Advanced Quantitative Methods in the Study of Political Behavior

University of Copenhagen
Department of Political Science

Televised ads revised

Replication of Krasno & Green (2008)

Authors
Eva B. Carstensen
Christian M. Bruhn
Jakob E. Hanghøj

Instructor
Gregory Eady

Number of words: 4,745
Number of characters: 28,470

December 13, 2020

Contents

1 Introduction

In their research paper *Do Televised Presidential Ads Increase Voter Turnout? Evidence from a Natural Experiment*, Jonathan Krasno & Donald Green (2008) study televised advertisements' effect on voter turnout in the 2000 US presidential election. Due to the US Electoral College, presidential candidates spend more resources in battleground states where the election outcome is affected by a swing in votes than in 'safe' states, where one party has dominated in past elections. This structure distorts advertising incentives across geographic areas, where advertising varies up to election day from no advertising at all in non-competitive states, such as California, to large levels of advertising in battleground states, such as Pennsylvania. But candidates are not able to precisely decide which states or districts to receive certain ads due to idiosyncrasies between states and media markets. The US Federal Communications Commission grants media companies local broadcast rights for a selection of counties in an area called a designated market area (DMA) (Spenkuch & Toniatti, 2018: 1982). The authors exploit this idiosyncrasy to set up a natural experiment within states, where they use the large variations in media exposure to estimate the effect of ads on turnout. Their overall finding is that televised advertisements in the final three weeks of the election had negligible effects on turnout.

In this paper, we wish to replicate Krasno & Green's research design to investigate whether the same results can be found in the 2008 US presidential election. We choose to focus on the 2008 election because it is the following open-seat presidential election. We replicate their state-level fixed effects approach to estimate within-state variation in states with more than one DMA. We replicate the results and graphs with new data, but we also add an extension to the original paper. We exploit the spill-over effect from competitive to non-competitive states caused by DMA locations and use an GRDD to compare neighboring counties from the same state but with different DMAs. We focus on a New Jersey county border where one county received ads from the Philadelphia DMA and the other received ads from the New York DMA. This extension enables us to estimate a local average treatment effect.

Krasno & Green put forward two empirical questions: 1) What is the net effect of the volume of presidential ads? 2) To what extent, if any, is the mobilizing influence of presidential ads masked by variations in their tone? Classifying whether the tone of the ad is either positive or negative, they find no proof that attack ads promote or diminish turnout. In our replication with data from the 2008 presidential election, we find similar negligible results.

1.1 Theoretical perspectives

The literature on televised advertisements' effect on turnout is somewhat ambiguous. Earlier studies find positive effects on both the intention to vote and that ads inform and mobilize voters (Hillygus, 2005; Freedman, Franz & Goldstein, 2004). And later studies have found similar confirming results to illustrate the effect on both intention and actual turnout (Law, 2020; Niebler, Neiheisel & Holleque, 2018; Gordon & Hartmann, 2012). However, similar to Krasno & Green, other studies find that ads have no effects on aggregate turnout (Keele & Titiunik, 2015; Spenkuch & Toniatti, 2018). The literature includes different elections, different measurements and different research strategies.

2 Data

2.1 Unit of analysis

Our unit of analysis is the DMA by state as in Krasno & Green's paper. DMAs do not necessarily follow state lines, so residents of a given state may be exposed to varying amounts of presidential advertising up to election day, as candidates invest their resources at the DMA level. Citizens of the same state can have different DMAs and citizens of different states can have the same DMA.

Because we investigate the differences between DMAs within states, we only include states with more than one DMA. In addition, we exclude states with no ads coded in the final three weeks of the election. Furthermore, we were not able to obtain county-level Senate election data for some states. Since senate election data is included in all the models, we exclude these states. These restrictions exclude 13 states including DC[1]. Lastly, we exclude Alaska because the electoral system is structured by districts that do not match the geographic boundaries of the DMAs.

This leaves us with 37 states and 150 DMAs, where Krasno & Green had 37 states and 75 DMAs. Since some DMAs cross state borders, we have 231 observations while Krasno & Green had 128 observations, which they term 'media zones' (Krasno & Green, 2008: 255). Each DMA comprises of several cases, so every unique combination of state and DMA counts as one observation.

Figure 1 is a replication of Figure 1 from Krasno & Green's paper which shows presidential advertising in the final three weeks of the 2000 election. The grey lines indicate DMA areas and the black lines indicate state borders. There are a few differences between their map and ours.[2] Krasno & Green did not include a DMA shapefile in their replication material, so we used a shapefile from Sood (2016).

[1] KY, WI, MA, MD, NY, AK, HI, UT, RI, CT, DC, DE, ND, SD

[2] Their map shows a DMA in Arizona as having GRPs, but their data set does not include any DMAs from Arizona. There is furthermore a slight difference in the exact DMA locations in Wyoming and Colorado.

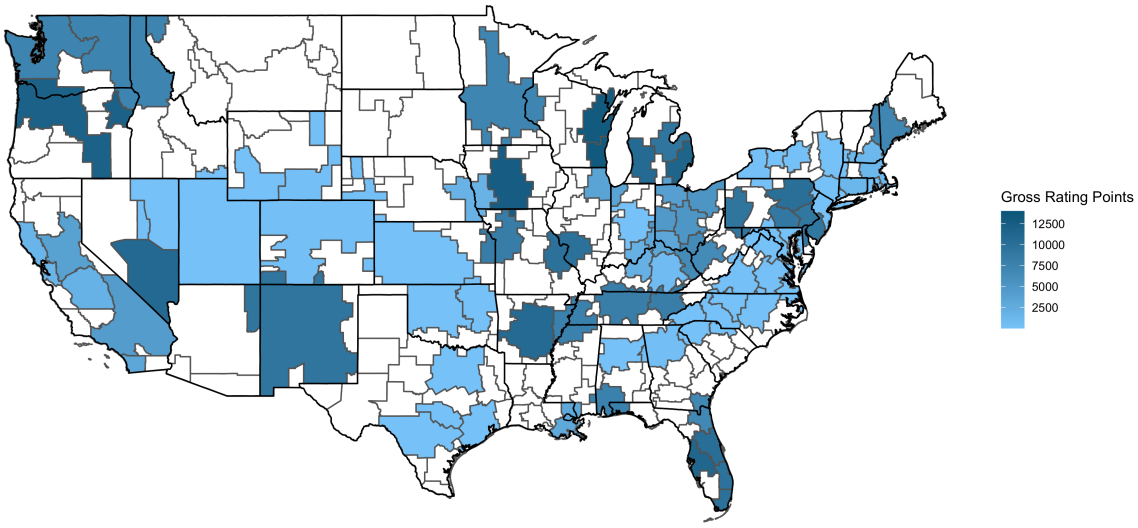

Figure 1: Advertising (in Gross Rating Points) by DMA 3 weeks up to the 2000 election

Figure 2 depicts total ads shown in the final three of the 2008 election. As Figure 2 reveals, our dataset has more DMAs than Krasno & Green's. As in Figure 1, we see a relationship between the volume of ads and the electoral value of each DMA. The map looks as expected based on a campaign logic to allocate resources to areas important for winning the election. Few ads were shown in historically Republican states like Idaho, while battleground states like Florida received a lot of attention. We see that several DMAs across state lines, also in battleground states. And most importantly for this analysis, there are large differences in ad exposure within the state. For example, the dataset contains 12 DMAs inside Ohio with varying levels of ad exposure.

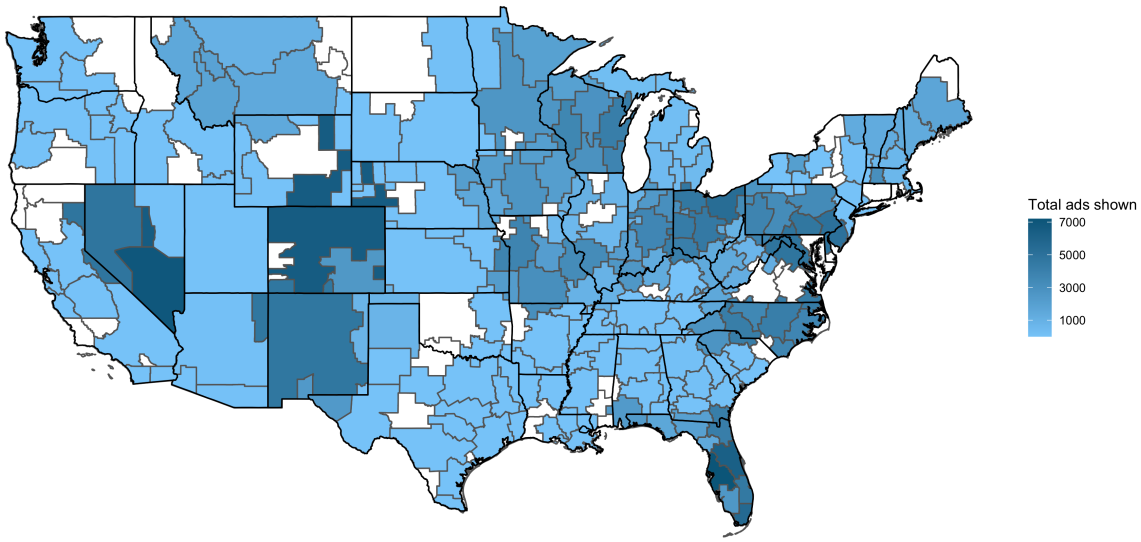

Figure 2: Advertising (in total ads) by DMA 3 weeks up to the 2008 election

2.2 Description of data

The dependent variable in Krasno & Green's paper is the turnout rate in 2000, measured as the proportion of voting-age citizens who vote. Their voting data is from the America Votes series and population data is from the 1990 and 2000 censuses (Krasno & Green, 2008: 248). In our analysis, the dependent variable is the turnout rate in the 2008 presidential election. We measure the turnout rate as Krasno & Green. Voting data is predominantly from the National Neighborhood Data Archive – NaNDA (Chenoweth et al., 2018). However, due to a lack of turnout data in some counties in the NaNDA, we supplement with turnout data from MIT Election Data and Science Lab (2018).

The independent variable in our analysis is media exposure. To measure media exposure, Krasno & Green use Gross Rating Points[3] (GRP) which is a measure of advertisement impact. The scope of media exposure is all presidential advertisements aired on broadcast and cable stations in the final three weeks of the 20008 election. Instead of GRPs, we measure media exposure as the total number of presidential advertisements but with the same scope as Krasno & Green. We do this because GRP-data is very expensive and almost impossible to obtain. However, the total number of ads is also accepted as a measure for media exposure. Krasno & Green actually run their models with total number of ads instead of GRPs and get the same results (Krasno & Green, 2008: 252). We purchased our ad data from the Wisconsin Advertising Project (Goldstein et al., 2011). Coders from the Wisconsin Advertising Project have evaluated if the tone of the ad is promotional, attacking, or contrasting (Krasno & Green, 2008: 251).

We use data on Senate and House election ads from the Wisconsin Advertising Project as control variables to prevent overestimation bias. Furthermore, we use the turnout rates for the recent Senate elections (2002 & 2006) and presidential elections (2000 & 2004) as control variables. The data on the turnout rate for the Senate election is obtained from the National Neighborhood Data Archive – NaNDA (Chenoweth et al., 2018). The data on the two recent presidential elections is obtained from MIT Election Data and Science Lab (2018) and County Intercensal Datasets: 2000-2010 from US Census Bureau (2016).

As Krasno & Green (2008), we add candidate appearances as a control variable because this is likely to increase turnout. The information about presidential and vice-presidential visits to each DMA is from Chen & Reeves (2011). Krasno & Green add an additional control variable on voter contact data. The voter contact data is from the National Annenberg Election Survey. We were able to obtain the National Annenberg Election Survey from 2008 on the state level. However, the Annenberg Public Policy Center would not allow us to get the county level data unless we sent a

[3]Calculated as a percent of the potential audience reached multiplied by the exposure frequency

hard copy mail to the US.

All of our data were obtained at the county level and aggregated up to state-level DMAs. The structure of the DMAs as they existed in 2008 was obtained from Sood (2016).

3 Research design and statistical models

The research design is structured as a natural experiment. The geographic boundaries of the DMAs allow for a somewhat random assignment of media exposure within states. The large variations in media exposure in the different DMAs are exploited to estimate the effect of televised ads on turnout.

3.1 Statistical model for Krasno & Green

In Krasno & Green (2008), the statistical model is outlined as:

$$T_i = b_0 + b_1 X_i + \sum_{2}^{J} b_j Z_{ij} + u_i \tag{1}$$

Where T_i denotes the voter turnout in media zone i, and X_i is the number of presidential ads aired in each media zone (GRPs). Z_{ij} represents a series of dummy variables for each state j, and u_i represents unobserved causes of voter turnout. We use the same estimation strategy for the 2008 election, but due to resource limitations, our X_i is total ads aired in each media zone and not GRPs.

As Krasno & Green, we test for heteroskedasticity by performing Breusch-Pagan tests in which the disturbance variance is predicted by the log of population and the voting rate times 100 minus the voting rate. We also find evidence of heteroskedasticity. Because of this, we run the OLS models with standard errors and robust standard errors (Krasno & Green, 2008: 254).

Krasno & Green extend their models with two specifications based on the effect of different tones of the ads: attacking, promoting, or contrasting the candidate. The first specification is an F-test to see if attack, promote, or contrast ads have distinct effects. The second specification is a model where turnout is a function of the difference between the number of GRPs devoted to ads attacking or promoting candidates. In our statistical model, we conduct the same specifications, but we substitute GRPs with the number of ads.

3.2 Geographic Regression Discontinuity Design

We demonstrate the robustness of our findings by estimating the effect of televised advertising on turnout in 2008 by using a Geographic Regression Discontinuity Design (GRDD). GRDD is useful when a geographic border divides two areas as-if randomly into treatment and control areas. The intuition is that people close to each side of the border are comparable since the division has separated them somewhat arbitrarily (Angrist & Pischke, 2015: 147).

Keele & Titiunik (2015) have collected individual-level data on voter turnout and geographic location of voters on each side of the border between the Philadelphia DMA and the New York DMA in the State of New Jersey as illustrated in Figure 3. There were approximately 5,000 presidential ads in the final three weeks before election day 2008 in the Philadelphia DMA. By contrast, there were 3 presidential ads in the same period in the New York DMA.

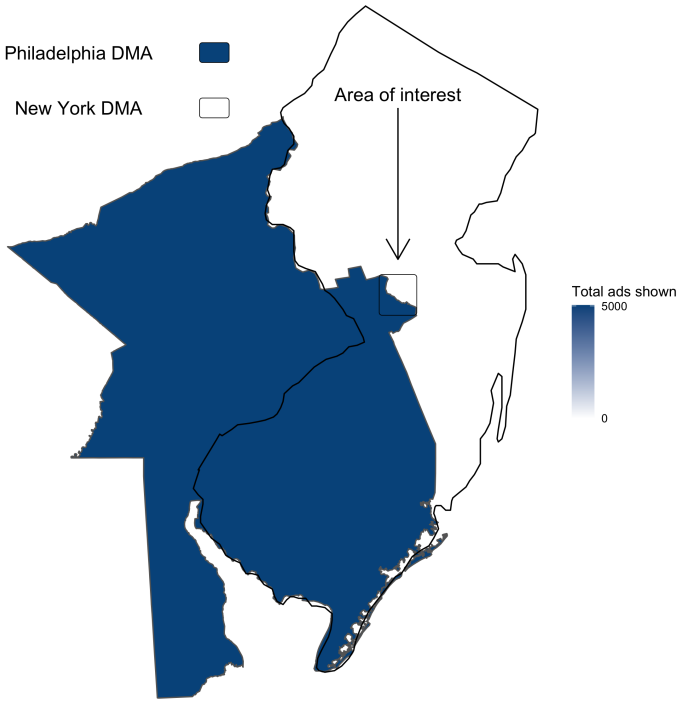

Figure 3: Close-up on the border between the Philadelphia and the New York DMA in New Jersey

The model for the GRDD estimation strategy is:

$$T_i = b_0 + b_1 D_i + b_2 X_i + u_i \tag{2}$$

This time X_i denotes the shortest distance from the geographic border between the two DMAs, and due to the GRDD, the state-level fixed effects are no longer relevant because we are looking at similar units that differ in whether or not they have been treated with ads (D_i).

4 Results

4.1 Replication of Krasno & Green

Table 1 presents our replication of Table 2 in Krasno & Green's paper. Our results are exactly the same for model 1 through model 5. The differences in model 6 and 7 are of an inconsiderable size. Krasno & Green's estimates are -0.07 in model 6 and -0,04 in model 7 with equivalent standard errors. We suspect that the small differences between the estimates are caused by our generation of the variables for candidate visits and contact rate. In section 4.2, we will elaborate more on the different models.

Table 1: Replication of Regression Estimates of the Effects of Presidential Advertisements on Voter Turnout '00 from Krasno & Green

	Turnout '00						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Gross Rating Points	0.25**	0.02	0.05	0.06	0.04	−0.06	−0.06
SE	(0.09)	(0.08)	(0.06)	(0.05)	(0.06)	(0.06)	(0.06)
Robust SE	(0.08)	(0.07)	(0.05)	(0.05)	(0.05)	(0.05)	(0.06)
Midterm turnout	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State fixed effects		Yes	Yes	No	Yes	Yes	No
Past presidential turnout			Yes	Yes	Yes	Yes	Yes
Control for other election ads					Yes	Yes	Yes
Control for candidate visits and contact rate						Yes	Yes
Observations	128	128	128	128	128	128	128
Adjusted R^2	0.60	0.86	0.94	0.88	0.93	0.94	0.88

Note: *p<0.05; **p<0.01; ***p<0.001

In their Table 3, Krasno & Green also present multiplicative heteroskedasticity regression estimates for the different models as a robustness test. Following instructions from Gregory Eady, we have not replicated these estimates.

4.2 Replication with data from the 2008 election

Table 2 presents our analysis of the 2008 presidential election. The models are based on the same specifications as Table 2 in Krasno & Green's paper.

Model 1 reports a naïve regression of media exposure on voter turnout including control for turnout in the recent midterm election. This model shows that ads in the final three weeks of the election had a statistically significant effect on turnout. An increase in total ads by 10,000 is associated with an increase in turnout by .1 percentage points. Introducing state-level fixed effects in model 2 makes the estimate insignificant. This means that when we only look within states, advertising

has no effect on turnout. The effect is still insignificant when we control for turnout in the two recent presidential elections. In model 4, we drop state-level fixed effects and the results are still statistically insignificant. However, the p-value is just below 0.1. In model 5, we control for ads in other elections and include fixed effects, and the results are still statistically insignificant. In model 6 and 7, where we add a control for candidate visits, the results are statistically insignificant, both with and without fixed effects.

All of our results from the 2008 presidential election are similar to what Krasno & Green find in the 2000 election. We reflect on this in section 5.

Table 2: Regression Estimates of the Effects of Presidential Advertisements on Voter Turnout '08

	Turnout '08						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Total ads (divided by 10,000)	0.10**	0.04	0.01	0.05	−0.001	−0.004	0.03
SE	(0.04)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)	(0.03)
Robust SE	(0.03)	(0.03)	(0.02)	(0.02)	(0.02)	(0.02)	(0.03)
Midterm turnout	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State fixed effects		Yes	Yes	No	Yes	Yes	No
Past presidential turnout			Yes	Yes	Yes	Yes	Yes
Control for other election ads					Yes	Yes	Yes
Control for candidate visits						Yes	Yes
Observations	231	231	231	231	231	231	231
Adjusted R^2	0.13	0.56	0.67	0.51	0.67	0.67	0.52

Note: *p<0.05; **p<0.01; ***p<0.001

Like Krasno & Green, we re-estimated the models in Table 2 using three separate ad tone variables: promotional ads, attack ads, and contrast ads. F-tests were conducted in order to ascertain whether allowing the three categories of ads to have different coefficients significantly improved the fit of the models. In none of these cases were the statistical tests significant.[4]

To improve our ability to detect a statistical difference between the effects of promotional and attack ads, we also conduct an estimation using the difference between promotional ads and attack ads. In general, our estimates were insignificant (except model 1 and 4) and negative, suggesting that, if anything, attack ads are slightly more likely to increase turnout than promotional ads.[4] This is the same conclusion Krasno & Green reached in their paper.

4.3 Estimates from Geographic Regression Discontinuity Design

To demonstrate the robustness of the results in Table 2, we estimate the local average treatment effect (LATE) on the border between the Philadelphia and New York DMA. We use a GRDD with

[4]The results are in our R script

optimal bandwidth[5] and different polynomial specifications.

Even though the density test indicates potential sorting close to the cut-off, we argue that the sorting issue and discontinuities in pre-treatment outcome are minimized. The areas on each side of the border are in the same legislative and school district, and the house prices are similar (Keele & Titiunik, 2015: 145). Therefore, we believe arguments to prefer living on one side of the DMA border compared to the other are few, and as a result, we believe that the assumptions have been met.

As it can be seen in Table 3, the treatment effect is insignificant for all four specifications.

Table 3: GRDD Estimates of the Effects of Presidential Advertisements on Voter Turnout '08

	Turnout '08			
	(A)	(B)	(C)	(D)
Local Average Treatment Effect	0.05	0.11	0.12	-0.002
SE	(0.07)	(0.06)	(0.07)	(0.14)
Optimal bandwidth	797	1,673	3,059	4,101
Polynomial order	Linear	2	3	4
Observations	2,044	5,895	10,085	6,710
Note:				*p<0.05

To illustrate the results, we have plotted all the models from Table 3 in Figure 4. The left side (negative numbers) of the border is the New York DMA which is the control DMA and the right side (positive numbers) is the Philadelphia DMA which is the treated DMA. As shown in the figures, the difference on each side of the border is small and insignificant for all models.

Overall, we find no evidence that televised ads impact voter turnout in the 2008 presidential election. Based on this, we find support for the conclusions in Krasno & Green's paper.

[5]We use the rdrobust-function to calculate the optimal bandwidth

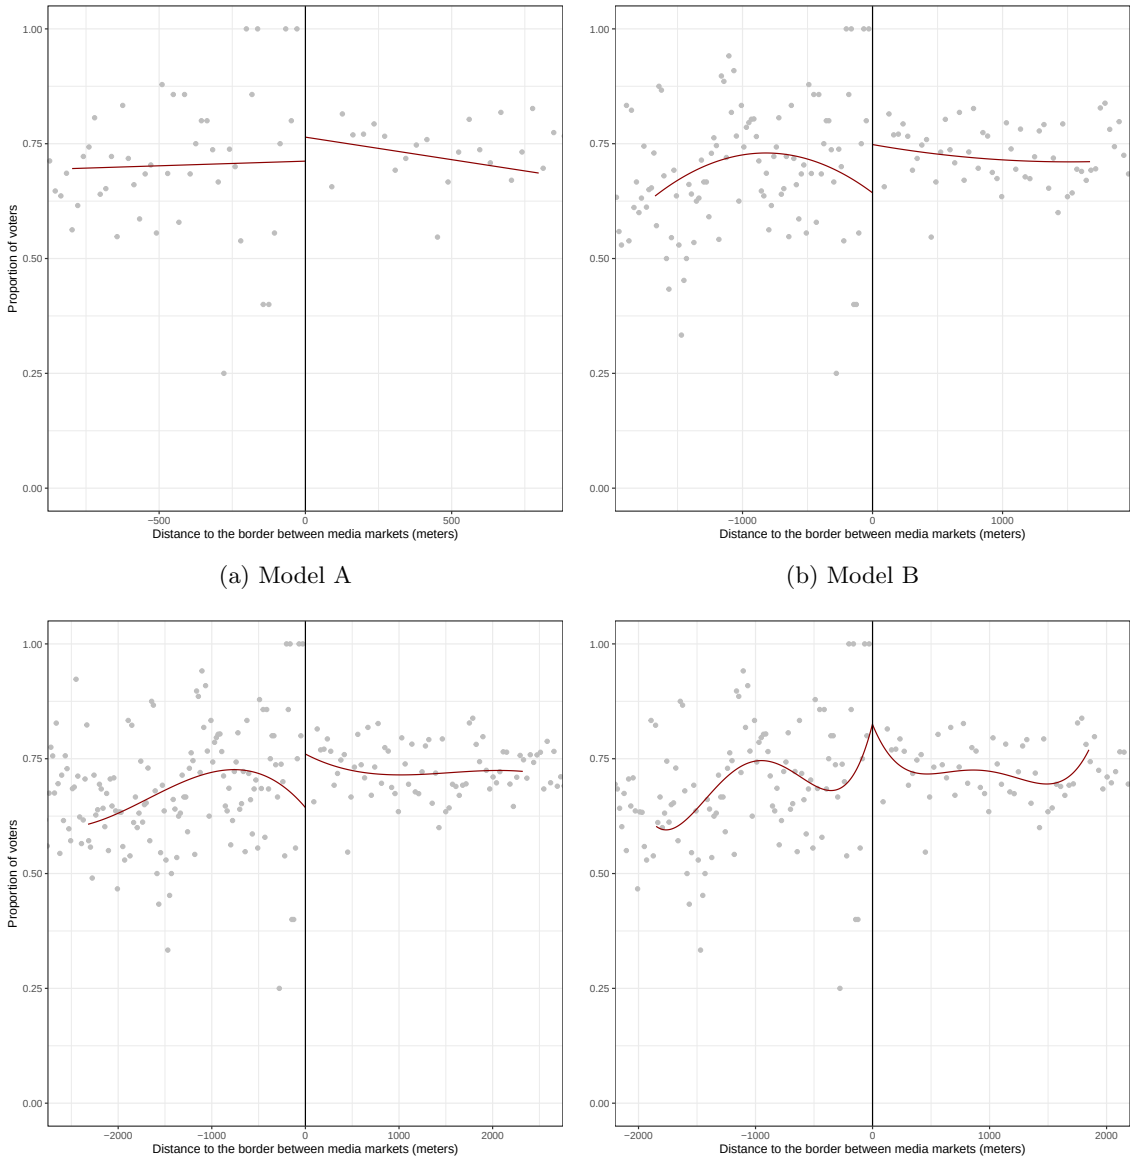

(a) Model A

(b) Model B

(c) Model C

(d) Model D

Figure 4: GRDD models with different specifications

5 Discussion

5.1 Similarities and dissimilarities between the studies

As mentioned in section 4.1 and 4.2, our results are similar to the findings in Krasno & Green (2008). We find no effect of televised ads on voter turnout in 2000 nor in 2008. This strengthens the inferences we can draw because the conclusions are not just driven by one election.

We find these similar results even though the two elections differ on several variables. The 2008 election had a higher voter turnout and a higher number of total ads. In the 2008 election, the turnout rate was 61.6 percent compared to 54.2 percent in 2000, and there was a record in spending

on televised ads (Woolley & Peters, 2016; Rutenberg, 2008). In our data, we see a larger volume of ads in the final three weeks of the 2008 election compared to 2000. Krasno & Green reported that the DMA with the most ads reached around 4,030 ads shown, while our dataset reveals that a DMA in Florida reached 7,166 ads shown (Krasno & Green, 2008: 249). This larger volume of total ads could have enhanced the chance of televised ads affecting voter turnout because it means that the treatment in some DMAs was stronger. The higher turnout rate could also enhance the chance of finding a positive, significant effect because of a higher possibility that some new groups have been mobilized by campaign means.

Another difference is the use of social media in the 2008 election in which Obama's campaign diligently used the emerging political online scene to advertise (Kiss, 2008). Social media advertising is not structured by DMAs but might be influenced by specific geographical conditions, so that battleground states still receive more social media ads. This is likely to have affected voter turnout in 2008. Thus, it could have been a relevant control variable in our design, especially if we wish to compare our results to the 2000 election. Another extension could be to look at the individual candidate's campaign. It could be that Republican voters are more affected by television than Democrats or the other way around. Furthermore, different voter groups might be more affected than others, i.e. older voters who generally watch more television than younger voters.

Based on this, the election in 2008 could be perceived as a most-likely case due to these favorable conditions for finding a significant effect of televised ads on turnout. In this respect, our insignificant findings in 2008 strengthen the causal inferences we can draw on the (lack of) effect of televised ads.

5.2 Data limitations

Conducting this replication, we have experienced some data limitations. First, it was not possible for us to obtain all the indicators identical to Krasno & Green. To fully compare our results, we would need equivalent data from the same sources with identical indicators. But due to data inaccessibility and differences in data setup between the two elections, this standard was difficult to achieve. This is best exemplified by our use of total ads instead of GRPs. GRPs are a more sophisticated measure, as total ads do not distinguish between ads shown in primetime and ads shown in the middle of the night. If we use different measurements for our independent variables, we could potentially find a different causal connection than Krasno & Green. On the other hand, other researchers also use different data sources and operationalizations to compare their results to Krasno & Green's.

Another potential problem is associated with the fact that our data consists of eight different

data sources which we have merged based on FIPS-code and DMA. Our method of merging has an inherent risk of data leakage which could be a problem as we are unaware of a potential bias in our results. Lastly, we have used different data sources for turnout. Even though we use the same method for calculating turnout as Krasno & Green, there is a risk associated with using data sources with potentially different collection methods. Furthermore, Krasno & Green were able to use updated census data on voters from 2000, whereas the data on voters in 2008 naturally is more inaccurate because census data is decennial.

Overall, our results are similar to Krasno & Green's, but because of data limitations, potential data leakage and different data sources, we are cautious when concluding on the precise comparability between the results of two studies.

5.3 Alternative models

The sharp RDD is ideal for our research design as it exploits abrupt changes in treatment status caused when the treatment is determined by a specific cut-off (Angrist & Pischke, 2015: 175). RDDs require seemingly 'mild' assumptions compared to other approaches, and some researchers believe that causal inferences from the RDD are the most credible (Lee & Lemieux, 2010: 282). In section 4.3, we estimated a LATE on the border between the Philadelphia and New York DMA. We found no significant effect even though citizens in the Philadelphia DMA received more than 5,000 ads compared to 3 ads for citizens in the New York DMA.

One important drawback to this strategy is that it is a LATE. We only know that televised ads did not affect people just near the cut-off in New Jersey. If we expand the cut-off, we create concerns about polynomial choice and sorting, and if we minimize the cut-off, we get less data (Angrist & Pischke, 2015: 162). To estimate more LATEs and increase generalizability, we could potentially estimate these local effects in all areas with a large discontinuity for citizens exposed to an either high or low level of ads. But this strategy was only possible because Keele & Titiunik had access to individual-level voter turnout data in an area where each side of the border was almost identical on different parameters. Due to data inaccessibility, it is difficult to broaden this strategy and estimate LATEs from several areas. It is furthermore difficult to choose ideal areas to run this estimation; i.e. if DMA and county boundaries are identical, one will not be able to isolate the two potentially different effects (Keele & Titiunik, 2015: 133). The RDD is an alternative way to estimate the effect, but the strategy requires very precise geographic data. We use the strategy as a supplement to Krasno & Green's fixed effects strategy, thus we interpret the insignificant results as another evidence for a null effect.

Krasno & Green emphasize the importance of running their models with state-level fixed effects,

but they use a unit of analysis including only 128 observations. That few observations can prove problematic, as they include 37 state dummy variables to the model, thus they have less than three observations for each estimated parameter. This few observations increase the risk for multicollinearity (Franz et al., 2008: 265). Multicollinearity can lead to large standard errors, whereby they risk false-negative results. In general, it is problematic to use a fixed-effects model where the unit of analysis includes few observations because it increases the possibility of estimating false nulls. Our unit of analysis includes 231 observations with 37 state dummy variables, leaving us with around six observations for each estimated parameter. This definitely reduces the risk of false-negative results but the observations are still few. An alternative model could look at the influence of ads over time to increase the number of observations. Franz et al. (2008) emphasize that a better approach to increase observations would be to switch the unit of analysis to a county-level model. The problem with this approach is that ads are assigned to DMAs, so if counties were the unit, some units would cluster and result in larger standard errors. Overall, we believe that the best way to increase observations is not to change the unit of analysis, but to look at several elections combined.

6 Conclusion

We replicated Krasno & Green's paper and found the same insignificant results. We furthermore tested if the same effect could be found in the following open-seat presidential election in 2008.

We found no evidence of televised ads' impact on voter turnout in the 2008 election. We extended their statistical model in order to demonstrate the robustness of these findings with a GRDD approach on a critical case with individual-level data provided by Keele & Titiunik (2015). We exploited the variation in total ads between the geographically linked Philadelphia and New York DMAs. Likewise, we found no significant results with this approach.

Our causal inference is strong because our findings are consistent with Krasno & Green's conclusions and robust to an alternative model. Moreover, the election in 2008 can be perceived as a most-likely case due to the high turnout, the high number of ads, and the emergence of social media ads. This strengthens the insignificant results' inference further. However, we are hesitant when comparing our results to Krasno & Green's due to data inaccessibility and the risk of data leakage.

An alternative model could be to run the GRDD model on a larger, more representative sample of the entire country to estimate several LATEs. However, this model is conditioned by obtaining large amounts of individual-level data. Another alternative model could be to examine the effect of ads over several elections, thereby increasing the number of observations and avoiding the risk of multicollinearity from using state-level fixed effects.

Bibliography

Angrist, Joshua D. & Jörn-Steffen Pischke. 2015. *Mastering 'Metrics: The Path from Cause to Effect*. Princeton, NJ: Princeton University Press.

Chen, Lanhee & Andrew Reeves. 2011. "Replication data for: Turning out the Base or Appealing to the Periphery? An Analysis of County-Level Candidate Appearances in the 2008 Presidential Campaign." *Harvard Dataverse* .
URL: *https://doi.org/10.7910/DVN/0XTAAU*

Chenoweth, Megan, Mao Li, Iris N. Gomez-Lopez & Ken Kollman. 2018. "Voter Registration, Turnout, and Partisanship by County, United States, 2004-2018." *National Neighborhood Data Archive* .
URL: *https://doi.org/10.3886/E125781V2*

Franz, Michael M., Paul Freedman, Ken Goldstein & Travis N. Ridout. 2008. "Understanding the Effect of Political Advertising on Voter Turnout: A Response to Krasno and Green." *The Journal of Politics* 70(1):262–268.

Freedman, Paul, Michael Franz & Kenneth Goldstein. 2004. "Campaign Advertising and Democratic Citizenship." *American Journal of Political Science* 48(4):723–741.

Goldstein, Kenneth, Sarah Niebler, Jacob Neiheisel & Matthew Holleque. 2011. "Presidential, Congressional, and Gubernatorial Advertising, 2008." *Wisconson Advertising Project* .
URL: *https://elections.wisc.edu/wisconsin-advertising-project/*

Gordon, Brett & Wesley Hartmann. 2012. "Advertising Effects in Presidential Elections." *Marketing Science* 32(1):19–35.

Hillygus, D. Sunshine. 2005. "Campaign Effects and the Dynamics of Turnout Intention in Election 2000." *Journal of Politics* 67(1):50–68.

Keele, Luke & Rocio Titiunik. 2015. "Geographic Boundaries as Regression Discontinuities." *Political Analysis* 23:127–155.

Kiss, Jemima. 2008. "Why everyone's a winner." *The Guardian* .
URL: *https://www.theguardian.com/media/2008/nov/10/obama-online-strategy*

Krasno, Jonathan S. & Donald P. Green. 2008. "Do Televised Presidential Ads Increase Voter Turnout? Evidence from a Natural Experiment." *The Journal of Politics* 70(1):245–261.

Law, Wilson. 2020. "Decomposing political advertising effects on vote choices." *Public Choice* pp. 1–23.

Lee, David S. & Thomas Lemieux. 2010. "Regression Discontinuity Designs in Economics." *Journal of Economic Literature* 48(2):281–355.

MIT Election Data and Science Lab. 2018. "County Presidential Election Returns 2000-2016." *Harvard Dataverse* .
URL: *https://doi.org/10.7910/DVN/VOQCHQ*

Niebler, Sarah, Jacob Neiheisel & Matthew Holleque. 2018. "By ground or by air? Voter mobilization during the United States' 2008 presidential campaign." *Journal of Elections, Public Opinion and Parties* 28(1):78–104.

Rutenberg, Jim. 2008. "Nearing Record, Obama's Ad Effort Swamps McCain." *New York Times* .
URL: *https://www.nytimes.com/2008/10/18/us/politics/18ads.html*

Sood, Gaurav. 2016. "Geographic Information on Designated Media Markets." *Harvard Dataverse* .
URL: *https://doi.org/10.7910/DVN/IVXEHT*

Spenkuch, Jörg & David Toniatti. 2018. "Political Advertising and Election Results*." *Quarterly Journal of Economics* 133(4).

US Census Bureau. 2016. "County Intercensal Datasets: 2000-2010." *US Census Bureau* .
URL: *https://www.census.gov/data/datasets/time-series/demo/popest/intercensal-2000-2010-counties.html*

Woolley, John & Gerhard Peters. 2016. "Voter Turnout in Presidential Elections." *The American Presidency Project* .
URL: *https://www.presidency.ucsb.edu/statistics/data/voter-turnout-in-presidential-elections*