



Contents lists available at ScienceDirect

Journal of Economic Behavior & Organization

journal homepage: www.elsevier.com/locate/jebo

Negative campaigning, fundraising, and voter turnout: A field experiment

Jared Barton^{a,*}, Marco Castillo^b, Ragan Petrie^b^a Martin V. Smith School of Business & Economics, One University Drive, California State University–Channel Islands, Barton, Camarillo, CA, United States^b Interdisciplinary Center for Economic Science and Department of Economics, George Mason University, Fairfax, VA, United States

ARTICLE INFO

Article history:

Received 31 March 2015
 Received in revised form 3 October 2015
 Accepted 15 October 2015
 Available online 26 October 2015

JEL classification:

D72
 C93
 M37

Keywords:

Voter turnout
 Negative campaigning
 Comparative advertising
 Fundraising
 Field experiment

ABSTRACT

Why do candidates risk alienating voters by engaging in negative campaigning? One answer may lie in the large empirical literature indicating that negative messages are more effective than positive messages in getting individuals to do many things, including voting and purchasing goods. Few contributions to this literature, however, gather data from a field environment with messages whose tone has been validated. We conduct field experiments in two elections for local office which test the effect of confirmed negative and positive letters sent to candidates' partisans on two measurable activities: donating to the candidate and turning out to vote. We find that message tone increases partisan support in ways that may help explain the persistence of negative campaigning. Negative messages are no better than positive messages at earning the candidates donations, but negative messages yield significantly higher rates of voter turnout among the candidates' partisans relative to positive messages. Positive messages, however, are not neutral relative to no message.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction

Negative campaigning in American politics is as old as the country (Felkner, 1966), despite the fact that large majorities of the current U.S. voting public report the belief that negative campaigning is unethical (86 percent), produces less ethical leaders (76 percent), and hurts democracy (81 percent) (Green, in press). While a large empirical literature in political science (Lau et al., 2007) finds a small but positive effect of negative campaigning on voter turnout, and the literature that has examined comparative advertising – of which negative campaigning is one type – has found comparative messages more effective at changing consumers' buying intentions (Grewal et al., 1997), there are few randomized experiments measuring individual behavior on this topic in naturally occurring settings, as such tests impose costs on those running for office. Outside of some notable exceptions (Arceneaux and Nickerson, 2010; Gottfried et al., 2009; Niven, 2006), previous studies frequently measured intentions rather than behavior, used laboratory experiments with synthetic candidates or products, or examined indirect evidence and required strong identification assumptions to reach their conclusions¹. In this paper, we present the

* Corresponding author. Tel.: +1 3012197326.

E-mail addresses: jared.barton@csuci.edu (J. Barton), mcastil8@gmu.edu (M. Castillo), rpetrie1@gmu.edu (R. Petrie).

¹ Over half of the studies reviewed by Lau et al. (2007) use non-experimental observational data. Of the experimental studies included in their analysis, few measure intended or actual voter turnout, and 18 of the 49 laboratory experiments use fictitious candidates, advertisements, or both (e.g., Carraro et al., 2010; Fridkin and Kenney, 2011; Wu and Dahmen, 2010). Only 6 of the 77 studies examined by Grewal et al. (1997) examined actual buying behavior.

results of a field experiment on negative campaigning with candidates running for office in a real political campaign. We find that negative messages do affect voters' behavior and are sometimes more effective than positive messages. However, consistent with informational theories of campaigning, we find that communication doesn't always increase voters' support for candidates.

The field experiment was designed to test the effect of externally-validated negative and positive messages on actual campaign outcomes. Working with two campaigns for local office, we sent either a negative or a positive letter to the candidates' partisans and measured its effect on campaign donations and their voter turnout. Positive letters highlighted a candidate's qualifications, while negative letters alerted voters to the opponent's undesirable qualities (from like-minded partisans' point of view). We compare these two treatments to each other, and to a control group that receives no letter. We used letters because they allow us to manipulate messages in a non-intrusive way that is carefully controlled, as there is no human interaction. All letters contained a contribution card and return envelope, stated the date of the election, and asked for voters' "support," but did not explicitly mention giving to or voting for either candidate. As the messages are delivered by letter, and not through direct personal contact, we know that nothing about the messages is correlated with the method of delivery or the receptivity of the subject. The advantage of targeting partisans is that it allows us to cautiously interpret voter turnout as a proxy for voter support, as partisans who turn out to vote are generally unlikely to support the opposition (Abramowitz et al., 1981; Phillips et al., 2008). We verify turnout with official voter records.

We pair this field experiment with a pre-experimental survey among partisans outside the district. We asked subjects from a population similar to our target population – same party voters but in another city – to rate the campaigns' messages along several dimensions (randomizing the order of the two messages, and also which candidate's messages the subject examined), including their open-ended impressions of each message, the tone of the message, how informative each message was, and their affect toward the sender. The survey has several purposes. First, it ensures that our manipulations are indeed as positive and negative as we claim. Previous field experiments utilizing negative messages or differences in message tone do not confirm that their manipulations are interpreted as they intend among voters similar to those they target. This leads to uncertainty as to whether voters view these messages as the researchers (or their coders) do. Our messages are validated: positive messages are viewed as positive and our negative messages as negative by partisan voters. This difference is strongly statistically significant (Wilcoxon signed-rank $z = -4.42$, $p > |z| = 0.000$), and is reflected in subjects' open-ended responses as well.

Second, the survey allows us to examine more deeply elements of positive and negative messages that may be drivers of behavior. Previous research suggests that negative campaigning (Brians and Wattenberg, 1996; Joslyn, 1986) and comparative advertising more generally is found to be more informative (Harmon et al., 1983; Chou et al., 1987) and memorable (Faber and Storey, 1984; Appleton-Knapp and Mantonakis, 2009) than positive or non-comparative advertising, and researchers have suggested this difference as a possible reason for a mobilizing effect of negative campaigns. Contrary to these findings, survey respondents rated the candidates' positive messages as more informative than their negative messages in our experiment (Wilcoxon signed-rank $z = -3.83$, $p > |z| = 0.000$), suggesting that any relative mobilizing effect of our negative messages is not due to greater informational content of the negative message.

We find that the negative messages are no better than positive messages at earning the candidates donations, but negative messages yield significantly higher rates of voter turnout among the candidates' partisans relative to positive messages. The donation rate in the positive treatment was 0.9 percent and was 0.7 percent in the negative treatment; these are not statistically different (p -value = 0.65)². However, negative message recipients are 3.8 percentage points more likely to vote (p -value = 0.024)³. We find this pattern of results (negative messages increase turnout relative to positive ones) in both districts, suggesting it is not something particular to the electoral environment or the specific race. Since the fundraising letter was sent five months prior to the election, we check the robustness of our turnout results with a placebo check. We compare the turnout of the voters in our sample in each of the previous four elections as a function of our treatments. There is no relationship between our treatment and past turnout behavior, indicating that the effect of a negative message on turnout in the current election is not spurious.

While comparing negative messages to positive ones allows us to consider relative mobilization (of money and votes), it is also important to examine the absolute levels of mobilization compared to having sent no message. Compared to the control group, we find that both messages stimulate financial contributions to the candidates, as candidates receive no unsolicited contributions from the control partisans. Relative to no message, our turnout findings are more nuanced. In one district, negative message recipients have higher turnout than the control group (though the difference is not statistically significant), while turnout for the positive message recipients is slightly lower than the control (and again not significantly different). In the other district, it is the turnout of negative message recipients which is nearly identical to the control group, while the positive messages led to significantly lower voter turnout relative to the control.

Though not as high profile as elections for federal office, local races are the most common elections in the United States (U.S. Dept. of Commerce Census Bureau, 1995) and provide opportunities to conduct experiments with common campaign tactics that candidates in larger races do not use as often, such as in-person canvassing (Barton et al., 2014). Our results,

² As we discuss below, however, our available sample was likely underpowered to detect differences in the donation rate.

³ This p -value applies when pooling the data across districts. For reasons discussed below, we only present our turnout results for each district separately. The pooled analysis is available upon request.

across two local races within the same county, provide important field evidence on the effects of negative campaigning on fundraising and voter turnout⁴.

Previous studies suggest several reasons for a mobilizing effect of negative campaigns—that they stimulate a visceral emotional response, that negative campaigns receive greater weight because they highlight potential losses to avoid (i.e., prospect theory) or explicitly invoke competition and competitive (partisan) behavior, or that negative ads have been found more informative than positive ones generally. Our results present initial evidence against some of these explanations. First, prior research finds that the emotional impact of an event or message on behavior fades considerably over a short time period (Adler et al., 1998; Grimm and Mengel, 2011). As several months transpire between voters' receipt of the messages and their turnout decision, it is unlikely that it is an immediate, emotional reaction to the information provided that drives turnout. Second, while relative voter mobilization would suggest that the negative message received greater weight in voters' minds, it is important to consider how behavior changes relative to the uncontacted control group. It is only the positive message that shows differences in turnout relative to the control group, suggesting that it is not only highlighting potential losses that causes the differences we see between positive and negative messages when considered in the absence of the control group.

Finally, we consider whether the differential effect of the messages is due to differences in information. We take advantage of two attributes of our experiment to address this possibility. Our pre-experimental survey explicitly asks which message the participants find more informative; contra previous findings, participants find the positive message more informative than the negative message, so it is unlikely that the negative message generates higher relative turnout because its final recipients found it more informative. We also use the presence of a control group that does not receive either letter to distinguish between tone and information. If the effect on voter turnout of getting a campaign letter is solely due to having received additional information, then whether information is presented in a positive or negative light should have the same ordered effects on outcomes relative to the control group in both districts. As we mentioned above (and develop further in the final section of the paper), we observe different orderings of voter turnout for the two treatments and the control between the two districts, despite observing the same ordering for their information content in the pre-experimental surveys, suggesting that the effect of tone is separable from the quantity of information provided.

Results from this experiment have application to several literatures. The literature examining the effect of negative campaigning on voting behavior is large (see Lau et al., 1999, 2007). There are, however, few randomized experiments in naturally occurring settings (with the exception of Arceneaux and Nickerson, 2010; Gottfried et al., 2009; Niven, 2006). Most studies rely on indirect evidence and require strong identification assumptions to reach their conclusions. Our paper, by design, can examine how negative campaigning by a candidate works in a natural setting.

The effect of negative messages in campaigns also speaks to the broader marketing literature on comparative advertising (e.g., Barone and Jewell, 2013; Dianoux et al., 2013; Lovett and Shachar, 2011; Yagci et al., 2009). Negative advertising against another brand (or candidate) is one type of comparative advertising (Pinkleton, 1997; Shiv et al., 1997; Collens, 2011), which is generally more effective than non-comparative advertising (Barry, 1993; Grewal et al., 1997), but is rarely if ever compared to not advertising at all. Consistent with results from the marketing literature, we find the negative (implicitly comparative) message to yield greater voter turnout than the positive message. That we find the positive, non-comparative message and not the negative message affecting voter turnout relative to no message suggests that examining differences in intentions or actions between exposure to two forms of advertising without examining the views or deeds of those left alone produces an incomplete picture of individual behavior.

We also contribute to two literatures in economics. First, we add to the empirical literature on advertising. Advertising is thought to work by providing information and through persuasion (see Bagwell, 2007; DellaVigna and Gentzkow, 2010, for reviews), but economists have paid less attention to comparative advertising specifically. Anderson and Renault (2009), Barigozzi et al. (2009), Emons and Fluet (2012) provide game theoretical models of comparative advertising in product markets; Anderson et al. (2012, 2013) provide empirical analysis in the over-the-counter analgesics industry. Our results add to the growing empirical literature on the topic.

Our results also relate to the large literature on contributing to public goods (Vesterlund, 2006, for a review). The policies of a government of a particular jurisdiction are similar to a public good. They are non-excludable and non-rival, and so in that respect participation by a candidate's partisans, both through contributing and voting, is analogous to contributing to a public good⁵. Each partisan would presumably prefer her party's policies be enacted, but her personal payoff is higher by free-riding on the actions of others. Our results are consistent with previous evidence on the power of asking (Andreoni and Rao, 2011). We find that those who are not asked to contribute free ride on the monetary contributions of others, though they do turn out to vote.

Additionally, we find that supporters are more likely to contribute their vote when asked using a negative message than when using a positive one. Voting (by a partisan supporter) is a meaningful contribution to the campaign in its own right. This finding is consistent with Augenblick and Cunha's (2015) result that campaign contributions are higher when the request

⁴ We are cautious to generalize our results to all elections, be they local, regional or national. Our study offers field evidence on negative campaigning, and we encourage replication of experimental results in other races and settings (Maniadis et al., 2014).

⁵ Augenblick and Cunha (2015) perform a public goods experiment in the field using political fundraising, and make a similar argument regarding the public good nature of policy.

is framed competitively, in their case referencing the average giving behavior of out-partisans, and in our case emphasizing the consequences of an own-partisan loss. Contrary to many previous findings on the importance of asking, it isn't always helpful: the positive message lowers voter turnout relative to the control group in one district.

The paper proceeds as follows. Section 2 provides background information and motivates our design. Section 3 describes our experimental design. Section 4 presents results from our pre-experimental survey, and Section 5 the results from the field experiment. We offer a discussion and conclusion in Section 6.

2. Background

Negative campaigning involves any attack against a candidate's opponent, rather than an argument for the candidate. It is a form of comparative advertising, as highlighting the undesirable traits of one's opponent is an implicit claim to be better (Pinkleton, 1997). In the marketing literature, comparative messages have been found to receive greater consumer attention, yield greater brand and message awareness, message processing, and favorable attitudes toward the sponsored brand than noncomparative ads (Grewal et al., 1997). It is possible that, in political contexts, negative comparisons serve a similar role. While some authors argue that negative campaigning reduces participation by voters of all persuasions (Ansolabehere et al., 1994; Ansolabehere and Iyengar, 1995), the balance of the current evidence suggests there is a mild *mobilizing* effect of negative campaigning on voter turnout. Lau et al. (1999, 2009) find that, across multiple studies, negative campaigning has a positive impact on actual voter turnout⁶.

Few of these studies, however, identify the effect of negative political messages on behavior cleanly. Over half of the studies Lau et al. collect use observational data, which requires strong assumptions necessary to infer causal relationships. Of the experimental studies included and conducted since Lau et al.'s analysis, few measure intended or actual voter turnout, and many use fictitious candidates, advertisements, or both. Only two of the studies Lau et al. review are field experiments conducted within an actual election (Arceneaux and Nickerson, 2010; Niven, 2006), and few authors since have tried to measure the effect of negative campaigning in an election (Gottfried et al., 2009).

Of the field experiments on negative advertising, Arceneaux and Nickerson (2010) worked with an independent organization in the 2004 election and varied whether voters received a campaign phone call with a negative or positively framed message regarding various policy outcomes. They find no turnout effects and insignificant candidate preference effects. Niven (2006) also uses messages from an independent organization in a mayoral contest – not messages sent from one candidate attacking another as we do in our field experiment – to test the effect of negative campaigning. He finds voters who receive the negative messages have higher turnout rates. Gottfried et al. (2009) find that positive messages for judges up for reelection lead to higher voter turnout relative to arguments against reelection and negative campaigns from past judicial elections in other states.

Because we send messages from the candidate himself, we can examine the effect of positive arguments for a candidate compared to negative arguments against an opponent, which none of the previous experiments on this topic have done. Our design is more akin to comparative advertising. Also, it reduces potential confounds because our messages are sent within an actual campaign, and we directly compare the effects of positive to negative messages sent by the same candidate within the same election.

Previous research suggests that a possible reason that negative campaigning works is because it stimulates a more immediate emotional response from voters (Finkel and Geer, 1998). If invoking an emotional response is the primary mechanism behind negative campaign messages, we would expect differences between messages to be stronger in fundraising but not in voter turnout, as previous research has found the impact of emotional states on behavior to diminish over time (Adler et al., 1998; Grimm and Mengel, 2011). We do find a difference in messages in voter turnout, but not fundraising, leading us to believe that the effect of emotional responses on outcomes is of smaller importance.

It is also suggested that negative campaigns stimulate voter participation because of differences in the quantity of information conveyed. In other contexts (though not in ours), negative political advertising has been found to contain more information than positive advertising (Brians and Wattenberg, 1996; Harmon et al., 1983; Chou et al., 1987; Joslyn, 1986). For this mechanism to work, however, two requirements need to hold. First, a negative message would need to contain more information than a positive one. Second, having more information would need to lead to higher turnout. There is no reason, a priori, for either requirement to hold. Indeed, we find no evidence from our study to support the first requirement, and for the second, while additional information allows individuals to have more precise posterior beliefs, the *direction* in which those beliefs change is unspecified (Greene, 2008).

Even if negative campaign messages are not more informative, they may receive more weight in voters' minds. A negative message may simply "stand out" against a general backdrop of positive information and life experience (Lau, 1985). Or, they may draw attention to possible costs or losses to avoid, which may receive more attention if voters are loss averse (Kahneman and Tversky, 1979; Jain et al., 2007; Miller and Krosnick, 2004)⁷. Our results do not support the hypothesis that

⁶ They do report that negative campaigning reduces intended turnout, which is, however, a less reliable dependent variable than actual voting (Traugott, 2008).

⁷ Miller and Krosnick (2004) conducted a field experiment with an abortion-rights organization that framed an upcoming policy debate in terms of an "opportunity for or a "threat to abortion rights. The threat generated the greatest rates and levels of giving; they attribute their results to loss aversion. Jain

negative messages are more memorable than positive ones. In one district, the largest effect on turnout corresponds to the positive rather than the negative message. Finally, negative messages, by drawing implicit or explicit comparisons, evoke a competition, which has been shown to matter for the provision of public goods in the lab (Bornstein and Ben-Yossef, 1994; Bornstein et al., 2002) and the field (Augenblick and Cunha, 2015; Erev et al., 1993).

The above explanations for higher turnout due to negative campaigning are applicable to all voters, but several are particularly applicable to partisans. Candidates' partisans likely have a greater emotional reaction to information about the opposing candidate than independent voters. Relative to independents, partisans likely see the opponent's negatives in greater relief, and view an opponent's victory as a genuine loss in terms of public policy. Finally, partisan elections naturally set up a contest between parties. Reminding voters of the opposition is much more a call to compete for partisans than it is for the independent. In other words, partisans should be particularly sensitive to the difference between negative and positive messages⁸.

In the environment we examine, we use positive and negative campaign messages, where the latter emphasizes that failing to act will result in the opposition's control of the county legislature, undoing past progress. Given the results surveyed in this section, we posit the following hypotheses:

H1. Negative messages produce higher contribution rates than positive messages.

H2. Negative messages produce higher voter turnout rates than positive messages.

We turn now to the experimental design.

3. Experimental design

We conducted this experiment in two local elections for county legislature during the 2010 general election. The county legislature has nine three-member districts; we conducted the experiment with two Democratic candidates in two different districts. In the first district ("District A"), only a single seat was up for election, while in the other district ("District B") two seats were up for election. District A was predominantly Republican; the average Democratic share of the two-party vote for the county legislature from 2004 to 2008 was roughly 40 percent. The Democrats fielded no candidates in the district in 2002. District B was predominantly Democrats; the average Democratic share of the two-party vote from 2002 through 2006 was about 60 percent. There were no Republican candidates in the district in 2008. Both candidates in our experiment are Democrats. Both candidates had run for the office previously: the District A candidate lost the general election in 2008, while the District B candidate lost the 2008 Democratic primary. In 2010, the District A candidate again lost the general election; the District B candidate won⁹.

Our experiment focuses on the candidates' attempt to mobilize funds and votes from partisan supporters in their respective districts; it is for this first reason that the candidates send the letter five months prior to the election. According to voter registration records, District A contained about 15,200 registered voters (8400 households), and District B had roughly 11,800 registered voters (7150 households). We used voters' participation in party primaries to construct a population of likely supporters. We used primary election activity as the indicator of partisanship¹⁰. First, we kept only those voters who had participated in at least one of the last three Democratic Party primary elections (and no Republican Party primary elections) from 2004 through 2008. This left 2152 voters (1611 households) in District A, and 2784 voters (2089 households). Next, we removed all likely Democrats where any member of their household had participated in at least one of the three Republican Party primary elections from 2004 through 2008, reducing the target population to 1886 voters in District A (1367 households) and 2619 voters (1944 households) in District B. Finally, the campaigns employed a private address verification system to remove voters who had left the district; we removed all voters that moved outside of their current city, leaving 1798 individuals (1296 households) in District A, and 2415 individuals (1788 households).

Households in the candidates' districts with at least one likely partisan supporter by the above criteria were randomly assigned to receive a negative letter, a positive letter or no letter. The District A candidate sent letters to partisans in 1037 targeted households; the District B candidate sent letters to partisans in 1432 targeted households¹¹. As some households

et al. (2007) examines interaction of "regulatory focus"—whether someone is promotion or prevention oriented—and comparative advertising, and find that promotion-oriented subjects' opinion toward the advertised brand is higher under a positive comparison (A is better than B at X) while prevention-oriented subjects' opinion is higher under a negative comparison (B is worse than A at X).

⁸ Karlan and List (2007) find a stronger reaction by Democrats in Republican-dominant states to a message from a liberal organization.

⁹ We provide here some additional political context for the interested reader. In 2008, the District A candidate was selected to represent the Democratic Party post-primary, while the District B candidate came in second in a four-way primary with 30% of the vote. In 2010, the District A candidate was unopposed in the Democratic primary, while the District B candidate was the top vote-getter of his primary (getting 37% of the vote, or the support of 53% of the voters, as each could vote for two candidates). The 2010 general election in this state included competitive (in terms of ex post result) gubernatorial and senatorial campaigns, as well as contested (but not competitive, ex post) campaigns for U.S. Congress, state senate, and state legislator. Neither candidate used this message in subsequent campaigning; the District B candidate did campaign among these voters again in a manner completely orthogonal to the treatments here. The District A candidate lost with 31% of the vote in a two-person field. The District B candidate prevailed with 33% of the votes cast; he was the only one of four candidates to receive a vote from a majority of voters in the district.

¹⁰ This state does not register voters by party.

¹¹ We assigned 80 percent of households into one of the treatment groups, leaving 20 percent in the control group in each district. As mentioned, the control group received no letter from the candidates. Both campaigns had resources to reach nearly all the target households, and allowed us to remove

with Democratic partisans contain more than one partisan, we randomly selected the addressed recipient from among the Democrats in the household. Table 1 presents summary statistics on the targeted individuals in the sample by district and treatment; note that, as partisans, they have participated frequently in elections and are predicted to do so in 2010¹².

We produced the candidates' mailings using the candidates' resources and shipped the mail to the candidates to send¹³. The authors and the candidates developed the candidates' letters, and confirmed our experimental manipulation among 24 like-minded partisans outside of the candidates' districts. Each candidate's letters have the same opening and concluding paragraphs; two middle paragraphs contain the content that differs between treatments. The letters are close in length for both candidates: the District A candidate's positive (negative) letter contains 270 (263) words. The District B candidate's positive (negative) letter contains 281 (262) words¹⁴.

The two treatments involve the candidates' making an argument for their attributes and positions or against those of their opponents, in contrast to previous field research using messages that cast events in a positive or negative light or an argument from a third party. Each candidate's positive message is something positive about the candidate sending the message. Their negative messages are something undesirable about their opponent and the opponent's party, not merely negative information about the circumstances of the electorate. While we would have preferred that the positive and negative information be symmetrical – that what each campaign highlighted as a weakness of the opponent plays to something that is a strength in the candidate – the actual positive and negative characteristics of our cooperating candidates and their opponents did not align, as is often the case in actual campaigns. The mail pieces use the candidates' names and not generic party labels (though they refer to their opponents in the third person), and the letter is sent by the candidate (with the candidate's return address). These differences make the messages both more relevant and more directly linked to the campaign. Also, they more precisely test the effect of campaign tone by the primary actor in the electoral contest – the candidate – on voter behavior. We turn now to the experimental results.

4. Pre-experimental survey

We begin with the results of the individual interviews. In late April 2010, we recruited 24 registered voters in northern Virginia who frequently participate in Democratic Party primary elections through the email list of faculty and staff of a large state university. We scheduled individual sessions with each voter-subject in a one-week period. Subjects were randomly assigned to inspect of one candidate's mail pieces¹⁵. Subjects in the interviews read both messages from one candidate, with the order randomized for each subject. We asked the subjects to rate the tone of the messages, their informational content, and their affect toward the sender of the messages. Surveys lasted about 20 min, and we paid subjects \$15 for their time.

The interviews serve two purposes. First, they validate our experimental manipulation among like-minded voters outside the candidates' districts. Second, they allow us to examine explanations for differences between positive and negative messages. In particular, we consider whether the subjects found the negative messages more informative, as Finkel and Geer (1998) posit this difference as a potential mechanism negative advertising's effect as discussed above.

The quantitative results of these interviews are reported in Table 2, and provide strong evidence validating our interpretation of the treatments. Survey participants view the positive letter as “very positive” (it is median and modal response) and the negative message as between “somewhat” and “very negative.” This difference is strongly statistically significant (Wilcoxon signed-rank $z = -4.42$, $p > |z| = 0.000$), and is reflected in subjects' open-ended responses.

20 percent. They did not think it prudent to contact any fewer potential supporters, and due to the size of the districts, there were no more additional supporters to put into either a treatment or control group. We performed initial power calculations for the districts pooled, which indicated that we would detect a 5 percent difference in turnout rates at $p = 0.05$ between the negative and positive groups with a power of 0.85. When looking at each district individually, however, the effect size needs to be roughly ten percentage points to have an 80% chance of finding it at $p = 0.05$.

¹² We confirm our household-level randomization procedure in the supporting information; there are no observable characteristics that predict household assignment to treatment. The variable “predicted likelihood to vote” comes from a probabilistic model of voting in the 2006 midterm election using voters' demographics and pre-2006 voting behavior (age and age squared in 2006, sex, whether the individual voted in the three previous elections, and whether they voted in a party primary in 2006). We applied the coefficients of this model to the voters' 2010 demographic characteristics and voting histories to estimate each voter's likelihood to vote in the 2010 midterm election, and verified this approach using the same estimation techniques on the 2002 data to predict voter turnout in 2006. Our prediction correlated strongly with behavior ($\rho = 0.68$). Note that we apply the model even to those subjects whose age precludes having voted in the last three elections (i.e., those aged 18–24). See Brox and Hoppe (2005) for a discussion of such methods and their predictive accuracy.

¹³ Readers may wonder if our involvement in the campaigns was ethical; subject-voters lack informed consent. In addition to obtaining institutional review board approval for the research at our university, it is worth noting that campaigns are already actively contacting voters multiple times in a campaign and already using experimental methods to study voter behavior in order to win elections (e.g., the Analyst Institute, an organization designed for this purpose). Like others, we partner with campaigns in order to advance the scholarship of voter behavior.

¹⁴ The supporting information contains all messages' text and a figure depicting one complete letter. We also vary the quality of the delivery mechanism. We examined other local candidates' mail pieces and consulted professional printers to design a “high quality” and “low quality” single-page letter with a contribution card and return envelope for each candidate. The high quality mailing envelope, return envelope, and the letterhead had the candidate's logo in color; the letter was printed on heavier, brighter paper with the candidate's signature in color ink. The low quality letter used no logos and was printed monochromatically on the lowest quality of printer paper. There were no discernible effects across these two treatments, so we pool the data across these treatments.

¹⁵ The supporting information contains the interview script.

Table 1
Sample descriptive statistics.

District A	Control			Positive letter			Negative letter		
	Obs	Mean	SD	Obs	Mean	SD	Obs	Mean	SD
Male	259	0.38	0.49	517	0.42	0.49	520	0.42	0.49
Age	256	52.4	15.7	510	53.4	15.9	515	51.8	15.7
Strong democrat	259	0.27	0.44	517	0.31	0.46	520	0.30	0.46
Weak democrat	259	0.73	0.44	517	0.69	0.46	520	0.70	0.46
Percent democrats in household	259	0.80	0.26	517	0.80	0.25	520	0.76	0.27
Voters in household	259	1.93	0.94	517	1.93	0.91	520	2.02	0.85
Previous donor household	259	0.10	0.30	517	0.14	0.35	520	0.12	0.33
Predicted likelihood to vote	256	0.66	0.25	510	0.69	0.25	515	0.69	0.24
Voted in 2008 general	259	0.93	0.26	517	0.91	0.28	520	0.92	0.27
Voted in 2006 general	259	0.65	0.48	517	0.71	0.46	520	0.71	0.45
Voted in 2004 general	259	0.78	0.41	517	0.81	0.40	520	0.82	0.39

District B	Control			Positive letter			Negative letter		
	Obs	Mean	SD	Obs	Mean	SD	Obs	Mean	SD
Male	356	0.37	0.48	716	0.41	0.49	716	0.36	0.48
Age	355	50.8	16.0	706	50.2	16.4	712	50.7	16.2
Strong democrat	356	0.37	0.48	716	0.36	0.48	716	0.36	0.48
Weak democrat	356	0.63	0.48	716	0.64	0.48	716	0.64	0.48
Percent democrats in household	356	0.84	0.25	716	0.82	0.25	716	0.82	0.25
Voters in household	356	1.84	0.98	716	1.82	0.93	716	1.86	0.95
Previous donor household	356	0.19	0.39	716	0.18	0.38	716	0.20	0.40
Predicted likelihood to vote	351	0.68	0.24	695	0.67	0.25	704	0.69	0.25
Voted in 2008 general	356	0.94	0.24	716	0.93	0.26	716	0.92	0.27
Voted in 2006 general	356	0.72	0.45	716	0.71	0.45	716	0.74	0.44
Voted in 2004 general	356	0.83	0.37	716	0.81	0.39	716	0.83	0.38

Strong democrats voted in at least two of the last democratic primaries and no other party primary. Weak democrats voted in one of the last three democratic primaries, or two of the last three and one non-republican primary
According to Board of Election officials, a birthday of "01/01/1900" in the voter registration records indicates a missing value.

District A	Control			Positive letter			Negative letter		
	Obs	Mean	SD	Obs	Mean	SD	Obs	Mean	SD
Male	259	0.38	0.49	517	0.42	0.49	520	0.42	0.49
Age	256	52.4	15.7	510	53.4	15.9	515	51.8	15.7
Strong democrat	259	0.27	0.44	517	0.31	0.46	520	0.30	0.46
Weak democrat	259	0.73	0.44	517	0.69	0.46	520	0.70	0.46
Percent democrats in household	259	0.80	0.26	517	0.80	0.25	520	0.76	0.27
Voters in household	259	1.93	0.94	517	1.93	0.91	520	2.02	0.85
Previous donor household	259	0.10	0.30	517	0.14	0.35	520	0.12	0.33
Predicted likelihood to vote	256	0.66	0.25	510	0.69	0.25	515	0.69	0.24
Voted in 2008 general	259	0.93	0.26	517	0.91	0.28	520	0.92	0.27
Voted in 2006 general	259	0.65	0.48	517	0.71	0.46	520	0.71	0.45
Voted in 2004 general	259	0.78	0.41	517	0.81	0.40	520	0.82	0.39

District B	Control			Positive Letter			Negative Letter		
	Obs	Mean	SD	Obs	Mean	SD	Obs	Mean	SD
Male	356	0.37	0.48	716	0.41	0.49	716	0.36	0.48
Age	355	50.8	16.0	706	50.2	16.4	712	50.7	16.2
Strong democrat	356	0.37	0.48	716	0.36	0.48	716	0.36	0.48
Weak democrat	356	0.63	0.48	716	0.64	0.48	716	0.64	0.48
Percent democrats in household	356	0.84	0.25	716	0.82	0.25	716	0.82	0.25
Voters in household	356	1.84	0.98	716	1.82	0.93	716	1.86	0.95
Previous donor household	356	0.19	0.39	716	0.18	0.38	716	0.20	0.40
Predicted likelihood to vote	351	0.68	0.24	695	0.67	0.25	704	0.69	0.25
Voted in 2008 general	356	0.94	0.24	716	0.93	0.26	716	0.92	0.27
Voted in 2006 general	356	0.72	0.45	716	0.71	0.45	716	0.74	0.44
Voted in 2004 general	356	0.83	0.37	716	0.81	0.39	716	0.83	0.38

Strong democrats voted in at least two of the last democratic primaries and no other party primary. Weak democrats voted in one of the last three democratic primaries, or two of the last three and one non-republican primary
According to Board of Election officials, a birthday of "01/01/1900" in the voter registration records indicates a missing value

Table 2
Subject evaluation of letter tone, information, and candidate likeability.

	Tone of letter [1 = very positive; 4 = very negative]		Negative		$p \geq z $
	Positive		Mean	SD	
District A	1.1	0.3	3.3	0.6	0.002
District B	1.1	0.4	3.2	0.6	0.002
Both districts	1.1	0.4	3.2	0.6	0.000
	Information in letter [1 = very; 4 = not at all]		Negative		$p \geq z $
	Positive		Mean	SD	
District A	1.5	0.5	2.9	0.7	0.002
District B	1.7	0.9	2.2	0.9	0.034
Both districts	1.6	0.7	2.5	0.9	0.000
	Likeability of candidate [1 = Much more; 4 = Much less]		Negative		$p \geq z $
	Positive		Mean	SD	
District A	1.5	0.6	3.1	0.7	0.002
District B	1.9	0.7	2.8	0.9	0.082
Both districts	1.7	0.7	2.9	0.8	0.001

Last column ($p > |z|$) reports the statistical significance of Wilcoxon signed-rank test between the positive and the negative letters

Participants responded to the framing effect without prompting. They described the positive letters as “positive”, also indicating that it “emphasizes qualifications” of the candidates. One participant stated that the positive letter was “selling himself.” Participants described the negative letter as “negative” and an “attack” and clearly had mixed feelings regarding the content. It focused one participant’s attention on “what they’ll do if we let them win.” Another said it highlighted “threats from the other party”, while another said it was “clearly designed to get blood boiling.” A participant who found the negative message distasteful still offered that it “forces you to do something.” A participant who liked the attack summarized it thus: “do you *know* what the *Republicans* are up to?”¹⁶ Both the quantitative and qualitative evidence provide support that partisan-minded readers perceive the tonal difference between letters, and some qualitative evidence suggesting an emotional or loss-avoiding reaction to the negative letter.

Participants also saw a difference in informational content between the letters. On average, subjects found the positive letter to be between “very informative” and “somewhat informative,” but the negative letter to be between “somewhat” and “not very informative.” This difference is statistically significant (Wilcoxon signed-rank $z = -3.83$, $p > |z| = 0.000$). While it is possible that the ultimate recipients get more information out of the negative than the positive letter, our survey participants do not, suggesting that voters in the field finding the negative messages *more* informative is likely not the cause of any differences we see in fundraising and voter turnout. Furthermore, differences in tone are not merely differences in information. Participants found the tone difference between the two letters to be larger than the information difference (Wilcoxon signed-rank $z = -3.97$, $p > |z| = 0.000$).

We also examined candidate affect. Because we do not observe candidate choice for the voters in the field, we wanted to assess the degree to which negative feelings might possibly influence candidate support. Consistent with the results of previous laboratory experiments (e.g., Jain and Posavac, 2004), the positive letter makes the candidates between “much more” and “a little bit more likeable” to survey participants, while the negative letter makes them between “a little less” and “much less likeable.” This difference is also statistically significant (Wilcoxon signed-rank $z = -3.32$, $p > |z| = 0.001$). While participants felt less favorably toward the candidate following the negative message, several volunteered that they would still likely vote for him. One put it thus: I “vote for [the Democrat] unless he’s a real doofus.” Another, who was disinclined to support candidates who go negative, said “I might vote for him, because he is a Democrat. But I would hold my nose.” Even if a candidate’s partisans find him less likeable after negative campaigning, they appear unlikely to switch sides. We turn now to the results from the field experiment.

5. Field experiment results

We conducted the experiment with the candidates in the first two weeks of June 2010. The authors produced both candidates’ letters using their campaign funds in late May 2010 and shipped the solicitations to the candidates. The candidates

¹⁶ These quotations are excerpts of answers made by survey participants to the interviewer (an author) to the question “what is your impression of Item X?” The use of “positive” and “negative” is unprompted by the question. The supporting information contains the interview script; interview notes are available upon request.

Table 3
Contribution rate, revenue per solicitation, and total contributions by district.

	Contribution rate			Negative			p-Value
	Positive			Obs	Mean	SD	
	Obs	Mean	SD	Obs	Mean	SD	
District A	516	0.014	0.116	519	0.012	0.107	0.77
District B	706	0.006	0.075	709	0.004	0.065	0.70
	Revenue per solicitation			Negative			p-Value
	Positive			Obs	Mean	SD	
	Obs	Mean	SD	Obs	Mean	SD	
District A	516	0.562	5.183	519	0.578	5.989	0.96
District B	706	0.283	4.407	709	0.106	1.624	0.32
	Total contributions			Negative			
	Positive			Obs	Mean	SD	
District A		\$290			\$300		
District B		\$200			\$75		

Last column (p-value) reports the statistical significance of two-sided t-test (unequal variances) between the positive and the negative letters

then mailed the solicitations to households. All letters were sent in the first two weeks of June. Candidates collected contributions over the next six weeks, and received no contributions from those solicited following the six week recording period. Following the election, we acquired voter turnout records for each district from the county board of elections. We consider each outcome below¹⁷.

5.1. Campaign contributions

Table 3 presents the contribution rates, revenue per solicitation (RPS), and total contributions received by district. In neither district do we find evidence that the negative letter stimulates greater rates of giving than the positive one. In District A, 1.4 percent of the positive message recipients donated to the candidate, while 1.2 percent of the negative message recipients did so (p-value = 0.775, two-sided t-test with unequal variances). Only 0.6 percent of District B’s positive message recipients donated, while 0.4 percent of the negative message recipients gave to the candidate (p-value = 0.701, two-sided t-test with unequal variances).

The fundraising data provide no evidence for the hypothesis that negative campaigning stimulates partisan financial support. We note, however, given how unresponsive donation behavior was to our treatments, that our lack of difference may stem more from a lack of power than from a lack of actual differences (as seen in, e.g., Miller and Krosnick, 2004)¹⁸. Both positive and negative letters, however, served to stimulate giving among Democrats, as neither candidate received donations from the control group during or after the period of the study.

5.2. Voter turnout

We now turn to the effect of messages on voter turnout. While receiving a letter is effective in getting partisans to give irrespective of tone, we will see that for voter turnout, tone does make a difference.

The main treatment effects can be seen in Fig. 1. This shows the rate of voter turnout by treatments and control across districts. In both districts, a negative message yields significantly larger voter turnout relative to a positive message. This is not, however, due to the mobilizing power of negative messages relative to the control group. In District A, the negative letter yields higher turnout than the control, while the positive letter yields slightly lower turnout (though not significantly). In District B, both letters lead to lower levels of voter turnout than the control; the negative letter just has a smaller (and statistically insignificant) negative impact on turnout.

We examine the robustness of these results in Table 4. The table shows linear probability model regressions of whether or not the voter turned out to vote on dummies for the two message treatments and the additional covariates we used in the regressions on contribution behavior¹⁹. Confirming the results from Fig. 1, we see that voter turnout is higher for those individuals who received a negative message compared to those who received a positive message. This holds in each district separately, with or without covariates. For example, in District A, a negative message yields a 4.9 percentage point increase

¹⁷ It is important to note that all effects here are intent-to-treat (ITT), as we cannot determine which households actually read the letter (see Perez-Truglia and Cruces, 2013). That said, for many forms of campaigning, including direct mail, radio, and television advertising, the ITT effect is the effect of interest, as campaigns cannot confirm actual receipt of the message.

¹⁸ We were uncertain as to how responsive donations would be to the treatments – particularly those discussed in footnote 11 above – given that we contacted active partisans of both candidates. Actual donation responses, as shown in Table 3, are far too low for us to generate any meaningful results given the reach of these campaigns. Our donation rate is not dissimilar from other political campaigning experiments, however (Augenblick and Cunha, 2015).

¹⁹ In control households with multiple partisans, we randomly selected one voter as the “recipient. Our results do not change if we include all partisans in control households.




Fig. 1. Voter turnout by treatment.

in voter turnout relative to a positive message (p -value = 0.064). In District B, this difference is 3.2 percentage points, but is not statistically significant (p -value = 0.116).

Negative messages consistently have a stronger mobilizing effect than positive messages²⁰. As we saw in Fig. 1, the effect of each message relative to the control is not consistent across districts. In District B, both messages reduce turnout relative to the control, and the positive message has the largest and most significant negative impact, with and without covariates. In District A, however, the negative message significantly increases turnout relative to the control (but only without controlling for covariates) but the positive message decreases turnout (although not significantly)²¹. We return to this finding in the discussion.

The fundraising letter was sent five months prior to the election, so as a robustness check on our voter turnout results, we run a “placebo” check. We regress the turnout of each of the previous four elections as a function of our treatments. We report the results of these analyses in Table 5. In panel A, we regress the turnout of all voters on our treatment variables alone. In panel B, we restrict the sample to only those voters eligible to vote in all elections examined (i.e., who were at least 18 in 2002)²². Panels C and D use the same restrictions on the sample as panels A and B, respectively, but include covariates as in Table 4. As our method for calculating the predicted likelihood to vote requires data from the previous five elections, we use instead control for likelihood to vote using the voter’s turnout in the three general elections preceding

²⁰ Nickerson (2008) demonstrated that affecting the voter turnout decision of one household member can influence the decision of others to vote. We examined non-recipients’ voter turnout in 2010 across treatments in the supporting information; the difference in voter turnout between the negative and positive letter treatments is positive, but is not statistically significant. There is no spillover behavior to analyze in the fundraising treatments. In one District B household, a different member than the addressee sent the return check. The sender, however, had previously indicated a desire to contribute to the candidate.

²¹ Because of this difference, we do not discuss the pooled results here, and thank an anonymous reviewer for raising this issue. Of course, the difference between the positive and negative letter treatments holds – and is more strongly significant – if we pool the two districts (3.8 percentage points, p -value = 0.024); these results are presented in columns 5 and 6.

²² As with the predicted likelihood to vote above, we include those subjects who were 18 in 2002, even though they would not have a previous voting history for the 1996–2000 election cycles.

Table 4
Voter turnout of letter recipients.

	(1) District A	(2) District A	(3) District B	(4) District B	(5) Both	(6) Both
Positive letter	−0.004 (0.037)	−0.030 (0.034)	−0.072** (0.030)	−0.070** (0.028)	−0.043* (0.024)	−0.052** (0.021)
Negative letter	0.053 (0.037)	0.019 (0.033)	−0.032 (0.030)	−0.038 (0.027)	0.003 (0.023)	−0.014 (0.021)
Positive = Negative? (<i>p</i> -Value, <i>F</i> -test)	0.062 [†]	0.064 [†]	0.116	0.155	0.016**	0.024**
Male		−0.040 (0.024)		0.007 (0.020)		−0.012 (0.016)
Age		−0.001 (0.001)		−0.000 (0.001)		−0.001 (0.001)
Strong democrat		−0.005 (0.028)		0.035 (0.023)		0.017 (0.018)
Percent democrats in household		0.007 (0.060)		0.058 (0.054)		0.033 (0.040)
Voters in Household		0.025 (0.017)		0.017 (0.014)		0.019* (0.011)
Previous democratic donor household		0.077** (0.033)		0.035 (0.024)		0.049** (0.019)
Predicted likelihood to vote		0.954*** (0.048)		0.849*** (0.043)		0.897*** (0.032)
District B binary						0.059*** (0.016)
Constant	0.595*** (0.031)	−0.048 (0.086)	0.711*** (0.024)	0.062 (0.076)	0.662*** (0.019)	−0.015 (0.058)
Observations	1294	1279	1771	1734	3065	3013
R-squared	0.003	0.242	0.003	0.230	0.002	0.236

Positive = Negative? reports the *p*-value of an *F*-test of the equality of coefficients between the positive and negative letters. Robust standard errors in parentheses.

*** *p* < 0.01.

** *p* < 0.05.

† *p* < 0.1.

each election analyzed²³. Across all elections from 2002 to 2008, there is no relationship between our treatment and past behavior, indicating that what we find in Table 4 is not spurious²⁴. The voter turnout effects we find from the treatment letters are robust.

6. Discussion and conclusions

Overall, we find mixed evidence on the effect of negative campaigning in the field. Unlike previous field experiments that framed fundraising in terms of policy “threats” and “opportunities” (Miller and Krosnick, 2004) or that primed competitive motivations for partisan fundraising (Augenblick and Cunha, 2015), we find no evidence that negative messages about the opposition spurred more giving than the positive message about the candidate. The positive message had higher donation rates and levels, though not statistically significantly so.

With respect to turnout, however, we do find a strong relative effect. This is consistent with previous research that compares positive and negative messages using observational and field experimental data. In our experiment, recipients of the negative message were roughly 4 percentage points more likely to go to the polls than recipients of the positive message. While it is tempting to conclude therefore that negative messages motivate a candidate’s core supporters, it is not the whole story. Unlike Niven (2006), our negative message did not have a statistically significant impact on turnout relative to the control group, while the positive message lowered voter turnout relative to the control group. This result runs contrary to evidence on partisan get-out-the-vote (GOTV) efforts (Nickerson et al., 2006) which find a small mobilizing effect. Unlike the GOTV efforts, targeted voters received the message several months prior to Election Day, and our letters do not contain any information about how or where to vote.

²³ We also use participation in the previous three Democratic elections to indicate partisanship in place of the calculated “strong Democrat” variable. We demonstrate in the first two columns of panels C and D that these changes do not affect our results in the 2010 general election.

²⁴ In the supporting information, we conduct an additional robustness check using data from 2002 through 2014. Using data on all recipients eligible to vote in elections from 2002 forward, we estimate the following regression: $V_{it} = \sum \alpha_i + \sum \beta * Negative_{it} + \sum \delta * Positive_{it} + \sum \gamma_t + \varepsilon_{it}$. This fixed-effect regression at the individual level allows us to examine the difference between the positive and negative letter recipients in years before and after the election of the experiment, controlling for any individual-level characteristics that are time invariant. We find a significant difference between the negative and positive letter recipients in 2010, and a weakly significant effect in 2014. This remains true even if we estimate the standard errors via bootstrapping, suggesting again that our result is not spurious.

Table 5
Robustness check: treatment differences in prior elections.

Panel A: Without covariates						
	2010	2008	2006	2004	2002	
Positive letter	−0.043*	−0.008	0.017	−0.004	0.005	
	(0.024)	(0.013)	(0.023)	(0.019)	(0.024)	
Negative letter	0.004	−0.013	0.035	0.010	0.010	
	(0.023)	(0.013)	(0.022)	(0.019)	(0.024)	
Observations	3065	3065	3065	3065	3065	
R-squared	0.002	0.000	0.001	0.000	0.000	
Panel B: Without covariates, restricted to 2002 age-eligible voters						
	2010	2008	2006	2004	2002	
Positive letter	−0.037	−0.006	0.015	0.000	0.012	
	(0.024)	(0.013)	(0.023)	(0.019)	(0.025)	
Negative letter	0.007	−0.010	0.033	0.016	0.013	
	(0.024)	(0.013)	(0.022)	(0.019)	(0.024)	
Observations	2976	2976	2976	2976	2976	
R-squared	0.002	0.000	0.001	0.000	0.000	
Panel C: With covariates						
	2010a	2010b	2008	2006	2004	2002
Positive letter	−0.052**	−0.049**	−0.007	0.014	−0.000	0.009
	(0.021)	(0.022)	(0.012)	(0.020)	(0.016)	(0.017)
Negative letter	−0.014	−0.001	−0.016	0.032	0.008	0.001
	(0.021)	(0.021)	(0.012)	(0.020)	(0.016)	(0.017)
Observations	3013	3035	3035	3035	3035	3035
R-squared	0.236	0.196	0.058	0.241	0.322	0.524
Panel D: With covariates, restricted to 2002 age-eligible voters						
	2010a	2010b	2008	2006	2004	2002
Positive letter	−0.048**	−0.044**	−0.005	0.009	0.001	0.011
	(0.022)	(0.022)	(0.012)	(0.020)	(0.016)	(0.017)
Negative letter	−0.011	0.001	−0.014	0.027	0.011	0.002
	(0.021)	(0.021)	(0.012)	(0.020)	(0.016)	(0.017)
Observations	2927	2946	2946	2946	2946	2946
R-squared	0.225	0.187	0.060	0.227	0.301	0.504

2010a refers to coefficients from regressions estimated as in Table 4 for the pooled districts; 2010b refers to coefficients from regressions estimated with binary variables for the last three general elections and the last three primary elections in place of “Predicted Likelihood to Vote” and “Strong Democrat”, respectively. Robust standard errors in parentheses.

*** $p < 0.01$.

** $p < 0.05$.

* $p < 0.1$.

Because our experiment uses a combination of methods (survey and field experiment), measurements of behavior (monetary and voting contribution), and comparisons (both between treatments and to a control group) to study negative campaigning in an actual election, we are able to consider several possible causal mechanisms behind our results. First, consider emotional response to the messages as an explanation. Subjects in the survey do admit to feeling riled and action-prone by the negative message relative to the positive one. But if an emotional reaction to the negative messages were the mechanism behind the relative effect, we would expect to see it primarily in the donation data, where affect could drive immediate behavior, but less so if at all in the turnout data, by which time the emotional reaction to the argument would likely have faded (Adler et al., 1998; Grimm and Mengel, 2011). This is not what we see. There is no significant difference in financial contribution between messages; only months later do we see a difference in voter turnout.

It is also difficult to attribute the difference in turnout between messages to the negative message receiving more weight in the partisans' minds, as it had no impact on their turnout relative to voters in the control group. It is, however, possible to explain our results as information transmission, though not of the kind posited by previous research that finds negative messages more informative. Recall that our survey respondents found the positive message more informative than the negative one. When comparing only the negative and positive messages, we cannot say whether partisans' turnout behavior was affected by a message, just what the difference in the two messages is. Compared to the control, however, we see that the negative message had no significant effect on turnout, while the positive message decreased turnout. This is entirely consistent with the survey result: the more informative message changes voter behavior, while the less informative one does not.

We do not, however, argue here that tone is merely information. We suggest, rather, that tone has an effect separable from how informative a message is. Consider the following thought experiment. Suppose that only having received information from our messages matters for voter turnout, and that the messages we employed can be ranked according to their information content, as indeed they can: the positive message has more information than the negative message, according to our survey respondents. We should expect that the ordering of the effect of messages is consistent across districts. Even

if the effect of information is positive in one district or negative in another, if information is the only variable at work, we should observe that the ordering of messages according to their impact – including the no-information control group – is the same across districts. If the two districts do not show the same ordering of turnout across the three groups, this implies that it is not just information that matters, but that tone has its own separate effect.

In all orderings, the message treatment that contains the most information should have the most extreme effect, either increasing or decreasing turnout, relative to the control. The effect of the less informative treatment should lie between the control and the more informative treatment. Now, if information is helpful in one district but hurtful in the other, the ordering of the magnitude always holds (the most informative has the largest effect, then the least informative, then the control), just in opposite directions (helpful information increases turnout, and hurtful information lowers it). Such orderings are consistent with information, and not tone itself, affecting voter turnout.

As we saw in Fig. 1, the effect of each message relative to no message is not consistent across districts. There is a clear ordering of voter turnout outcomes (t) in District B, with $t_{\text{Bcontrol}} \geq t_{\text{Bnegative}} > t_{\text{Bpositive}}$. In District A, the ordering is $t_{\text{Anegative}} > t_{\text{Acontrol}} \geq t_{\text{Apositive}}$. The results from District B are consistent with a pure information story, that is, information affects turnout, and positive messages having more information than negative ones. If this were the case, then we would expect a similar ordering of turnout across treatments and control in District A based on the assumptions above. We do not see this in the data. Indeed, using a Cuzick (1985) nonparametric trend test²⁵, we reject the null for District B (p -value = 0.014) but not in District A (p -value = 0.599). That is, voter turnout can be ordered as highest in the control, then in the negative message treatment, then in the positive message treatment in District B, but there is no such order in the data from District A.

We also test whether the alternative ordering, where the negative message has the largest effect relative to the control, and the positive message has an intermediate effect, is significant. We find no evidence of this ordering in District B (p -value = 0.563) but do find evidence for it in District A (p -value = 0.090). Taking the results from both tests together, we can conclude that there is statistical evidence for the positive message having the largest effect relative to the control in District B and the negative message having the largest effect in District A. These results are not consistent with only information mattering. Tone is important independent of the informational content.

Our turnout results confirm that negative messages do not diminish turnout relative to positive ones; at least among partisans, it is the negative message that yielded significantly higher turnout, though not relative to receiving no message. This result is consistent with naturally occurring and laboratory evidence, even if the explanations for a mobilizing effect of negative campaigns do not readily explain our findings. That the levels of turnout across the control and treatment groups differs by district presents an additional puzzle. In particular, in District A, both letters lead to higher turnout relative to the control (albeit not significantly so), while in District B, the positive letter significantly lowers turnout.

One possible explanation for this difference is the political particulars of each candidate. As both messages in District B yield lower turnout than the control, perhaps the candidate himself turns off voters, while the negative message reminds these partisans of the stakes of not supporting their own party. In District A, however, both messages yield (slightly) higher turnout, as most of the get-out-the-vote literature would anticipate, with the negative message simply outperforming the positive message. This implies that the District B candidate, though, is more unpopular among his own partisans compared to the District A candidate. This is certainly possible, but seems unlikely, as the District B candidate won his primary (and placed a close second in the 2008 primary), and went on to win more votes than any other candidate in the district during the general election.

Another possibility is that it is not merely the candidate, but both the candidate and political context that matter for messages. We noted above that District A leans heavily Republican, while District B is more Democratic, but not nearly as lopsided. Perhaps both letters in District A remind these Republican-surrounded Democrats of the importance of supporting fellow partisans for office, with the negative message making that importance clearer, while in “safe” District B, the positive message merely indicates that a perfectly electable Democrat is running – as he’ll prevail, no need to turn out here – while the negative message removes some, but not all, of this all-will-be-well affect, potentially triggering feelings of loss aversion. While such an interpretation is plausible, we withhold judgment as to the best explanation for between-district results, as it both requires assumptions that are not readily testable with our data, and tries to craft an explanation for a sample size of two districts. We believe that observing such differences across more than two districts would be best before offering additional speculation on these particular districts.

In addition to providing some evidence on the causal mechanism behind negative advertising, there are practical implications of these findings. Candidates and political parties depend on contributions, for which they generally have to ask. Our results suggest that these groups face a tradeoff when making purely positive arguments: the act of asking positively yields monetary contributions but reduces “political contributions” at the ballot box. No such tradeoff exists in our data when asking with a negative appeal. As campaigns need to fundraise in order to operate, but do not want to demobilize their own supporters, it appears that going negative is the best way to ask for political donations. Thus, despite voters’ stated preferences for positive messages, it should not surprise partisan supporters of any stripe that their favored candidates go on the attack when asking for their support.

²⁵ The null hypothesis of the test is that there is no order in voter turnout across the control and treatment groups, against the alternative hypothesis that there is a significant ordering of the data with the positive message having the largest difference from the control.

Acknowledgements

We thank Thomas Stratmann, Kevin Arceneaux, and several anonymous referees, as well as seminar participants the 2012 Public Choice Society Conference and the 2013 North American Economic Science Association Conference. We cannot thank our two cooperating candidates enough, who executed our protocol willingly and without complaint. All errors are, of course, our own.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at <http://dx.doi.org/10.1016/j.jebo.2015.10.007>.

References

- Abramowitz, A., McGlennon, J., Rapoport, R., 1981. A note on strategic voting in a primary election. *J. Polit.* 43 (3), 899–904.
- Adler, R.S., Rosen, B., Silverstein, E.M., 1998. Emotions in negotiation: how to manage fear and anger. *Negotiat. J.* 14 (2), 161–179.
- Anderson, S., Renault, R., 2009. Comparative advertising: disclosing horizontal match information. *RAND J. Econ.* 40, 558–581.
- Anderson, S., Ciliberto, F., Liaukonyte, J., 2013. Information content of advertising: empirical evidence from the OTC analgesic industry. *Int. J. Ind. Org.* 31 (5), 355–367.
- Anderson, S., Ciliberto, F., Liaukonyte, J., Renault, R., 2012. Push-Me Pull-You: Advertising in the OTC Analgesics Industry. (http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2047106).
- Andreoni, J., Rao, J., 2011. The power of asking: how communication affects selfishness empathy, and altruism. *J. Publ. Econ.* 95, 513–520.
- Ansolabehere, S., Iyengar, S., Simon, A., Valentino, N., 1994. Does attack advertising demobilize the electorate? *Am. Polit. Sci. Rev.* 88 (4), 829–838.
- Ansolabehere, S., Iyengar, S., 1995. *Going Negative: How Attack Ads Shrink and Polarize the Electorate*. The Free Press, New York, NY.
- Appleton-Knapp, S., Mantonakis, A., 2009. Implications of the relationship between retrieval strength and storage strength in a comparative advertising context. *Adv. Consum. Res.* 8, 179–184.
- Arceneaux, K., Nickerson, D.W., 2010. Comparing positive and negative campaign messages: evidence from two field experiments. *Am. Polit. Res.* 38 (1), 54–83.
- Augenblick, N., Cunha, J., 2015. Using competition to elicit cooperation in a political public goods game: a field experiment. *Econ. Inq.* 53 (1), 574–588.
- Bagwell, K., 2007. The economic analysis of advertising. Mark Armstrong and Rob Porter Handbook of Industrial Organization, vol. 3. North-Holland, Amsterdam, pp. 1701–1844 (2007).
- Barigozzi, F., Garella, P., Peitz, M., 2009. With a little help from my enemy: comparative advertising as a signal of quality. *J. Econ. Manage. Strat.* 18, 1071–1094.
- Barone, M.J., Jewell, R.D., 2013. The innovator's license: a latitude to deviate from category norms. *J. Mark.* 77 (January), 120–134.
- Barry, T., 1993. Comparative advertising: what have we learned in two decades? *J. Advert. Res.* 33 (2), 19–29.
- Barton, J., Castillo, M., Petrie, R., 2014. What persuades voters? A field experiment in political campaigning. *Econ. J.* 124, F293–F326.
- Bornstein, G., Ben-Yossef, M., 1994. Cooperation in intergroup and single-group social dilemmas. *J. Exp. Soc. Psychol.* 30, 52–67.
- Bornstein, G., Gneezy, U., Nagel, R., 2002. The effect of intergroup competition on intragroup coordination: an experimental study. *Games Econ. Behav.* 41, 1–25.
- Brians, C., Wattenberg, M., 1996. Campaign Issue knowledge and salience: comparing reception from TV commercials TV news, and newspapers. *Am. J. Polit. Sci.* 40 (1), 172–193.
- Brox, B., Hoppe, R., 2005. Predicting voter turnout: testing new tools. In: Paper Presented at the Midwest Political Science Association Annual Meeting, Chicago, IL, p. 2005 (4/7).
- Carraro, L., Gawronski, B., Castelli, L., 2010. Losing on all fronts: the effects of negative versus positive person-based campaigns on implicit and explicit evaluations of political candidates. *Br. J. Soc. Psychol.* 49, 453–470.
- Chou, L., Franke, G.R., White, G.B., 1987. The information content of comparative magazine ads: a longitudinal analysis. *J. Q.* 64 (1), 119–124.
- Collens, J.D., 2011. Comparative advertising: strategy in U.S house elections. In: Paper Presented at the Pre-Conference of the Political Communication Section of the American Political Science Association, Seattle, WA.
- DellaVigna, S., Gentzkow, M., 2010. Persuasion: empirical evidence. *Annu. Rev. Econ.* 2, 643–669.
- Dianoux, C., Herrmann, J., Zeitoun, H., 2013. Comparative advertising: citing the leading brand and its price. *J. Consum. Mark.* 30 (4), 345–354.
- Emons, W., Fluet, C., 2012. Non-comparative versus comparative advertising of quality. *Int. J. Ind. Org.* 30 (2012), 352–360.
- Erev, I., Bornstein, G., Galili, R., 1993. Constructive intergroup competition as a solution to the free rider problem: a field experiment. *J. Exp. Soc. Psychol.* 29, 463–478.
- Faber, R.J., Storey, M.C., 1984. Recall of information from political advertising. *J. Advert.* 13 (3), 39–44.
- Felkner, B., 1966. *Dirty Politics*. W.W. Norton, New York, NY.
- Finkel, S., Geer, J., 1998. A spot check: casting doubt on the demobilizing effect of attack advertising. *Am. J. Polit. Sci.* 42 (2), 573–595.
- Fridkin, K., Kenney, P., 2011. Variability in citizens' reactions to different types of negative campaigns. *Am. J. Polit. Sci.* 55 (2), 307–325.
- Gottfried, J.A., Ben-Porath, E.N., Gibson, J.L., Hall Jamieson, K., 2009. The effects of judicial campaign messages on voter mobilization: an experimental study. In: Paper Delivered at the Annual Conference of the American Political Science Association, September 2009, Toronto.
- Green, J.C., in press. The Dimensions of Disgust: Citizen Attitudes and Codes of Campaign Conduct. White paper issued by Institute for Global Ethics available at (<http://www.globalethics.org/resources/The-Dimensions-of-Disgust/13/>).
- Greene, W.H., 2008. *Econometric Analysis*, sixth ed. Prentice Hall, New Jersey.
- Grewal, D., Kavanoor, S., Fern, E.F., Costley, C., Barnes, J., 1997. Comparative versus noncomparative advertising: a meta-analysis. *J. Mark.* 61 (4), 1–15.
- Grimm, V., Mengel, F., 2011. Let me sleep on it: delay reduces rejection rates in ultimatum games. *Econ. Lett.* 111, 113–115.
- Harmon, R.R., Razzouk, N.Y., Stern, B.L., 1983. The information content of comparative magazine advertisements. *J. Advert.* 12 (4), 10–19.
- Jain, S.P., Lindsey, C., Agrawal, N., Maheswaran, D., 2007. For better or for worse? Valenced comparative frames and regulatory focus. *J. Consum. Res.* 34, 57–66.
- Jain, S.P., Posavac, S.S., 2004. Valenced comparisons. *J. Mark. Res.* 41, 46–58.
- Joslyn, R.A., 1986. Political advertising and the meaning of elections. In: Kaid, Linda L., Dan, Nimmo, Sanders, K.R. (Eds.), *New Perspectives on Political Advertising*. Southern Illinois University Press, Carbondale, IL.
- Kahneman, D., Tversky, A., 1979. Prospect theory An analysis of decision under risk. *Econometrica* 47, 263–291.
- Karlan, D., List, J.A., 2007. Does price matter in charitable giving? Evidence from a large-scale natural field experiment. *Am. Econ. Rev.* 97 (5), 1774–1793.
- Lau, R.R., 1985. Two explanations for negativity effects in political behavior. *Am. J. Polit. Sci.* 29 (1), 119–138.
- Lau, R.R., Sigelman, L., Rovner, I.B., 2007. The effects of negative political campaigns: a meta-analytic reassessment. *J. Polit.* 69 (4), 1176–1209.
- Lau, R.R., Sigelman, L., Heldman, C., Babbitt, P., 1999. The effects of negative political campaigns: a meta-analytic review. *Am. Polit. Sci. Rev.* 93, 851–875.

- Lovett, M.J., Shachar, R., 2011. The seeds of negativity, knowledge and money. *Mark. Sci.* 30 (3), 430–446.
- Maniadis, Z., Tufano, F., List, J.A., 2014. One swallow doesn't make a summer: new evidence on anchoring effects. *Am. Econ. Rev.* 104 (1), 277–290.
- Miller, J.M., Krosnick, J.A., 2004. Threat as a motivator of political activism: a field experiment. *Polit. Psychol.* 25 (4), 507–523.
- Nickerson, D.W., 2008. Is voting contagious? Evidence from two field experiments. *Am. Polit. Sci. Rev.* 102 (1), 49–57.
- Nickerson, D.W., Friedrichs, R.D., King, D.C., 2006. Partisan mobilization campaigns in the field results from a statewide turnout experiment in michigan. *Polit. Res. Q.* 59 (1), 85–97.
- Niven, D., 2006. A field experiment on the effects of negative campaign mail on voter turnout in a municipal election. *Polit. Res. Q.* 59, 203–210.
- Perez-Truglia, R., Cruces, G., 2013. Social Incentives in Contributions Field Experiment Evidence from the 2012 U. S. Presidential Campaigns, Available at (scholar.harvard.edu/rtruglia/research).
- Phillips, J.M., Urbany, J.E., Reynolds, T.J., 2008. Confirmation and the effects of valenced political advertising: a field experiment. *J. Consum. Res.* 34, 794–807.
- Pinkleton, B., 1997. The effects of negative comparative political advertising on candidate evaluations and advertising evaluations an exploration. *J. Advert.* 26, 19–29.
- Shiv, B., Edell, J.A., Payne, J.W., 1997. Factors affecting the impact of negatively and positively framed ad messages. *J. Consum. Res.* 24 (3), 285–294.
- Traugott, M.W., 2008. Validation studies. In: Wolfgang, Donsbach, Traugott, Michael W. (Eds.), *The SAGE Handbook of Public Opinion Research*. SAGE Publications, Thousand Oaks, CA, pp. 408–416.
- U.S. Dept. of Commerce Census Bureau, 1995. 1992 Census of Governments, Vol. 1 No. 2 Popularly Elected Officials, (<http://www.census.gov/prod/2/gov/gc/gc92-1-2.pdf>).
- Vesterlund, L., 2006. In: Steinberg, Richard, Powell, Walter W. (Eds.), *Why Do People Give? in The Nonprofit Sector*, second ed. Yale Press, New Haven, CT, p. 2006.
- Yagci, M.I., Biswas, A., Dutta, S., 2009. Effects of comparative advertising format on consumer responses, the moderating effects of brand image and attribute relevance. *J. Bus. Res.* 62, 768–774.
- Wu, H.D., Dahmen, N., 2010. Web sponsorship and campaign effects assessing the difference between positive and negative web sites. *J. Polit. Mark.* 9, 314–329.