 2013. "Agnostic Notes on Regression Adjustments to Experimental Data: Reexamining Freedman's Critique." *The Annals of Applied Statistics* 7(1):295–318.
- MacKuen, Michael B., Robert S. Erikson and James A. Stimson. 1989. "Macropartisanship." *American Political Science Review* 83(4):1125–1142.
- Mullinix, Kevin J., Thomas J. Leeper, James N. Druckman and Jeremy Freese. 2015. "The Generalizability of Survey Experiments." *Journal of Experimental Political Science* 2:109–138.
- Petrocik, John R. 1996. "Issue Ownership in Presidential Elections, with a 1980 Case Study." *American Journal of Political Science* 40(3):825–850.
- Rahn, Wendy M. 1993. "The Role of Partisan Stereotypes in Information Processing about Political Candidates." *American Journal of Political Science* 37(2):472–496.
- Schickler, Eric and Donald P. Green. 1997. "The Stability of Party Identification in Western Democracies: Results from Eight Panel Surveys." *Comparative Political Studies* 30(4):450–483.
- Tomz, Michael and Paul M. Sniderman. 2005. "Brand Names and the Organization of Mass Belief Systems." *Unpublished Manuscript* .
- Valentino, Nicholas A., Michael W. Traugott and Vincent L. Hutchings. 2002. "Group Cues and Ideological Constraint: A Replication of Political Advertising Effects Studies in the Lab and in the Field." *Political Communication* (19):29–48.