

The Persuasive Effects of Direct Mail: A Regression Discontinuity Based Approach

Alan S. Gerber, Daniel P. Kessler, Marc Meredith

Alan S. Gerber is George C. and Dorathea S. Dilley Professor of Political Science, and Faculty in Residence, Institution for Social and Policy Studies, Yale University, 77 Prospect Street, PO Box 208209, New Haven, CT, 06520-8209 and Faculty Research Fellow, National Bureau of Economic Research (alan.gerber@yale.edu).

Daniel P. Kessler is the David S. and Ann M. Barlow Professor of Management, Professor of Law, Senior Fellow at the Hoover Institution, Stanford University, 434 Galvez Mall, Stanford, CA 94305 and Research Associate, National Bureau of Economic Research (fkessler@stanford.edu).

Marc Meredith is Assistant Professor, Department of Political Science, University of Pennsylvania, 208 S. 37th Street, Philadelphia, PA 19104-6215 (marcmere@sas.upenn.edu).

We thank Daniel Butler, Daniel Hopkins, Greg Huber, Gabriel Lenz, Neil Malhotra, Janelle Wong, and seminar participants at Columbia, Haas, MIT, Stanford, Wharton, and 2008 Midwest Political Science Association Meetings for useful comments and suggestions. Supplementary material available on the Journal of Politics webpage. Replication data available at <http://www.sas.upenn.edu/~marcmere/replicationdata/Persuasive Mail Replication Data.zip>.

Abstract

During the contest for Kansas attorney general in 2006, an organization sent out 6 pieces of mail criticizing the incumbent's conduct in office. We exploit a discontinuity in the rule used to select which households received the mailings to identify the causal effect of mail on vote choice and voter turnout. We find these mailings had both a statistically and politically significant effect on the challenger's vote share. Our estimates suggest that a ten percentage point increase in the amount of mail sent to a precinct increased the challenger's vote share by approximately three percentage points. Furthermore, our results suggest that the mechanism for this increase was persuasion rather than mobilization.

Keywords: political communication, regression discontinuity, voter targeting

Understanding how partisan campaign activity affects the behavior of voters is a central issue in the study of political economy. If campaign activity has an independent effect on election outcomes, then the policy preferences of election-oriented politicians may not perfectly reflect those of the median voter (Baron 1989, Grossman and Helpman 1994). Estimates of the effect of campaign activity are also essential to predicting the impact of reforms to the political process such as restrictions on campaign spending.

Yet, despite an extensive empirical literature devoted to this topic, the magnitude of the impact of campaign activity on voters is still unresolved. For several reasons simple correlations between campaign activity and vote share do not represent a causal effect. Campaign activity and vote share are both outcomes of a complex process that depends on many aspects of candidates and elections that are difficult to measure. For example, if more able candidates attract both more campaign resources and more votes, and candidate ability is not fully observable, then regression estimates of the effect of campaign spending on vote share will inevitably reflect some combination of the true causal effect and unobserved heterogeneity.

In this paper, we use a regression discontinuity (RD) based approach to identify the effects of campaign activity on turnout and vote share. Although previous papers have sought to address the endogeneity of campaign activity, none has used RD to do so. During the contest for Kansas attorney general in 2006, an organization sent out 6 pieces of mail criticizing the Republican incumbent's conduct in office. We obtained a complete record of which households received the mailings as well as the algorithm used to select the universe of households that received the mail. We also obtained precinct-level candidate vote totals, the lowest level of aggregation at which candidate choice is

observed, and individual level voter turnout records. We exploit our knowledge of the selection rule to isolate a discontinuity in the targeting algorithm which resulted in substantially different amounts of mail in otherwise similar precincts. Our identification strategy compares precinct-level vote shares and individual-level turnout decisions in similar precincts that received substantially different amounts of mail as result of this discontinuity. We find that the 6 piece mail campaign caused a sizable increase in the vote share of the Democratic challenger, but had no effect on turnout.

Previous Literature

A number of empirical strategies have been used to deal with the problem of the endogeneity when estimating the effect of campaign activity on political outcomes. One approach seeks to find quasi-experimental variation in observational data. For example, Levitt (1994) identifies the effect of campaign spending on outcomes by looking at how variation in campaign spending relates to election outcomes in cases where the same candidates face each other on repeated occasions. A number of other papers posit instrumental variables that induce variation in spending but are independent of the characteristics of candidates and elections (e.g., Green and Krasno 1988; Gerber 1998; Erikson and Palfrey 2000). Each of these studies is vulnerable on the familiar grounds that the modeling details or exclusion restrictions are not valid.

Another approach is to use randomized field experiments to assess how voters respond to campaign activity. Most of this work has focused on estimating the effects of campaign activity on turnout (e.g., Gerber and Green 2000; Green and Gerber 2008). Although randomized experiments are also well suited to study voter choice from a theoretical standpoint, there are practical difficulties that limit their use. First, the secret

ballot means that individual voters' candidate choices, unlike their turnout, are not a matter of public record. As a result, studies must either randomize mailings at the household-level and measure the effect on vote share using post election surveys (e.g., Gerber 2004; Arceneaux 2007), or randomize campaign activity at the precinct level (e.g. Gerber 2004). Studies of the former type tend to be small due to the cost of administering post-election surveys and vulnerable to bias due to high rates of survey non-response. Studies of the former type are limited by the fact that few campaigns, particularly those in competitive electoral environments, are willing to remove a substantial number of precincts from the campaign's communications efforts. Moreover, the clustered nature of such experiments can substantially increase standard errors (Arceneaux and Nickerson 2009), and therefore increase the cost required to implement experiments with sufficient statistical power.

In this paper, we use a regression discontinuity (RD) based approach to avoid some of these limitations of previous approaches. The exclusion restrictions utilized in RD analyses are often more theoretically justified than other observational instrumental variables. Moreover, while field experiments require fewer identifying assumptions than RD, RD also enjoys several potential advantages in this context. First, RD is less obtrusive. Campaigns may be reluctant to alter their campaign plans to produce control groups, while RD merely requires the campaign to keep track of the rules (and cutoffs) used to determine the campaign targets. This suggests that RD might be applied to a larger and potentially more representative sample of campaigns. Second, RD can be applied historically if campaign records can be obtained. Many campaigns are currently being conducted and treatment assignment rules are held secret. If the selection rules for

the mailings and other campaign activity are preserved, these campaigns can be a source of valuable information about the effects of campaign activity once there is no longer a need to maintain secrecy. This holds for all contests, including competitive races, where political actors might be especially resistant to setting aside control groups.

Data

Our study focuses on the effects of a 6 piece direct mail campaign in the Kansas state attorney general race in the 2006 midterm election. This election featured a Republican incumbent against a Democratic challenger. The mail, which was sponsored by the advocacy group Kansans for Consumer Privacy Protection, informed constituents about the group's concerns regarding the incumbent's conduct in office, suggesting that the incumbent was violating citizens' privacy by "snooping" around peoples' medical records rather than fighting crime. The mailings, which featured pictures of a variety of dogs to illustrate the snooping theme, are included in the supplemental appendix. The mailings were sent every two or three days in the final two weeks before the election to a selected set of households. Households received either one set of mailings or no mailings, regardless of how many registered voters resided there. We consider a voter to have received mail if anyone in his or her household received mail.

We use two sources of data to estimate the effect of these mailings on voter behavior. First, we use precinct-level election returns compiled by the Kansas Secretary of State's office from the state attorney general, gubernatorial, and secretary of state races from the 2002 and 2006 election. Second, we match to this file to precinct-level summaries of an individual-registrant level voter file obtained directly from the mail vendor. The voter file contains all of the variables used by the vendor to select

households to receive mail and the turnout outcome for each household's registered voter in the 2002 and 2006 election. We use these data to construct precinct-level demographic measures of partisan identification, gender, and age. The mail vendor also provided us with information on the exact targeting rule used to select households to receive mail. The vendor sent mail to a voter according to a function of three groups of variables: the background characteristics of each of the registered voters in the household, including party registration, gender, voting history, time since registration, and other demographic characteristics; responses to a phone survey designed to exclude from the mailing registered voters who stated that they were committed to voting for a specific candidate; and characteristics of the census block group in which the household resides.¹ We observe the variables in the first and third group, but do not observe responses to the phone survey. We use these data to construct precinct-level measures of the share of voters in the precinct receiving the mailings.

Simple least-squares estimates of the Democratic challenger's vote share on the proportion of registered voters in a precinct receiving mail are likely to be inconsistent estimates of effect of mail on voting behavior. The probability of receiving mail is positively correlated with voters' background characteristics, phone survey responses, and census characteristics, all of which might also be correlated with voting behavior. However, the targeting rule contains a quasi-experimental element which we exploit for identification in the upcoming sections. Specifically, registrants were significantly more likely to receive mail if they resided in a census block in which more than 5.08131 percent of the households have incomes greater than \$150,000. We refer to this condition as the *income threshold*. We construct a variable $forcing_b$ for each census block b , which

equals the percentage of households in the census block with incomes greater than \$150,000 minus the income threshold. A census block satisfies the income-threshold if it has a value of $forcing_b$ greater than zero.

Only a subset of the registrants' mail status was affected by the income threshold. The income threshold specifically affected the probability that registered Democrats and Libertarians received mail; the targeting rule used for registered Republicans and Independents did not depend on the income threshold. We use knowledge of the targeting formula to construct a *mail eligibility* variable; that is, an indicator for whether Democratic and Libertarian identifiers would receive mail if they lived in a census-block that satisfied the income threshold.²

Table 1 reports descriptive statistics for a selected set of these characteristics. The Secretary of State's office reports election results for 2,711 precincts in the targeted Congressional Districts. After eliminating precincts because the precincts used in the vendor's voter file did not match those used by the Secretary of State's office, because precinct boundaries changed over time, or because individual vote history was unavailable, we are left with a final sample of 1,731 precincts.³

Study Design

The presence of a variable that affects the assignment of mail in a discontinuous manner suggests that we may potentially employ a fuzzy regression discontinuity design to estimate the effect of mail on election outcomes (Hahn, Todd, and Vand der Klaauw, 2001). Briefly, a fuzzy regression discontinuity design is analogous to an instrumental variable (IV) regression where the instrumental variable is whether the forcing variable (e.g., the variable that discontinuously affect treatment status) is just above relative to just

below the discontinuous threshold. The first stage regression estimates how the treatment concentration changes for observations with values of the forcing variable above the threshold. The second stage regression then estimates the effect of this variation in the treatment concentration on outcomes. The exclusion restriction necessary to ensure the consistency of the IV estimate is that, conditional on the observables, observations above and below the discontinuous threshold are similar but for the difference in treatment concentration. To enhance the validity of this exclusion restriction, the analysis is often restricted to observations with values of the forcing variable in a narrow bandwidth around the discontinuous threshold.

Our ideal design would compare electoral outcomes in census blocks with values of $forcing_b$ just above zero with electoral outcomes in census blocks with values of $forcing_b$ just below zero. Unfortunately, the lowest level of aggregation that we can observe election results are precincts, which have boundaries drawn independently of census blocks. As a result, many of the precincts in our data contain registrants from multiple census blocks. This is problematic because the standard methods used to implement regression discontinuity designs are based around treatment status being affected by a single forcing variable (Imbens and Lemieux 2007).

We use two approaches to deal with this issue. First, we focus our analysis on a subset of precincts in our data where we can impose a single precinct-level forcing variable and perform standard regression discontinuity design analysis. Specifically, we restrict our analysis to precincts p in which a majority of registrants reside in a single census block b and set $forcing_p = forcing_b$. Second, we generate a specification that accommodates multiple forcing variables within a precinct that is motivated by

regression discontinuity design. Specifically, we compare precincts that have a given percentage of mail eligible voters living in census blocks just above the income threshold with precincts that have a similar percentage of mail eligible voters living in census blocks just below the income threshold.

Each approach has relative advantages and disadvantages. The analysis imposing a single forcing variable allows us to perform the standard graphical analysis used in the RD literature. It also allows for the use of the local linear regression techniques highlighted by Imbens and Lemieux (2007). However, we are able to incorporate additional information using the specification that accommodates multiple forcing variables, which provides us with significantly more statistical power. Finding similar results using both approaches helps assuage concerns about either approach individually.

Single Forcing Variable Analysis

In this section, we focus on the subset of precincts in which a majority of registrants reside in a single census block. 1,169 of the 1,731 precincts in our dataset satisfy this property. We assign $forcing_p$ in precinct p equal to $forcing_b$ in the majority census block group b . Because multiple precincts can have a majority of registrants from the same census block, we cluster standard errors by primary census block.

Figure 1a displays how the share of registrants in a precinct receiving mail varies with the value of $forcing_p$. It illustrates a discontinuous jump in the share of registrants in mail household in precincts above the income threshold. As Table 2 indicates, about 17.0 percent of registrants received mail in precincts with values of $forcing_p$ between zero and 1.57.⁴ In comparison, only about 7.6 percent of registrants received mail in precincts with values of $forcing_p$ between -1.57 and zero.

Figure 1b shows how the share of voters supporting the Democratic attorney general and gubernatorial candidates in 2006 varies with the value of $forcing_p$ in the precinct. It shows that the Democratic attorney general candidate's vote share in 2006 was greater in precincts above the income threshold. As Table 2 indicates, the Democratic attorney general candidate's votes share was 56.2 percent in precincts with values of $forcing_p$ between zero and 1.57. In comparison, the Democratic attorney general candidate's vote share in 2006 in precincts with values of $forcing_p$ between -1.57 and zero was 53.5 percent. The single-difference estimate of the effect of mail is thus

$$D = 56.2 - 53.5 = 2.7 \text{ percentage points (standard error 2.3).}$$

Because we are using a quasi-experiment rather than a randomized experiment, we cannot be certain that having a value of $forcing_p$ above the income threshold is orthogonal to unobservable determinants of vote share. We therefore develop an estimation framework that attempts to isolate the effect of mail from any pre-treatment differences in those precincts containing census blocks just above versus just below the income threshold. One way to control for pre-treatment differences is by making the same comparison for untreated races held on same day. Figure 1b shows the Democratic gubernatorial candidate's vote share was similar in precincts above and below the income threshold. Table 2 indicates that in 2006 the Democratic gubernatorial candidate received 56.0 percent of the vote in precincts with values of $forcing_p$ between zero and 1.57, as compared to 55.7 percent in precincts with values of $forcing_p$ between -1.57 and zero. If we assume that the difference in Democratic performance observed in the governor race is also what would be observed in the attorney general race absent any mailings, the difference-in-difference (DD) estimate of the causal effect of mail is:

$$DD = ((55.3 - 53.1) - (56.0 - 55.7)) = 2.5 \text{ percentage points (standard error 1.2).}^5$$

The DD estimate allows voter preferences for Democrats in the precincts just above and just below the income threshold to differ. However, it assumes that the underlying preferences for Democratic attorney general and gubernatorial candidates are the same. This assumption may be problematic if precincts just above the income threshold have a differential preference for a Democratic attorney general relative to a Democratic governor. To investigate this possibility, Figure 1c compares the Democratic vote share of the attorney general relative to the governor in 2002 and 2006 in precincts above and below the income threshold. Figure 1c shows some evidence of a relative preference for a Democratic attorney general candidate above the income threshold in 2002. Table 2 indicates that the Democratic attorney-general and gubernatorial candidates received 47.1 and 51.8 percent of vote respectively in precincts with values of $forcing_p$ between zero and 1.57, and 46.3 and 52.1 percent of the vote respectively in precincts with values of $forcing_p$ between -1.57 and zero. The resulting pre-treatment difference-in-difference (DD') estimate for 2002 is:

$$DD' = ((47.1 - 46.3) - (51.8 - 52.1)) = 1.2 \text{ percentage points (standard error 1.6).}$$

If we assume that the preference for the Democratic attorney-general candidate relative to the governor observed in 2002 is also what would be observed in 2006 absent any mailings, the difference-in-difference-difference (DDD) estimate of the causal effect of mail is:

$$DDD = DD - DD' = ((55.3 - 53.1) - (56.0 - 55.7)) - ((47.1 - 46.3) - (51.8 - 52.1)) = 1.3 \text{ percentage points (standard error 1.7).}^6$$

This effect is quite substantial in political terms. Given that the difference in the share of households receiving mail in the two sorts of precincts is about 9.5 percentage points ($= 0.170 - 0.076$), this implies the estimated average treatment-on-the-treated effect of mail is about 29.0, 26.1, and 13.8 percentage points using the difference, DD, and DDD estimates respectively, although only the DD estimate is statistically significant at conventional levels.

Table 3 suggests that the estimates of the effect of mail from Table 2 are not unique to a particular “control” office. This might occur, for example, if the relationship between income and preference for Democratic candidates were stronger in up-ballot races such as governor as compared to down-ballot elections like attorney-general or secretary of state. In both the 2002 and 2006 elections, the vote share of the Democratic secretary of state candidate was about 1.5 percentage points smaller in precincts just above versus just below the income threshold. This suggests replicating the above analysis using vote shares from the election for secretary of state, rather than vote shares from the election for governor, would increase our estimated effect sizes. This is notable because the exact same two candidates competed for Secretary of State in 2002 and 2006.

Table 3 also shows that registrants’ observable characteristics look similar in precincts above and below the income threshold. Registrants’ partisan identification is nearly identical with registrants from precincts above the income threshold being 0.5 percentage points (standard error 2.0) more likely to be registered Republican and 0.1 percentage points (standard error 1.5) less likely to be Democratic. Registrants from precincts above the income threshold are a bit older, and slightly more likely to be male or reside in a mail eligible household. Given the relative balance, it is not surprising that

unreported regressions show similar results to what is observed in Table 2 if we include these observable characteristics as controls.

Two concerns with the above analysis are that there is no justification for the selection of the bandwidth around the income threshold and that we did not control for any differential effect of $forcing_p$ across the different elections. To address these concerns we also apply the local linear regression methods highlighted by Imbens and Lemieux (2007) to our data. The local linear method assumes that the function mapping a precinct's value of $forcing_p$ into electoral outcomes is approximately locally linear. The realization of this function for any given value of $forcing_p$ is estimated by running a regression of electoral outcomes on $forcing_p$ using only data with values of $forcing_p$ in narrow bandwidth around the given value. The idea is to compare our estimates of the function for values of $forcing_p$ just above and just below the income threshold. In practice, this is done by estimating the value of the function at the income threshold twice; once only using data from precincts with values of $forcing_p$ below zero, and once only using data from precincts with values of $forcing_p$ above zero. The difference between these two estimates of the function value at zero is the estimated effect resulting from the treatment change at the income threshold.

A critical component of the above analysis is selecting the bandwidth around the income threshold used to estimate the values of function when $forcing_p$ equals zero. We use the cross-validation procedure proposed by Ludwig and Miller (2005) that is detailed in the supplemental appendix. Briefly, the cross-validation procedure selects the bandwidth that performs the best in an out-of-sample predictive test. This procedure selects a bandwidth of 1.57 when using a rectangle kernel.

Table 4 presents the point estimates and standard errors for these estimates.⁷ It indicates that there is about an 11.1 percentage point difference in mail concentration above the income threshold relative to below. The difference point estimate of 1.6 percentage points (standard error = 4.2 percentage points) is slightly smaller than we observed in Table 2. In contrast, the DD point estimate of 5.3 percentage points (standard error = 2.4 percentage points) and the DDD point estimate of 6.2 percentage points (standard error = 3.5 percentage points) are larger than we observed in Table 2.⁸

The point estimates reported in Table 2 and Table 4 are consistent with their being a sizeable effect of mail on election outcomes. However, they are estimated with sufficient uncertainty that, for the most part, we cannot rule out that they are different from zero at standard levels of statistical significance. This highlights a limitation of the single forcing variable analysis. Our analysis in Table 2 and Table 4 is limited to 104 precincts in 78 census blocks in which a majority of registrants reside in a census block with $|forcing_b| < 1.57$. However, 382 precincts in 121 census blocks have at least one mail eligible voter living in a census block with a value of forcing in this range. Thus, our analysis in this section is ignoring a substantial amount of information. In the next section, we develop an alternative approach that allows us to incorporate this information in our analysis while keeping our specification in the spirit of a RD design.

Econometric Model and Results

Model

The results above suggest that receipt of mail may have had large effect on voter behavior in the 2006 Kansas attorney-general race. In this section we develop an econometric model to formalize the assumptions necessary to identify this effect. We

model the share of votes received by the Democratic candidate in race $r =$ (attorney general, governor) or (attorney general, secretary of state) and precinct $p = (1, \dots, P)$ at year $t = (2002, 2006)$, S_{rpt} . Vote share S_{rpt} depends on (race \times year)-specific constant terms⁹; the share of voters in the precinct who received mail, m_p ; a vector function $g_p(\Omega)$ of the characteristics Ω of voters and census blocks contained in a precinct; interactions between race, year and m_p and $g_p()$; and an error term ε_{rpt} :

$$\begin{aligned}
S_{rpt} = & \alpha_0 + \alpha_1 m_p + g_p(\Omega) \alpha_2 + \\
& AG_{rt} (\beta_0 + \beta_1 m_p + g_p(\Omega) \beta_2) + \\
& YEAR2006_t (\delta_0 + \delta_1 m_p + g_p(\Omega) \delta_2) \\
& AG_{rt} * YEAR2006_t (\phi_0 + \phi_1 m_p + g_p(\Omega) \phi_2) + \varepsilon_{rpt}
\end{aligned} \tag{1}$$

For each of the N_p registered voters in a precinct, the matrix Ω contains 5 variables: whether their household is mail eligible, the voter's partisan identification, the voter's gender, the voter's age, and the value of $forcing_b$ in the voter's census block. The vector function $g_p(\Omega)$ returns 24 variables defined at the precinct level: the share of voters who are mail eligible, registered Democratic, registered Republican, female, age 18 to 29, age 30 to 44, age 45 to 64 X_p ; a cubic polynomial in the precinct-level approximation of the

average level of the census-block income, with $I_p = \frac{\sum_{\text{census blocks } b} R_{pb} * forcing_b}{\sum_b R_{pb}}$ where R_{pb} is the

number of registered voters in precinct p who live in census block b ; interaction terms between X_p and this cubic polynomial; and three additional variables $g1_p$, $g2_p$, and $g3_p$. The variables $g1_p$, $g2_p$, and $g3_p$ partition the precinct level variable *mail eligible* share into sections defined by census-block income to capture nonlinearities in the relationship between census-block income and mail eligibility. They measure the share of voters

from high- and moderate-income census blocks who are mail eligible (the omitted category is of voters from low-income census blocks who are mail eligible), as well as mail eligible voters who were missing census data:¹⁰

$g1_p$ = share of voters who are from a census block with $forcing_b > s$ AND are mail eligible,

$g2_p$ = share of voters who are from a census block with $-s < forcing_b < s$ AND are mail eligible,

and

$g3_p$ = share of voters who are missing census data AND are mail eligible,

where s is a bandwidth term.

For the reasons discussed above, simple OLS estimates of (1) are likely to be inconsistent: m_p explicitly depends on $forcing_p$ and implicitly on ε_{rpt} (because of the unobserved phone survey). We therefore write m_p as a function of a constant term θ_0 , a vector function $h_p(\cdot)$, and an error term ω_p :

$$m_p = \theta_0 + h_p(\Omega)\theta_2 + \omega_p.$$

To estimate (1) by instrumental variables, we impose the exclusion restriction that mail-eligible voters from census blocks where $forcing_b$ is just above zero have the same propensity to vote for a Democratic candidate as mail-eligible voters from census blocks where $forcing_b$ is just below zero. We implement this identification strategy with the following specification of $h(\cdot)$:

$$h(\Omega) = [g(\Omega) \mid h1_p],$$

where

$h1_p$ = share of voters who are from a census block group with $0 < forcing_b < s$ AND are mail eligible.

The vector functions $g(\cdot)$ and $h(\cdot)$ are specified to make this exclusion restriction as weak as possible. By including a cubic polynomial of I_p in $g(\cdot)$, we attempt to capture any possible direct effects of census-block income on voters' propensity to choose Democratic candidates. Interacting the cubic polynomial of I_p with X_p in $g(\cdot)$ allows the effects of the covariates, most importantly mail eligibility, to vary with the income-level of the precinct.¹¹

The inclusion of gI_p and $g2_p$ in $g(\cdot)$ allows for the possibility that mail eligibility and census-block income interact to affect the probability of voting Democratic in some way that is not captured by the precinct-level interactions between where I_p and X_p . By including gI_p and $g2_p$, identification is reduced to the condition that after controlling for X_p , the cubic polynomial of I_p , and the interaction of X_p and the cubic polynomial of I_p , the interactive effect between census-block income and mail eligibility is the same in census blocks with values of $forcing_b$ between $-s$ and s . The exclusion of hI_p from $g(\cdot)$ only assumes that mail-eligible voters from census blocks with values of $forcing_b$ between 0 and s have the same propensity to vote Democratic, conditional on observables, as mail-eligible voters from census blocks with values of $forcing_b$ between $-s$ and 0. We vary the size of the window s to investigate the sensitivity of our results to the scope of our exclusion restriction.

We take a conservative approach to estimating standard errors when estimating equation (1) to account for both between and within precinct correlation of the error term ε_{ipt} . The between precinct correlation occurs because the income of a single census block affects the concentration of mail in multiple precincts, while the within precinct correlation results from using four observations from each precinct. To account for the

resulting non-independence of ε_{rpt} , we find the census block with a value of $forcing_b$ between $-s$ and s that has the most registrants in precinct p and cluster the standard errors for all precincts that share this maximal census block. If a precinct has no registrants with a value of $forcing_b$ between $-s$ and s , we cluster the observation within the precinct.

Results

Table 5 presents estimates of the effect of the proportion of registered voters receiving mail on Democratic vote shares for attorney general and the control race in 2002 and 2006 using a bandwidth of 1.57. The table reports estimates of four parameters: the effect of mail on the vote share in the control race in 2002, α_1 ; the differential effect of mail on the attorney general race in 2002, β_1 ; the differential effect of mail on the control race in 2006, δ_1 ; and the differential effect of mail on the attorney general race in 2006, φ_1 . These four parameters can be combined to construct the three different estimates of the effects of mail. The value of $\alpha_1 + \beta_1 + \delta_1 + \varphi_1$ is the IV analogue to the difference estimate of the effect of mail in Table 2. Similarly, $\beta_1 + \varphi_1$ is the IV analogue to the DD estimate of the effect of mail in Table 2. Finally, φ_1 is the IV analogue of the DDD estimate from Table 2.

Table 5 shows robust evidence that the receipt of mail increased the Democratic attorney-general candidate's vote share in 2006. Our baseline difference estimate reported in column (1) is 0.151 (standard error 0.145), indicating that a 10 percentage point increase in mail increases the vote share by 1.5 percentage points. Our DD and DDD estimates increase to 0.260 (standard error 0.145) and 0.270 (standard error 0.144) respectively. The nearly identical DD and DDD estimates indicate the increased

performance of the Democratic attorney general relative to the governor was unique to 2006, and that these candidates performed almost equally in 2002.

The remaining columns of Table 5 present a number of additional specifications to show the robustness of the baseline result. Column (2) uses the Secretary of State race as the baseline instead of the governor race. The DD estimate of 0.219 (standard error 0.162) and the DDD estimate of 0.243 (standard error 0.110) are similar to the estimates when the governor race is used as the baseline. Column (3) truncates the sample to include only those precincts with moderate values of I_p , which allows us to estimate the income polynomial and interactions with the covariates using data in the same range that we identify the effects of mail. All of our estimates of the effects of mail increase using this specification.¹² Column (4) drops all of the covariates from equation (1) except $g2_p$ (e.g., the percent mail eligible and living in census blocks with values of $forcing_b$ between -1.57 and 1.57). We find a smaller difference estimate and larger DD and DDD estimates. Column (5) compares the estimate in column (4) to an estimate that uses the share of the precinct living in census blocks with values of $forcing_b$ between -1.57 and 1.57 irrespective of mail eligibility status. This estimate is the continuous treatment analog to the results reported in Table 2. We find similar patterns to what we observed in Table 2, except that the DDD estimate is larger.

Figure 3a, Figure 3b, and Figure 3c illustrate the stability of the results reported in our baseline specification as we vary the bandwidth s . Figure 3a presents the point estimates and 90 percent confidence intervals on the difference estimates. The figure indicates that we consistently find a point estimate of about 0.2 across bandwidths and that these estimate generally a bit shy of significance at the 10 percent level, two-tailed.

Figure 3b and Figure 3c present the DD and DDD estimates and 90 confidence intervals respectively. The DD results tend to be slightly larger than the difference results, particularly when the governor is used as the benchmark race. When the governor race is used as the benchmark the estimates are consistently statistically significant at the 10 percent level, two-tailed, while the results using the secretary of state race as the benchmark are only significant at this level using some larger bandwidths. Nearly identical point estimates are found when using the DDD specification, albeit with smaller standard errors making the results significant at the 90 percent level, two tailed, using either race as the benchmark for almost any bandwidth.

Identifying the Mechanism for the Change in Vote Share

In the previous section, we showed that a direct mail campaign that criticized the Republican incumbent attorney general increased voters' propensity to vote for the Democratic challenger for that office relative to their propensity to vote for the Democratic candidate for governor. There could be two possible reasons for such a causal effect. The first is that receiving the direct mail persuaded individuals who were already going to turnout to switch for whom they voted. The second is that receiving the direct mail persuaded individuals who supported the Democratic attorney general candidate to turnout to vote. In this section, we investigate whether there is any evidence for the second possibility.

Investigating the effect of mail on turnout is more straightforward than estimating the effect of mail on vote share because we observe at the individual level both whether an individual received mail and whether s/he turned out to vote. Figure 3a shows how the percentage of registrants receiving mail varies with the value of *forcing_b* and the

household mail eligibility status. Table 6 shows that 75.0 percent of the mail eligible registrants in census blocks with values of $forcing_b$ between 0 and 1.57 received mail, compared to only 9.3 percent of mail eligible voters in census blocks with values of $forcing_b$ between -1.57 and 0. In contrast, there is almost no difference in the amount of mail received by not mail eligible households in census blocks above and below the income threshold.

There is no evidence that receiving mail substantially increased turnout rates. Figure 3b shows a slight increase in the turnout rates of mail eligible registrants living in census blocks above the income threshold. However, Figure 3c also shows that the turnout rates of registrants who were not mail eligible also increased above the income threshold. Table 6 indicates that while mail eligible registrants' turnout rates in 2006 were 1.8 percentage points (standard error 1.6) higher in census blocks with values of $forcing_b$ between 0 and 1.57 received mail compared to census block with values of $forcing_b$ between -1.57 and 0, registrants not mail eligible turnout rates were 2.0 percentage points higher (standard error 2.1). This implies a DD effect of the mail threshold on turnout rates of -0.2 percentage points (standard error 2.2). Comparing this to the patterns observed pre-treatment in 2002 implies a DDD effect of 0.5 percentage points (standard error 1.2).

While we cannot rule out that the mailings had a small mobilizing effect, we can rule out a mobilizing effect that can explain the findings in the previous section. There, a 10 percent increase in the proportion of registered voters receiving mail led to an increase in vote share of approximately 3 percentage points. Even a mobilizing effect equal to 5 percentage points, an upper bound of the 95 percent confidence interval of the estimate of

the difference effect in column (3), can explain only a very small portion of the estimated effect of mail on vote share. At least for this specific mailing, the effect on vote share appears to come through persuasion of individuals who were already going to vote.¹³

Conclusion

To explore the importance of campaign activity on political outcomes, we estimate the effect of a series of direct mailings on voter turnout and candidate choice in the 2006 Kansas attorney general race. We find that mailings sent criticizing the Republican incumbent had both a statistically and politically significant effect on his vote share. Our estimates suggest that a ten percentage point increase in the amount of mail sent to a precinct increased the Democratic challenger's vote share by about three percentage points. Furthermore, we find no evidence that these mailings affected turnout. As a result, we conclude that these mailings persuaded individuals who were already going to turnout to switch for whom they voted.¹⁴

These effects are quite large. By comparison, in a meta-analysis of field experiments, Green and Gerber (2008) find that several pieces of direct mail increase turnout by about one percentage point. There is only limited evidence on the persuasive effects of direct mail campaigns on precinct-level vote shares, but previous studies find much smaller effects than those reported in tables 2 and 4 (Gerber 2004).¹⁵ A number of factors might account for this difference. First, these are estimates and one standard deviation is approximately 40%-50% of the size of the measured effect. Although we reject the hypothesis that the true effect is zero, we can not reject the hypothesis that the true effect is materially lower than the point estimate. Second, the particular race we study is a down-ballot race; it was not the primary race mobilizing voters to the polls.

Direct mail likely has a larger potential effect in such environment than in a presidential race where voters are much better informed about the issues. Third, the influence of the mailings may have been affected by the fact that the mailings themselves became news. Many of the largest Kansas newspapers, including the Kansas City Star, Lawrence Journal World, and Topeka Capital Journal ran stories about the mailings. It is possible that there exist complementarities between receipt and the coverage of the mailings. Fourth, the local-average treatment effect we estimate only applies to so-called "mail eligible" voters -- namely, those who were predetermined by the vendor to be particularly susceptible to persuasion; mail almost surely would be less effective in the population at large. Finally, the magnitude of our estimate might also be evidence of spillover effects, or social interaction among voters. Our treatment effects are estimated at the precinct level and may reflect changes in households that were not send mail but live in a precinct where the mailings were concentrated.

The loss of statistical power from aggregation makes it difficult to precisely estimate the effect of a single mailing. As a result, a more general understanding of the persuasive effects of campaign communication will likely require a meta-analysis of the estimates from many races. Although our study is of a single direct mail campaign, the method we propose for analyzing campaign communications can be applied to other targeting formula with geographic discontinuities. Because this method can be applied retroactively, it may permit scholars to analyze elections that have already occurred. Moreover, RD is likely to be acceptable to a broader set of campaigns than randomized experimentation, which requires setting aside randomized control groups. In addition, ethical and legal constraints may prevent scholars from conducting field experiments

using many forms of political communication used by campaigns. Thus, we believe the analytical methods we develop in this paper will complement field experiments for learning how and when voter behavior is affected campaign activity.

References

- Arceneaux, Kevin. 2007. "I'm Asking for Your Support: The Effects of Personally Delivered Campaign Messages on Voting Decisions and Opinion Formation." *Quarterly Journal of Political Science* 2: 43-65.
- Arceneaux, Kevin and David W. Nickerson. 2009. "Modeling Certainty with Clustered Data: A Comparison of Methods." *Political Analysis* 17(2): 177-190.
- Baron, David P. 1989. "Service-Induced Campaign Contributions and the Electoral Equilibrium." *Quarterly Journal of Economics* 104(1): 45-72.
- Erikson, Robert S. and Thomas R. Palfrey. 2000. "Equilibria in Campaign Spending Games: Theory and Evidence." *American Political Science Review* 94(3): 595-609.
- Gerber, Alan. 1998. "Estimating the Effect of Campaign Spending on Senate Election Outcomes Using Instrumental Variables." *American Political Science Review* 92(2): 401-411
- Gerber, Alan S. 2004. "Does Campaign Spending Work?: Field Experiments Provide Evidence and Suggest New Theory." *American Behavioral Scientist* 47(5):541-574.
- Gerber, Alan S. and Donald P. Green. 2000. "The Effects of Canvassing, Telephone Calls, and Direct Mail on Voter Turnout: A Field Experiment." *American Political Science Review* 94(3): 653-663.

- Green, Donald P. and Alan S. Gerber. 2008. *Get Out the Vote!: How to Increase Voter Turnout*. 2nd Ed. Washington, DC: Brooking Institution.
- Green, Donald. P. and Johnthan S. Krasno. 1988. "Salvation for the Spendthrift Incumbent." *American Journal of Political Science* 32(4): 884-907.
- Grossman, Gene and Ethan Helpman. 1994. "Protection for Sale." *American Economic Review* 84(4): 833-850.
- Hahn, Jinyong, Petra Todd, and Wilbert Vand der Klaauw. 2001. "Identification and Estimation of Treatment Effects with a Regression Discontinuity Design." *Econometrica* 69(1): 201-209.
- Imbens, Guido W. and Thomas Lemieux. 2007. "Regression Discontinuity Designs: A Guide to Practice." *Journal of Econometrics* 142(2): 615-635.
- Levitt, Steven D. 1994. "Using Repeat Challengers to Estimate the Effect of Campaign Spending on Election Outcomes in the House". *Journal of Political Economy*, 102(4): 77-98.
- Ludwig, Jens and Douglas L. Miller. 2005. Does Head Start Improve Children's Life Chances? Evidence From a Regression Discontinuity Design. *NBER Working Paper* No. 11702.

¹ Census block groups typically contain 600 to 3000 people, with an optimal size of 1500.

It is the lowest level of aggregation for which much census data are publically available.

For brevity, we refer to these as census blocks in the paper.

² Phone survey responses do not affect the mail eligibility variable because we do not observe all responses. Therefore, not all mail eligible households in census blocks above the income threshold receive mail.

³ Of the excluded precincts, 293, covering 8.1 percent of registrants, are dropped because the precincts in the vendor's file did not match those used by the Secretary of State's office. 383 more precincts, covering 25.8 percent of registrants, are excluded because precinct boundaries change between 2002 and 2006. An additional 109 precincts, covering 2.8 of registrants, are excluded because we don't have either 2002 or 2006 vote history. Finally in 265 precincts we need to consolidate two or more precincts into a single observation in order to match between the vendor's file and the Secretary of State's files. These 265 precincts are collapsed down to 70 observations.

⁴ We justify the bandwidth of 1.57 later in this section. In the subsequent section, we show our results are robust to the choice of bandwidth.

⁵ A statistical model of vote outcomes that can support this assumption is that the vote share for a Democratic candidate in a precinct equals an office invariant normal Democratic vote plus an additive race specific i.i.d random valence shock.

⁶ A statistical model of vote outcomes that can support this assumption is that the vote share for a Democratic candidate in a precinct equals a time invariant and office specific normal Democratic vote plus an additive race specific i.i.d random valence shock.

⁷ Figures A4-A7 in the supplemental appendix display the point estimates of the local linear function mapping $forcing_p$ to electoral outcomes.

⁸ We ran a number of variants of these regressions. These include specifications that constrain slopes to be equal on either side of the income threshold, include covariates, and use alternative bandwidths. The difference estimates are sensitive to specification. The DD and DDD estimates are less sensitive to specification, although the point estimates decrease using larger bandwidths.

⁹ α_0 , β_0 , δ_0 , and φ_0 are each five-element vectors: a constant term and dummy variables for the home county and home media markets of the Republican and Democratic candidates. We include fixed effects for the home counties and media markets to account for hometown candidate preferences that may cause differentials within precincts in Democratic vote shares across the two races.

¹⁰ About 9 percent of the sample was missing census data. A mail eligible registrant with missing demographic data received mail.

¹¹ Results are robust to the inclusion of higher-order polynomials.

¹² We truncate the sample by finding the 5th and 95th percentile of I_p in precincts with positive values of $g2_p$. In this case, this results in bounds of -4.11 and 2.60 on I_p .

¹³ In the supplemental appendix we also present estimates of precinct-level mail share on precinct-level turnout, which are estimated using a model similar to equation (1). While we would not conclude that mail had a statistically significant effect on turnout in the aggregate analysis, we also would not be able to rule out many effect sizes, illustrating the statistical costs of using aggregated data.

¹⁴ One concern is that other campaign activity may also be affected by the income threshold. It is illegal for the sponsoring group of issue ads to coordinate its targeting with other campaign actors. The vendor performing the targeting assured us that no coordination occurred. Thus, it is unlikely that other campaign activity was affected by whether a household resided in a census block above the income threshold.

¹⁵ Arceneaux (2007) finds estimates of a similar magnitude on post-contact surveys following randomized persuasive canvassing and phone calls in a local primary election.

Table 1
Descriptive Statistics, Voting Precincts, 2002 and 2006

Variable	Mean	Std. Dev.
Percentage of voters:		
voting for Democratic Attorney General in 2006	0.533	0.093
voting for Democratic Governor in 2006	0.554	0.099
voting for Democratic Secretary of State in 2006	0.283	0.107
voting for Democratic Attorney General in 2002	0.483	0.108
voting for Democratic Governor in 2002	0.532	0.105
voting for Democratic Secretary of State in 2002	0.282	0.101
Percentage of registrants:		
registered Democrat	0.256	0.088
registered Republican	0.516	0.124
female	0.534	0.031
age 18 - 29	0.160	0.056
age 30 - 44	0.219	0.047
age 45 - 64	0.375	0.059
mail eligible household	0.270	0.089
mailed household	0.099	0.074
Number of voters (unweighted)	203.61	252.75

Note: N = 1731 precincts. Observations weighted by total votes in 2006 Attorney General Race.

Table 2: Democratic Candidates' Vote Shares in 2006 and 2002 in Precincts Just Above versus Just Below the Income Threshold for Receiving Mail in 2006

	Share Mailed	Vote Share for Democrat in 2006		Vote Share for Democrat in 2002	
		AG	GOV	AG	GOV
	(1)	(2)	(3)	(4)	(5)
Precincts with $0 < Forcing_p < 1.57$	0.171	0.562	0.560	0.471	0.518
Precincts with $-1.57 < Forcing_p < 0$	0.076	0.535	0.557	0.463	0.521
Difference in Vote Shares		0.027 (0.023)	0.003 (0.025)	0.008 (0.027)	-0.003 (0.025)
Difference-in-Difference in Vote Shares		0.025 (0.012)		0.012 (0.016)	
Difference-in-Difference-in-Difference in Vote Shares			0.013 (0.017)		

Note: 42 precincts with $0 < Forcing_p < 1.57$ in 31 census blocks, 62 precincts with $-1.57 < Forcing_p < 0$ in 47 census blocks. Standard errors clustered by census block.

Table 3: Other Characteristics of Precincts Just Above and Just Below Income Threshold for Receiving Mail in 2006

Value of Forcing in Precinct	$0 < Forcing_p < 1.57$	$-1.57 < Forcing_p < 0$	Difference
voting for Democratic Sec'y of State in 2006	0.273	0.288	-0.015 (0.020)
voting for Democratic Sec'y of State in 2002	0.262	0.277	-0.015 (0.018)
registered Democrat	0.236	0.236	-0.001 (0.013)
registered Republican	0.541	0.536	0.005 (0.020)
Female	0.521	0.531	-0.010 (0.006)
age 18 – 29	0.146	0.154	-0.007 (0.011)
age 30 – 44	0.214	0.229	-0.015 (0.013)
age 45 – 64	0.416	0.393	0.023 (0.016)
mail eligible household	0.225	0.264	-0.038 (0.026)

Note: 42 precincts with $0 < Forcing_p < 1.57$ in 31 census blocks, 62 precincts with $-1.57 < Forcing_p < 0$ in 47 census blocks. Standard errors clustered by census block.

Table 4: Fitted Democratic Candidates' Vote Shares in 2006 and 2002 from Local Linear Regression at the Income Threshold

	Share Mailed	Vote Share for Democrat in 2006		Vote Share for Democrat in 2002	
		AG	GOV	AG	GOV
	(1)	(2)	(3)	(4)	(5)
Value of Local Linear Regression Function when $Forcing_p = 0+$	0.168	0.558	0.545	0.479	0.538
Value of Local Linear Regression Function when $Forcing_p = 0-$	0.057	0.542	0.582	0.494	0.544
Difference in Vote Shares		0.016 (0.042)	-0.037 (0.044)	-0.015 (0.052)	-0.005 (0.045)
Difference-in-Difference in Vote Shares		0.053 (0.024)		-0.010 (0.029)	
Difference-in-Difference-in-Difference in Vote Shares			0.062 (0.035)		

Note: Estimated using a rectangle kernel with a bandwidth of 1.57. Standard errors clustered by census block.

Table 5: IV Effect of Mail on Vote Share

	(1)	(2)	(3)	(4)	(5)
% registered voters receiving mail in 2006 (α_1)	0.030 (0.143)	0.044 (0.089)	0.083 (0.140)	-0.228 (0.216)	-0.034 (0.258)
% registered voters receiving mail in 2006* attorney general race (β_1)	-0.010 (0.105)	-0.024 (0.152)	-0.003 (0.117)	0.040 (0.115)	0.040 (0.139)
% registered voters receiving mail in 2006* year=2006 (δ_1)	-0.139 (0.115)	-0.113 (0.082)	-0.221 (0.116)	-0.100 (0.121)	-0.131 (0.153)
% registered voters receiving mail in 2006* year=2006 * attorney general (φ_1)	0.270 (0.144)	0.243 (0.110)	0.324 (0.151)	0.323 (0.146)	0.302 (0.164)
Difference Estimate ($\alpha_1 + \beta_1 + \delta_1 + \varphi_1$)	0.151 (0.140)	0.151 (0.139)	0.184 (0.153)	0.035 (0.173)	0.176 (0.204)
Difference-in-Difference Estimate ($\beta_1 + \varphi_1$)	0.260 (0.145)	0.219 (0.162)	0.321 (0.167)	0.363 (0.122)	0.342 (0.124)
Control race	Governor	SS	Governor	Governor	Governor
Number of Observations	1731	1731	1060	382	477
Covariates	Yes	Yes	Yes	No	No
Instruments:					
% $0 < forcing_b < s$ & Mail Eligible	X	X	X	X	
% $0 < forcing_b < s$					X
F-Statistic on First Stage Instruments	57.33	57.39	45.91	30.77	19.09

Note: All estimates using bandwidth $s = 1.57$. Standard errors clustered as described in text.

Table 6: Turnout Rates in 2006 and 2002 in Census Block Groups
Just Above versus Just Below the Income Threshold for Receiving Mail in 2006

	Share Mailed		Turnout Rate in 2006		Turnout Rate in 2002	
	Mail Eligible	Not Mail Eligible	Mail Eligible	Not Mail Eligible	Mail Eligible	Not Mail Eligible
	(1)	(2)	(3)	(4)	(5)	(6)
Census Blocks Groups with $0 < Forcing_b < 1.57$	0.750	0.031	0.692	0.602	0.671	0.574
Census Blocks Groups with $-1.57 < Forcing_b < 0$	0.093	0.039	0.675	0.582	0.670	0.566
Difference in Turnout Rates			0.018 (0.016)	0.020 (0.021)	0.001 (0.019)	0.007 (0.022)
Difference-in-Difference in Turnout Rates			-0.002 (0.022)		-0.007 (0.022)	
Difference-in-Difference-in-Difference in Turnout Rates					0.005 (0.012)	

Note: 6297 mail eligible and 21146 not mail eligible in 47 census blocks b with $0 < Forcing_b < 1.57$. 12317 mail eligible and 34071 not mail eligible in 89 census block groups b with $-1.57 < Forcing_b < 0$. Standard errors clustered by census block.




Figure 1a: Mail Concentration in 2006 by Majority Census Block Group Income

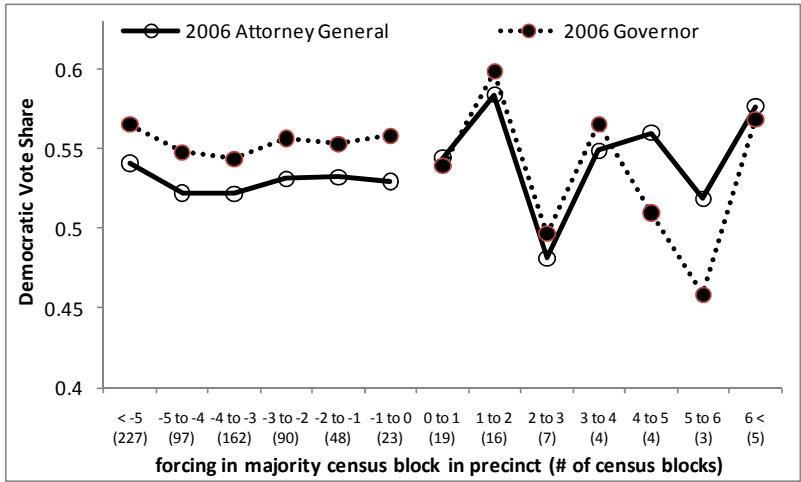


Figure 1b: Democratic Vote Share in 2006 by Majority Census Block Group Income




Figure 1c: Difference in Democratic Attorney General and Governor's Vote Share in 2002 and 2006 by Majority Census Block Group Income




Figure 2a: Difference Estimate on Vote Share by Bandwidth




Figure 2b: Difference-in-Difference Estimate on Vote Share by Bandwidth

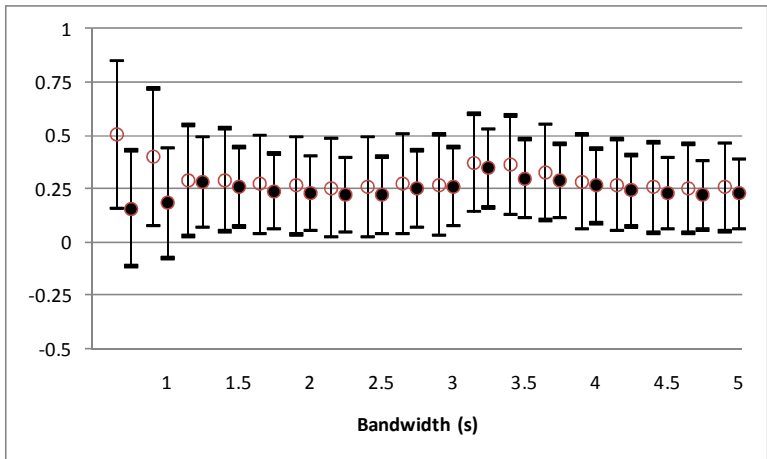


Figure 2c: Difference-in-Difference-Difference Estimate on Vote Share by Bandwidth

White Circles Relative to Governor, Black Circles Relative to Secretary of State. Black bars represent 90 percent confidence intervals

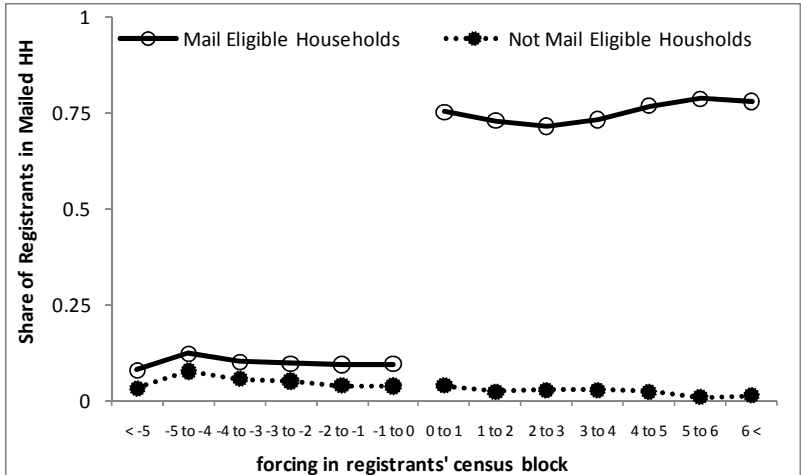


Figure 3a: Mail Concentration in 2006 by Mail Eligibility Status and Census Block Income




Figure 3b: Turnout in Mail Eligible Households in 2002 and 2006 by Census Block Income




Figure 3c: Turnout in Not Mail Eligible Households in 2002 and 2006 by Census Block Income