Flight to Safety: 2020 Democratic Primary Election Results and COVID-19

James Bisbee
Princeton University

Dan Honig
Johns Hopkins University

This version: April 1, 2020

First public version: March 25, 2020

Abstract

What is the impact of anxiety on vote choice? Building on a well-documented phenomenon in finance, we posit that voters will exhibit a "flight to safety" by turning toward establishment candidates. We test this theory in the context of the Democratic primary election of 2020 by examining changes in the vote shares of Bernie Sanders, a candidate promising disruptive change. We use the outbreak of the novel coronavirus across both space and time to identify a causal effect of anxiety on voting. By comparing counties with and without reported cases in their local media market, before and after the outbreak of the virus, we show that COVID-19 anxiety resulted in diminished support for Sanders as compared to his support in the 2016 election. Our findings contribute empirical evidence to an as-yet underappreciated preference for "safe" candidates in times of social anxiety.

1 Introduction

As COVID-19 began to dominate the headlines of US Newspapers in March 2020, it displaced coverage of the Democratic primary election. In that campaign, one of the two leading candidates had been running on a platform centered on universal health care. One might reasonably imagine that a growing pandemic would lead to a surge in support for that candidate. It did not.

Bernie Sanders did not merely fail to surge as the novel coronavirus came to dominate the headlines in 2020 – if anything, his campaign faded. We explore whether COVID-19 had anything to do with Sanders' apparent decline in support.

The novel coronavirus appeared in the middle of the Democratic primary season. On Super Tuesday, COVID-19 cases were in the news, but public concern in the US was modest. There were only 112 total cases in the US, and the bulk of TV and print news content was focused on the elections. And why shouldn't it be? President Trump, after all, had said just a few days prior that COVID-19 was soon "going to disappear" (Leonhardt, 2020).

President Trump's projection was, unfortunately, incorrect. Just 14 days later, when voters in 3 states cast their ballots in March 17th Democratic primaries, President Trump had already declared a national state of emergency; most of the nation's schools were closed; the stock market had lost over 20% of its value; and many people had begun staying home to practice "social distancing".

In this paper, we ask whether the anxiety created by the novel coronavirus hurt the electoral prospects of Bernie Sanders, the more anti-establishment candidate. Our question is motivated by a well-documented financial phenomenon that has, as yet, not been applied to voting behavior – namely, a "flight to safety". We registered a preanalysis plan prior to analyzing any data, and prior to the primaries of March 17th. Our plan outlines the hypotheses and empirical specifications we employ.

Examining a primary election between two challengers – two non-incumbents not associated with the administration or party in power – allows us to disentangle a "flight to safety" from any assessment of incumbents in response to a crisis. This paper provides what we believe is the first well-identified quantitative analysis of the impact of anxiety on vote choice, distinct from voter evaluations of current office-holders. We explore whether anxiety generated by the unexpected outbreak of COVID-19 impacted voting decisions. Empirically, we compare how counties voted before and after the virus was widespread, in areas where the virus was relatively prevalent and where it was not.

We find that the novel coronavirus disproportionately hurt Bernie Sanders. We show that where the virus emerged prior to the primary election, vote shares for Sanders fell. That this result obtains despite what should be a policy platform whose appeal increases with the pandemic leads us to conclude that the power of the "flight to safety" in the context of voting is an important, but as yet unaccounted for, phenomenon in the voting literature.

2 Theory & Context

Scholars of financial markets and market analysts often discuss markets' "flight to safety" (e.g. Adrian, Crump and Vogt 2019; Inghelbrecht et al. 2013). As market outcomes become more uncertain, risk appetite falls. Anxiety drives market players to reduce their level of risk. In the context of investing, this behavior typically involves shifting assets towards more liquid and Government-insured assets, which are perceived as safer.

While there is a literature on voters' response to terrorism, (e.g. Getmansky and Zeitzoff 2014; Montalvo 2011) there is little research on the effects of anxiety more broadly and whether anxiety shifts votes shifting towards candidates perceived as less risky.[1] Existing studies of crisis voting largely focus on the retrospective evaluation of incumbents in the context of adverse shocks, be they security-related (Gutiérrez, 2014), economy-related (e.g. Nezi 2012; Remmer 1991; Abramson et al. 2007), or broadly about the performance of incumbents in crises (e.g. Smith 1998).

Typically, political scientists rely on either rational actor models or cognitive frameworks to predict vote choice (i.e., Canes-Wrone, Herron and Shotts 2001; Green and Palmquist 1994; Maskin and Tirole 2004). A rational actor model might predict that Bernie Sanders' platform emphasizing universal healthcare access should win the day by appealing to a timely concern of voters. The "flight to safety" perspective generates the opposite prediction – namely that the Sanders campaign would suffer from the increased anxiety generated by the outbreak of the novel coronavirus. By examining the effect of anxiety on the choice between two aspirants for President not part of the administration in power at the time of the anxiety-inducing crisis, this paper provides insight on whether a more general "flight to safety" occurs in voting independent of any attribution of responsibility to the candidates for the crisis itself.

In the case of the 2020 democratic primary, Joe Biden represented safety and Bernie Sanders a disruption of "political as usual". Biden portrayed himself as representing continuity and the security of the known – an "Obama-Biden Democrat", as Biden himself put it in a campaign speech (Fegenheimer and Glueck, 2020). Sanders, in contrast, promised to "change the power the structure in America" (Stewart 2020), portraying himself as a candidate who (in the words of his 2020 campaign spokesman) pushed against "the limits of politics as usual" (Eilperin 2020). Voters apparently understood these divergent appeals, with exit polls in a number of states indicating that Sanders won a majority of those voters for whom the most important quality in a candidate was "Can Bring Needed Change", while Biden was preferred by those for who most valued "Can Unite the Country".[2]

[1] One notable exception is Campante, Depetris-Chauvin and Durante 2020, who examine candidates' strategic manipulation of Ebola-induced fear of immigrants in the 2014 US midterm elections and find results complementary to this paper's. Other papers have explored

[2] See exit polls as reported by CNN, https://edition.cnn.com/election/2020/entrance-and-exit-polls/STATE NAME/democratic, e.g. those from Michigan and Washington. In some states – e.g. California – the candidates won a plurality, but not the majority, of those who felt the most important quality was change and unity respectively.

We hypothesize that growing anxiety due to the outbreak of the novel coronavirus reduces the appeal of a disruptive outsider like Sanders. We predict that a political flight to safety will manifest in decreased votes for Sanders where voting occurs after a COVID-19 infection is identified in a Designated Market Area (DMA), all else equal. Those living in places where positive COVID-19 tests occurred are likely to have experienced more anxiety than those for whom infection was a more distant possibility, at least during the period we examine.[3] We use these twin sources of variation in anxiety induced by the disease – i.e., cross-sectional variation due to differences in exposure and temporal variation in the timing of the outbreak – to empirically estimate the effect of COVID-19 on Democratic primary vote choice.

We emphasize that if the anxiety mechanism we describe does *not* obtain, voters might be *more* supportive of Sanders due to his policy platform, making this a particularly hard test for the theory. That is, Sanders' emphasis on universal healthcare should appeal to voters who are exposed to the novel coronavirus and face a more acute need for care. Similarly, Sanders' more expansive protections for working class voters should grow more appealing as the spectre of recession and job losses grew. Given that Sanders' policy platform should be more attractive following COVID-19's emergence, we believe our empirics constitute a hard test of the motivating theory.

3 Data and Methods

We combine several data sources to measure our outcome variable, explanatory variable, and controls.

Outcome Variable

Our outcome variable is the change in the county-level vote share for Bernie Sanders between 2016 and 2020. The 2016 data was obtained from `https://www.nytimes.com/elections/2016/results/primaries/[STATE]`. The 2020 data was obtained from the "State Results" tab on the `https://www.nytimes.com/interactive/2020/03/17/us/elections/results-primary-elections-florida-illinois-arizona.html` page at noon on March 18th. At the time of writing, over 97% of counties had 100% reporting.

Throughout our paper, we refer to the "start-date" of the outbreak as either after March 1st, after March 3rd, or after March 10th. These dates are chosen such that the three waves of primary elections in March fall into either treatment or control, as defined in Table 1. We further exploit the timing of elections for robustness checks and placebo tests in our Supporting Information.

[3] As national media coverage of the outbreak became ubiquitous, our ability to leverage cross-sectional variation declines. We discuss and test these SUTVA assumptions below and in our Supporting Information.

Start Date	Control	Treatment
March 1st	Feb	ST, March 10th, & March 17th
March 3rd	Feb & ST	March 10th & March 17th
March 10th	Feb, ST, & March 10th	March 17th

Table 1: Treatment and control elections by outbreak "start date". February (Feb) primaries include IA, NV, and SC. Super Tuesday (ST) primaries include AL, AR, CA, CO, ME, MN, NC, OK, TN, TX, UT, and VA. March 10th primaries include ID, MI, MS, ND, and WA. March 17th primaries include AZ, FL, and IL. (Ohio's was postponed due to the outbreak.) MA, VT, and NH are excluded as they do not aggregate votes by county in reporting totals.

Explanatory Variable

We use data from two separate sources for county-level COVID-19 infection data. The first is the github account for Johns Hopkins University CSSE Coronavirus Resource Center `https://github.com/CSSEGISandData/COVID-19/blob/master/csse_covid_19_data/csse_covid_19_time_series/time_series_19-covid-Confirmed.csv`. The second is from a non-profit website developed by a variety of academics and professionals called 1Point3Acres `https://coronavirus.1point3acres.com/#stat`. We scraped these data in the evening of March 16th, 2020 using the `rSelenium` package for `R`. At the time of writing, the JHU data coverage only extends through March 9th while our March 16th scrape of 1Point3Acres is, to the best of our understanding, accurate for that date. Maps of the geographic distribution of the outbreak by DMA on March 2nd (the eve of Super Tuesday), March 9th (the eve of the second round of multiple state primaries), and March 16th (the eve of Arizona, Illinois, and Florida) are presented in Figure 1.

Controls

We obtain a rich set of pre-treatment county-level controls from the five year averages of the American Community Survey (2018). These county-level controls are:

- total population
- % of the population that is rural
- % of the population that is white
- % of the population with a bachelor's degree
- the county's old-age dependency ratio (retirees to workers)
- Share of households that are headed by a woman without a husband present
- % of the population that speaks only English

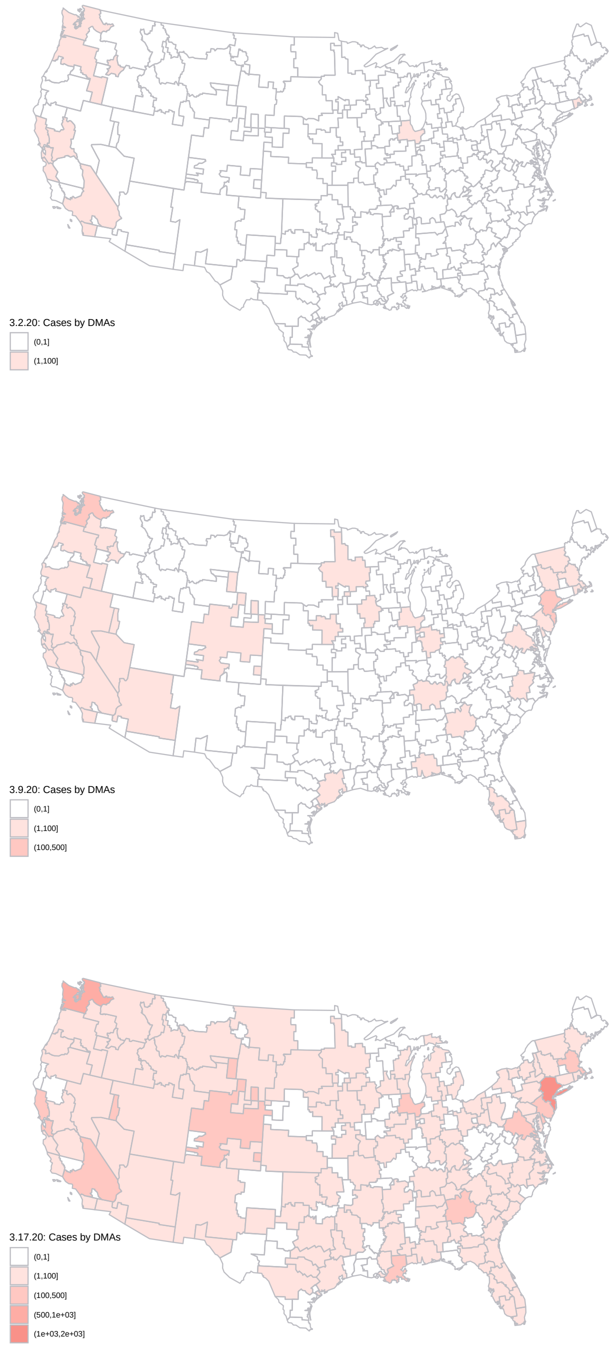

Figure 1: Maps of the geographic and temporal variation in the spread of the virus between March 2nd and March 17th. March 2nd and March 9th data are from Johns Hopkins University. March 16th data are from 1Point3Acres.

- % of population between 18 and 64 that is below the poverty level

- % of the population employed in manufacturing
- the county-level unemployment rate
- the county-level labor force participation rate
- the median household income

Methods

We are interested in identifying the causal effect of exposure to the novel coronavirus on Democratic primary voters' decisions. While the outbreak of COVID-19 was an exogenous shock to voter anxiety, it is confounded in two ways. First, the timing of treatment is colinear with other explanations for changing electoral fortunes, such as the decision by several primary candidates to drop out (Staff, 2020), signaling a consolidation of party support behind Biden (Yglesias and Beauchamp, 2020). A simple before-after comparison of election returns would be unable to disentagle our "flight to safety" theory from a coincidental shift in electoral momentum.

Second, we might expect that older voters are more dissuaded from appearing at the polls following the appearance of COVID-19 due to the increased risks of exposure. Insofar as younger voters are relatively more supportive of Sanders, this would bias our results in a conservative direction, making it harder to identify a negative relationship between exposure and Sanders' vote share.[4]

We posit that anxiety due to the disease is a function of both temporal and geographical variation, allowing us to address these confounds. We define exposure as binary variable taking on the value of 1 if a county c resides in a designated market area (DMA) with confirmed cases of COVID-19 on the eve of their primary election date, and 0 otherwise, denoted with $COVID_c$. To absorb county-specific characteristics that are time-invariant, we estimate the difference in the change in Bernie Sanders' vote share from 2016 to 2020, between exposed and unexposed counties, or $E[\Delta Y_1 - \Delta Y_0]$, where ΔY is the county-level change in Sanders' primary vote share, and the 1 and 0 subscripts represented treated and control counties, respectively.[5] We also include fixed effects for the date of the election, thus

[4]There is also the possibility of a selection bias which would obtain of exposed counties were more anti-Sanders to begin with. We predict Sanders' 2016 voteshare as a function of exposure and find, if anything, these counties are *more* pro-Sanders. To the degree that there is selection bias, it works against our results.

[5]By defining the outbreak as happening after Super Tuesday when the field narrowed to a contest between Sanders and Biden the 2020 Sanders vote share is more directly comparable to the 2016 contest, effectively a two-candidate contest between Sanders and Hillary Clinton. To the extent that this comparison fails we expect it will bias results downwards rather than increase the chance of spurious statistically significant findings; when the field narrowed after Super Tuesday that increased the Sanders vote share and thus biases against us finding decreased support for Sanders, all else equal. Furthermore we also acknowledge that several states (including CO, ME, MN, UT, and WA) switched from caucuses to primaries between 2016 and 2020. While Bernie does better on average in caucus states, all five of these are included in the pre-period in our main specification, ensuring that they do not drive our results.

restricting comparisons to differences in a given election "wave" and thus partially controlling for anything that may have varied between election dates (e.g. number of candidates in the race). This helps ensure that our findings are indeed picking up on growing anxiety around COVID-19 associated with local knowledge (via media coverage) of cases. Our simplest specification, as per our PAP, takes the form:

$$\Delta Y_c = \beta_0 + \beta_1 COVID_c + \gamma \mathbf{X} + \lambda + \epsilon_c \tag{1}$$

where $\mathbf{X}$ is the vector of county-level controls summarized above, and λ are date of election fixed effects.

However, since counties with earlier exposure to COVID-19 are disproportionately more densely populated coastal areas, this specification risks dissimilarities between treatment and control counties. To address this bias, we also leverage the staggered timing of both vote date and exposure to estimate a pseudo difference-in-differences (DID) specification where we compare the difference in the outcome between treated and control groups prior to the outbreak to the difference in these groups following the outbreak, thus allowing us to identify if the anxiety induced by a COVID case was greater following the outbreak date than before.

This design is complicated by the fact that, unlike standard DID settings, we do not observe outcomes in the pre and post period for every unit, precluding our ability to measure $E[Y_{i,t=1} - Y_{i,t=0}]$ at the county level. Instead, we must assume that those counties who voted in the pre period but would go on to be exposed to COVID-19 are valid counter-factuals for those counties that were exposed to COVID-19 and voted in the post period. Similarly, we must assume that the control counties that voted in the pre period (i.e., those that did not experience the COVID-19 outbreak in the post period) are valid counterfactuals for the control counties that voted in the post period.

We augment our conditional independence assumption (CIA) with matching and balancing strategies to ensure we are comparing otherwise similar counties who differ only in the timing of their exposure to COVID-19 - that is, comparing counties that were ultimately exposed to COVID-19 but differ only in whether that exposure occurred before or after the election. With a rich set of pre-treatment covariates we obtain good balance using either nearest neighbor matching (based on minimized Mahalanobis distance), or covariate balanced propensity score weights (CBPS). Exposure is as-if randomly assigned to counties conditional on the observables we control for, match on, and balance over.

One final concern that we believe grows more problematic as the virus spreads is the Stable Unit Treatment Value Assumption, or SUTVA. Substantively, this assumption requires that our control counties are not affected by treatment spillovers from treated counties. Our treatment exposure is defined at the DMA-level, based on the assumption that the salience of the disease is elevated via local media markets which report on more geographically proximate cases. We believe this is sensible for the beginning of March, when the virus was just beginning to spread across the United States. However, by the time of the March 17th elections, national media outlets (e.g. cable news, newspapers, news websites, and online social media such as Facebook in sharing news (Roose and Dance, 2020)) had shifted coverage to focus almost exclusively on the outbreak as the crisis worsened. Thus many of our notionally "control" counties experienced substantial levels of anxiety despite not residing in a DMA with confirmed cases of the virus, with "control" counties becoming decreasingly valid counterfactuals for counties in a DMA where a COVID-19 case had been diagnosed with each passing week.

4 Results

Our main results are summarized in Table 2, in which treatment is defined at the DMA as all confirmed cases of COVID-19 on March 9th, 2020 as reported in the Johns Hopkins University data as of March 21st, 2020. The first two columns present the coefficients on a binary measure of exposure (1 if any cases were recorded in the DMA, 0 otherwise), and a continuous count of the number of confirmed cases as of March 9th, 2020. Clustered standard errors at the DMA-election are presented in parentheses. The coefficients represent a standard deviation change in the change in support for Bernie Sanders' between 2020 and 2016 associated with either moving from 0 to 1 on the binary measure of COVID exposure, or with a one standard deviation increase in the number of confirmed cases of COVID-19 on the continuous measure of exposure. The results indicate that counties that voted after Super Tuesday (March 3rd) and which were exposed to the novel coronavirus were less likely to support Sanders as compared to counties that voted prior to March 10th and counties that voted on or after March 10th but did not reside in a DMA with any reported cases. According to Columns 1 and 2, being exposed to the virus corresponds to an estimated 0.36 standard deviation decline in support for Sanders as compared to his 2016 vote share, over and above the decline in Sanders vote share in matched counties in the control group. This corresponds to an expected change of approximately 4.1 percentage points less support for Sanders compared to the 2016 vote share he enjoyed in the average county. This result is reinforced at the intensive margin, as illustrated by the negative and significant coefficient in column 2, suggesting that a standard deviation increase in the number of confirmed cases (roughly 99 new cases in the DMA) corresponds to a 0.12 standard deviation decline in support for the Sanders' campaign, or roughly 1.5 percentage points, relative to 2016.

However, the results in columns 1 and 2 rely on the assumption that insulated and exposed counties are valid counterfactuals for each other after controlling for a variety of demographic and economic county-level factors. In columns 3 and 4, we reduce our reliance on this assumption by employing a nearest-neighbors matching strategy in which we identify the most similar control county for each treated county in our dataset based on the same county-level covariates. We use Mahalanobis distance measures to summarize the difference across our twelve county-level covariates and choose the county that is most similar to each treated county in terms of this distance measure.[6]

Substantively, this approach strengthens our conditional independence claim that we

[6] We achieve good balance, as demonstrated in the balance tests reported in the Supporting Information, specifically Figure 8.

Table 2: Main Results: Change in Sanders Support $\sim$ Exposure

	Dependent variable: Δ Sanders Vote Share					
	Basic		Matching		Weighting	
	Bin.	Cont.	Bin.	Cont.	Bin.	Cont.
	(1)	(2)	(3)	(4)	(5)	(6)
Treatbin	−0.361		−0.600**		−0.472**	
	(0.231)		(0.251)		(0.226)	
Treatcont		−0.102***		−0.300***		−0.119***
		(0.012)		(0.036)		(0.006)
Tot pop	0.036	0.032	0.131**	0.109*	−0.015	−0.016
	(0.030)	(0.028)	(0.065)	(0.062)	(0.011)	(0.011)
Old age dep ratio	0.024	0.017	0.047	0.055	0.098**	0.096**
	(0.046)	(0.045)	(0.087)	(0.091)	(0.045)	(0.044)
Bachelor's degree	−0.071	−0.066	0.156	0.194*	−0.030	−0.024
	(0.054)	(0.055)	(0.115)	(0.110)	(0.060)	(0.060)
Female HH no husband	0.380***	0.383***	0.346***	0.329***	0.363***	0.357***
	(0.056)	(0.056)	(0.115)	(0.099)	(0.068)	(0.068)
Md inc HH	0.084	0.087	−0.136	−0.121	0.027	0.027
	(0.064)	(0.064)	(0.184)	(0.177)	(0.056)	(0.055)
Manufacturing	0.125***	0.129***	0.080	0.085	0.059	0.062
	(0.048)	(0.048)	(0.078)	(0.067)	(0.047)	(0.047)
Speak only english	−0.309***	−0.306***	−0.195*	−0.190**	−0.360***	−0.360***
	(0.051)	(0.051)	(0.101)	(0.094)	(0.041)	(0.042)
Below poverty level	0.021	0.020	−0.106	−0.110	−0.015	−0.017
	(0.039)	(0.039)	(0.094)	(0.089)	(0.050)	(0.050)
White	−0.069	−0.067	0.012	−0.006	−0.132**	−0.134**
	(0.066)	(0.066)	(0.098)	(0.093)	(0.062)	(0.063)
LFPR	−0.073*	−0.072*	−0.163	−0.122	−0.059	−0.061
	(0.041)	(0.040)	(0.145)	(0.160)	(0.055)	(0.055)
Unem rate	0.040	0.043	0.147**	0.150**	0.034	0.037
	(0.041)	(0.041)	(0.074)	(0.070)	(0.036)	(0.036)
Rural	0.017	0.027	−0.136	−0.085	−0.066	−0.064
	(0.037)	(0.037)	(0.096)	(0.102)	(0.059)	(0.058)
Constant			0.300***	0.000	−0.083	−0.116
			(0.108)	(0.110)	(0.071)	(0.072)
Observations	1,657	1,657	234	234	1,657	1,657
R^2	0.420	0.424	0.389	0.388	0.455	0.455
Election FE	Y	Y	N	N	N	N

Note: *p<0.1; **p<0.05; ***p<0.01
Robust standard errors clustered at the DMA-election indicated in parentheses.

are comparing otherwise similar counties that differ only in the timing of their exposure and the number of cases experienced, which are both exogenous events. As indicated in columns 3 and 4 of Table 2, this matching strategy strengthens our conclusions, suggesting that exposure to the pandemic reduces support for the Sanders' campaign by almost 60% of a standard deviation for the binary measure (column 3), and a third of a standard deviation for the continuous measure (column 4). These suggest that the substantive impact on COVID-19 exposure is non-trivial, accounting for approximately 7 percentage points slippage for Sanders vote share between 2016 and 2020.

However, matching strategies such as the method we implement require us to jettison a substantial number of observations. As indicated at the bottom of Table 2, we rely on less than 15% of our total observations to draw these conclusions, choosing only those control observations that are most similar to the treated according to the Mahalanobis distance measure across the 12-dimensional covariate space. As a final test, we instead employ a weighting strategy that re-weights the control observations to best approximate the treated observations, without throwing any information away. Specifically, we implement the optimal weighting method of Zubizarreta (2015), achieve good balance across all observables, as summarized in Table 3 in our Supporting Information.

Columns 5 and 6 summarize the weighted estimates, suggesting that exposure to the novel coronavirus predicts a decline in support quite similar to the unmatched regression in columns 1 and 2 - with Sanders' support declining slightly less than half a standard deviation for the binary treatment measure, and 0.12 standard deviations in response to a standard deviation increase in cases. Notably, our predictive power increases meaningfully from an R^2 of just under 0.40 to over 0.45 with the weighting method employed in columns 5 and 6.

Differences-in-Differences

The preceding results exploit temporal variation in exposure, but operationalize this variation in cross-sectional statistical analyses. In the following section, we instead turn to a difference-in-differences specification in which we compare the difference between treated and control counties prior to the outbreak with the difference in Sanders support among these groups of counties following the outbreak.

Figure 2 plots the simple averages of treated (blue) and control (red) groups prior to (left) and following (right) the outbreak of the virus. Based solely on this simple difference-in-differences, one might draw several conclusions. First, there appears to be a decline in support for Sanders among both treated and control counties following the outbreak of the novel coronavirus. Second, there is some evidence suggesting that the counties that were exposed to the virus and voted after the outbreak shifted more strongly against Sanders than those counties that were not exposed. Simple bivariate regressions across groups suggest that there is no correlation between the number of cases (logged, x-axis) and the change in Sanders vote share between 2016 and 2020 in the pre-outbreak period (March 3rd and earlier). Conversely, there is a clear negative correlation following the outbreak.

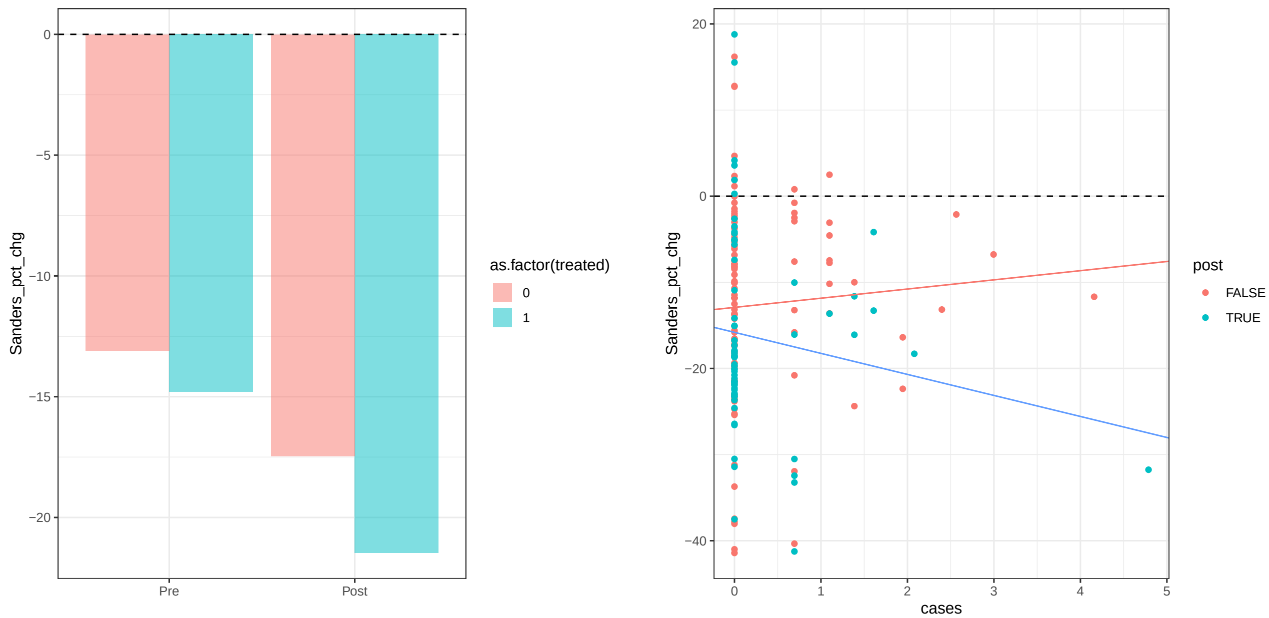

Figure 2: Descriptive differences between treated and control voting behavior before and after the outbreak, defined as starting on March 4th. Left panel groups counties by whether they were exposed as of March 9th, right panel plots the logged cases as of March 9th by whether the county voted prior to, or following, the outbreak.

These plots are descriptive, and are not meant to support well-identified inferential conclusions. As such, we turn to our conditional difference-in-difference regression specifications. We examine both the basic conditional results as well as the matched and weighted results using different dates for the beginning of "treatment" in Figure 3. When we set the treatment period to March 1st and include the exposed counties voting in Super Tuesday among our treated group, we find significant evidence that exposure leads to declining support for Bernie Sanders. However, this effect declines over time, with the result attenuating to a null when we define the outbreak starting after Super Tuesday and even some suggestive evidence that the virus actually benefited Sanders among the counties voting on March 17th. We suspect that these patterns reflect a broadening of the national coverage of the outbreak, prompting SUTVA violations when we define exposure at the DMA. We test this suspicion in our Supporting Information, and find that redefining the unit of exposure at the state level recovers our main results. This suggests that the "flight to safety" occurs throughout, but that as time passes information about diagnoses induces panic not just within the DMA but state-wide.

The Supplementary Information also includes additional analyses specified in the pre-analysis plan, including robustness checks, exploring sensitivity to shifting definitions of treatment, choices of matching strategy, and balancing weights. We also adjust the geographic definition of treatment to the county, the DMA, and the state, and compare our results across different outbreak dates (thus allowing a detailed examination of whether e.g. party consolidation of support behind Biden explains these results - we believe it does not).

These tests confirm the main findings described above, with stronger results if we define treatment using deaths due to COVID-19 instead of confirmed cases. We also find that differential turnout is not a likely explanation of these findings, with turnout not being suppressed until the March 17th election and no evidence that age was particularly determinative. Finally, we run a placebo test by permuting treatment. This test, in concert with our comparisons between different election dates and the election date fixed effects we include in our main specifications above in Table 2, strongly suggest that our findings are not being driven by any secular Democratic party elite consolidation behind Biden over time.

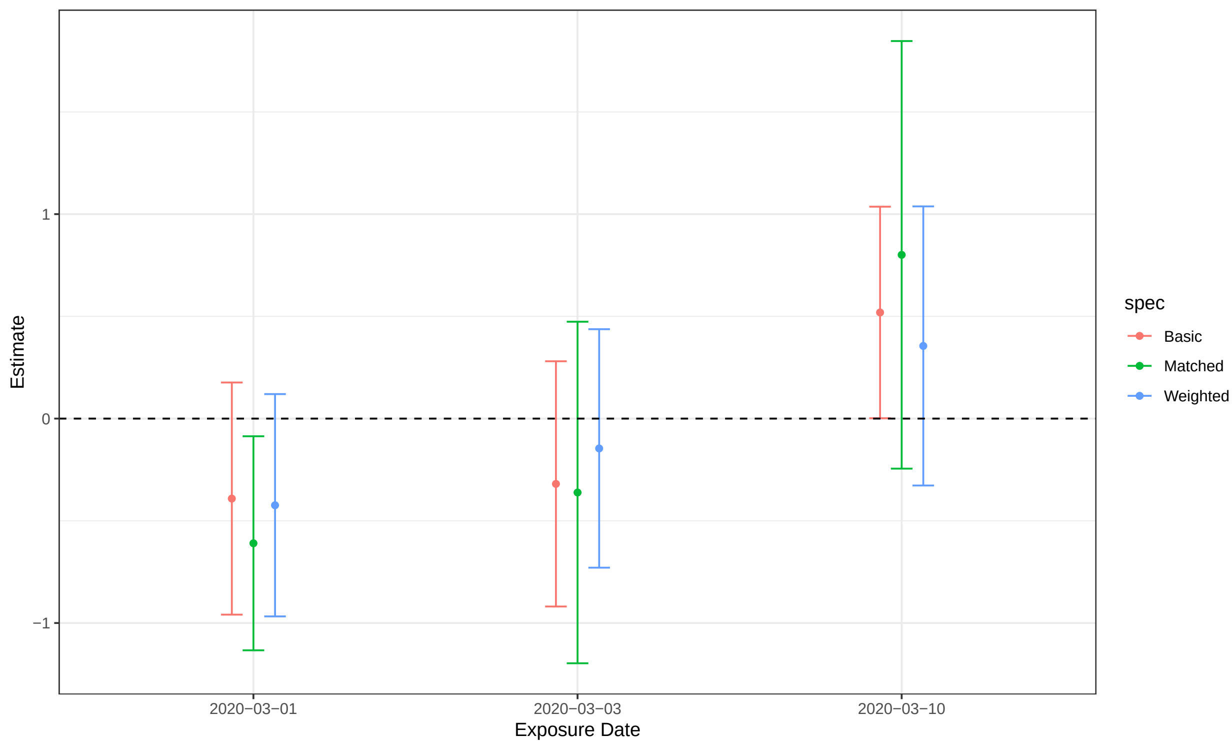

Figure 3: Diff-in-diff estimates for different start-dates of the outbreak (x-axis).

5 Discussion

These findings explore the substantive political effects of the novel coronavirus. They also help us understand the effects of anxiety on voting. These findings complement those of Campante, Depetris-Chauvin and Durante 2020, who find Ebola-induced fear had substantial electoral consequences in the 2014 US midterm elections. As they put it, "emotional reactions associated with fear can have a strong electoral impact".

The size of the COVID-19 effect is not large enough to have made Bernie Sandersthe front-runner in the absence of the novel coronavirus' appearance; it is, however, a far from trivial effect. We find in our primary specification (Table 2) that COVID-19 exposure in a local media market depresses vote shares for the relatively more extreme candidate by up to 7 percentage points.

These findings suggest that psychological factors may warrant consideration in assessing what conditions are necessary for a free and fair election. Those considering the timing of elections may wish to consider the potential for anxiety in considering dates for elections both Democratic primaries and potentially November 2020's general election as well. Additionally, it is possible that actors (foreign or domestic) may wish to induce a "flight to safety" in an effort to influence electoral outcomes. Under certain conditions those wishing to manipulate the vote towards more status quo candidates need not go anywhere near voting machines or registration lists; they simply need to propagate anxiety and fear.

References

Abramson, Paul R, John H Aldrich, Jill Rickershauser and David W Rohde. 2007. "Fear in the voting booth: The 2004 presidential election." *Political Behavior* 29(2):197–220.

Adrian, Tobias, Richard K Crump and Erik Vogt. 2019. "Nonlinearity and Flight-to-Safety in the Risk-Return Trade-Off for Stocks and Bonds." *The Journal of Finance* 74(4):1931–1973.

Blackwell, Matthew, Stefano Iacus, Gary King and Giuseppe Porro. 2009. "cem: Coarsened exact matching in Stata." *The Stata Journal* 9(4):524–546.

Campante, Filipe, Emilio Depetris-Chauvin and Ruben Durante. 2020. "The Virus of Fear: The Political Impact of Ebola in the US.".

Canes-Wrone, Brandice, Michael C Herron and Kenneth W Shotts. 2001. "Leadership and pandering: A theory of executive policymaking." *American Journal of Political Science* pp. 532–550.

Eilperin, Juliet. 2020. "Bernie Sanders's climate record in Congress: Lots of advocacy, no compromise." *Washington Post* .
URL: *https://www.washingtonpost.com/climate-environment/bernie-sanderss-climate-record-in-congress-lots-of-advocacy-no-compromise/2020/03/12/7aa2772a-5fd7-11ea-b29b-9db42f7803a7$_s$tory.html*

Fegenheimer, Matt and Katie Glueck. 2020. "Joe Biden Had a Big Night. He Needs Another in 72 Hours." *New York Times* .
URL: *https://www.nytimes.com/2020/02/29/us/politics/joe-biden-super-tuesday.html*

Getmansky, Anna and Thomas Zeitzoff. 2014. "Terrorism and voting: The effect of rocket threat on voting in Israeli elections." *American Political Science Review* 108(3):588–604.

Green, Donald Philip and Bradley Palmquist. 1994. "How stable is party identification?" *Political behavior* 16(4):437–466.

Gutiérrez, Sandra Jessica Ley. 2014. Citizens in fear: political participation and voting behavior in the midst of violence PhD thesis tesis de doctorado, Durham, Duke University, 2014), en especial 1-10 . . . .

Imai, Kosuke and Marc Ratkovic. 2014. "Covariate balancing propensity score." *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 76(1):243–263.

Inghelbrecht, Koen, Geert Bekaert, Lieven Baele and Min Wei. 2013. "Flights to safety." *NBER WORKING PAPER SERIES* 19095:1–51.

Leonhardt, David. 2020. "A Complete List of Trump's Attempts to Play Down Coronavirus." *New York Times* .
URL: *https://www.nytimes.com/2020/03/15/opinion/trump-coronavirus.html*

Maskin, Eric and Jean Tirole. 2004. "The politician and the judge: Accountability in government." *American Economic Review* 94(4):1034–1054.

Montalvo, Jose G. 2011. "Voting after the bombings: A natural experiment on the effect of terrorist attacks on democratic elections." *Review of Economics and Statistics* 93(4):1146–1154.

Nezi, Roula. 2012. "Economic voting under the economic crisis: Evidence from Greece." *Electoral studies* 31(3):498–505.

Remmer, Karen L. 1991. "The political impact of economic crisis in Latin America in the 1980s." *American Political Science Review* 85(3):777–800.

Roose, Kevin and Gabriel J.X. Dance. 2020. "The Coronavirus Revives Facebook as a News Powerhouse." *New York Times* .
URL: *https://www.nytimes.com/2020/03/23/technology/coronavirus-facebook-news.html*

Smith, Alastair. 1998. "International crises and domestic politics." *American Political Science Review* 92(3):623–638.

Staff, FiveThirtyEight. 2020. "Who Will Win the 2020 Democratic Primary?" *FiveThirtyEight.com* .
URL: *https://projects.fivethirtyeight.com/2020-primary-forecast/?cid=rrpromo*

Stewart, Cameron. 2020. "Now it's Bernie v Biden." *The Australian* .
URL: *https://www.theaustralian.com.au/inquirer/super-tuesday-now-its-bernie-v-biden/news-story/11aff0ef309b7ce6be787d4457ab9510*

Yglesias, Matthew and Zack Beauchamp. 2020. "4 winners and 3 losers on Super Tuesday." *Vox.com* .
URL: *https://www.vox.com/2020/3/3/21163826/super-tuesday-winners-losers-biden-sanders*

Zubizarreta, José R. 2015. "Stable weights that balance covariates for estimation with incomplete outcome data." *Journal of the American Statistical Association* 110(511):910–922.

Supplementary Information

"Party Decides" Placebo Tests

The main results suggest that exposure to the novel coronavirus results in a greater decline in support for a Sanders presidency than what we observe in relatively insulated counties or those that voted prior to the outbreak. However, even with our matching and weighting strategies to argue that the outbreak is as-good-as-randomly assigned conditional on observables, there remains a concern with regards to timing. Specifically, our definition of "exposure" is defined as any county residing within a DMA that had confirmed cases of the virus as of March 9th, 2020. Effectively, this definition risks conflating other contemporaneous changes in the political landscape that occurred between Super Tuesday (March 3rd), and the 10 states that voted afterwards (7 on March 10th, 3 on March 17th). Specifically, this period saw the Democratic party rally around the establishment candidacy of Joseph Biden as several candidates dropped out of the race and endorsed Biden.

To confirm our results are not simply picking temporal variation and the momentum shift that occurred on Super Tuesday, we run a placebo test in which we permute our explanatory variable while keeping our definitions of pre and post exposure at March 3rd. If our main results are driven by the "party decides" phenomenon, we should still find a significant negative relationship between Sanders' declining vote share and our permuted treatment. We bootstrap sample our data, each time drawing a permuted explanatory variable, and re-estimate our main specifications. As illustrated in Figure 4, our results are noisily estimated nulls.

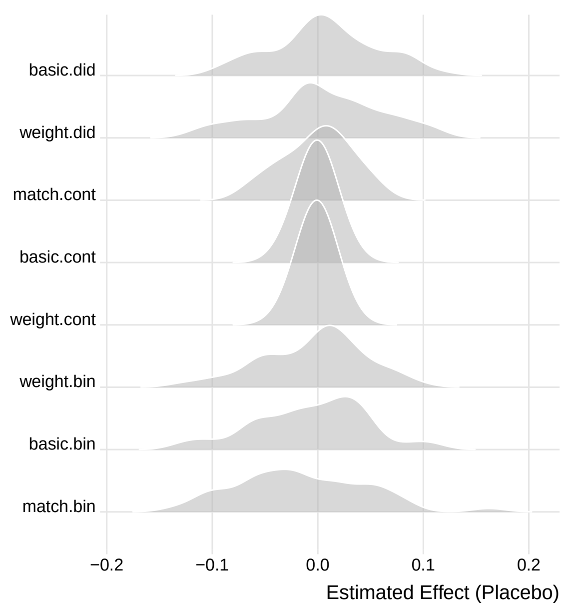

Figure 4: Placebo test bootstrapped estimates generated by permuting the COVID-19 cases while keeping the exposure start date starting after Super Tuesday.

The results summarized above use outbreak dates to separate treated and control elections as per Table 1, meaning that all elections prior and including a given cutoff are defined as control, and all elections following the cutoff are treated. We also re-run our analyses by conducting a series of pairwise comparisons in which one election is defined as control and the other is defined as treated. Doing so allows us to identify where (and more precisely, when) our effects obtain. We treat all primary elections prior to Super Tuesday as one group in order to include multiple states in each treatment and control condition. Figure 5 summarizes these results for every specification at our disposal. The Democratic party consolidated support behind Biden ahead of Super Tuesday. As Figure 5 demonstrates, the results do not depend on comparing the period before Super Tuesday to the period after, and thus are not collinear with a "party consolidation" effect, though we cannot rule out that such an effect may also contribute to the findings in the panel comparing Super Tuesday to pre-Super Tuesday voting states.

We can also reverse the temporal sequencing of these results, creating placebo tests for our conclusions. We treat the later election as the control, and the earlier as the treated, and estimate the effect of future COVID-19 cases on vote choice. Our results are summarized in Figure 6

Turnout and Age

Our main results suggest that Bernie Sanders was hurt by the outbreak of the virus, although the effect attenuated over time. We argue that this is consistent with our theorized mechanism of an electoral "flight to safety". However, an alternative mechanism might be that the outbreak differentially reduced turnout among different voting groups. One plausible scenario might be that those most threatened by exposure might be less likely to turn out. If this group is also more likely to support Sanders, there is an alternative explanation for the effects we document. Of course, Sanders' popularity among young voters is well-documented, while the elderly are most threatened by the virus. As such, if this mechanism is operating, it should be the case that older voters are *less* likely to turn out, and that therefore we should see an *increase* in support for Sanders from younger voters.

We examine this alternative mechanism by replacing the change in Sanders' vote share with the change in county-level turnout. As illustrated in Figure 7, there is little evidence to suggest that such an age-based dynamic is at play. As illustrated, it appears that turnout wasn't suppressed by the virus until after the March 10th elections, when it reduced turnout for Illinois, Arizona, and Florida. (Ohio chose to delay its primaries due to concerns about the virus.) Furthermore, the marginal effects of exposure across different populations of older voters suggests that COVID-19 was *less* depressing to turnout in counties with more elderly constituents, although the interaction effects are not statistically significant. Taken together, these findings suggest that turnout was not appreciably influenced by COVID-19 exposure as of March 3rd (95% confidence intervals for the marginal effects always include zero) and that, although there is some evidence of heterogeneity across counties by share of the population

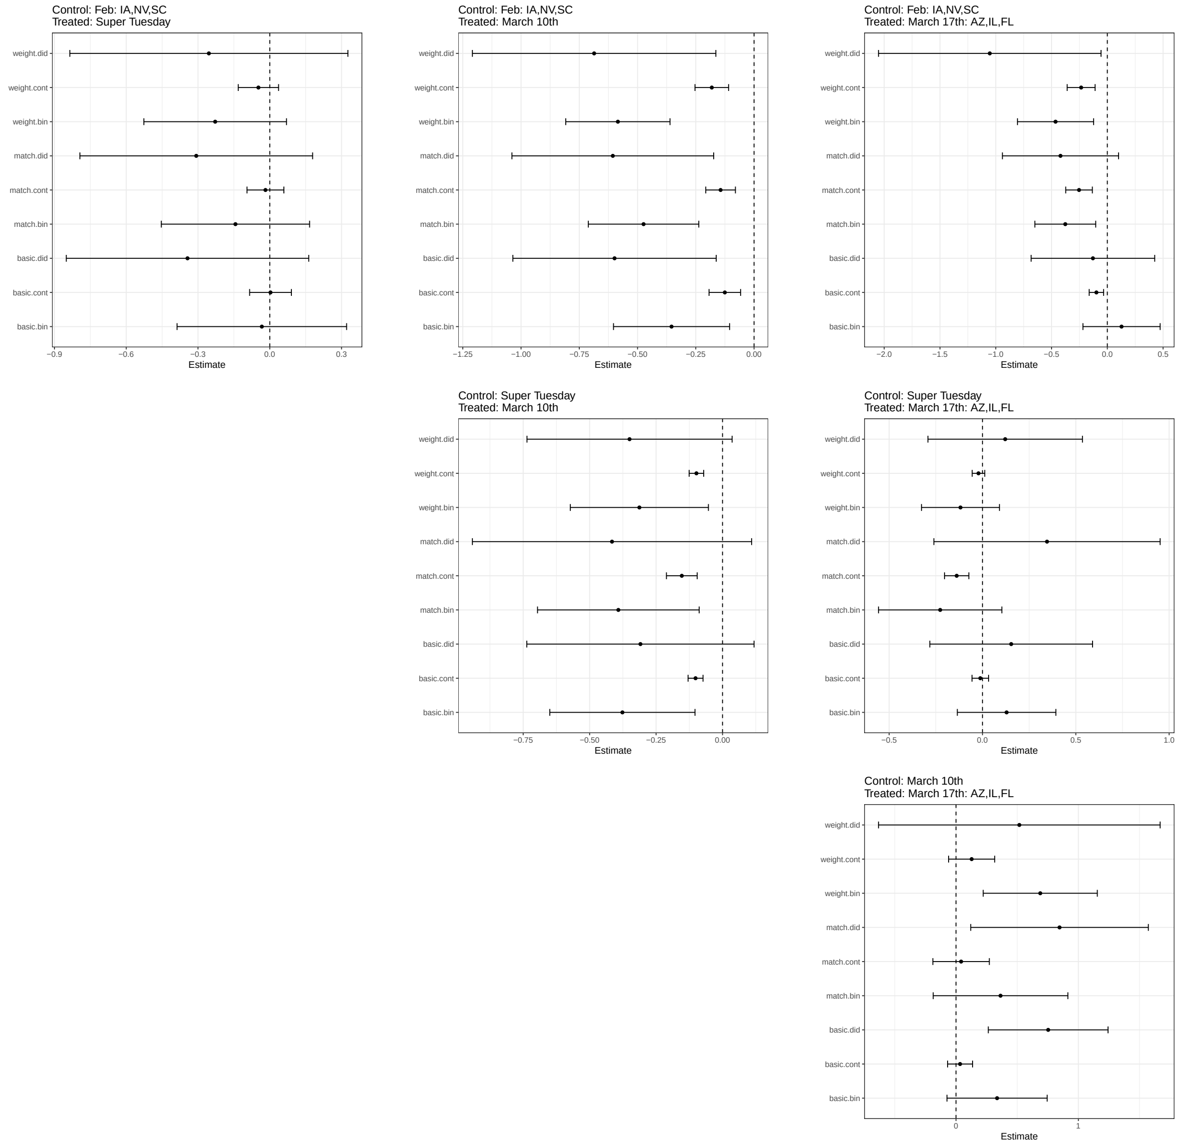

Figure 5: Pairwise election comparisons by control (rows) and treatment (columns) given in plot titles.

older than 64 years of age, these interaction coefficients are also insignificant. As such, we are confident in our conclusion that the reduction in Sanders' support is attributable to more than simply shifting turnout dynamics across the period of analysis.

Balance and Weighting Robustness

We achieve good balance on both the matching and weighting strategies employed in the body of our paper. Figure 8 plots the improvements to balance on observables between treated and control units generated by our choice of nearest-neighbor matching using mini-

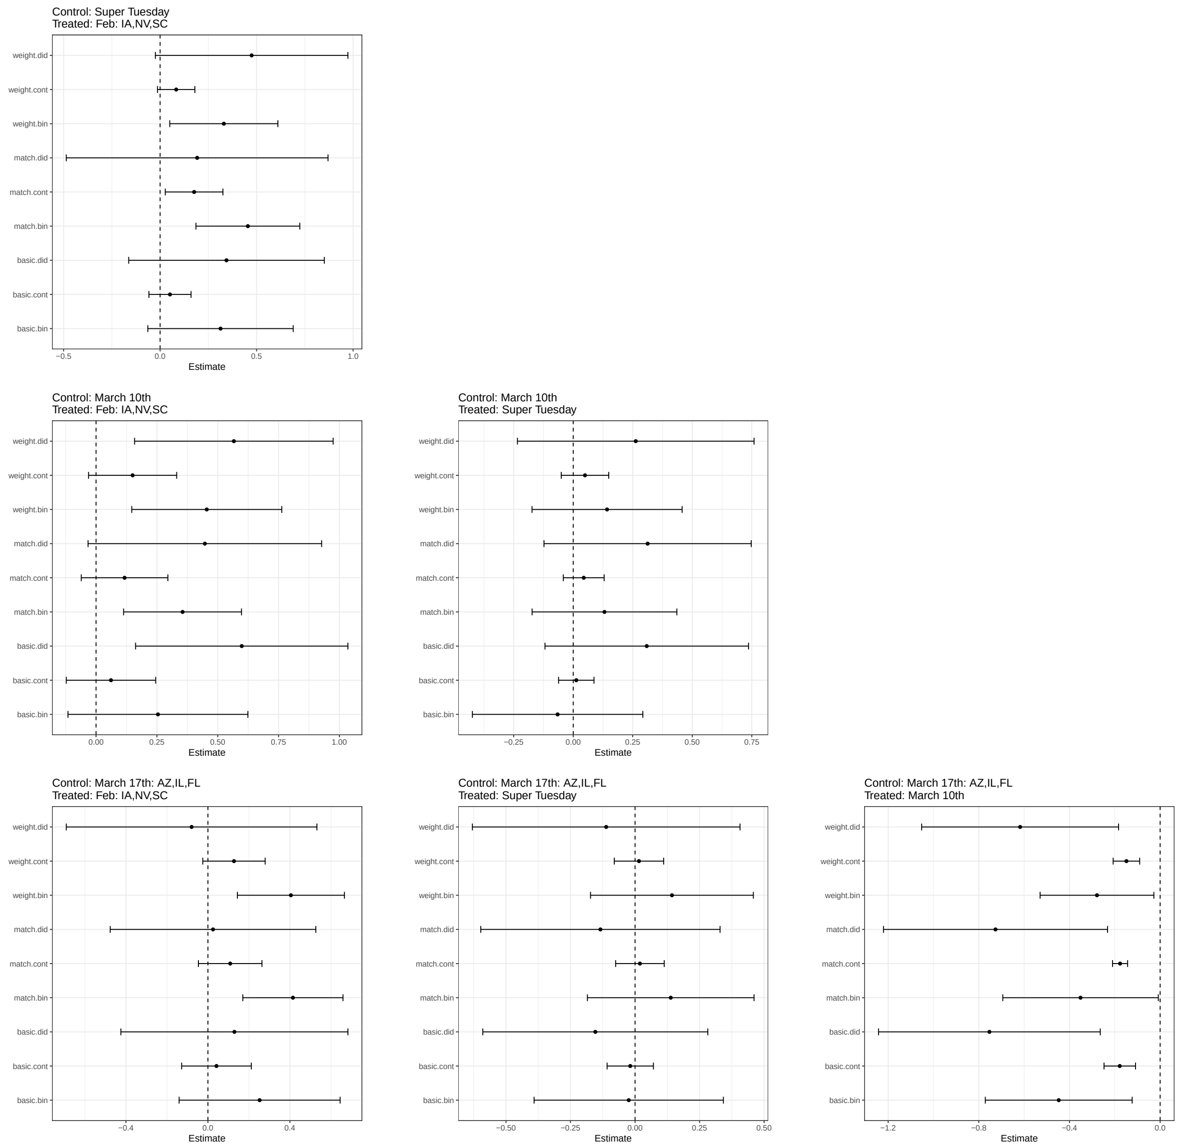

Figure 6: Pairwise election comparisons by control (rows) and treatment (columns) given in plot titles.

mized Mahalanobis distance. And Table 3 summarizes the differences in treated and control covariates prior to, and following the `optmatch` weights. In both cases, we successfully adjust our data to better reflect the distribution of observables in an experimental context in which treatment is randomly assigned.

We also confirm the robustness of our main findings to different choices about the matching strategy and the balancing weights. Specifically, we re-estimate our main findings replacing the `optweight` method of Zubizarreta (2015) with covariate balancing propensity scores (CBPS, Imai and Ratkovic 2014), and replacing the nearest neighbor matching strategy with coarsened exact matching (CEM, Blackwell et al. 2009). The former robustness check yields substantively and statistically similar findings to our main results, as illustrated

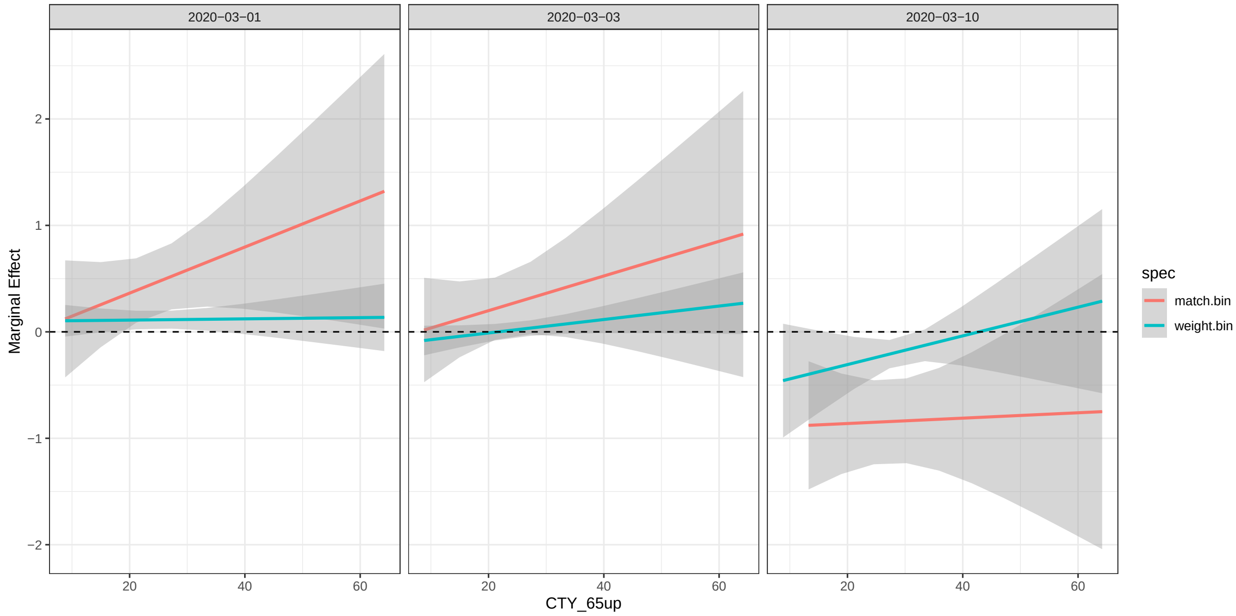

Figure 7: Marginal effects of exposure on turnout (y-axes) across counties with smaller and larger proportions of their population older than 64 years of age (x-axes) by outbreak onset date (panels). None of the marginal effects are themselves significant at conventional levels.

Table 3: Weighting Balance Checks

	Covs	Diff_Unm	Bal_Test_Unm	Diff_Match	Bal_Test_Match
1	County Pop	0.320	Not Balanced, >0.05	0	Balanced, <0.05
2	Old-Age Dep Ratio	0.240	Not Balanced, >0.05	0	Balanced, <0.05
3	% Bachelor's	0.370	Not Balanced, >0.05	0	Balanced, <0.05
4	Female HH, No Hub	-0.410	Not Balanced, >0.05	0	Balanced, <0.05
5	Median HH Inc	0.370	Not Balanced, >0.05	0	Balanced, <0.05
6	% Manuf	-0.690	Not Balanced, >0.05	0	Balanced, <0.05
7	% Speak English	-0.300	Not Balanced, >0.05	0	Balanced, <0.05
8	% Below Pov	-0.160	Not Balanced, >0.05	0	Balanced, <0.05
9	% White	0.250	Not Balanced, >0.05	0	Balanced, <0.05
10	LFPR	-0.100	Not Balanced, >0.05	0	Balanced, <0.05
11	Unemp Rate	0.130	Not Balanced, >0.05	0	Balanced, <0.05
12	%Rural	-0.770	Not Balanced, >0.05	0	Balanced, <0.05

in Tables 4 and 5.

Moving from nearest neighbor matching based on Mahalanobis distance to the CEM method requires us to reduce the number of county-level covariates we use for matching. This is due to the default parameter settings yielding only two matched observations, precluding our ability to estimate treatment effects. We reduce our set of covariates to select the following six across which we can obtain reasonably good performance on our balance tests while also obtaining enough observations for statistical inference:

Table 4: Main results using CBPS instead of optweights

	Dependent variable:					
	bsc.bin	bsc.cont	mtc.bin	mtc.cont	wgt.bin	wgt.cont
	(1)	(2)	(3)	(4)	(5)	(6)
Treatbin	−0.361		−0.600**		−0.496**	
	(0.238)		(0.258)		(0.225)	
Treatcont		−0.102***		−0.300***		−0.106***
		(0.013)		(0.038)		(0.011)
Tot pop	0.036	0.032	0.131**	0.109*	0.017	0.013
	(0.030)	(0.028)	(0.065)	(0.063)	(0.026)	(0.021)
Old age dep ratio	0.024	0.017	0.047	0.055	0.065*	0.047
	(0.046)	(0.046)	(0.087)	(0.088)	(0.035)	(0.030)
Bachelor's degree	−0.071	−0.066	0.156	0.194*	0.012	0.049
	(0.054)	(0.055)	(0.114)	(0.106)	(0.078)	(0.078)
Female HH no husband	0.380***	0.383***	0.346***	0.329***	0.354***	0.317***
	(0.053)	(0.054)	(0.115)	(0.096)	(0.081)	(0.075)
Md inc HH	0.084	0.087	−0.136	−0.121	−0.096	−0.081
	(0.061)	(0.061)	(0.183)	(0.173)	(0.087)	(0.082)
Manufacturing	0.125***	0.129***	0.080	0.085	−0.011	0.002
	(0.048)	(0.048)	(0.077)	(0.064)	(0.068)	(0.055)
Speak only english	−0.309***	−0.306***	−0.195**	−0.190**	−0.211***	−0.215***
	(0.050)	(0.050)	(0.097)	(0.088)	(0.066)	(0.065)
Below poverty level	0.021	0.020	−0.106	−0.110	−0.070	−0.084*
	(0.040)	(0.040)	(0.093)	(0.083)	(0.052)	(0.049)
White	−0.069	−0.067	0.012	−0.006	−0.111*	−0.130**
	(0.064)	(0.064)	(0.091)	(0.084)	(0.059)	(0.063)
LFPR	−0.073**	−0.072**	−0.163	−0.122	−0.008	−0.023
	(0.036)	(0.036)	(0.143)	(0.158)	(0.067)	(0.070)
Unem rate	0.040	0.043	0.147*	0.150**	0.138**	0.165***
	(0.040)	(0.040)	(0.077)	(0.073)	(0.057)	(0.060)
Rural	0.017	0.027	−0.136	−0.085	−0.153***	−0.131*
	(0.038)	(0.039)	(0.091)	(0.090)	(0.057)	(0.069)
Constant			0.300***	0.000	−0.107	−0.285*
			(0.107)	(0.124)	(0.092)	(0.163)
Observations	1,657	1,657	234	234	1,657	1,657
R^2	0.420	0.424	0.389	0.388	0.406	0.418
Note:						*p<0.1; **p<0.05; ***p<0.01

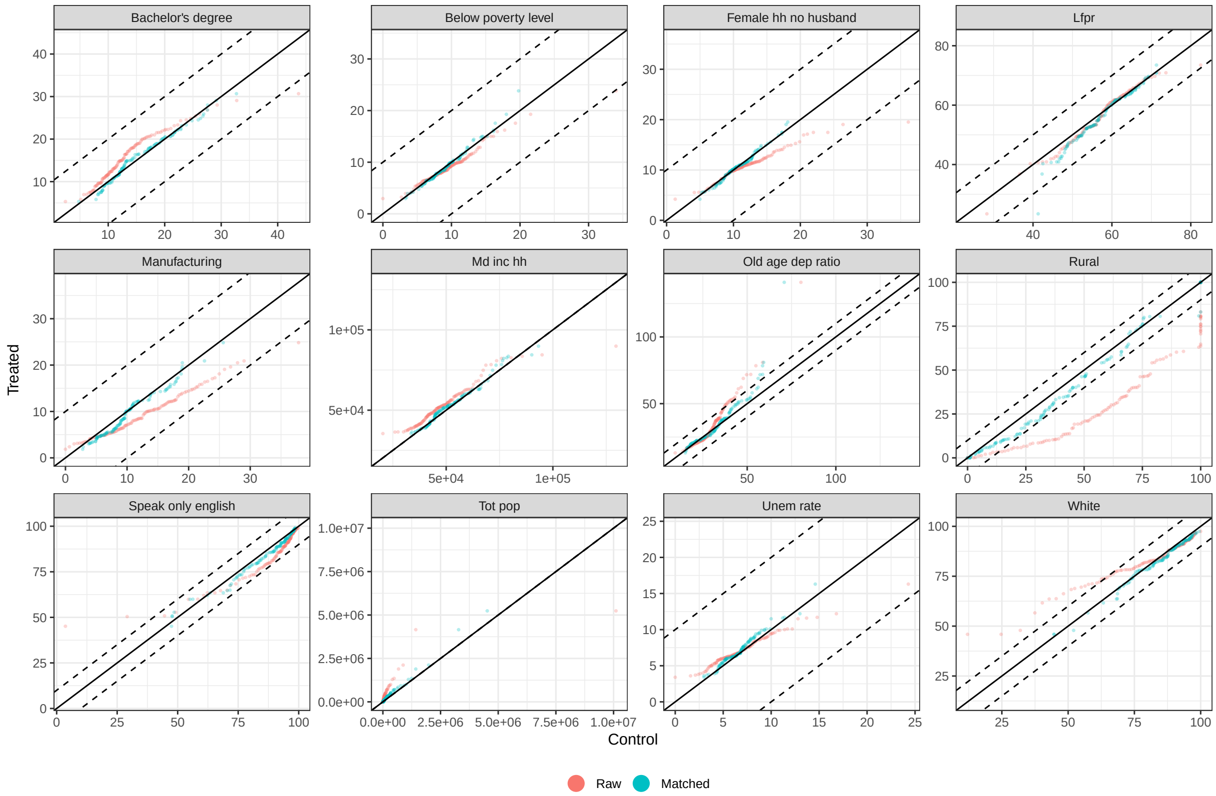

Figure 8: Balance of treated and control covariates before (red) and after (blue) matching. 45 degree line indicates perfect match.

Table 5: Balance results for CBPS

	Covs	Diff_Unm	Bal_Test_Unm	Diff_Match	Bal_Test_Match
1	CTY_tot_pop	0.320	Not Balanced, >0.05	0	Balanced, <0.05
2	CTY_Old_age_dep_ratio	0.240	Not Balanced, >0.05	0.110	Not Balanced, >0.05
3	CTY_Bachelor_s_degree	0.370	Not Balanced, >0.05	0.010	Balanced, <0.05
4	CTY_Female_hher_no_husbandhh	-0.410	Not Balanced, >0.05	-0.050	Balanced, <0.05
5	CTY_Md_inc_hhs	0.370	Not Balanced, >0.05	0.010	Balanced, <0.05
6	CTY_Manufactur	-0.690	Not Balanced, >0.05	-0.020	Balanced, <0.05
7	CTY_Speak_only_English	-0.300	Not Balanced, >0.05	0.010	Balanced, <0.05
8	CTY_Below_poverty_level_AGE_18_64	-0.160	Not Balanced, >0.05	-0.030	Balanced, <0.05
9	CTY_White	0.250	Not Balanced, >0.05	0.010	Balanced, <0.05
10	CTY_Labor_Force_Part_Rate_pop_16_over	-0.100	Not Balanced, >0.05	-0.070	Not Balanced, >0.05
11	CTY_Unem_rate_pop_16_over	0.140	Not Balanced, >0.05	0	Balanced, <0.05
12	CTY_POPPCT_RURAL	-0.770	Not Balanced, >0.05	0	Balanced, <0.05

- % 65 and older
- % with bachelor's degree
- Median household income
- % speak only English
- County unemployment rate

- % White

These choices reduce the number of total observations to 152 but yield substantively and statistically similar results to our main findings, as illustrated in Table 6. The balance test results are visualized in Figure 9.

Table 6: Main Results Estimated with CEM instead of nearest neighbor matching

	Dependent variable:					
	bsc.bin	bsc.cont	mtc.bin	mtc.cont	wgt.bin	wgt.cont
	(1)	(2)	(3)	(4)	(5)	(6)
Treatbin	−0.392		−0.924**		−0.531**	
	(0.262)		(0.408)		(0.252)	
Treatcont		−0.107***		−0.257***		−0.114***
		(0.014)		(0.029)		(0.011)
Bachelor's degree	−0.178***	−0.181***	−0.080	−0.118	−0.031	−0.011
	(0.042)	(0.041)	(0.100)	(0.114)	(0.065)	(0.061)
Md inc HH	−0.018	−0.015	−0.123	−0.064	−0.069	−0.044
	(0.049)	(0.049)	(0.093)	(0.118)	(0.066)	(0.070)
Speak only english	−0.301***	−0.293***	−0.071	−0.090	−0.329***	−0.320***
	(0.051)	(0.051)	(0.091)	(0.097)	(0.063)	(0.059)
White	−0.292***	−0.290***	−0.230**	−0.243**	−0.313***	−0.309***
	(0.059)	(0.059)	(0.104)	(0.105)	(0.055)	(0.055)
Unem rate	0.117**	0.118**	0.181**	0.171**	0.164***	0.180***
	(0.049)	(0.048)	(0.072)	(0.075)	(0.060)	(0.061)
Old age dep ratio	−0.042	−0.048	−0.092	−0.075	0.009	0.013
	(0.042)	(0.041)	(0.078)	(0.076)	(0.044)	(0.031)
Constant			0.175	−0.009	−0.011	−0.213
			(0.118)	(0.128)	(0.077)	(0.156)
Observations	1,657	1,657	219	219	1,657	1,657
R^2	0.363	0.366	0.315	0.258	0.373	0.386
Note:						*p<0.1; **p<0.05; ***p<0.01

With counties nested within DMAs and states, our data facilitate multilevel models as an alternative to standard linear regression analyses, as well as allowing for more rigorous fixed effects at the DMA or state level. Per our PAP we implement both in examining the results of March 17th in Figure 10 which summarizes the impact of different fixed effects / mixed effects, producing noisier but still negative estimates for most checks.

Diff-in-Diff Over Time

We summarize the descriptive difference-in-differences for March 1st, March 3rd, and March 10th in the plots below. As illustrated, prior to Super Tuesday, counties that would become exposed were more supportive of Sanders than counties that would remain in control

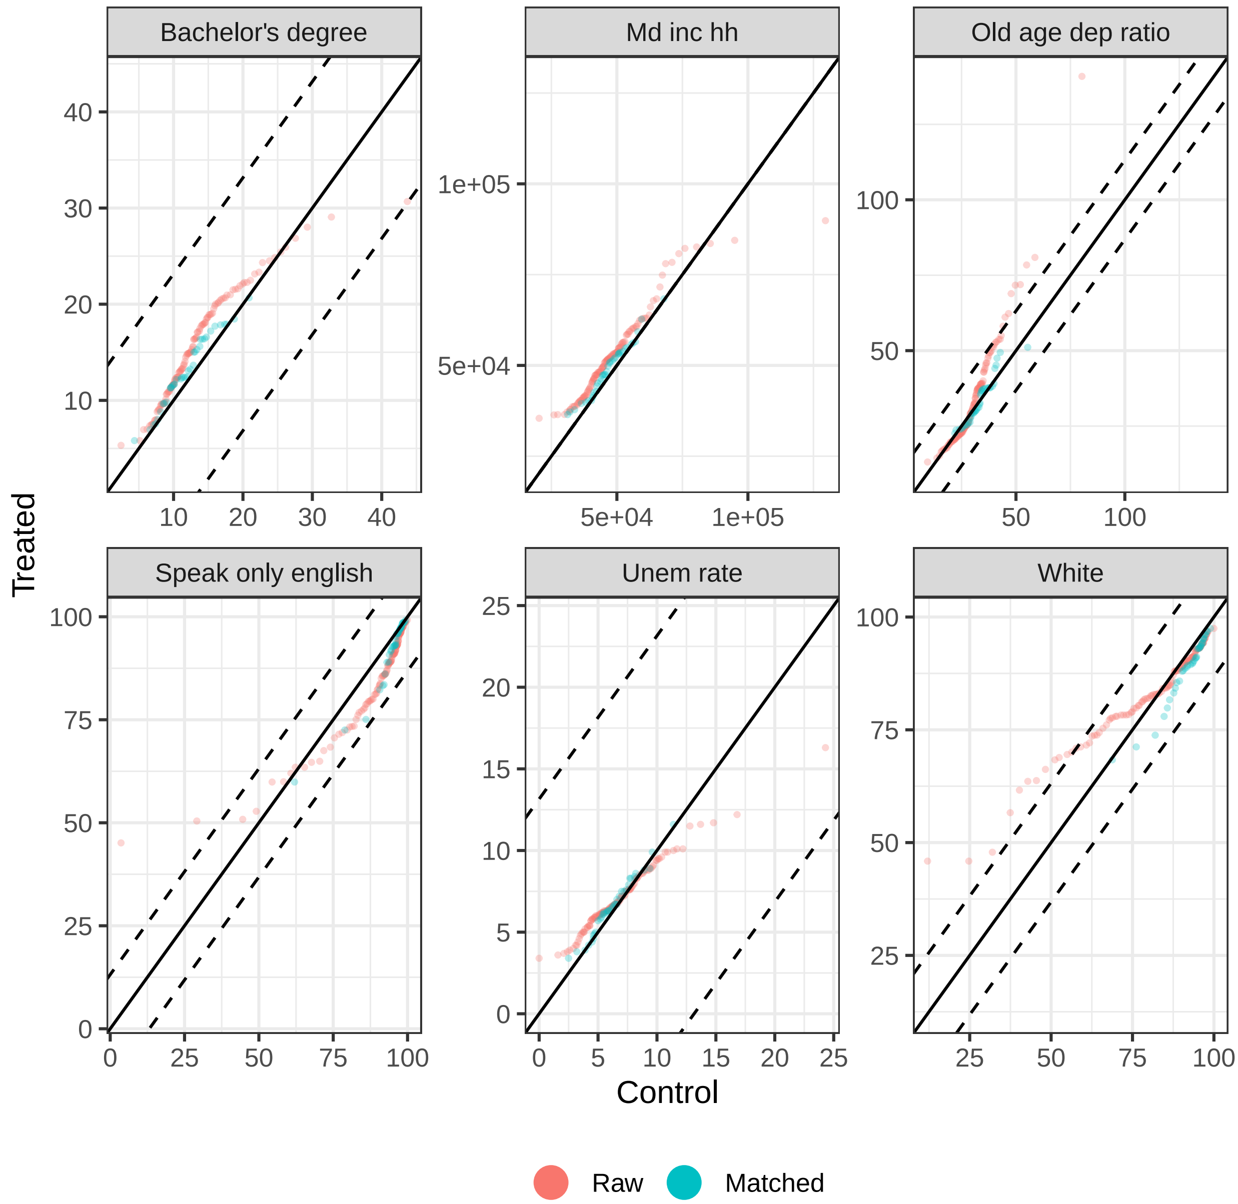

Figure 9: Balance performance across 6 covariates using CEM (Blackwell et al., 2009).

over the period of analysis (although the general shift was still away from Sanders relative to 2016). In addition, the March 17th primary voters in exposed counties also shifted less away from Sanders compared to the control counties.

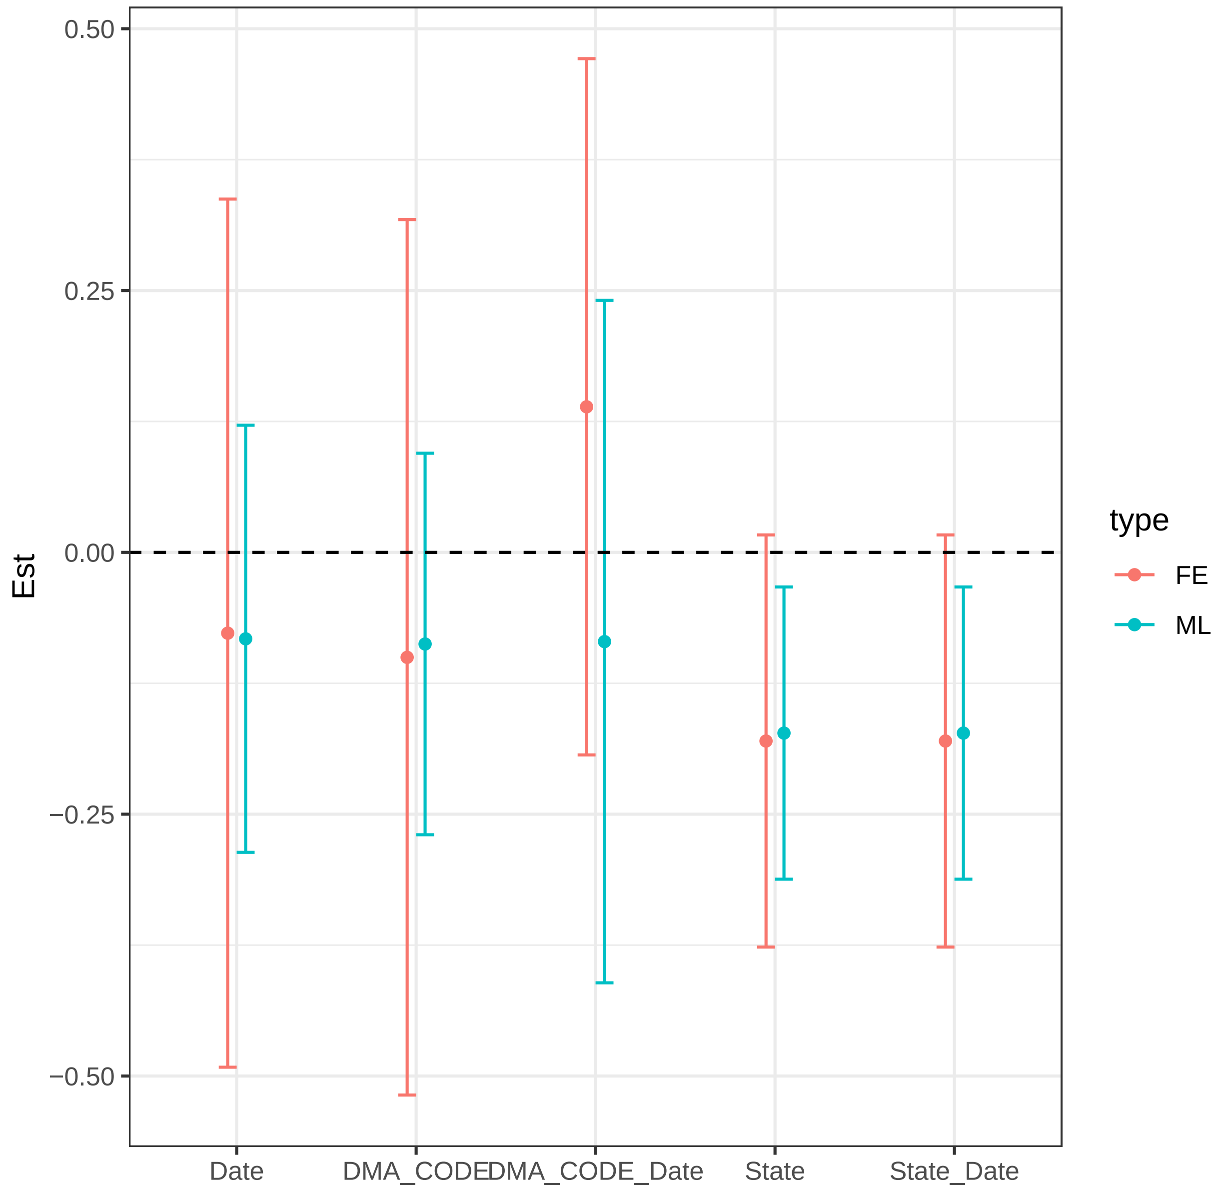

Figure 10: Estimates subject to different choices of fixed (red) and mixed (blue) effects.

Geographic Units of Treatment

Our main findings define treatment as a function of the local media market in which a county resides. The intuition is that, as the virus was initially spreading, local media was more likely to report on the virus when cases appeared in their market. However, by March 17th news about the virus was a constant fixture on national stations, suggesting that the DMA would no longer be an appropriate border by which to define exposure to the fear and uncertainty generated by the outbreak. To the extent that larger units grew more salient as

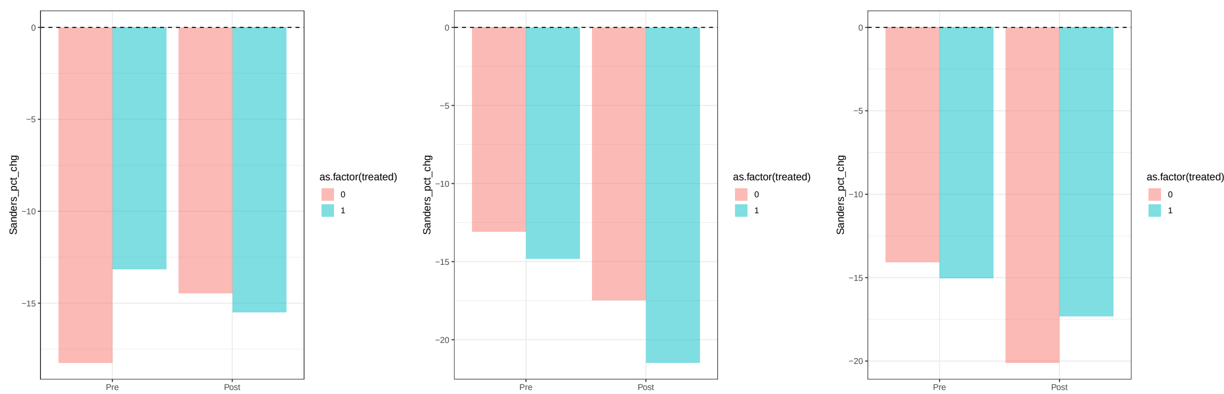

Figure 11: Descriptive DID data for outbreak dates of March 1st, March 3rd, and March 10th

overall coverage of the outbreak increased, we compare the estimates generated by defining the virus at the county, the DMA, and the state in Figure 12. As illustrated, aggregating at smaller geographic units produces null to positive results when defining the outbreak as starting after March 10th. Conversely, the negative findings persist when defining treatment assignment at the state level, although this unit appears too large for the earlier days of the outbreak when local news sources would be more appropriate for transmitting information and uncertainty.

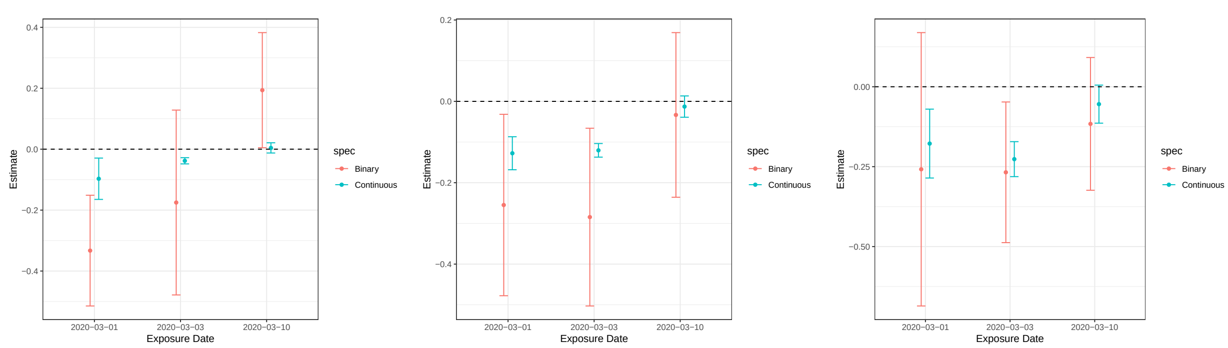

Figure 12: Matching estimates for impact of exposure on the change in Sanders' vote share when treatment is defined at the level of the county (left plot), the DMA (center plot), or at the state (right plot).

Deaths

The main results use confirmed cases of the virus to define treatment. However, we also have data on deaths due to COVID-19. We reproduce our main table, replacing confirmed cases with deaths, and find even stronger results, as illustrated in Table 7.

Table 7: Relationship between Sanders vote share and COVID-19 Deaths

	Dependent variable: Δ Sanders Vote Share 2020-2016								
	Basic		Matched		Weighted		Diff-in-Diff		
	Bin	Cont	Bin	Cont	Bin	Cont	Basic	Match	Weight
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
treatBin	-0.922^{***}		-0.995^{***}		-0.922^{***}				
	(0.191)		(0.250)		(0.185)				
treatCont		-0.124^{***}		-0.424^{***}		-0.124^{***}			
		(0.010)		(0.078)		(0.011)			
treatGroup							-0.113	0.424	-0.030
							(0.092)	(0.344)	(0.094)
post							-0.203^{*}	-0.116	-0.161
							(0.116)	(0.305)	(0.136)
Tot Pop	0.021	0.023	0.097	0.089	0.001	0.001	0.028	0.106	0.008
	(0.025)	(0.025)	(0.145)	(0.143)	(0.024)	(0.025)	(0.028)	(0.148)	(0.028)
Old Age Dep	0.024	0.018	0.111	0.228	0.058	0.057	0.023	0.109	0.063^{*}
	(0.049)	(0.049)	(0.264)	(0.200)	(0.037)	(0.037)	(0.050)	(0.264)	(0.037)
Bach Deg	-0.065	-0.063	0.125	0.066	-0.042	-0.040	-0.070	0.147	-0.053
	(0.051)	(0.051)	(0.227)	(0.217)	(0.061)	(0.061)	(0.053)	(0.240)	(0.065)
Female HH	0.369^{***}	0.383^{***}	0.176	0.216^{*}	0.321^{***}	0.319^{***}	0.366^{***}	0.173	0.303^{***}
	(0.054)	(0.053)	(0.157)	(0.126)	(0.083)	(0.083)	(0.054)	(0.161)	(0.089)
Med HH Inc	0.079	0.084	-0.172	0.060	0.053	0.058	0.082	-0.223	0.054
	(0.065)	(0.065)	(0.281)	(0.239)	(0.062)	(0.061)	(0.064)	(0.297)	(0.062)
% Manuf	0.120^{**}	0.127^{**}	0.009	0.085	0.089^{*}	0.092^{*}	0.121^{**}	0.018	0.094^{*}
	(0.050)	(0.050)	(0.113)	(0.097)	(0.054)	(0.053)	(0.050)	(0.114)	(0.052)
Speak English	-0.321^{***}	-0.320^{***}	-0.273^{*}	-0.223	-0.416^{***}	-0.415^{***}	-0.301^{***}	-0.272^{*}	-0.398^{***}
	(0.043)	(0.043)	(0.139)	(0.152)	(0.055)	(0.055)	(0.048)	(0.143)	(0.064)
Below Pov	0.005	-0.002	-0.010	0.054	0.041	0.039	0.011	-0.032	0.042
	(0.040)	(0.041)	(0.222)	(0.187)	(0.045)	(0.046)	(0.039)	(0.227)	(0.047)
% White	-0.116^{*}	-0.113^{*}	0.078	0.097	-0.083	-0.084	-0.096	0.073	-0.071
	(0.066)	(0.066)	(0.172)	(0.175)	(0.101)	(0.101)	(0.064)	(0.175)	(0.100)
LFPR	-0.079^{*}	-0.074^{*}	-0.113	-0.022	-0.063	-0.064	-0.082^{**}	-0.099	-0.066
	(0.043)	(0.043)	(0.221)	(0.200)	(0.049)	(0.049)	(0.041)	(0.217)	(0.049)
Unemp Rate	0.036	0.038	0.049	0.077	0.048	0.051	0.054	0.057	0.066
	(0.043)	(0.042)	(0.065)	(0.058)	(0.047)	(0.047)	(0.041)	(0.066)	(0.047)
% Rural	0.022	0.034	-0.147	-0.110	-0.006	-0.007	0.011	-0.159	-0.030
	(0.035)	(0.035)	(0.121)	(0.124)	(0.049)	(0.049)	(0.037)	(0.124)	(0.056)
trtGrp:post							-0.677^{***}	-1.326^{***}	-0.791^{***}
							(0.210)	(0.415)	(0.219)
Constant	0.027	0.000	0.498^{***}	0.000	-0.059	-0.086	0.091	0.521^{**}	-0.015
	(0.063)	(0.062)	(0.155)	(0.170)	(0.085)	(0.086)	(0.087)	(0.240)	(0.100)
Observations	1,657	1,657	96	96	1,657	1,657	1,657	96	1,657
R^2	0.406	0.398	0.394	0.306	0.349	0.338	0.414	0.400	0.355
Note:									*p<0.1; **p<0.05; ***p<0.01

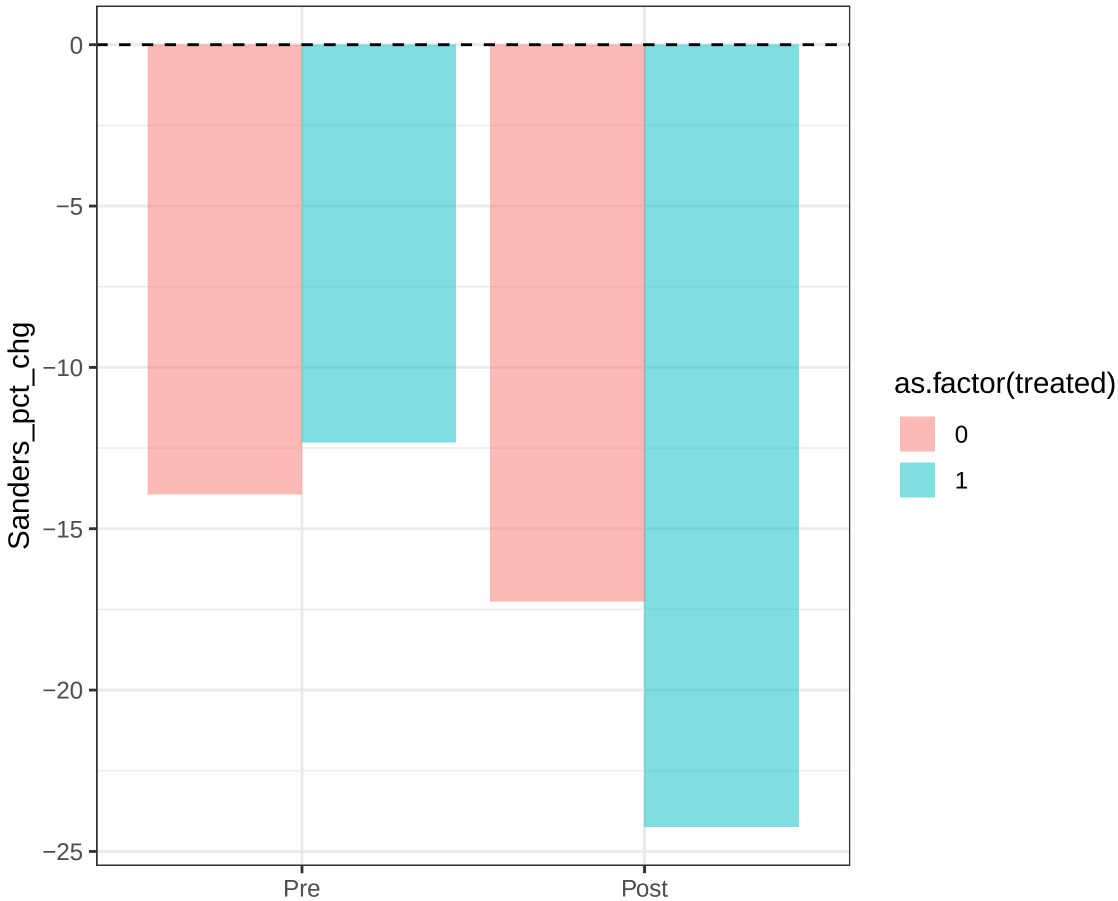

Figure 13: Change in support for Sanders (y-axis) between 2016 and 2020 by exposure to COVID-19 deaths in the DMA prior to and following the outbreak (dated to after March 3rd, x-axis).

The associated descriptive statistics are presented in Figure 13.