

Targeted Campaign Appeals and the Value of Ambiguity*

Eitan D. Hersh and Brian F. Schaffner[†]

September 24, 2012

Abstract

Political campaigns increasingly micro-target. Given detailed knowledge of voters' identities, campaigns try to persuade voters by pandering to these identities. Through multiple survey experiments, we examine the persuasiveness of group-directed pandering. We ask: Do group-members respond more favorably to appeals geared to them, or do they prefer broad-based appeals? Do voters not in a group penalize candidates who appeal to a group? Answers to these questions help us grapple with the evolving relationship between voters and candidates in a rapidly changing information environment. Our results suggest that voters rarely prefer targeted pandering to general messages and that "mistargeted" voters penalize candidates enough to erase the positive returns to targeting. Theoretically, targeting may allow candidates to quietly promise particularistic benefits to narrow audiences, thereby altering the nature of political representation, but voters seem to prefer being solicited based on broad principles and collective benefits.

*For helpful advice, the authors thank Stephen Ansolabehere, Anthony Fowler, Seth Hill, Dan Hopkins, Timothy Ryan, and the anonymous reviewers. We are also grateful for feedback received by members of the American Politics Research Group at UMass, participants in the Political Psychology and Behavior Workshop at Harvard, and a panel at the 2011 Annual Meeting of the American Political Science Association. We thank John Hird and the Department of Political Science at UMass and the Center for the Study of American Politics at Yale for supporting the CCES modules and Mechanical Turk experiment on which the paper is based. The data and syntax files necessary to reproduce the findings in this paper are available at <http://dvn.iq.harvard.edu/dvn/dv/schaffner>.

[†]Eitan Hersh is Assistant Professor of Political Science, Yale University. Brian Schaffner is Associate Professor of Political Science, University of Massachusetts-Amherst.

The “big data” revolution is underway, and political campaigns are investing resources in individual-level, targeted appeals. Such appeals raise a number of questions for normative theorists and social scientists. Targeting elicits concerns about privacy and the attenuation of information flows (Sunstein 2007), but there are also questions about how a “surgical” approach to mobilization affects a democracy. Does targeting technology allow politicians to make secret promises to voters? Does targeting enable politicians to restrict the set of constituents they pay attention to? Closer to political science research on public opinion and political behavior, there are questions about how voters’ judgments are affected when candidates attempt to prime and persuade, engage in “cheap talk” pandering, and send cues about their personal beliefs and policy positions. The possibility that campaigns can narrowcast messages to different groups based on voters’ individual characteristics makes these questions especially salient in the Information Age.

Neither researchers nor campaigns yet know very much about how well targeting works at persuading voters (Grossmann 2009). While a growing body of research examines the extent to which candidates target their messages to particular subgroups (e.g. Hillygus and Shields 2008; Nteta and Schaffner n.d.) as well as the effects of campaigns in general (Brady, Johnston, and Sides 2006; Iyengar and Simon 2000), there has been less attention paid to whether targeted pandering actually works. The effects of targeted messages are difficult to understand from observational studies on account of selection bias and endogeneity (Arece-neaux 2010). Moreover, when campaigns send hidden messages to voters, by definition they do not want these messages to be publicized.

Because of the inherent difficulties in studying targeting, we turn to experiments to test the persuasive effects of targeted messages. Over a three-year period, we tested more than a dozen experimental conditions in which we describe to survey respondents a fictional candidate who is running for Congress. We randomize whether the candidate claims to be supporting the interests of a sub-group, like born-again Christians or Latinos, or whether the

candidate offers a generic appeal. We measure how voters evaluate these messages depending on whether they had previously identified as a member of the targeted group.

After testing a number of variations on our experiment, we find that voters who are members of particular groups do not generally reward candidates for pandering to their groups, but voters who receive a message targeted to a group of which they are not a member consistently penalize the candidate for “mistargeting.” Mistargeting is important because even in contemporary elections, any list of voters that might be used for targeting contains enough individuals who are not actually members of the intended group (i.e. false positives) that targeting may generally result in net losses to campaigns.

Our study speaks to a wide range of political science research, from classic “cheap talk” models (e.g. Farrell and Gibbons 1989, Banks 1990) to empirical analyses on how candidate traits and positions affect voters (e.g. Rahn 1993, McDermott 1997, Hayes 2005); from studies of campaign strategy (e.g. Hillygus and Shields 2008) to new research by Hopkins (2011a) finding that whether a voter hears a Spanish or an English message about immigration policy will affect his or her judgment on that policy. Most of all, our evidence builds on research about the benefits incurred to candidates of all kinds from being ambiguous rather than specific (e.g. Tomz and Van Houweling 2009; Norton, Frost and Ariely 2007; Jensen 2009; Sides 2006; Shepsle 1972). As Norton, Frost and Ariely find in a study of potential romantic partners, “more information about any one person leads, on average, to less liking for that person.” We find the same principle at play with respect to voters evaluating political candidates. This research helps us grapple with the evolving relationship between voters and candidates in a rapidly changing information environment. If we can learn more about the effectiveness of targeted group appeals, we will be better equipped to answer the normative questions about democracy in an era of abundant and personalized information.

At the outset, we emphasize three points of clarification. First, this is a study that explores the persuasive effects of targeted messages, not the mobilizing effects. Much of political

campaign targeting may be designed to demonstrate to particular groups the importance of voting in an upcoming election by invoking the group’s connection with a candidate or party. In this brief analysis, we confine ourselves to the persuasive effects and leave the study of mobilizing effects to others (e.g. Green and Gerber 2008). Second, as we try to simulate aspects of real-world targeting, we conflate political science concepts such as targeting, priming, pandering, and signaling. This conflation is intentional, as the targeting that campaigns engage in involves all of these effects simultaneously. Thirdly, as with all studies based on survey experiments, our results come from an artificial environment. We will address concerns about external validity below, but here we emphasize that our aim is not to simulate the environment of a real-world campaign in a survey experiment, but rather to capture how voters may react to certain kinds of messages similar to ones that campaigns give them.

The Emergence of Targeted Appeals

During the past two decades, candidates have become increasingly adept at voter *targeting* – identifying voter characteristics and sending tailored messages to voters based on those characteristics. Hillygus and Shields (2008) describe the evolution of candidate strategies from an emphasis on broad appeals delivered via mass media to fine-tuned appeals transmitted to specific subgroups over targeted media. That candidates craft different messages for targeted media than for mass media (as evidenced by Hillygus and Shields) suggests that candidates believe these messages will have differing effects. In a campaign message that will reach all voters, candidates may prefer to make broad appeals that will alienate the smallest share of the population. When messages can be hidden from all but the intended recipients, candidates might craft messages that are more pointed and perhaps more effective.

However, no research to date indicates whether this strategy actually works. Targeting voters is easier said than done. For one thing, voters may not be responsive to directed appeals. They might prefer more inclusive messages. Furthermore, when campaigns compose

lists of voters they think will be positively responsive to appeals, oftentimes such target lists include a large number of false positives - voters who appear to campaigns to be responsive to an appeal but actually do not possess the characteristics that the politicians are hoping to target. Thus, these two phenomena might cut against a politician's ability to target.

A well-developed body of field-experimental research has shown that targeted appeals can increase turnout, though usually only by a modest amount (e.g. Green and Gerber 2008). However, only a handful of studies have examined whether targeted appeals are successful in *persuading* voters to support a particular candidate (e.g. Arceneaux 2007). Most scholarship on the influence of campaign messages has focused primarily on the effect of television advertising rather than targeted appeals (Iyengar and Simon 2000). Television advertisements have been shown to increase citizens' knowledge about the candidates, but the notion that they can persuade voters to support a particular candidate is a source of some debate (Brader 2005; Johnston, Hall, and Jamieson 2004; Huber and Arceneaux 2007; Gerber, et al. 2011). Unlike broadcast TV ads, targeted appeals are often difficult for political scientists to observe, and causal inference with observational data is anyway challenging. For example, relying on citizens to report their contact with campaigns is problematic since individuals tend to have unreliable memories of such contact (Bradburn, Rips, and Shevell 1987; Ansolabehere and Simon 1999; Vavreck 2007).

However, whether or not targeting "works" is important because if candidates can successfully target messages to particular groups, this could reshape the relationship between voters and politicians. On the positive side, when politicians know more about voters, they can pay attention to citizen preferences with a level of specificity that was not before possible. Targeting, in this sense, means politicians know more about what each voter cares about. On the negative side, when politicians know more they may find it desirable to *pander* to those groups by making narrowly tailored promises.¹ Of course, politicians have always been able to pander to their donors and personal associates, but if voter targeting "works," then

politicians can cut up the electorate into narrow sub-divisions and provide enough particularized benefits to each group to put together a winning coalition. Perhaps more troubling is that such narrow coalition-building would not be transparent to other voters. A coalition of voters built quietly by promising narrow benefits to narrow audiences implies a very different kind of democracy than a coalition of voters built on broad principles and collective benefits.² Thus, the effectiveness of targeting is not merely a logistical question of interest to campaign operatives; it can influence the nature of representation in a democracy.

Theory: Why Might a Pander Succeed or Fail?

Our experiments help to answer three questions about the effectiveness of targeting. First, *are targeted panders more effective than general appeals?* Narrow messages transmitted to sub-groups may not be any more effective than broad-based appeals. For example, Latinos tend to demonstrate less of a sense of group identity than African-Americans (Rodriguez 2000, Jones-Correa and Leal 1996). In messaging to Latinos, a targeted appeal may be no more successful than a general appeal to “workers” or the “middle class.”

Second, *are appeals to particular groups more or less effective depending on the candidate’s political party?* Iyengar and Valentino (2000) show that “campaign advertising is most effective when the sponsoring candidates pitch their messages to the traditional strength of their parties (p.127).” For example, a Democrat making an appeal to unions may be more credible than a Republican making a similar appeal. On the other hand, a candidate might have more success with directed appeals when the targeted group is outside the party’s traditional coalition, insomuch as the appeal may provide new information to voters.

Third, *will individuals who are not members of the targeted group penalize a candidate for appealing to that group?* Since candidates tend to send different types of messages in targeted media than they do in mass media, they may believe that narrower (group-based) appeals would cost them support among non-group members. These mis-targeted appeals

may lead voters to think the candidate will prioritize an out-group over people like them.

We offer three broad theoretical frameworks for the conditions under which targeting may affect voters: one informational, one strategic, and one psychological. The *informational* model is the most straightforward. If targeted pandering influences voters' support for candidates, it likely does so by providing citizens with new information about the candidates' priorities. Party labels hold information for citizens seeking to make inferences about a candidate's policies (Rahn 1993, Schaffner and Streb 2002), but group-based panders, when they work, may provide additional clarity to these assumptions. For example, born-again Christians may view their agenda as being supported by some, but not all, Republican politicians; thus, a Republican party label may not be a sufficient condition for support. If a Republican campaign panders to born-again, it may send a signal that the politician will be particularly responsive to this group on social issues.

A *strategic* model emphasizes the limitations for targeted messages to transmit credible information. A problem with targeted pandering is that such messages might easily be dismissed as cheap talk. Individuals tend to take messages more seriously if the sender must pay a cost for sending those messages. Broadcast promises to a specific group are costly because the candidate is publicly on record and because non-group members may react negatively.³ Targeting, on the other hand, allows candidates to reduce the potential costs of group-based pandering by hiding the message from citizens who are not part of the group. But since a targeted pander is less costly, group members may view it as less credible.

In the analysis below, we find that individuals who see a targeted message that is not intended for them react negatively to that message. Under a strategic framework, this may happen because citizens hearing a pander for a group to which they do not belong assume the message was broadcast widely. Presumably the voter would not know that he or she had been mis-targeted but rather understands the candidate to be stating a genuine priority. Because this commitment is not only credible but also seeks to advantage a group the voter

is not a member of, mis-targeted voters penalize candidates. And because the message's intended audience may question the message's credibility, candidates may often lose out by trying to target voters with sub-group specific messages.

A final theoretical framework is *psychological*. Like the strategic framework, the psychological model stresses the limits of the informational value in targeted messages. Innovative experiments by Norton, Frost, and Ariely (2007) find that when singles using online dating websites learn more specific information about potential romantic partners, they tend to like the person less. Tomz and Van Houweling (2009) find similar results in the political realm: when fictional politicians are ambiguous about their policy preferences, survey respondents reward those candidates with expressions of support. Sides (2006) finds campaign television advertisements often discuss issues while remaining ambiguous about the candidates' positions. There are several plausible reasons why individuals may favor ambiguity (i.e. they are risk-takers or are themselves ambiguous about their own views), but central to these studies is the idea that people respond to ambiguity with optimism. When voters know less about a candidate, they project their own favored values upon that candidate.

If ambiguity engenders optimism, then targeted pandering may be a suboptimal strategy for reasons unrelated to cheap talk. Suppose that when voters receive a generic message from a candidate, they project their own values upon the candidate and thus view the candidate favorably. Then, when a person of a particular group is targeted based on the group characteristic, the voter learns nothing new, not because he thinks the candidate is pandering but because he already was optimistic about the candidate's view. Conversely, if a candidate mis-targets, the voter now updates his view about the true nature of the candidate. Thus, under this framework too, the candidate may find no gain in support among targeted voters and a loss from voters who are mis-targeted.

The clear-cut expectation from all three theoretical frameworks is that voters who are mistargeted will penalize the candidate, because they a.) learn the candidate does not

represent them, b.) believe the candidate's out-group message is sincere, and/or c.) realize the candidate is not as similar to them as they may have imagined. As for in-group targeting, the frameworks offer differing predictions. Under the strategic and psychological models, the candidate does not gain from targeting in-group members. On the other hand, the information model predicts that the voter may reward the candidate if the voter learns new details about the candidate's priorities.

For both in-group and out-groups, multiple theories can explain the same outcomes. How can we hope to sort out the theories? We offer some direct tests of mechanisms. In particular, we ask voters to evaluate the ideological disposition of the candidate. By comparing their evaluation to their own ideological position, we are able to test whether voters project their ideology onto the ambiguous candidate and whether voters update their ideological placement of the candidate following a pandering message. Our tests of mechanisms lend most support to the strategic framework. However, while our tests favor the strategic framework and while the three mechanisms together provide some intuition for how voters may react to targeted messages, the task of decisively sorting out the mechanisms is left to future work. The main contribution of this paper is to gain new insights about whether targeting works better or worse than wholesale messaging, whether the effect of targeting depends on the group targeted and the party of the candidate, and whether mis-targeted voters penalize a candidate. Though we cannot entirely parse the theoretical foundations, our results suggest that a simple informational model, whereby voters learn from appeals and then update their beliefs, is not sufficient to explain how voters react to targeted pandering messages. Consistent with Tomz and Van Houweling (2007) and others, candidates may benefit from remaining ambiguous rather than informing voters about their actual priorities.

Experimental Design

Our experiments were designed as follows. In the 2010 and 2011 Cooperative Congressional Election Study (CCES) surveys,⁴ respondents were presented with a fictional candidate, Williams. In the 2011 study, 2,596 respondents were shown a mock campaign mailer and were asked to evaluate the candidate. An example of one of the treatment conditions is displayed in Figure 1. In the 2010 study, 2,492 respondents were simply told about the candidate in text format. The text is available in the online appendix. In both versions, respondents could place themselves on a scale ranging from “very unlikely to vote for him” on one extreme (coded 0) to “very likely to vote for him” on the other end of the scale (coded 100). The average rating across all conditions was a 35 in both studies, and ratings ranged from 0 to 100 with a standard deviation of 31-33.⁵ The online appendix provides summary statistics on the evaluation measure from each of our experiments.

The experiments were designed to randomize two pieces of information that respondents were given about fictional candidate Williams. First, half of the respondents were told that Williams was a Republican while the other half were told that he was a Democrat. Second, respondents were told a particular group on whose behalf Williams vowed to work. In the 2010 study, respondents were randomly selected to one of four different groups—“the middle class,” “Latinos,” “Unions,” or “religious conservatives.” In the 2011 study, the groups were “constituents,” “Latinos,” “Gun Owners,” and “Born-Again Christians.”

Any number of groups might have been chosen for the experimental conditions, thus our selection of groups merits attention. In the 2010 study, our baseline, or control group, was “the middle class.” The “middle class” is broadly defined in American politics - most Americans identify themselves as part of this group.⁶ Thus, we expect that most respondents in this condition would view themselves as a target of this appeal or, at the very least, would view the appeal as addressing a general audience. In the 2011 study, out of concern that “middle class” could be perceived as too targeted for a baseline condition, we changed the

baseline group simply to “constituents.”

For the treatment conditions, we chose groups carefully to maximize the believability of the appeal while also generating variance on the extent to which the appeal might appear to come from an atypical source. For example, a Republican appeal to religious conservatives or born-again Christians might be expected while the same appeal from a Democratic candidate might be considered more unusual, but still believable. Likewise, while Democrats generally perform better among union members than Republicans, Republicans have made appeals for union votes and often make inroads among these voters.

We designed our treatments to mimic the way real campaigns might pander to particular groups. After talking to campaign consultants, we discovered that microtargeted messages are often quite overt. Instead of crafting subtle, “dog whistle”-style messages, campaigns often prioritize boldness over subtlety in direct mail ads because they want the ads to make a lasting impression. Consider Figure 2, which displays a direct mail advertisement sent to Latinos by Jon Corzine’s 2005 gubernatorial campaign in New Jersey. The Corzine mailer is indicative of direct mail that is sent to Latinos. The use of Spanish language aside, the message in Corzine’s mailer is quite similar to ours. Or consider data from a prominent Democratic media consultant who recently donated a database of direct mail pieces to American University (Nteta and Schaffner, n.d.). The database offers additional examples of straight-forward pandering messages like the ones we simulate: “For 25 Years, A Fighter for Latino Workers and Their Families,” “Para Nuestras Familias,” “The Latino Voice in Government.” For ads targeted to union households, the headlining messages read, “Standing Up for Maine’s Working Families,” “A Voice for Working America,” and “Fighting for Illinois Fire Fighters.” Democratic ads targeted to Republican voters include messages like “Faith in Action - Tim Kaine: Faithful, Moral Leadership for Virginia” and “Protecting Our Sporting Heritage.” Messages geared toward senior citizens have language such as “Leadership for Our Seniors,” and “Who will protect our senior citizens?”

From Corzine’s campaign, we also have an example of a real-world generic appeal. Corzine sent a mail piece to a broad swath of targeted voters that simply showed a picture of Corzine and the following words: “A more affordable New Jersey. That’s Jon Corzine’s plan.” On the inside of that postcard, the campaign offered basic biographical information and general bullet points about Corzine’s policy goals. Of course, our control conditions are designed to provide a baseline for understanding the targeted appeals. As such, they are quite pared down. Nevertheless, very generic messages (like those we employ in our control conditions) are used widely in non-targeted appeals. All together, our experiment follows very closely with this class of targeted pandering.

Prior to the experiment portion of the survey, respondents were asked standard demographic questions that identified them as part or not part of the groups we targeted. Respondents who identified as Hispanic (7% of the 2010 sample; 11% of the 2011 sample) are taken to be potential targets of a Latino appeal. Respondents identifying as a union member or as living in the household of a union member (13% of the 2010 sample) are taken to be potential targets for a union appeal. Respondents identifying as born-again Christians (32% of each sample) are taken to be potential targets for an appeal to “religious conservatives” or “born-again Christians.”^{7 8} In the 2011 study, we also asked respondents whether they felt a sense of “linked fate” with each of a series of groups.⁹ Both the demographic questions and linked fate questions were asked earlier in the survey and were separated from the randomized candidate evaluation treatments.

Before turning to the results, we draw attention to two features of our experiments. First, we made several tweaks between the two experiments for the sake of robustness, including differences in question wording, treatment design (i.e. text vs. graphical), baseline groups and treatment groups. We present results from all these conditions, and the key lessons drawn from the results are the same. Below, we discuss results from another follow-up study we conducted using Amazon.com’s Mechanical Turk that adds additional robustness checks

to the design: we embed the treatment in a series of non-targeted campaign messages and we distract the respondent between the treatment and the evaluation.

Second, consider the extent to which our design realistically portrays the interaction between voters and candidates. In many down-ballot races, such as Congressional races or in-state races, it is utterly plausible that citizens go to the voting booth knowing little else about a candidate than his/her name, party affiliation, and perhaps a snippet of information about the candidate's priorities. The information that our experimental conditions provide to respondents about fictional candidate Williams is admittedly sparse, but for better or worse, it is not a far-fetched portrayal of a low-information election environment.

A separate question of external validity is whether the mock messages realistically simulate actual campaign communications. We have shown that our experiments replicate one form of group-based targeting, but they do not replicate other forms of targeting. Our experiments also do not replicate all the various direct and indirect exchanges that occur between candidates and voters. The survey environment is artificial. However, our goal is not to measure the average effect of campaign messages but rather to understand how voters react to certain types of messages and how the reactions vary based on whether the voter is or is not a member of a targeted group. We expect voters' reactions to be stronger in the survey than in the real-world since we have captured their attention and we ask for their immediate reaction to our messages (Chong and Druckman 2010). If voters do not react positively even in this environment, we suspect they will not react positively in the real world, where the treatments are more diluted by other aspects of the campaign. But, of course, more research will be necessary to understand the full real-world implications.

The Advantages and Penalties for Targeted Pandering

The main results of our experiments are in Figure 3 and Figure 4.¹⁰ In Figure 3, we restrict the analysis to voters who earlier in the survey had identified as a member of one of

our targeted groups. The figure shows the average difference between evaluations of Williams when Williams is targeting the group from when Williams vows to work just on behalf of the “middle class” (in 2010) or “constituents” (in 2011). For example, if group members receiving the group-specific treatment rate the candidate at 60 (out of 100) and the group members receiving the baseline condition rate the candidate at 40, this would appear in the figure as a value of 20. Positive values indicate the group appeal was preferred; negative values indicate that the baseline appeal was preferred.

When Williams is presented as a Democrat, in only one of our five experiments do in-group members evaluate Williams more favorably than if he had just offered a generic appeal. Born-again Christians liked Democratic candidate Williams better when the candidate vowed to work for born-again Christians than when he vowed to work on behalf of constituents. But none of the other groups favored the targeted appeal. When Williams is presented as a Republican, he gets some traction with the targeted appeal in two of the five treatments (in addition to the large but statistically insignificant effect for Latinos in 2011). Worth noting is the differences in variance across treatments in Figure 3: The confidence intervals are much wider in the Latino treatments than the Christian treatments not only because of the smaller number of Latinos in our sample, but also because Latinos had much more variance in their reactions to the directed appeal than did born-again Christians. Thus, even though the point estimates for Latinos targeted by Republicans are relatively large, the distribution of responses is quite disperse.

Consistent with our expectations, the in-group treatments are modest. However, it appears that our fictitious candidate is more successful when the appeal is directed to a group not typically associated with the candidate’s party base. The one instance in which Democratic Williams gains support is when he targets born-again Christians. Similarly, Republican Williams seems, on average, to do better with Latinos. We will test mechanisms below, but at first glance, the results lend support to the idea that voters update their views more when

the message provides information beyond what is contained in the party label.

Compare the weak and inconsistent results in Figure 3 with the uniformly strong results in Figure 4. Here, we study respondents who are not part of the targeted groups. These results simulate a situation in which a voter is “mistargeted” with an appeal. Across all treatment conditions, respondents who were randomly assigned to an out-group treatment preferred Williams much less than similar respondents who were assigned to a baseline condition.

There is evidently a serious risk for candidates crafting group-based appeals. If enough group-specific messages reach individuals who are not part of the intended group, then a significant backlash could occur. But how high would the error rate need to be to make targeted pandering ineffective? Clearly, for the appeals that garnered no positive reaction among in-group members, any backlash from mistargeted voters results in a net loss for the candidate. What about the groups that did increase support for the candidate given that they received a targeted appeal? In the 2010 Republican/religious treatment, the candidate received nearly a 20-point bump from born-again Christians, but when the same pander reached citizens who were not born-again, the candidate was penalized by about 30 points. Thus, to gain a net positive effect, a Republican candidate would have to ensure that for every three born-again Christians that receive the appeal, fewer than two non born-again Christians receive it. Similar calculations can be made for other conditions.

In the real world of campaigns, the number of false-positives on targeting lists is not trivial. Consider evidence from the 2008 CCES, which was matched to a national campaign database maintained by the Democratic firm Catalist (Ansolabehere and Hersh, 2012). Catalist uses name-matching tools and Census geography to predict voters’ races. Because Catalist’s data was matched to survey records, we can study the consistency between a campaign’s prediction of race and self-reported race in the survey. Among people whom Catalist predicts to be Hispanic, 25% indicated on the CCES that they are not Hispanic. Thus, for every three Hispanics successfully targeted, one non-Hispanic is mis-targeted. Given the

availability of Census data on the location of racial groups and given the homogeneity of Hispanic surnames, it is fair to assume that predicting which voters are Hispanic is much easier than predicting which are born-again Christians. Thus, based on the ten experimental conditions reported in Figures 3 and 4, it seems like it would rarely be worthwhile for a candidate to offer targeted messages instead of general messages.

The lack of consistent in-group effects for pandering may be surprising, but part of the reason for these limited effects may be that simply belonging to a group does not necessarily make one feel a strong identity with that group (Gay and Hochschild n.d.), and there is little reason to expect those with weak group identities to respond strongly to group-based appeals. For example, many U.S. Latinos do not feel a particularly strong identity with other Latinos (Rodriguez 2000, Jones-Correa and Leal 1996). Thus, a pledge by a candidate to work on behalf of Latinos may be no more persuasive than an appeal to work on behalf of all constituents. On the other hand, some non-group members may feel an affinity to a group that they are not a part of. For example, religious people who are not born-again Christians may feel a sense of linked fate with born-again Christians if they identify as part of a larger community of people with similar values. If they do feel this linked fate, they may not penalize a candidate who mistargets them.¹¹

Thus, for another approach to these experiments, in the 2011 study we asked respondents how much they think that what happens to various groups in this country affect what happens in their lives.¹² Respondents placed themselves on a 7-point scale for this standard “linked fate” question. We plot the relationship between a person’s connection to the group and his or her reaction to the targeted treatment versus the baseline condition (we include all respondents in this analysis, not just those that were part of a group). The candidate evaluation is on an 100-point scale. As previously mentioned, we have one treatment - an appeal to gun owners - that we do not report in Figures 3 and 4 because a miscommunication led to there being no question on the survey asking about gun ownership. However, we did

ask about linked fate with gun owners, so we can now report results from that treatment condition as well.

The results appear in Figure 5. As should be expected, a person’s reaction to a candidate who vows to work on behalf of “constituents” does not noticeably vary according to his or her linked fate with particular groups. However, when the candidate offers a targeted message, voters’ reactions vary greatly based on their connection to the group. Relative to the control conditions, respondents with low levels of linked fate with Latinos and born-again Christians were very unlikely to vote for the candidate who panders to those groups. Respondents with the highest levels of gun-owner linked fate and born-again linked fate prefer the targeted treatment to the baseline condition.

The results in Figure 5 provide some additional context for understanding just how difficult it is for campaigns to engage in targeted pandering. Pandering to born-again Christians is only effective when a respondent falls in one of the top two categories of linked fate with that group. While identifying as a born-again Christian is certainly associated with taking on higher levels of linked fate with born agains, such an identification is not determinative. In fact, only half of the self-identified born-again Christians in our sample selected one of the two highest values of linked fate with that group, meaning that the other half did not identify enough with born again Christians to make a targeted pander to the group an effective strategy relative to a general appeal. Of course, it is even harder for candidates to predict an attitudinal disposition liked “linked fate” than to predict demographics.

Mechanisms

To begin to sort out the mechanisms at the root of the effects found above, we asked subjects in our 2011 study to place the hypothetical candidate on an ideological scale after receiving the treatment. The scale ranged from “very liberal” (coded 1) to “very conservative” (coded 7). Before the experiment, respondents were also asked where they would

place themselves on the same scale. We use these ideology measures to search for evidence of behavior consistent with the theoretical frameworks discussed above.

The first question about mechanisms is whether respondents who receive the control condition project their own values on a candidate who offers only ambiguous messages. Does a candidate who provides only his party and a promise to represent “constituents” come across as conservative to conservatives and as liberal to liberals? Figure 6 plots respondents’ own self-reported ideology against their estimate of Williams’ ideology, given that Williams offers an ambiguous appeal (the control condition). We see no evidence of projection. For the Republican candidate, respondents across the ideological spectrum place Williams as a moderate conservative. For the Democratic candidate, respondents place Williams as moderate to liberal, and it is respondents who are most conservative who think of Williams as most liberal, not the other way around. From these results, we do not see evidence for the psychological model we described above. It seems that respondents are not optimistic that the ambiguous candidate is already close to them ideologically. Thus, we do not believe that the reason for a limited positive effect of the treatment on in-group members is that in-group members project their positions on the generic candidate in the control condition.

Figure 7 helps to address the informational and strategic theoretical frameworks. For each in-group and out-group, we estimate the average ideological position of the respondents, and the perceived ideological position of Williams among those same respondents. For each group, the figure shows the distance between the respondents’ ideology and their perception of the candidate’s ideology.

For in-groups, we see that born-again Christians evaluated the pandering Republican and the baseline Republican at the same ideological location, which was also the same location that they placed themselves. However, they updated their position much more for the Democratic panderer. Conversely, Latinos were more affected by the Republican appeal than the Democratic one. What jumps out of this figure, though, is that the in-group

updating is much less consistent than the out-group updating. In all four out-group plots in Figure 7, respondents shifted the ideological placement of the candidate away from their own position under the pandering condition.

The ideological shifts in Figure 7 suggest that in-group respondents are indeed learning some relevant policy information from the candidate who panders, especially when the targeted group is more distant from the candidate's party base. At the same time, that there are more instances of ideological movement for out-group respondents than in-group respondents lends support to the strategic model. The data are consistent with the idea that out-group respondents believe the pandering messages to be genuine but that in-group members discount some pandering as cheap talk.

In sum, Figure 6 demonstrates a lack of optimistic projection on the part of respondents, and Figure 7 demonstrates that in-group respondents may have discounted the messages as cheap talk as compared to out-group respondents. However, in some cases in-group respondents did indeed learn ideologically-relevant information from the panders, which correspond to modest positive evaluations of the pandering candidates in Figure 3. These tests of mechanisms only begin to sort out the causal chain. It is possible that the ideological dimension that is tested in these figures is not the primary basis by which information, projection, and reaction to cheap talk impact the respondents' judgement of the candidates.¹³ Furthermore, it is unclear whether cheap talk considerations attenuate the informational effect or whether there is limited learning among in-group members for other reasons. This investigation of mechanisms thus provides an initial set of tests, but future research should seek to extend our understanding of why targeted pandering so strongly affects the evaluations of out-group respondents but is less influential for in-group respondents.

Additional Robustness Check

In June 2012, we conducted a final experiment using Amazon.com’s Mechanical Turk to determine whether our findings would be robust in an experimental design that provided more information about the candidate and included a distraction task.¹⁴ A detailed description of the experiment appears in the online appendix, but several changes are worth noting. First, we give more biographical information about the fictional candidate in a paragraph that describes Williams’ career, family life, and political experience. Second, after reading the description of the candidate, respondents scroll through three pages, each showing a different advertisement. The first and third ads are the same in all conditions, and they offer generic messages. Only the second ad varies by experimental condition.

Third, instead of the treatment and control messages stating that Williams pledges to “represent the interests of [GROUP] in Congress,” we use language stating that Williams “understands our values” (control condition), “understands our Christian values” (treatment group 1), or “understands our Latino values” (treatment group 2). Fourth, after viewing the mailers, respondents answered seven unrelated questions, mostly dealing with local and national economic conditions. These questions were included to distract the respondents from the ads they saw and allow the salience of the ads to decay, even if only for a couple of minutes. After these questions, we ask the respondents to evaluate Williams.¹⁵

As with the CCES experiments, the out-groups in the Mechanical Turk experiment penalized the candidate for mis-targeting. There was, however, one exception to this: the Democratic candidate appealing to Latinos did not generate a statistically significant negative reaction among non-Latinos. The failure of a penalty to materialize in this case may be due to the under-representation of ideological conservatives and moderates in the MTurk sample. Similar to the 2011 CCES, the in-group targeting worked better when the Democratic candidate targeted born-again Christians and the Republican candidate targeted Latinos. The candidate made no gains by targeting groups closer to their party base. The full

set of results from this follow-up experiment are available in the online appendix.

Conclusion

The study of campaign effects is notoriously difficult, particularly with regard to analyzing the influence of targeted campaign messages and persuasive effects (Arceneaux 2010). By utilizing survey experiments, we are able to confront concerns about endogeneity and selection bias in order to gain a better understanding of the effects that targeted messages have on their intended audiences and on unintended audiences. In general, our findings speak to the limitations of targeted group-based pandering. In most cases, our hypothetical candidate did no better among group members by appealing directly to that group’s identity. Furthermore, these narrower appeals come with risks, since they lead to diminished support among non-group members who may easily be mis-targeted. If these results are applied to the real world of targeting, they suggest a puzzle: given the small payoffs and high risks, why do candidates engage in targeted pandering at all?

We point to several possible reasons why candidates may target narrow groups. First, not all groups were unresponsive to this pandering. In some of our trials, appeals to religious groups or Latinos seemed to have some promise for our candidate, as long as those messages were not transmitted to the general public. Second, we re-emphasize that our study focuses on persuasion rather than mobilization. Many scholars have argued that the main goal of targeted messages is to mobilize supporters rather than persuade uncommitted voters (e.g. Holbrook and McClurg 2005). Potentially, messages like the ones we tested could have increased turnout rather than support for a candidate. Third, our appeals were vague and not specifically policy-oriented, which might have served to mute their effect. And finally, while our experiments aim to explore an interaction between voters and candidates that is difficult to explore observationally, the survey environment is nevertheless artificial and so more work will be required to understand the full applicability of our findings to real-world

campaign efforts.

Hillygus and Shields (2008) argue that targeting poses a threat to democracy because many of the claims and promises made by the candidates can be hidden from the view of a majority of the public. Of course, candidates do not publicize their targeted messages, so it can be challenging to determine whether targeted appeals would really turn off the broader electorate. Our study serves to test this claim. Regardless of whether a targeted appeal generates more support for a candidate among the intended group, the candidate in our experiment loses support when his appeal is seen by those outside of the group. This pattern underscores the importance of transparency in campaigns—if candidates’ targeted messages were available for all citizens to see, it is likely that candidates would find it too costly to make promises to narrow groups. Yet, with more precise and accurate data on the electorate, campaigns will increasingly have the ability to target voters without penalty.

Even so, the capability of narrowcasting to voters while avoiding “false-positives” may not yield an election system in which politicians make narrow promises to narrow groups through voter targeting. Our experiments suggest that voters on average do not favor candidates who pander to their group identities; many group-identified voters seem to prefer an ambiguous message as much as a specific one. In response to the legitimate normative concern about the inevitable use of personal data in electioneering, the voters themselves thus might be the saving grace: the targeting of narrow messages is not merely determined by technical expertise, but by a responsive electorate. Our study suggests that while candidates may want to transmit pandering messages to selective audiences, the audiences may not be interested.

Figure 1: Screenshot From One Condition of the 2011 CCES Survey Experiment



Williams is a Republican running for the U.S. House of Representatives in 2012. Williams recently sent the mailer shown below to some voters in his district.



If you lived in Williams's district, how inclined would you be to vote for him.

Very unlikely to vote for him Very likely to vote for him

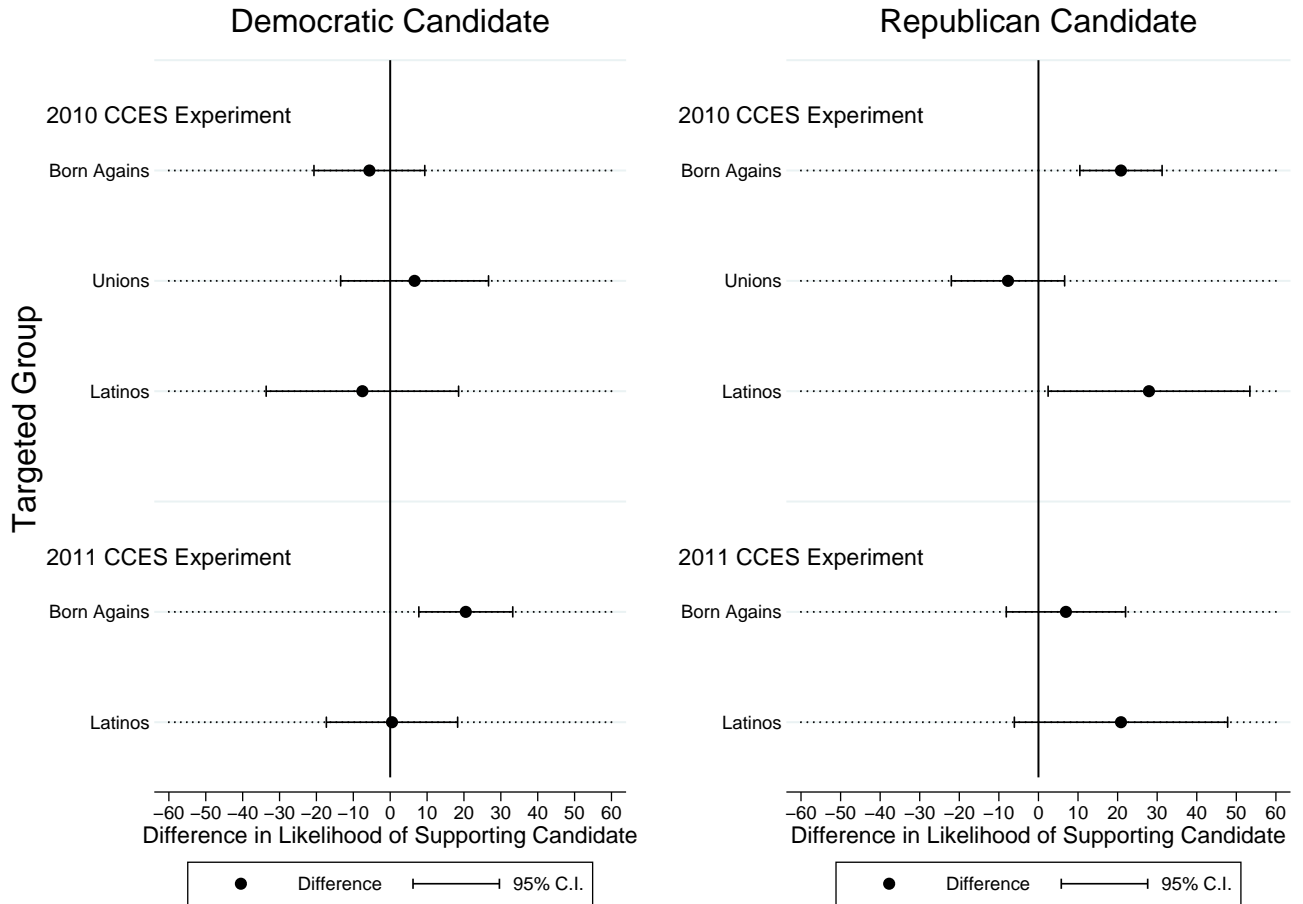
Not sure

Figure 2: Front of mailer sent to Latino targeted voters in 2005 Corzine Campaign




Note: Mailer is reproduced with permission.

Figure 3: Candidate Support Given Targeted Message Minus Candidate Support Given Baseline Message, Among Respondents Belonging to Targeted Groups.



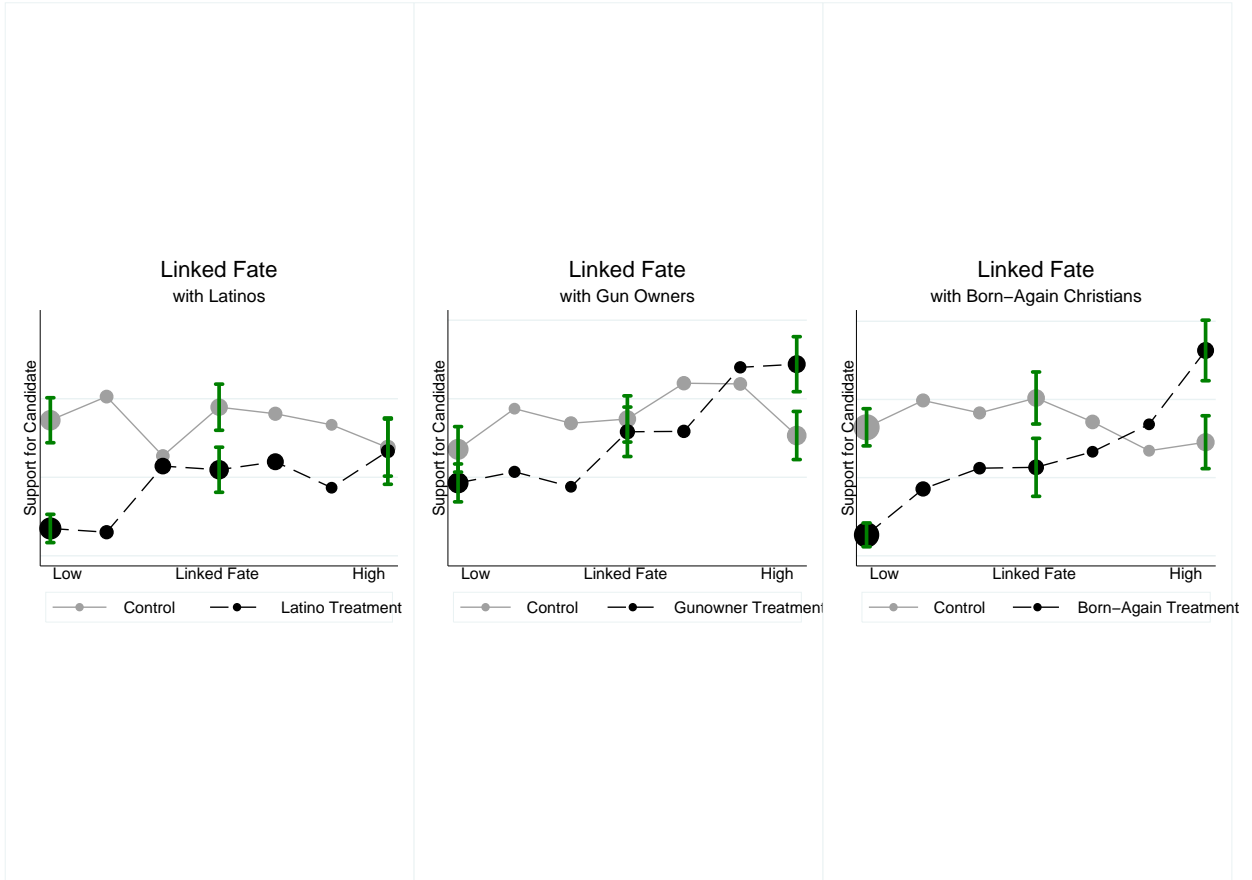
Note: 95% confidence intervals are displayed for the difference between respondents in treatment conditions and respondents in baseline conditions. Positive values indicate respondents prefer targeted message to baseline message.

Figure 4: Candidate Support Given Targeted Message Minus Candidate Support Given Baseline Message, Among Respondents NOT Belonging to Targeted Groups.




Note: 95% confidence intervals are displayed for the difference between respondents in treatment conditions and respondents in baseline conditions. Positive values indicate respondents prefer targeted message to baseline message.

Figure 5: Candidate Evaluations by the Appeal that was Randomly Assigned and by Respondents' Pre-Treatment Report about their Linked Fate with the Group



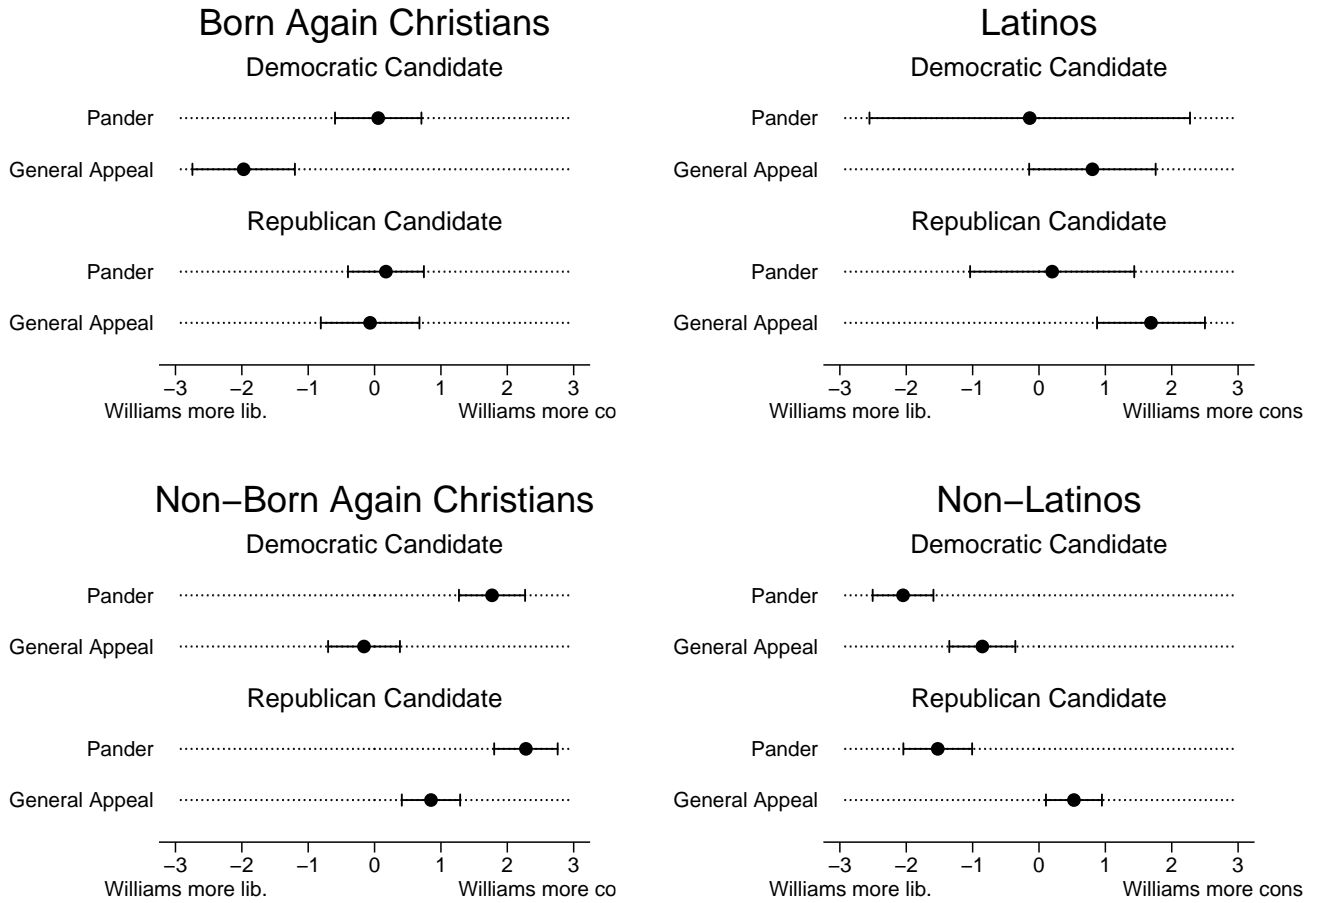
Note: 95% confidence intervals are displayed for selected groupings to ease interpretation. Dots in plot are sized in proportion to the number of respondents in each cell. Candidate evaluations are on an 100-point scale, but the mean evaluations displayed range from a value of 15 to a value of 63. All subplots are displayed on the same scale.

Figure 6: Are Voters Optimistic in the Face of Generic Appeals? Candidate Ideology versus Voter Ideology in Control Condition



Note: 95% confidence intervals are displayed.

Figure 7: Differences between Ideological Self-Placement of Respondents and Their Ideological Placement of Treatment and Control Candidates



Note: 95% confidence intervals are displayed. The ideological self-placement of each group is centered at zero in each subplot.

Notes

¹*Targeting* is the process of identifying groups of voters based on some characteristic and sending messages to just those groups. *Pandering* is the process of making promises to narrow groups; pandering may be accomplished via targeted messages, but it need not be so.

²In electoral systems with multi-member districts, parties may pander to their constituent groups without regard for the median voter (e.g. Cox 1990, Meyerson 1993). At a normative level, the kind of targeted pandering analyzed here is different because it lacks any transparency.

³For a formal model of candidate credibility and pandering, see Kartik and McAfee (2007).

⁴For details about the CCES sampling procedure, visit <http://projects.iq.harvard.edu/cces/home>. For a comparison of opt-in Internet sampling with telephone and mail studies, see Ansolabehere and Schaffner (2012). This study was conducted in the context of a larger political survey that may prime respondents to weight their political attitudes more than in their normal environs.

⁵We also conducted a trial set of experiments on Amazon.com’s Mechanical Turk program in early 2011 and found a similar set of results to those from the three experiments we present in this paper.

⁶A 2011 *Washington Post* poll asked respondents whether they considered themselves upper, upper-middle, middle, lower-middle, or lower class, and 86% considered themselves in a middle-class category.

⁷Due to a miscommunication, we have no in-group identifier compatible with our “gun owner” treatment.

⁸Results on union appeals are robust to an alternative specification in which only union members (rather than all respondents in union households) are considered to be in the targeted group. Results on appeals to religious conservatives are robust to an alternative specification in which questions about religiosity (e.g. frequency of prayer, church attendance) are utilized instead of identification as a born-again Christian.

⁹Specifically, the question asked “How much do you think what happens to the following groups in this country will affect what happens in your life?” Respondents could place themselves on a 7-point scale.

¹⁰All estimates are calculated applying the sampling weights.

¹¹See Tesler and Sears’ (2010) discussion of White racial liberals in the 2008 Presidential election.

¹²These questions were asked early in the questionnaire, before respondents encountered our experiment.

¹³At the suggestion of a reviewer, we performed an alternative test of the psychological model, whereby we asked respondents in a follow-up experiment which groups they thought the candidate cared about. For respondents in the control condition, we see no differences in these evaluations, suggesting an absence of projection.

¹⁴Amazon.com’s Mechanical Turk (MTurk) website has become an affordable and accessible venue for

recruiting subjects into social science surveys. MTurk has been found to be as good as or better than traditional convenience samples (Buhrmester, Kwang and Gosling (2011)) and Berinsky et al. (forthcoming) note that MTurk respondents “appear to respond to canonical experimental stimuli in a manner consistent with prior research” (p. 17).

¹⁵In the MTurk follow-up experiment, we did not provide respondents with the option of reporting they are “not sure” how to evaluate Williams. This decision was in part budgetary and in part to test yet another alteration on the design.

References

- Ansolabehere, S. and B.F. Schaffner. N.d. “Re-Examining the Validity of Different Survey Modes for Measuring Public Opinion in the US: Findings From a 2010 Multi-Mode Comparison.” .
- Ansolabehere, S.D., S. Iyengar and A. Simon. 1999. “Replicating experiments using aggregate and survey data.” *The American Political Science Review* 93(4):901–909.
- Ansolabehere, Stephen and Eitan Hersh. 2007. “Validation: What Big Data Reveal About Survey Misreporting and the Real Electorate.” *Political Analysis* Forthcoming.
- Arceneaux, K. 2010. “The Benefits of Experimental Methods for the Study of Campaign Effects.” *Political Communication* 27(2):199–215.
- Arceneaux, Kevin. 2007. “I’m Asking for Your Support.” *Quarterly Journal of Political Science* 2(1):43–65.
- Banks, Jeffrey S. 1989. “A Model of Electoral Competition with Incomplete Information.” *Journal of Economic Theory* 50(2):309–325.
- Berinsky, Adam J., Gregory A. Huber and Gabriel S. Lenz. Forthcoming. “Evaluating Online Labor Markets for Experimental Research.” *Political Analysis* .
- Bradburn, N.M., L.J. Rips and S.K. Shevell. 1987. “Answering autobiographical questions: The impact of memory and inference on surveys.” *Science* 236(4798):157.
- Brader, Ted. 2005. “Striking a Responsive Chord: How Political Ads Motivate and Persuade Voters by Appealing to Emotions.” *American Journal of Political Science* 49(2):388–405.
- Brady, H.E., R. Johnston and J. Sides. 2006. *The study of political campaigns*. University of Michigan Press pp. 1–26.

- Buhrmester, M., T. Kwang and S.D. Gosling. 2011. "Amazon's Mechanical Turk." *Perspectives on Psychological Science* 6(1):3.
- Chong, Dennis and James N. Druckman. 2010. "Dynamic Public Opinion: Communication Effects Over Time." *American Political Science Review* 104(4):663–680.
- Cox, Gary W. 1990. "Centripetal and Centrifugal Incentives in Electoral Systems." *American Journal of Political Science* 34(4):903–933.
- Farrell, Joseph and Robert Gibbons. 1989. "Cheap Talk with Two Audiences." *American Economic Review* 79(5):1214–1223.
- Gay, C. and J. Hochschild. N.d. "Is Racial Linked Fate Unique? Comparing Race, Ethnicity, Class, Gender, and Religion." . Forthcoming.
- Gerber, A.S., J.G. Gimpel, D.P. Green and D.R. Shaw. 2011. "How Large and Long-lasting Are the Persuasive Effects of Televised Campaign Ads? Results from a Randomized Field Experiment." *American Political Science Review* 1(1):1–16.
- Green, Donald P. and Alan S. Gerber. 2008. *Get Out the Vote: How to Increase Voter Turnout*. Brookings Institution Press.
- Grossmann, Matt. 2009. "Do the Strategists Know Something We Don't Know? Campaign Decisions in American Elections." *The Forum* 7(3).
- Hayes, Danny. 2005. "Candidate Qualities through a Partisan Lens: A Theory of Trait Ownership." *American Journal of Political Science* 49(4):908–923.
- Hillygus, D.S. and T.G. Shields. 2008. *The persuadable voter: Wedge issues in presidential campaigns*. Princeton University Press.

- Holbrook, Thomas and Scott McClurg. 2005. "The Mobilization of Core Supporters." *American Journal of Political Science* 49(4):689–703.
- Hopkins, Daniel J. 2011. "The Upside to Accents: Language, Skin Tone, and Attitudes Toward Immigration." Georgetown University Working Paper.
- Huber, Gregory and Kevin Arceneaux. 2007. "Identifying the Persuasive Effects of Presidential Advertising." *American Journal of Political Science* 51(4):957–977.
- Iyengar, S. and A.F. Simon. 2000. "New perspectives and evidence on political communication and campaign effects." *Annual Review of Psychology* 51(1):149–169.
- Iyengar, S. and N.A. Valentino. 2000. "Who says what? Source credibility as a mediator of campaign advertising." *Elements of reason: Cognition, choice, and the bounds of rationality* pp. 108–129.
- Jensen, Thomas. 2009. "Projection Effects and Strategic Ambiguity in Electoral Competition." *Public Choice* 141(1):213–232.
- Johnston, Richard, Michael G. Hagen and Kathleen Hall Jamieson. 2004. *The 2000 Presidential Election and the Foundations of Party Politics*. Cambridge University Press.
- Jones-Correa, Michael and David L. Leal. 1996. "Becoming Hispanic? Secondary Panethnic Identification Among Latin American-Origin Populations in the United States." *Hispanic Journal of Behavioral Sciences* 18(2):214–254.
- Kartik, N. and R.P. McAfee. 2007. "Signaling character in electoral competition." *The American economic review* 97(3):852–870.
- McDermott, Monika L. 1997. "Voting Cues in Low-Information Elections." *American Journal of Political Science* 41(1):270–283.

- Meyerson, Roger B. 1993. "Incentives to Cultivate Favored Minorities Under Alternative Electoral Systems." *American Political Science Review* 87(4):856–869.
- Norton, Michael I., Jeana H. Frost and Dan Ariely. 2007. "Less is More." *Journal of Personality and Social Psychology* 92(1):97–105.
- Nteta, Tatishe and Brian F. Schaffner. Forthcoming. "Substance and Symbolism: Campaign Appeals to Minority Voters in U.S. Senate Elections." *Political Communication* .
- Rahn, Wendy M. 1993. "The Role of Partisan Stereotypes in Information Processing about Political Candidates." *American Journal of Political Science* 37(2):472–496.
- Rodriguez, Clara. 2000. *Changing Race*. Temple University Press.
- Schaffner, Brian F. and Matthew J. Streb. 2002. "The Partisan Heuristic in Low-Information Elections." *Public Opinion Quarterly* 66(4):559–581.
- Shepsle, Kenneth A. 1972. "The Strategy of Ambiguity: Uncertainty and Electoral Competition." *American Political Science Review* 66(2):558–568.
- Sides, John. 2006. "The Origin of Campaign Agendas." *British Journal of Political Science* 36(3):407–436.
- Sunstein, C.R. 2007. *Republic. com 2.0*. Princeton University Press.
- Tesler, Michael and David O. Sears. 2010. *Obama's Race*. University of Chicago Press.
- Tomz, Michael and Robert P. Van Houweling. 2009. "The Electoral Implications of Candidate Ambiguity." *American Political Science Review* 103(1):83–98.
- Vavreck, Lynn. 2007. "The Exaggerated Effects of Advertising on Turnout: The Dangers of Self-Reports." *Quarterly Journal of Political Science* 2(4):325–343.