

Flight to Safety: Covid-Induced Changes in the Intensity of Status Quo Preference and Voting Behavior

James Bisbee
NEW YORK UNIVERSITY CSMAP

Dan Honig*
UNIVERSITY COLLEGE LONDON

Authors' Final Version, June 17, 2021

Please cite as Bisbee, James and Dan Honig. (2022) Flight to Safety: Covid-Induced Changes in the Intensity of Status Quo Preference and Voting Behavior. *American Political Science Review*, 116 (1), 70-86

Abstract

The relationship between anxiety and investor behavior is well known enough to warrant its own aphorism: a “flight to safety”. We posit that anxiety alters the intensity of voters’ preference for the status quo, inducing a political flight to safety towards establishment candidates that operates independently of preferences for leadership, retrospective voting, or the familiarity of a particular incumbent. Leveraging the outbreak of the novel coronavirus across space and time during the Democratic primary election of 2020, we identify a causal effect of the outbreak on voting, with Biden benefiting between 7 and 15 percentage points at Sanders’ expense. A survey experiment in which participants exposed to an anxiety-inducing prompt choose the less disruptive hypothetical candidate provides further evidence of our theorized flight to safety among US-based respondents. Evidence from 2020 French municipal and US House primary elections suggests a Covid-induced flight to safety generalizes to benefit mainstream candidates across a variety of settings. We find little support for alternative pathways, bolstering our claim that the findings suggest an as-yet underappreciated preference for “safe” candidates in times of anxiety.

10,828 Words
8,526 Appendix

1 Introduction

Emotions influence political behavior. Anxiety has been shown to stimulate preferences for protective policies (Albertson and Gadarian, 2015); increase support for conservative male candidates (Holman, Merolla and Zechmeister, 2016; Holman et al., 2019) and incumbents (Morgenstern and Zechmeister, 2001); and generally prompt voters to consider more carefully (Marcus, Neuman and MacKuen, 2000; MacKuen et al., 2007). In this paper, we build upon and extend existing theory by proposing that voters engage in a general “flight to safety” when faced with anxiety-inducing exogenous shocks. Substantively, our framework predicts that voters move toward the status quo under times of threat. This flight to safety is broader than a simple preference for incumbents, particular candidate attributes, or candidates’ policy platforms.

Our argument carries provocative implications for democratic accountability and governance. If and where a flight to safety operates, candidates who offer to preserve the status quo will find an easier path to (re-)election, regardless of whether their policy platforms are favored by the electorate on substantive grounds. In addition, strategic actors may attempt to manipulate anxiety in order to pursue political outcomes that would otherwise be unattainable. And where the prospect of change is itself a potential source of anxiety – as psychology suggests it often may be (Jost and Hunyady, 2003; Paterson and Cary, 2002) – a flight to safety may be a feedback mechanism which adds friction to efforts to disrupt the status quo, making radical change more difficult.

We test our argument using a variety of empirical contexts and methods. Our primary analysis uses a staggered primary election whose timing was decided independently of the Covid-19 pandemic: the 2020 Democratic primary election in the United States.¹ We

¹We pre-specified this hypothesis and our empirical specifications on March 15th, 2020, registering a pre-analysis plan prior to analyzing any data, and prior to the primaries of March 17th. This registration is at <https://osf.io/undv4>. We discuss one key departure from this plan below.

show that our hypothesized political flight to safety is meaningfully predictive of the anti-establishment candidate Bernie Sanders’ electoral fortunes relative to those of his primary rival, Joe Biden. To isolate our theorized causal mechanism, we field a survey experiment in which we randomize the policy platform and status quo qualities of two hypothetical candidates. We assign half of our respondents to a treatment condition that emphasizes the anxiety-inducing qualities of the Covid-19 pandemic, and the other half to a reassuring frame that emphasizes progress made on finding a vaccine. Again, we find evidence consistent with a political flight to safety that is independent of the policy platforms of the two hypothetical candidates. Finally, we test the generalizability of our findings by documenting similar patterns in the 2020 primary elections for the House of Representatives with more anti-establishment candidates garnering disproportionately less support when their elections were held after the spread of Covid-19. Similarly, we examine the fortunes of more mainstream and more anti-establishment French political parties in their 2020 municipal elections, finding that more anti-establishment parties lost support in the June 2020 round of elections relative to their fortunes in early March.

Across these settings we find consistent evidence of an electoral penalty for non-mainstream candidates and parties. We consider and reject plausible alternative mechanisms, providing suggestive evidence that our interpretation of the findings – a “flight to safety” – is most consistent with the empirical evidence. These results are substantively quite large, with the magnitude of the vote shift ranging from 2 to 15 percentage points across our empirical contexts. We conclude that the political flight to safety is an important, but not yet fully accounted for, phenomenon in voting behavior.

2 Theory & Empirical Contexts

Market analysts refer to a “flight to safety” to describe the behavior of investors in the face of uncertainty (Inghelbrecht et al., 2013; Adrian, Crump and Vogt, 2019). As market outcomes become more uncertain, investors shift towards more liquid and government-insured assets, which are perceived as safer (Cohn et al., 2015). We hypothesize that a similar flight to safety operates in the market for political candidates.

We center our intuition on a spatial model of voting with valence (Downs et al., 1957; Krehbiel, 1998; Ansolabehere and Snyder, 2000), given in Equation 1. In this framework, voter i ’s utility for supporting candidate j on issue k is some combination of policy preferences and the candidate’s valence attributes. The policy component is the classic spatial model of voting where utility is declining in the distance between the voter’s preferred policy (θ_k^*) and the platform of candidate j ($\theta_{j,k}$), where the shape of the decline is parameterized with α .² The valence component V_j represents non-policy candidate qualities like charisma, leadership, or – in our framework – safety.

$$u_{i,j,k} = \underbrace{-(1-w)(\theta_{j,k} - \theta_k^*)^\alpha}_{\text{Policy}} + \underbrace{wV_j}_{\text{Valence}} \quad (1)$$

Existing political science work on anxiety focuses attention on different parts of this model. One branch of the literature suggests anxiety effects political behavior by shifting the voter’s policy preferences θ_k^* towards “protective policies”, benefiting conservative candidates whose platforms ($\theta_{j,k}$) are more likely to emphasize these dimensions (e.g. Clifford and Jerit 2018; Stenner 2005; Albertson and Gadarian 2015). Empirical work in this vein shows that crises such as terrorist attacks (Getmansky and Zeitzoff, 2014) and public health threats (Campante, Depetris-Chauvin and Durante, 2020) disproportionately benefit conservatives.

²A more general formulation of Euclidean distance, such as $F(\|\theta - \theta^*\|)$ as in Carroll et al. (2013), complicates the discussion of our contribution without improving the intuition.

A related body of work focuses on risk aversion and how heterogeneity in risk appetite influences voting behavior (Kam and Simas, 2010, 2012; Ehrlich and Maestas, 2010; Eckel, El-Gamal and Wilson, 2009). Here, the shape of the utility function itself, denoted with α in Equation 1, carries implications about how much uncertainty a voter is willing to stomach (Berinsky and Lewis, 2007). Risk aversion models can generate substantively similar predictions to a flight to safety in which, for example, incumbent parties benefit as more risk-averse voters choose “the devil they know” (Morgenstern and Zechmeister, 2001). However, this theory treats the shape of an individual’s utility function as innate, implying that risk appetites cannot change in response to external events.³

In contrast, a different branch of the literature argues that crises influence vote choice not through the policy preferences or risk appetites of the first component of Equation 1, but rather through the valence component. These studies theorize that voters care less about the specific position of a candidate along a given policy dimension, as all political actors may share basic goals during a crisis (e.g. keeping citizens safe, Merolla and Zechmeister 2009). Instead, anxious voters prioritize candidate attributes like “leadership” and “strength”, leading to a bias toward conservative male candidates (Holman, Merolla and Zechmeister, 2016; Holman et al., 2019).⁴ These perspectives generate predictions of anxiety-induced changes to voting behavior through the relative weight w assigned to the candidate’s valence V_j .

Our work builds on this model, treating the safety of the status quo as a valence

³There is a growing body of work in economics that documents how risk appetites change in response to external shocks like natural disasters (Hanaoka, Shigeoka and Watanabe, 2018; Cameron and Shah, 2015; Callen, 2015; Cassar, Healy and Von Kessler, 2017), extreme weather (Eckel, El-Gamal and Wilson, 2009; Kahsay and Osberghaus, 2016), and violence (Voors et al., 2012; Callen et al., 2014). One potentially fruitful area of future political science research is testing whether and how an endogenous risk appetite translates into politically consequential outcomes.

⁴Our claim that anxiety influences vote choice has echoes of, but is distinct from, Affective Intelligence Theory (AIT). AIT posits that anxiety influences whether voters make choices based on (often partisan) heuristics or more cognitively taxing rational choice considerations (Marcus, Neuman and MacKuen, 2000; MacKuen et al., 2007). Our flight to safety framework is distinct in that it predicts a specific direction of the anxiety-induced rational choice calculus – one biased toward status quo safety. Whether and how this directed preference interacts with the heightened attention theorized by AIT is beyond the scope of our paper, but is an interesting avenue for future work.

term whose importance increases in response to anxiety. We posit that valence qualities such as leadership and strength are components of the broader safety that anxious voters seek. But we argue that the political flight to safety provides a more general understanding of how anxiety influences vote choice, while providing a more precise prediction about the direction of the effect. Specifically, it is not that ideologically conservative candidates benefit during crises but rather that more radical candidates on both ends of the spectrum are penalized. Furthermore, while a political flight to safety accommodates existing work on risk aversion and the incumbency advantage (Shepsle, 1972; Eckles et al., 2014), it extends the understanding to allow for a status quo bias even in races between two challengers. Finally, by shifting focus from policy preferences to the weight placed on certain valence attributes, our framework can make sense of electoral outcomes in which voters appear to vote against their self-interest.

On this latter point, a political flight to safety framework connects with System-Justification Theory (SJT), which argues that there is an inherent need to “defend, bolster, and justify aspects of existing social, economic, and political systems” (Jost, 2019). SJT argues that individuals are generally disposed to justify the existing social order to reduce uncertainty, and that this need is stronger among those facing greater uncertainty or experiencing larger feelings of powerlessness (Van der Toorn et al., 2015). While this perspective is wholly outside the spatial model of voting framework that unifies the literature summarized above, it does provide important insights on why the disadvantaged support the status quo (Jost et al., 2003). Consistent with the findings of Van der Toorn et al. (2015), a political flight to safety predicts that anxiety can place greater weight on a voter’s preference for the status quo.

Empirical Contexts

A central challenge in gathering empirical support for our claim is one of observational equivalence. Measuring support for incumbents versus challengers risks conflating retrospective evaluations with the strength of status quo preference. Crises may also adjust the policy positions taken by political candidates or the policy preferences of voters.⁵ With so many moving parts, it is a challenge to empirically isolate the channel of theoretical interest to us: anxiety-induced status quo preference.

We test our theoretical intuition in several empirical contexts. Our main analysis examines the impact of Covid-19 on the electoral fortunes of Bernie Sanders and Joe Biden, the two leading candidates for the Democratic Party’s presidential nomination in 2020. 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.⁶

We assume that, in the context of the 2020 democratic primaries, Joe Biden embodied safety while Bernie Sanders embodied disruption. Evidence in support of this assumption is overwhelming, starting with the candidates themselves. Biden portrayed himself as representing continuity and the security of the known – Biden described himself as an “Obama-Biden Democrat” (Fegenheimer and Glueck, 2020). “The heart of his [Biden’s] pitch, when he delivered it clearly, was status quo ante, back to normal”, as one journalist succinctly put it (Debenedetti, 2020). Sanders, by contrast, promised to “change the power structure in America” (Stewart 2020), portraying himself as a candidate who (in the words

⁵Our claim is not that crises *do not* have the potential to alter candidates’ policies or to move voters’ policy preferences, rather that there is an additional status quo preference which receives greater weight under conditions of anxiety.

⁶That both candidates are men holds constant the possibility that gendered beliefs regarding the valence qualities of candidates is driving response to perceived threat (Holman, Merolla and Zechmeister, 2016; Holman et al., 2019).

of his 2020 campaign spokesman) pushed against “the limits of politics as usual” (Eilperin 2020). In some sense both candidates agreed that Biden was the mainstream alternative, and Sanders the more radical and thus riskier choice. Voters apparently understood these divergent appeals, with exit polls in a number of states indicating that Sanders won a majority of those voters who preferred a candidate that “Can Bring Needed Change”, while Biden was preferred by those who sought a candidate that “Can Unite the Country”.⁷

We believe this empirical setting provides a hard test of the motivating theory in some important senses. Surveys of US primary voters in 2016 suggest that Democratic voters in general and Sanders supporters in particular were particularly unlikely to engage in system justification (Azevedo, Jost and Rothmund, 2017), and as such might be particularly unlikely to shift votes due to a need to preserve the status quo in response to an anxiety-provoking pandemic. Additionally we believe that Sanders’ policy platform – with its focus on health and worker protections – would likely be more attractive following Covid-19’s emergence, all else equal. A shift away from Sanders thus suggests that it is a change in the weight placed on voters’ status quo preference, rather than a change in voters’ policy preferences, that lies behind any observed flight to safety.

This is not the only possible interpretation of how Biden and Sanders’ relative attractiveness as candidates might be altered by Covid’s emergence. First, Biden’s prior service as Vice President may have led some primary voters to prefer his leadership qualities in the face of crisis (Merolla and Zechmeister, 2009). We conduct a survey experiment which isolates the anxiety mechanism, finding causally identified evidence of a political flight to safety that obtains in the absence of the leadership dimensions of safety, or any other confounds.⁸

⁷See exit polls as reported by CNN, [https://edition.cnn.com/election/2020/entrance-and-exit-polls/STATE NAME/democratic](https://edition.cnn.com/election/2020/entrance-and-exit-polls/STATE%20NAME/democratic), e.g. those from Michigan and Washington. Additional evidence supporting our characterization of Sanders as the disruptive candidate can be found in Azevedo, Jost and Rothmund (2017) which finds that support for Sanders in 2016 was highest among individuals with the lowest levels of system justification. Pew data from 2020 further reaffirms our claim by showing that the vast majority of respondents placed Bernie Sanders at the far-left of an ideological scale, while placing Biden toward the center. Please refer to the Supporting Information for these results.

⁸The full details of the experiment are summarized in the Supporting Information.

Second, it is possible that the economic consequences of the pandemic made Biden’s platform more appealing for voters concerned about the country’s long run economic health. In the Supporting Information we test whether negative views of the economy at the Congressional District level are correlated with the spread of the pandemic, finding little evidence in support of this alternative policy pathway. In addition, we examine the electoral fortunes of primary candidates for the House of Representatives, an analysis that includes anti-establishment candidates at both extremes of the political spectrum. In aggregate we observe a flight to safety away from non-mainstream candidates on both sides of the political spectrum, suggesting it is not merely candidates’ platforms that drive our results.

Third, inasmuch as Biden and Sanders were contesting a primary election, it is possible that some segment of the primary electorate voted strategically, choosing not their preferred candidate but the most “electable” candidate in the general election (Abramson et al., 1992; Rickershauser and Aldrich, 2007; Simas, 2017). In this setting, an anxiety-inducing crisis raises the stakes of the general election and pushes more primary voters to behave strategically. The flight to safety theory provides an explanation for the relative appeal of mainstream and anti-establishment candidates, and we expect this calculus to apply to a strategic voter’s calculation of electability as much as it applies to a sincere voter’s utility function. Consistent with this expectation, we document a similar flight to safety in French municipal elections and in our survey experiment, which are not subject to these strategic calculations of primary voters.

3 Data & Methods: 2020 Democratic Primary

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

Outcome Variable

Our outcome variable is the two-way vote share for Sanders in the 2020 Democratic primary election, aggregated to the county level. We obtained these data from David Leip’s Atlas of the United States, updated on May 30th, 2020.⁹

Explanatory Variable

We use data on the county-level spread of the pandemic obtained from The New York Times via their publicly available GitHub (<https://github.com/nytimes/covid-19-data>).¹⁰ We aggregate these cases to the designated market area (DMA) to reflect our expectation that exposure is most salient within media markets in which local news channels report on cases. However, we recognize that the geographic variation across these units becomes an increasingly poor proxy for anxiety as time progresses and the country shuts down. We use observable proxies for DMA-level anxiety (social distancing behaviors¹¹ and Google searches¹²) to confirm that our reliance on this source of geographic variation is plausible for the first three weeks of March, 2020.

Controls

We obtain a rich set of pre-treatment county-level controls from the five year averages of the American Community Survey (2018), as well as 2016 Democratic primary election

⁹<https://uselectionatlas.org/RESULTS/national.php?year=2020&f=0&off=0&elect=1>

¹⁰We recognize that these data are themselves incomplete and that there will remain a debate over the true toll of the pandemic for months if not years to come. We believe that these sources of measurement error are orthogonal to our empirical strategies since these data reflect the information environment about the severity of the pandemic for those who went to the polls.

¹¹Obtained from Cuebiq, <https://www.cuebiq.com/visitation-insights-covid19/>.

¹²<https://trends.google.com/trends/?geo=US>. Data obtained using the `gtrendsR` package for R.

data also obtained from David Leip’s Atlas of the United States.¹³ We list the full set of controls in the Supporting Information. In addition, we include indicators for both the election format (primary or caucus) and for whether the state switched from a caucus to a primary between 2016 and 2020.

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 four ways, visualized in Figure 1. The first two confounds challenge the causal claims we make. The second two confounds threaten our theorized mechanism of anxiety-induced changes to the intensity of voters’ preference for the status quo. We discuss each threat in turn, and organize our results along these pathways.

First, if areas that were already anti-Sanders were also those most exposed to the outbreak, our results would pick up a spurious selection effect, indicated by the “selection” pathway in Figure 1. We control for the county’s support for Sanders in 2016 in the main specifications. Furthermore, we predict Sanders’ 2016 vote share as a function of exposure and find, if anything, these counties are *more* pro-Sanders. Finally, insofar as it is possible that Sanders’ 2016 vote share may be a poor proxy for his support in 2020, we demonstrate that in the months prior to the 2020 Democratic primary election, areas that would become more exposed were also those that were more favorable toward Sanders based on weekly polling data at the Congressional District-level between 2019 and 2020.¹⁴ These checks confirm that, to the extent that there is selection bias, it works against our results.

Second, the timing of the outbreak is colinear with other explanations for chang-

¹³ACS 2018 data from <https://www.census.gov/programs-surveys/acs/news/data-releases/2018/release.html>.

¹⁴We use the weekly Nationscape survey with data measured at the Congressional District for this analysis (Tausanovitch et al., 2019).

ing electoral fortunes, such as the decision by several primary candidates to drop out (FiveThirtyEight, 2020), signaling a consolidation of party support behind Biden (Yglesias and Beauchamp, 2020). If later-voting Sanders supporters no longer saw their votes as pivotal, their exposure to health risks might be enough to differentially keep them home. A simple before-after comparison of election returns would be unable to disentangle our flight to safety theory from a coincidental shift in electoral momentum, and is indicated by the “spurious” confounder in Figure 1. We appeal to both the cross-sectional variation of exposure within states, as well as to a series of placebo tests, to defend our results against this concern. Furthermore, we test whether turnout in counties that favored Sanders in 2016 was noticeably lower following Super Tuesday, finding little support for this alternative explanation in our data.

Third, 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, illustrated by the “turnout” pathway in Figure 1. Insofar as younger voters are relatively more supportive of Sanders, this would also work against our results by making it harder to identify a negative relationship between exposure and Sanders’ vote share. We test for differential turnout by average age and confirm that this pathway is not supported in the data.

Fourth, as discussed in the theory section, the disease could also influence voter policy preferences directly, rather than simply altering the intensity of status quo preference. For example, the pandemic would plausibly increase demand for health care and unemployment insurance. These demands would suggest increased support for the Sanders campaign, making the bias work against our theory. Conversely, the pandemic’s economic consequences might prompt voters to care more about each candidate’s economic platform, with Biden’s being the more preferred by Wall Street. These stories constitute an alternative pathway, indicated by “Preferences” in the causal diagram illustrated in Figure 1. To isolate the theorized mechanism of a political flight to safety, and to make a causal claim about the

relationship between the pandemic and support for Sanders, we employ the following methods.

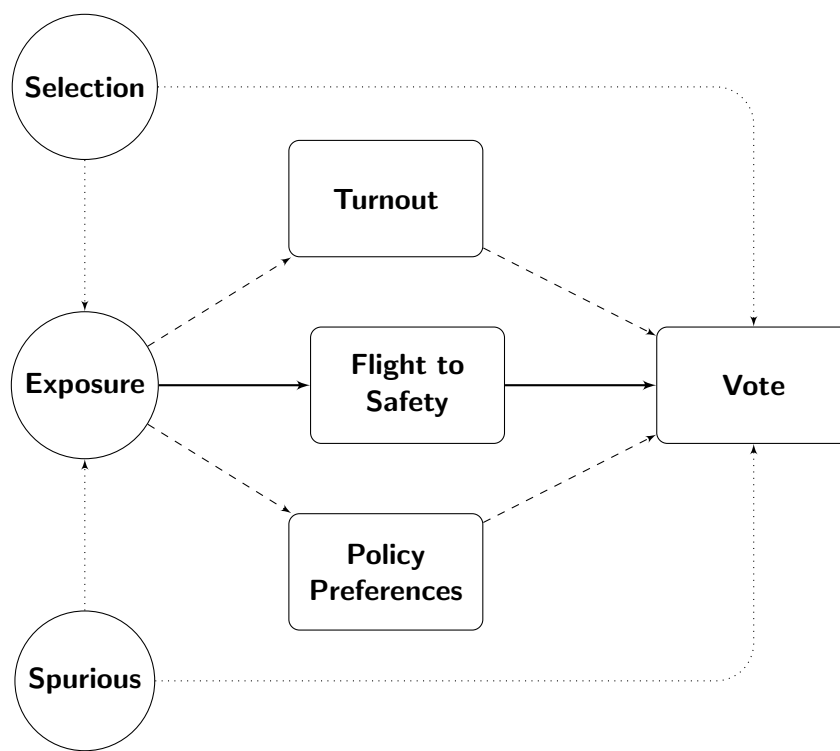


Figure 1: Causal DAG illustrating alternative pathways (dashed lines through turnout and policy preferences) and selection bias / OVB (dotted lines). Selection bias occurs if counties that were more exposed were also more likely to reduce support for Sanders. We use controls, weighting, matching, and generalized diff-in-diff strategies to account for possible selection bias. Furthermore, we demonstrate that exposed counties were more supportive of Sanders in 2016, and viewed him more favorably in 2019 and 2020, suggesting that this source of bias should work against our findings if it persists. Omitted variable bias (“spurious”) occurs if we attribute declining Sanders support to the pandemic when it was in fact due to the Democratic Party coalescing around Biden. We use placebo tests to show that this is unlikely. To adjudicate between the plausible mechanisms by which the virus influences support for Sanders, we first provide evidence of increased anxiety in areas where the virus first spread, suggesting that the theorized anxiety channel is open. We then test the alternative pathway of turnout directly, finding that older voters did not vote in greater numbers in response to the virus. Finally, we argue that the policy preferences pathway should bias against our findings (i.e., Sanders’ policy platform is theoretically more attractive given its emphasis on health care reform and improving the social safety net).

Our identifying assumption relies on the fact that states did not reschedule their primary election dates until after the period we analyze.¹⁵ The timing of the primary election

¹⁵The only exception is Ohio who canceled the primary election originally scheduled for March 17th. In

date across the outbreak of Covid-19 in the United States is plausibly independent of the outcome, allowing us to make comparisons between counties that voted earlier or later in the outbreak but had otherwise similar exposure trajectories. We make these comparisons using three different methods which we pre-registered on March 15th, prior to analyzing any observational data.¹⁶

Our first method is a linear regression predicting Sanders' vote share as a function of exposure to the virus, which we define as the number of cases reported in the designated market area (DMA) in which a county is located. We assign Covid-19 cases to counties according to the date of their election, resulting in a nominally cross-sectional dataset with rows indexing counties. We control for observable characteristics of these counties in a variety of ways. The simplest approach is to include these covariates as controls in the linear regression. In addition, we use matching and balancing strategies to ensure we are comparing otherwise similar counties who differ only in the timing of their exposure to Covid-19 when they went to the polls. We obtain good balance on a rich set of pre-treatment covariates using either nearest neighbor matching (based on minimized Mahalanobis distance) (Stuart, 2010), or covariate balanced propensity score weights (CBPS, Imai and Ratkovic 2014).¹⁷ Finally, we implement DMA fixed effects to force the comparison to be between counties that share the same information environment but are located in states that voted at different times. Figure 2 highlights the counties that reside in DMAs that cut across state borders between states that voted earlier (light gray) and those that voted later (dark gray). This fixed effects

our Supporting Information, we run an exhaustive set of pairwise comparisons to confirm our main results hold if we omit the March 17th elections from the data.

¹⁶Our PAP originally envisioned the diff-in-diff being a comparison between Sander's vote share in 2016 and 2020. However, as the primaries unfolded, this comparison was revealed to be intractable for two important reasons. First, the 2016 contest was effectively a two candidate race between Sanders and Hilary Clinton. The few alternative candidates who appeared on the ballot in 2016 were never as viable as the many candidates that competed in 2020. Second, turnout alone differed dramatically between 2016 and 2020. As an example, there were roughly 204,000 votes cast in the Minnesota 2016 primary, while Sanders alone received approximately 220,000 votes in the 2020 primary and didn't even win. Despite the intuitive appeal of comparing the change in Sanders' county-level support between 2016 and 2020, closer inspection suggests these two quantities are vastly different, precluding an apples-to-apples comparison, and justifying our decisions to abandon this part of our pre-analysis plan.

¹⁷Balance results are included in our Supporting Information.

specification isolates the timing of primary elections as the source of identifying variation, but at the cost of dropping counties residing in DMAs that are wholly located within one state.¹⁸

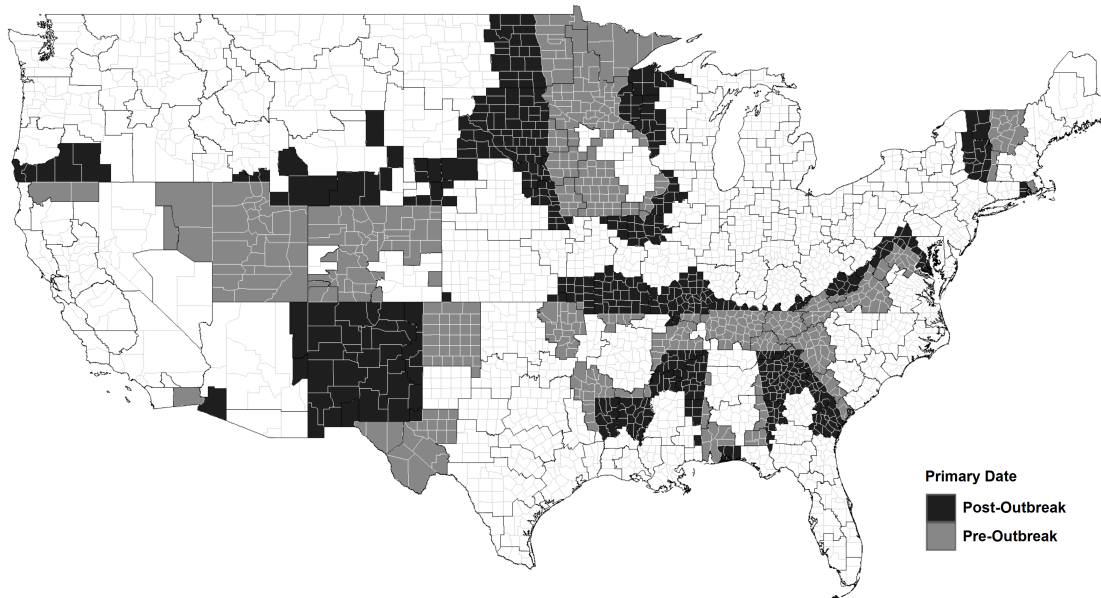


Figure 2: Visually isolating the identifying variation. Counties that share the same DMA but reside in states that voted earlier or later are shaded in gray and black, respectively.

Our second approach re-oriens our dataset as a panel dataset, recognizing that we only observe each county once. To implement a difference-in-differences estimator, we categorize every county as either exposed or insulated based on the number of cases in their DMA *as of March 17th*. We calculate the difference in the support for Sanders among these counties that voted prior to March 17th, and compare this difference to that measured between exposed and insulated counties that voted on March 17th. The identifying assumption underlying this method is that, by differencing out the support for Sanders among counties that would and would not be exposed as of March 17th, we remove any potential selection bias.

¹⁸As such, our causal estimand is “local” in the sense that we identify a relationship among a subset of counties that differ from those that we are forced to drop. Specifically, our results are based on counties that are either located near state borders or are more rural (or both). This shortcoming is only an issue in our regression models that use DMA fixed effects, and is not a problem for the diff-in-diff analyses. Furthermore, as we show in the SI, the counties we drop were more supportive of Sanders in 2016, suggesting that any bias introduced by this sample restriction is conservative.

We augment this strategy via trajectory balancing weights (Hazlett and Xu, 2018) which re-weight the control units such that they more closely approximate the treated units based on pre-treatment covariates. We balance on the the daily number of Covid-19 cases by county, matching all exposed counties that voted on March 17th to those counties that voted prior to March 17th based on the full list of controls given above, *and their daily exposure to the virus*. This approach ensures that we compare an exposed county who voted on March 17th to an as-close-to-identical-as-possible county who voted earlier, both in terms of its pre-election demographic, economic, and social characteristics, as well as full time series vector of daily cases through late April 2020. More theoretically, this method means that we are identifying the effect of exposure using the exogeneity of the pandemic as it interacts with the independently determined primary election calendar.

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 2020, when the virus was just beginning to spread across the United States. However, by the end of March, national media outlets (e.g. cable news, newspapers, news websites, and online social media such as Facebook, as discussed in [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 likely experienced substantial levels of anxiety despite not residing in a DMA with confirmed cases of the virus.

We test the SUTVA assumption by predicting week-by-week variation in two observable behaviors that we believe proxy for anxiety. The first is mobility data which we interpret as a proxy for social distancing behaviors. The second is Google search data for the term

“coronavirus”. We show that geographic variation in DMA-level cases is strongly predictive of variation in these proxies in the first three weeks of March, but that – in line with violations of SUTVA – exposed and insulated areas show no differences after March 23rd. In addition, we include an exhaustive series of pairwise comparisons in which we define one primary election as treated, and another as control, in our Supporting Information which serve as a series of placebo tests for the validity of our identifying variation in first three weeks of March, 2020.

4 Results

Our first set of results are summarized in Table 1 which presents the coefficient estimates returned by a regression of Sanders two-way vote share (Sanders / Sanders + Biden) on an indicator that takes the value 1 if the county is located in a DMA with one or more confirmed cases of Covid-19 at the election date. The first three columns present the coefficients on exposure to the pandemic using the full data running from February through April. The second three columns subset the data to focus only on the elections that occurred between March 1st and April 7th – the dates after which the pandemic started and before Sanders dropped out of the race. Clustered standard errors at the DMA are presented in parentheses. Columns 1 and 4 present the results of a basic regression using only the controls with DMA fixed effects. Columns 2 and 5 implement the nearest neighbor matching strategy using Mahalanobis distance (Stuart, 2010). Columns 3 and 6 apply weights generated using covariate balanced propensity scores (CBPS, Imai and Ratkovic 2014).

The results indicate that counties with 1 or more confirmed cases of the virus in their DMA at the time they went to the polls were significantly less likely to support Sanders than counties without confirmed cases. The coefficients themselves are standard deviations of the two-way vote share for Sanders (1SD = 14.9 percentage points), suggesting a statistically and

substantively significant relationship between exposure and voting behavior, albeit one whose magnitude is exaggerated when we include election results prior to the outbreak. The most conservative estimates generated in the March subset find an effect size of approximately 9 percentage points.

Table 1: Sanders two-way vote share \sim Exposure

	<i>Dependent variable: Sanders 2020 Two-Way Vote Share</i>					
	Full Sample			March & April 7th		
	(1) Basic	(2) Match	(3) Weight	(4) Basic	(5) Match	(6) Weight
Exposure Dummy	-1.321* (0.552)	-1.291* (0.632)	-1.348*** (0.380)	-0.885* (0.382)	-0.810† (0.417)	-0.885** (0.325)
Turnout 2020	0.175*** (0.026)	0.192*** (0.057)	0.137*** (0.038)	0.197*** (0.026)	0.214*** (0.056)	0.169*** (0.041)
Sanders 2016	0.027*** (0.004)	0.035*** (0.008)	0.032*** (0.006)	0.026*** (0.004)	0.033*** (0.006)	0.034*** (0.006)
Caucus switch	0.628 (0.959)	1.220*** (0.333)	0.874 (0.557)	0.205 (0.771)	0.490 (0.790)	0.764 (0.465)
Caucus Dummy	1.227*** (0.221)	0.932* (0.467)	0.764† (0.417)			
County Controls	Y	Y	Y	Y	Y	Y
DMA FEs	Y	Y	Y	Y	Y	Y
Observations	1,882	666	1,882	1,710	666	1,710
R ²	0.830	0.882	0.883	0.860	0.894	0.898

Notes: DMA-cluster robust standard errors in parentheses. The caucus indicator does not have sufficient variation to be included in the March and April subset (columns 4 - 6). † $p < 0.10$; * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

Differences-in-Differences

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

The left panel of Figure 3 plots the simple share of the voting age population that supported Sanders among exposed (dark gray) and insulated (light gray) counties prior to (left) and following (right) the outbreak of the virus (defined as beginning on March 10th). These descriptive plots highlight some important patterns in the voters’ response to Covid-19. First, there appears to be a secular decline in support for Sanders among both exposed and insulated counties following the outbreak of the novel coronavirus (left panel). 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 (right panel). These patterns suggest that Sanders enjoyed greater support in areas that were more affected by the virus (younger, more urban areas), but that this support eroded as the virus spread.

To test this descriptive intuition, we estimate a diff-in-diff regression of the following form:

$$Y_c = \beta_0 + \beta_1 \text{Exp}_{c, \text{Mar}17} + \beta_2 \mathbb{I}[\text{Vote Mar}17] + \beta_3 \text{Exp} \times \mathbb{I}[\text{Vote}] + \epsilon_c \quad (2)$$

and plot the marginal effects in Figure 4. As illustrated, there is evidence consistent with the descriptive plots above – counties that were exposed as of March 17th were more supportive of Sanders overall, but were significantly more supportive when they voted prior to March 17th. The interaction coefficient on this specification is statistically significant and of commensurate magnitude to that presented in Table 1 (estimate = -6.75 percentage points; SE = 2.14).

To confirm the sensitivity of the diff-in-diff finding to the choice of outbreak date, we regress support for Sanders on exposure week-by-week and plot these coefficients in Figure 5. As illustrated, there is clear evidence of a positive relationship between exposure as of March 17th and support for Sanders up until March 10th. On March 17th we see a significant negative relationship when weighting counties using CBPS.

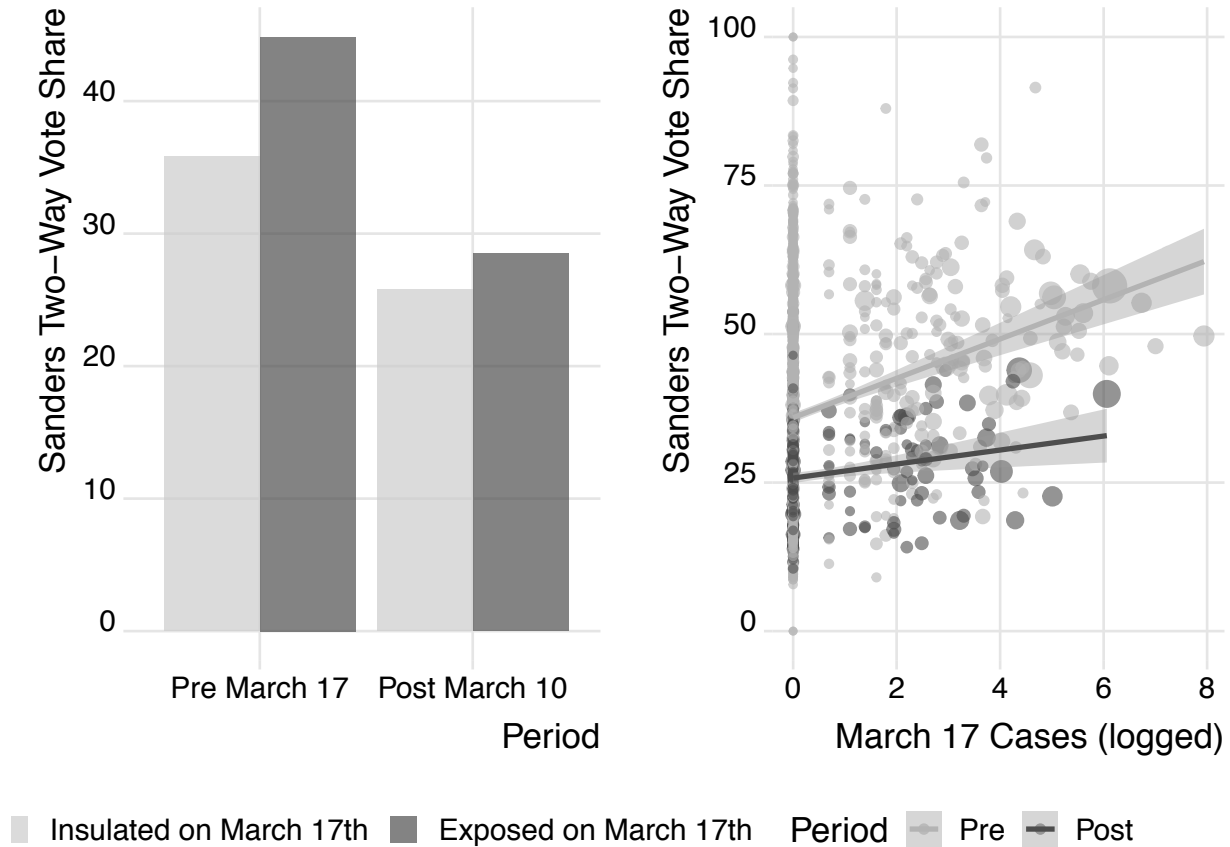


Figure 3: Descriptive differences between exposed and insulated voting behavior before and after the outbreak, defined as starting on March 10th. Left panel groups counties by whether they were exposed as of March 17th, right panel plots the logged cases as of March 17th by whether the county voted prior to, or following, the outbreak. Points sized to reflect total turnout.

These plots are not meant to support well-identified inferential conclusions, although the correlations we document consistently point toward support for Sanders slipping in exposed areas that we would otherwise expect to be quite supportive. These plots do clarify the sources of identifying variation we rely on in the results that follow. Specifically, by comparing counties based on their number of cases as of March 17th, and dividing them into groups that voted prior to that date and those that voted on March 17th, we emphasize the selection problems in this exercise. There is clear evidence that areas more affected by the pandemic were also those more naturally inclined to support Sanders. The goal therefore is to re-weight the data to provide the most appropriate counterfactual for our notionally

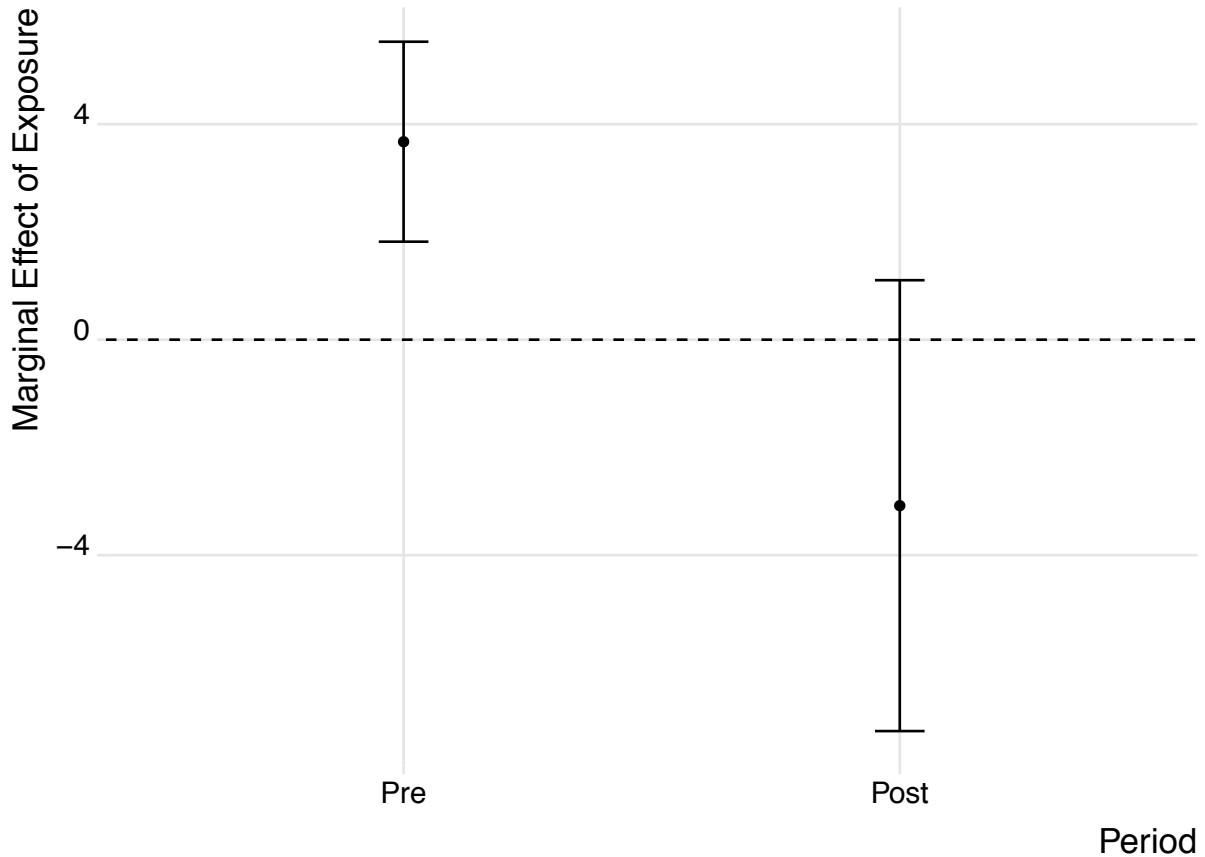


Figure 4: Marginal effects plot of the diff-in-diff results. Coefficients in the pre-period capture the difference between exposed and insulated counties as of March 17, but that voted prior to March 17th. Coefficients in the post period reflect the difference in exposed and insulated counties who voted on March 17th.

“treated” counties which (1) had cases and (2) voted on March 17th.

To do so we turn to trajectory balancing, a method that re-weights control units to appear as similar to treated groups as possible over the entire pre-treatment period (Hazlett and Xu, 2018). We use a jackknife approach in which we drop one exposed county at a time, re-estimate the trajectory balancing weights, and calculate the difference in Sanders two-way vote share between exposed and weighted control counties. We repeat this process three times, corresponding to the March election dates of March 3, March 10th, and March 17th. We plot the estimates as densities in Figure 6. Again we see evidence suggesting that exposed counties were insignificantly more supportive of Sanders as of March 3rd, as



Figure 5: Week-by-week regressions of Sanders’ two-way vote share on a dummy indicator for whether the county was exposed to the virus on March 17th. Lighter points and bars are coefficients estimated using the basic regression, including all controls. Darker points and bars are coefficients estimated applying the CBPS weights to balance exposed and insulated counties.

compared to otherwise similar counties that voted in February. However, by March 10th and then most strikingly by March 17th, this relationship is reversed. Substantively, these results suggest that counties that voted on March 17th and were located in a DMA with at least 1 confirmed Covid-19 case were approximately 15 percentage points less supportive of Sanders on average than similar counties that voted prior to March 17th, commensurate to the coefficients estimated in the baseline naive regressions in Table 1.

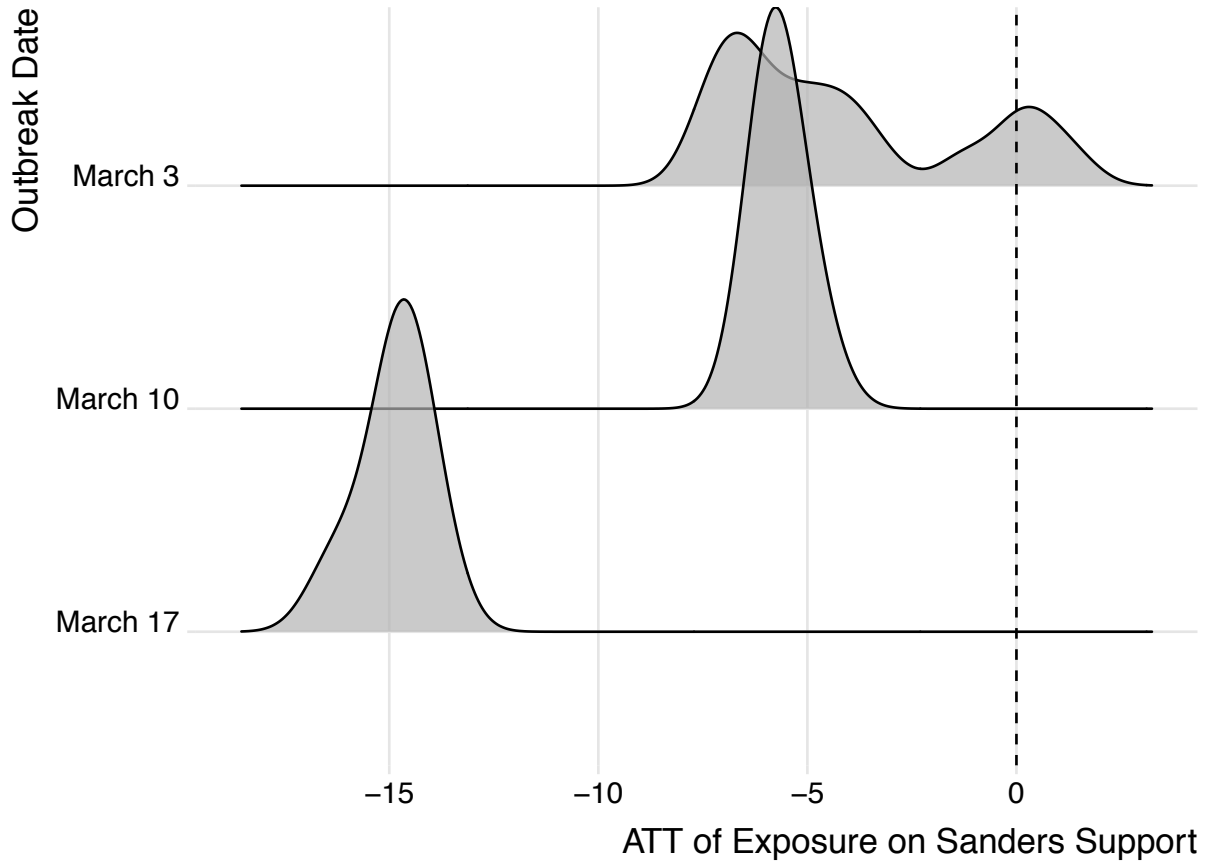


Figure 6: Jackknife estimates generated by dropping each exposed county one at a time, estimating trajectory balancing weights, and then calculating the difference between the exposed support for Sanders and the weighted control support. Y-axis indicates whether we compare exposed counties who voted on March 3rd, March 10th, or March 17th.

Mechanisms: the Spread of Anxiety

The results summarized above are consistent with our theorized flight-to-safety mechanism in which the Covid-19 pandemic alters the relative appeal of mainstream and anti-establishment candidates. The methods employed above combine the exogeneity of the pandemic with the orthogonal primary election dates, in so doing endeavoring to purge our estimates of confounding selection effects.

Our theorized mechanism rests on the assumption that differences in the exposure to the pandemic cause differences in anxiety, which then generate differences in observed

vote shares for Sanders. To test whether anxiety is in fact responding to the outbreak, we predict daily cross-sectional variation in Google searches about the virus. We view increased interest in the virus represented by Google search data is a proxy for increased anxiety about its health risks. We obtain these measures at the DMA level daily from December 30th, 2019 to April 30th, 2020.

In Figure 7, we show that search traffic is significantly correlated with DMA-level variation in cases, particularly in the week following Super Tuesday. The top words that were associated with searches for “coronavirus” are displayed above the coefficients. As illustrated, not only are those living in exposed areas more likely to search for “coronavirus” than those living in unexposed areas, they are also pairing their search terms with other anxiety-associated words. These plots provide growing evidence of “saturation” over time in the sense that geographic heterogeneity is no longer meaningful when everyone is equally anxious. Specifically, we note that the significant positive relationship between Google searches for “coronavirus” and DMA-level cases disappears after March 10th, meaning that the search profile in areas with many cases was no different from the search profile in areas with few cases. These patterns suggest that Covid-induced anxiety becomes so widespread by mid-March that we are no longer able to use geographic variation to identify the effect of the pandemic.¹⁹

Considering Alternative Explanations

It is possible that it is not anxiety that induces the shift toward Biden, but either differences in turnout which covary with anxiety or a concurrent “party decides” phenomenon. In the Supporting Information, we explore these dynamics, finding no support for the notion that they explain the empirical patterns we observe.

¹⁹In the Supporting Information, we run a similar analysis on a different proxy for anxiety – mobility – finding substantively similar patterns.




Figure 7: Daily (x-axis) coefficients (y-axis) on the relationship between DMA-level cases and DMA-level Google searches for “coronavirus”, including state random effects. Vertical dashed bars indicate primary dates. Top 3 related search phrases given above each bar.

First, we look for evidence of an alternative pathway in which the pandemic’s differential suppression of turnout drives the results. There are two versions of this alternative story. The first focuses on the differences in health risks by age and posits that those most threatened by exposure might be less likely to turnout. If this group is also more likely to support Sanders, it would suggest 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.) Nevertheless, we look for differences in turnout by county age demographics and find little support for this alternative turnout explanation.

The second turnout story is that the Sanders campaign was effectively finished after Biden’s convincing victory in South Carolina on February 29th. In this scenario, a number

of would-be Sanders voters were planning on casting what were effectively protest votes, and those in Covid-exposed areas didn't bother since the "costs" of doing so were higher. If this were the case, we should expect to see a decline in turnout in more pro-Sanders counties following the South Carolina primary. Again, we test this claim in our Supporting Information, finding little evidence to support it.

We additionally consider whether our results might be spuriously driven by a "party consolidation" effect. This concern is motivated by the theory that "the Party Decides" – that is, the possibility that party elites hold decisive power over the candidate selection process (Cohen et al., 2009), these elites exercised their power in favor of Biden, and that these elites did so just as the pandemic was spreading. If the Democratic Party decided to back Biden just as Covid-19 spread across the country, it might appear that the virus caused Sanders' decline when it was actually mere coincidence. This concern is somewhat mitigated by the observation that not only would the party decides explanation have to occur at the same time as the outbreak, it would have to be correlated with the pandemic's geographic distribution to account for our results. Furthermore, to the extent that the Democratic Party did "decide" on Biden, it arguably did so prior to Super Tuesday, meaning that its impact would be influencing primary vote shares in both the pre and post periods of the diff-in-diff analyses.²⁰ Nevertheless, we test this alternative explanation using a variety of placebo tests in the Supporting Information, including a permutation test in which we break the geographic distribution of the outbreak and examine whether the temporal variation still predicts a decline in Sanders' support (it doesn't); all alternative assignments of vote dates to treated and control conditions; and tests of whether more anti-Sanders areas were more exposed to Covid-19 (they weren't).

²⁰The notable drop-outs that signaled the Democratic Party coalescing around Biden happened on Sunday, March 1st (Tom Steyer and Pete Buttigieg), Monday March 2nd (Amy Klobuchar). In addition, former competitors including Buttigieg, Klobuchar, and O'Rourke all endorsed Biden on March 2nd (Drezner, 2020).

5 Survey Experimental Evidence

Even with the exhaustive checks on our main results summarized above and in our Supporting Information, convincingly identifying our theory is a challenge using observational data. There are a number of differences between any two candidates which might explain differential reaction to them in light of the pandemic. In order to isolate the anxiety mechanism, we fielded a survey experiment on Amazon’s Mechanical Turk platform between May 14th and May 20th, 2020 in which we randomly assigned respondents to read a summary of a potential future course for the pandemic and then choose between two hypothetical challengers for executive office. We randomly assigned respondents to read either an optimistic assessment of the pandemic (anxiety-relieving) or a pessimistic assessment (anxiety-inducing), both of which were based on real media accounts and expert assessments as of May 10th, 2020. The optimistic assessment highlighted the progress made toward creating a vaccine, the potential use of existing drugs to combat the virus, and statistical evidence suggesting that most of the country had the worst part of the pandemic behind them. The pessimistic assessment painted a bleak picture of the possibility of a second wave of infections, a longer wait before vaccines became available, and suggested that existing reports likely undercounted the number of cases and deaths to date.²¹

We randomly varied the candidates along 4 dimensions: age (45 or 48), occupation (accountant or lawyer), education (law school or local college), policy platform (healthcare or education), and – our primary object of interest – anti-establishment or mainstream candidate. The anti-establishment candidate was described as an individual who “seeks fundamental transformation of the economic, social, and political order. He believes that the system is broken, and the time for radical change is now.” The mainstream candidate was described as an individual who “believes in strengthening existing economic, social, and political institutions. He believes that we must come together and re-invest in our system,

²¹The text of the survey is provided in the Supporting Information.

strengthening its foundations to support generations to come.”

Our main quantity of interest is whether exposure to the pessimistic treatment reduced support for the anti-establishment candidate, corresponding to our theorized flight to safety mechanism. Table 2 illustrates the results, finding that reading a pessimistic description of the pandemic reduces support for the anti-establishment candidate, in line with our theoretical intuition.

Table 2: Suvey Experiment Results

	<i>Dependent variable:</i>			
	Favor Anti-Establishment Candidate			
	Bivariate	Demographic Controls	Politician Controls	Attentive Subset
	(1)	(2)	(3)	(4)
Anxiety Prime	-0.074 [†] (0.039)	-0.078* (0.038)	-0.082* (0.038)	-0.182** (0.061)
Democrat		-0.133* (0.061)	-0.133* (0.060)	-0.076 (0.098)
Lean Dem		-0.047 (0.072)	-0.054 (0.071)	0.032 (0.112)
Independent		-0.071 (0.063)	-0.070 (0.062)	0.083 (0.104)
Lean GOP		-0.311*** (0.080)	-0.329*** (0.079)	-0.308* (0.123)
Republican		-0.383*** (0.074)	-0.385*** (0.073)	-0.244* (0.118)
Strong GOP		-0.294*** (0.081)	-0.296*** (0.080)	-0.121 (0.131)
Demographic Controls	N	Y	Y	Y
Politician Controls	N	N	Y	Y
Observations	654	654	654	269
Mean Outcome	0.46	0.46	0.46	0.47
R ²	0.006	0.095	0.117	0.147

Notes: Robust standard errors presented in parentheses. Anxiety prime emphasizes the possibility of a second wave of infections, a longer wait before vaccines become available, and suggested that existing reports likely undercounted the number of cases and deaths. The “Attentive” subset are those who spent more than 3 minutes completing our short survey. † $p < 0.10$; * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

We also find that these relationships dominate even when we randomly vary the policy

platform adopted by our hypothetical candidates. Specifically, the increased preference for the mainstream candidate holds even when the comparison is with an anti-establishment candidate running on a platform centered on health care reform. Any variation in other qualities (e.g. leadership or executive experience) that may be inferred from the different occupations or educational histories of our candidates are randomly assigned, and in any case have no detectable effect on voting intent. Furthermore, the experimental setting is about two hypothetical challengers running for executive office, controlling for the potential strategic voting behavior of a primary setting (Abramson et al., 1992). As such we interpret these results as a direct test of mainstream vs. anti-establishment preference in response to anxiety that obtains independent of specific candidate qualities such as leadership or strategic assessments of candidate electability in a general election.

6 Generalizability

Our analysis thus far combines observational evidence from the 2020 Democratic primaries with a survey experiment fielded among U.S.-based respondents. Our findings consistently support our argument that voters exhibit a flight-to-safety during periods of heightened anxiety. However, does the effect generalize to other offices? And is there evidence that a flight to safety operates outside of the U.S.?

To address these questions, we conduct two additional observational studies. The first looks at primary elections for members of the House of Representatives in 2020 to show that more extreme candidates suffered in states with later elections and in congressional districts with greater exposure to the pandemic. The second looks at 2020 French municipal elections to demonstrate that similar patterns obtain in contexts outside the United States.

U.S. House Primaries

We classify a candidate for a House seat as “extreme” if they are endorsed either by the Justice Democrats movement (an anti-establishment leftist coalition), or by the Tea Party movement (an anti-establishment rightist coalition), yielding 8 Democratic Party candidates and 62 Republican party candidates that we consider anti-establishment. We estimate an interacted specification in which we predict the candidate’s electoral support as a function of whether they are an extreme candidate, the number of Covid-19 cases or deaths in their Congressional District, and the interaction. Formally,

$$y_{c,d} = \gamma_d + \beta_1 \text{anti-estab}_c + \beta_2 \text{covid}_d + \beta_3 \text{anti-estab} \times \text{covid} + \varepsilon_{c,d} \quad (3)$$

where subscripts d represent Congressional Districts and c indicates the candidates. We include district fixed effects and cluster standard errors at the district. The results, summarized in Table 3, indicate that anti-establishment candidates received less support at the ballot box in areas more exposed to the pandemic, as seen by the negative and statistically significant interaction terms. These patterns hold whether we predict variation in their contest-specific vote share, or the logged votes they received after controlling for total votes cast. The patterns also obtain when we replace the logged Covid-19 cases predictor with logged deaths, although they are weaker.²²

2020 French Municipal Elections

We conclude our analysis by testing whether a similar pattern holds outside of the United States. We obtain department-level data on Covid-19 cases in France, which we

²²Note that the insignificant positive coefficient on the Justice Democrats is based on a smaller sample (8 candidates total), only two of which competed in primary races on March 3rd that were plausibly prior to the Covid-19 outbreak. Furthermore, these candidates competed in the same race – California’s 6th Congressional District.

Table 3: Anti-establishment vote share as a function of Covid-19 exposure

	<i>Dependent variable: Vote Share</i>		
	Anti-Estab.	Justice Dems	Tea Party
	(1)	(2)	(3)
Anti-Estab.	0.342*** (0.043)	0.134** (0.050)	0.364*** (0.043)
Cases (ln)	-0.044 (0.150)	-0.014 (0.151)	-0.031 (0.151)
Anti-Estab. X Cases (ln)	-0.016** (0.006)	0.014 (0.012)	-0.020** (0.006)
District FE	Y	Y	Y
Observations	2,019	2,019	2,019
# Anti-Estab. Candidates	70	8	62
R ²	0.516	0.506	0.515

District-clustered SEs in parentheses. †p<0.1; *p<0.05; **p<0.01; ***p<0.001

match with election returns from the two-wave municipal election cycle of 2020. These elections occurred in two waves in the first half of 2020, first on March 15th just prior to the widespread outbreak of the pandemic, and again on June 28th.²³ We consider all centrist parties (LREM, LMDM, LUDI, LUC, LDVC), as well as the largest parties on the center-left (LSOC) and center-right (LRR), mainstream parties. We plot the descriptive change in electoral fortunes of these parties between March and June in Figure 8, highlighting the electoral penalty suffered by less mainstream parties. These descriptive patterns are highly statistically significant in a difference-in-differences specification in which we interact mainstream status with election wave, suggesting an almost 10 percentage point shift in relative electoral fortunes (results included in the Supporting Information). There is also evidence of greater effects where COVID was more widespread, suggesting both temporal

²³The latter round was postponed due to the pandemic, complicating the causal argument used in the context of the U.S. Democratic primary elections. We present these results as correlations consistent with our main results.

and geographic heterogeneity consistent with our other findings.

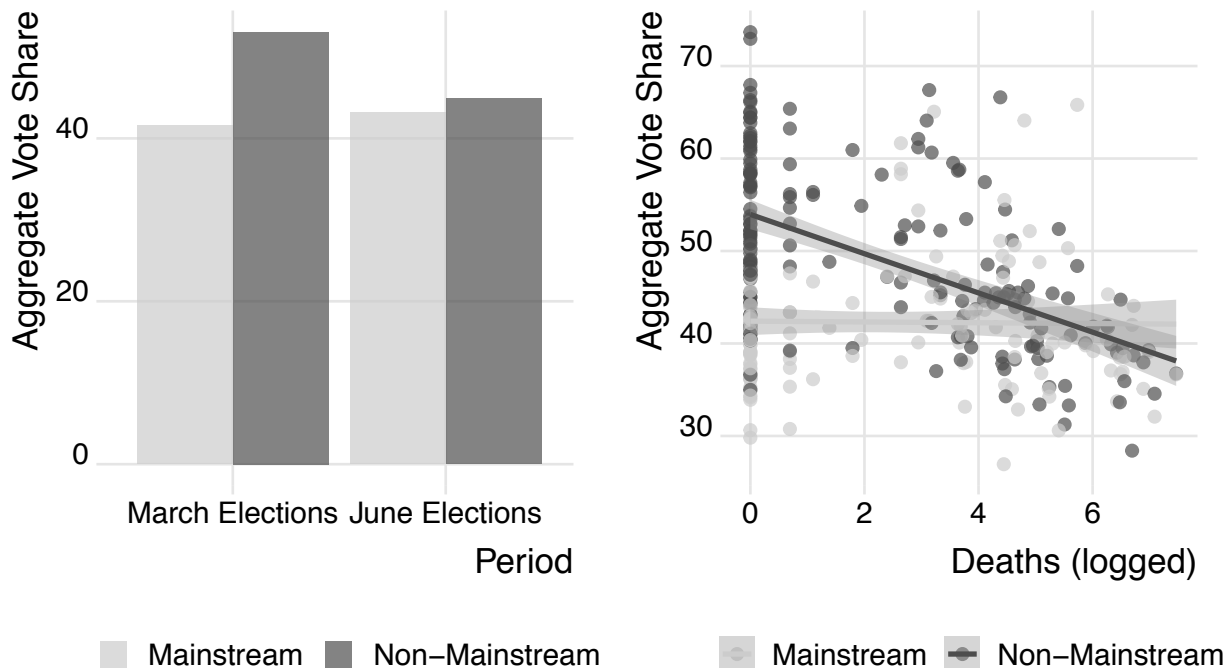


Figure 8: Descriptive evidence from France, indicating that mainstream parties benefited (gray), and non-mainstream parties lost (black), between the March and June municipal-level elections. Left panel plots average aggregate vote share (y-axis) by party (shaded bars) and period (x-axis). Right panel plots average aggregate vote share (y-axis) by logged Covid-19 deaths (x-axis) and party (shaded points).

7 Discussion

In this paper, we build on a rich literature to develop a general understanding of how anxiety influences vote choice. Our framework – which we refer to as a political flight to safety – predicts that shocks to voter anxiety improve the electoral chances of mainstream candidates.

We provide evidence of our claim across four empirical contexts related to Covid-19. Our main analysis describes a causal effect of the pandemic on Bernie Sanders’ declining fortunes in the Democratic primary election of 2020. We supplement this result with similar

evidence in the House of Representatives, in which anti-establishment candidates disproportionately lost where exposure to the virus was greater. We also show that this pattern travels outside of the United States, finding that less mainstream parties in French municipal elections were penalized at the ballot box between the first (March 2020) and second (June 2020) rounds of voting. Finally, we fielded a survey experiment in which we experimentally manipulated the anxiety-inducing qualities of a prime about Covid-19, before asking respondents to indicate their preference for a more mainstream or more anti-establishment candidate. Across all contexts, we find consistent evidence that anxiety prompts a preference for the status quo, regardless of other attributes including policy positions, experience, and office.

The magnitude of the effects we summarize are non-trivial, ranging from a 2 percentage point penalty against anti-establishment candidates for the House of Representatives, to 7 percentage points in our survey experiment, to 10 percentage points in French municipal elections, to between 7 and 15 percentage points in the Democratic primary election between Bernie Sanders and Joe Biden. These effect sizes are substantially larger than those from existing work on political campaigns themselves, where coefficient magnitudes greater than 1 percentage point are rare (Kalla and Broockman, 2018).

There are some reasons to believe we capture an unusually large treatment ‘dosage’ for our underlying anxiety mechanism in the sense that Covid-19 is a particularly large shock. Future research might explore the extent to which other sources of anxiety induce a political flight to safety and the magnitude of the flight to safety these sources produce. Further work might also explore whether the flight to safety is conditioned by, among other factors, the closeness of elections and thus the likelihood that an anti-establishment candidate – itself a potential source of anxiety for some voters – might actually be elected.

Our intuition accommodates existing research documenting a positive relationship between anxiety and voter preference for leadership (Merolla and Zechmeister, 2009), protection

(Albertson and Gadarian, 2015), and conservatism Stenner (2005). However, a political flight to safety is at once both more general and more precise. While it operates through the valence (rather than policy preference) component of voters' vote choice functions, it is not a specific candidate quality such as leadership or experience that is rewarded, nor a directional right-wing preference that is activated. Building on existing theory, this paper provides evidence that the flight to safety operates even where there is no incumbent and where retrospective evaluations do not influence voter behavior. A broad valence preference for the safety of the familiar is intensified by voters' anxiety – a preference that is always present, but receives more weight under conditions of anxiety. In this sense, our framework connects with work on system justification which predicts that even those disadvantaged under the status quo will defend it in order to reduce uncertainty, particularly when anxiety is higher (Jost et al., 2003; Van der Toorn et al., 2015).

The implications of voters' anxiety-induced preference for the status quo are far reaching, particularly if the political strength of radical movements is itself a source of public anxiety. A political flight to safety may operate as a brake on radical change, carrying provocative implications for anti-establishment candidates, parties, and even democratic accountability. We leave these extensions to future research. What this paper can say with confidence is that, as it does in financial markets, anxiety prompts a flight to safety in the market for political candidates. The COVID flight to safety is associated with quite substantial effect sizes, and is a force warranting further attention by scholars, political actors, and civically concerned citizens.

References

- Abramson, Paul R, John H Aldrich, Phil Paolino and David W Rohde. 1992. “Sophisticated” voting in the 1988 presidential primaries.” *The American Political Science Review* pp. 55–69.
- 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.
- Albertson, Bethany and Shana Kushner Gadarian. 2015. *Anxious politics: Democratic citizenship in a threatening world*. Cambridge University Press.
- Ansolabehere, Stephen and James M Snyder. 2000. “Valence politics and equilibrium in spatial election models.” *Public Choice* 103(3-4):327–336.
- Azevedo, Flávio, John T Jost and Tobias Rothmund. 2017. ““Making America great again”: System justification in the US presidential election of 2016.” *Translational Issues in Psychological Science* 3(3):231.
- Berinsky, Adam J and Jeffrey B Lewis. 2007. “An estimate of risk aversion in the US electorate.” *Quarterly Journal of Political Science* 2(2):139–154.
- Callen, Michael. 2015. “Catastrophes and time preference: Evidence from the Indian Ocean Earthquake.” *Journal of Economic Behavior & Organization* 118:199–214.
- Callen, Michael, Mohammad Isaqzadeh, James D Long and Charles Sprenger. 2014. “Violence and risk preference: Experimental evidence from Afghanistan.” *American Economic Review* 104(1):123–48.
- Cameron, Lisa and Manisha Shah. 2015. “Risk-taking behavior in the wake of natural disasters.” *Journal of Human Resources* 50(2):484–515.

- Campante, Filipe R, Emilio Depetris-Chauvin and Ruben Durante. 2020. “The Virus of fear: The political impact of Ebola in the US.”
- Carroll, Royce, Jeffrey B Lewis, James Lo, Keith T Poole and Howard Rosenthal. 2013. “The structure of utility in spatial models of voting.” *American Journal of Political Science* 57(4):1008–1028.
- Cassar, Alessandra, Andrew Healy and Carl Von Kessler. 2017. “Trust, risk, and time preferences after a natural disaster: experimental evidence from Thailand.” *World Development* 94:90–105.
- Clifford, Scott and Jennifer Jerit. 2018. “Disgust, anxiety, and political learning in the face of threat.” *American Journal of Political Science* 62(2):266–279.
- Cohen, Marty, David Karol, Hans Noel and John Zaller. 2009. *The party decides: Presidential nominations before and after reform*. University of Chicago Press.
- Cohn, Alain, Jan Engelmann, Ernst Fehr and Michel André Maréchal. 2015. “Evidence for countercyclical risk aversion: An experiment with financial professionals.” *American Economic Review* 105(2):860–85.
- Debenedetti, Gabriel. 2020. “Biden Is Planning an FDR-Size Presidency.” *Intelligencer* .
URL: <https://nymag.com/intelligencer/2020/05/joe-biden-presidential-plans.html>
- Downs, Anthony et al. 1957. “An economic theory of democracy.”
- Drezner, Daniel. 2020. “Perspective—Can the Democratic Party still decide?”
URL: <https://www.washingtonpost.com/outlook/2020/03/03/can-democratic-party-still-decide/>
- Eckel, Catherine C, Mahmoud A El-Gamal and Rick K Wilson. 2009. “Risk loving after the storm: A Bayesian-Network study of Hurricane Katrina evacuees.” *Journal of Economic Behavior & Organization* 69(2):110–124.

- Eckles, David L, Cindy D Kam, Cherie L Maestas and Brian F Schaffner. 2014. "Risk attitudes and the incumbency advantage." *Political Behavior* 36(4):731–749.
- Ehrlich, Sean and Cherie Maestas. 2010. "Risk orientation, risk exposure, and policy opinions: The case of free trade." *Political Psychology* 31(5):657–684.
- 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_story.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>
- FiveThirtyEight. 2020. "Who Will Win the 2020 Democratic Primary?" *FiveThirtyEight.com* .
URL: <https://projects.fivethirtyeight.com/2020-primary-forecast/?cid=rrpromo>
- 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.
- Hanaoka, Chie, Hitoshi Shigeoka and Yasutora Watanabe. 2018. "Do risk preferences change? evidence from the great east japan earthquake." *American Economic Journal: Applied Economics* 10(2):298–330.
- Hazlett, Chad and Yiqing Xu. 2018. "Trajectory balancing: A general reweighting approach to causal inference with time-series cross-sectional data." *Available at SSRN 3214231* .
- Holman, Mirya R, Jennifer L Merolla and Elizabeth J Zechmeister. 2016. "Terrorist threat, male stereotypes, and candidate evaluations." *Political Research Quarterly* 69(1):134–147.

- Holman, Mirya R, Jennifer L Merolla, Elizabeth J Zechmeister and Ding Wang. 2019. "Terrorism, gender, and the 2016 US presidential election." *Electoral Studies* 61:102033.
- 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.
- Jost, John and Orsolya Hunyady. 2003. "The psychology of system justification and the palliative function of ideology." *European review of social psychology* 13(1):111–153.
- Jost, John T. 2019. "A quarter century of system justification theory: Questions, answers, criticisms, and societal applications." *British Journal of Social Psychology* 58(2):263–314.
- Jost, John T, Jack Glaser, Arie W Kruglanski and Frank J Sulloway. 2003. "Political conservatism as motivated social cognition." *Psychological bulletin* 129(3):339.
- Kahsay, Goytom Abraha and Daniel Osberghaus. 2016. "Extreme weather and risk preference: Panel evidence from Germany." *ZEW Discussion Papers* 16.
- Kalla, Joshua L and David E Broockman. 2018. "The minimal persuasive effects of campaign contact in general elections: Evidence from 49 field experiments." *American Political Science Review* 112(1):148–166.
- Kam, Cindy D and Elizabeth N Simas. 2010. "Risk orientations and policy frames." *The Journal of Politics* 72(2):381–396.
- Kam, Cindy D and Elizabeth N Simas. 2012. "Risk attitudes, candidate characteristics, and vote choice." *Public Opinion Quarterly* 76(4):747–760.
- Krehbiel, Keith. 1998. *Pivotal politics: A theory of US lawmaking*. University of Chicago Press.

- MacKuen, Michael, George E Marcus, W Russell Neuman and Luke Keele. 2007. "The third way: The theory of affective intelligence and American democracy." *The affect effect: Dynamics of emotion in political thinking and behavior* pp. 124–151.
- Marcus, George E, W Russell Neuman and Michael MacKuen. 2000. *Affective intelligence and political judgment*. University of Chicago Press.
- Merolla, Jennifer L and Elizabeth J Zechmeister. 2009. "Terrorist threat, leadership, and the vote: Evidence from three experiments." *Political Behavior* 31(4):575.
- Morgenstern, Scott and Elizabeth Zechmeister. 2001. "Better the devil you know than the saint you don't? Risk propensity and vote choice in Mexico." *Journal of Politics* 63(1):93–119.
- Paterson, Janice M and Jane Cary. 2002. "Organizational justice, change anxiety, and acceptance of downsizing: Preliminary tests of an AET-based model." *Motivation and emotion* 26(1):83–103.
- Rickershauser, Jill and John H Aldrich. 2007. "'It's the electability, stupid' - or maybe not? Electability, substance, and strategic voting in presidential primaries." *Electoral Studies* 26(2):371–380.
- 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>
- Shepsle, Kenneth A. 1972. "The strategy of ambiguity: Uncertainty and electoral competition." *American Political Science Review* 66(2):555–568.
- Simas, Elizabeth N. 2017. "The effects of electability on US primary voters." *Journal of Elections, Public Opinion and Parties* 27(3):274–290.
- Stenner, Karen. 2005. *The authoritarian dynamic*. Cambridge University Press.

- 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>
- Stuart, Elizabeth A. 2010. “Matching methods for causal inference: A review and a look forward.” *Statistical science: a review journal of the Institute of Mathematical Statistics* 25(1):1.
- Tausanovitch, Chris, Lynn Vavreck, Tyler Reny, Alex Russell Hayes and Aaron Rudkin. 2019. “Democracy fund+ UCLA nationscape methodology and representativeness assessment.” .
- Van der Toorn, Jojanneke, Matthew Feinberg, John T Jost, Aaron C Kay, Tom R Tyler, Robb Willer and Caroline Wilmuth. 2015. “A sense of powerlessness fosters system justification: Implications for the legitimation of authority, hierarchy, and government.” *Political Psychology* 36(1):93–110.
- Voors, Maarten J, Eleonora EM Nillesen, Philip Verwimp, Erwin H Bulte, Robert Lensink and Daan P Van Soest. 2012. “Violent conflict and behavior: a field experiment in Burundi.” *American Economic Review* 102(2):941–64.
- 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>

SUPPORTING INFORMATION
Flight to Safety:
Covid-Induced Changes in the Intensity of Status Quo
Preference and Voting Behavior

James Bisbee
NEW YORK UNIVERSITY CSMAP

Dan Honig
UNIVERSITY COLLEGE LONDON

Authors' Final Version, June 17, 2021

Contents

A	Controls	3
B	Trajectory Balancing Weights	3
C	“Localness” of the DMA Fixed Effects	5
D	Considering Alternative Explanations	6
	D.1 Evidence of Anxiety: Robustness	6
	D.2 “Party Decides” Placebo Tests	7
	D.3 Turnout	9
	D.4 Selection Effects	11
	D.5 Economic Policy Preferences	11
E	Balance and Weighting Robustness	13
F	Survey Experiment	17

F.1 Ethical Considerations	22
F.2 Survey Analysis	23
G Sanders v Biden Ideological Placement	24
H Generalizability	25
H.1 Democratic Party House Primary	25
H.2 France Diff-in-Diff	25

A Controls

We obtain a rich set of county-level controls from the American Community Survey 5-year averages, collected in 2018. These data are publicly available from ACS from <https://www.census.gov/programs-surveys/acs/news/data-releases/2018/release.html>. Our main specifications include:

- % of the population with less than a high school education
- % of the population with a college degree or higher
- % of the population younger than 30
- % of the population older than 60
- % of population between 18 and 64 that is below the poverty level
- Share of households that are headed by a woman without a husband present
- the county-level unemployment rate
- the county-level labor force participation rate
- % of the population employed in manufacturing
- the median household income
- % of the population that is rural
- % of the population that speaks only English
- % of the population that is white
- % of the population that is black
- total population (logged)
- Turnout in 2020 primary
- Support for Sanders in 2016

B Trajectory Balancing Weights

To bolster our assumption that we are capturing exogenous variation in exposure to Covid-19 at the time of the primary election, we use trajectory balancing methods to reweight our control counties to match our treated counties (Hazlett and Xu, 2018). Unlike the motivating use case for this method where the outcome is measured multiple times

per unit, we only observe a county’s vote choice once. As such, we target the method to balance treated and control counties on their pre-treatment covariates as well as on their full history of exposure. The results of this weighting procedure are summarized in Figure 1. By matching on the full history of Covid-19 cases, we bolster our claim that the comparison is between those counties who voted on March 17th and those who voted earlier but were otherwise identical.




Figure 1: Trajectory balancing example for counties that voted on March 17th with one or more confirmed case in their DMA. We implement trajectory balancing (Hazlett and Xu, 2018) to re-weight the control counties that voted prior to March 17th (in red) to more closely resemble the treated counties (thick gray) in terms of the trajectory of the outbreak and pre-treatment covariates. The comparison of Sanders support is thus between those treated counties that voted on March 17th, and the re-weighted control counties given by the dotted black line that voted earlier.

C “Localness” of the DMA Fixed Effects

Our simplest regression specification uses DMA fixed effects to compare counties in the same media environment who voted at different times. Since we assign the exposure measure at the DMA-level, these fixed effects mean that counties located within DMAs that don’t cross state lines are dropped, since their treatment measure is collinear with the fixed effect. The benefit of this estimation strategy is that we compare counties in the same area who happen to fall on two sides of a state border, one of which voted earlier than the other. The downside is that we rely on a subset of counties that differ from the those that are dropped, although this problem only obtains for the first set of results that rely on these fixed effect specifications. Table 1 summarizes these differences, highlighting that the counties driving the first set of results differ systematically from those that are dropped from the analysis. These covariates are ordered by the t-statistics of a t-test difference in means, in which negative statistics mean that the counties we drop have smaller values for the covariates than those we keep. As illustrated, while there are important differences between these counties, again we note that these differences largely work against our main findings. In particular, we note that the counties we keep were more supportive of Sanders in 2016 (46.2 versus 39.9), suggesting that any selection bias introduced by our sample works against our main findings of a penalty for Sanders in 2020.

	Variable	avg_0	avg_1	p_value	t_value
1	Speak only English	86.5	92.4	0	-10
2	% Sanders16	39.9	46.2	0	-8.9
3	Manufacturing	11.7	13.6	0	-6.4
4	% Rural	54.6	60.5	0	-4.3
5	Labor Force Part Rate	57.3	58.8	0	-4.2
6	% White	81.7	84.5	0	-3.9
7	% CollUp	29.1	30.2	0.02	-2.4
8	% 60Up	24.8	24.8	0.9	-0.1
9	Md HH Income	49,621.4	49,362.3	0.6	0.5
10	Old age dep ratio	31.1	30.6	0.3	1
11	% Below poverty level	9.2	9	0.2	1.2
12	% Turnout 2020	8.8	8.4	0.1	1.5
13	% LT30yo	37.4	37	0.1	1.9
14	% Black or African American	10.4	8.6	0.01	2.5
15	Unem rate pop 16 over	6.7	6.3	0	3.3
16	% Female hher no hus	11.7	10.9	0	4.1
17	Total Pop	158,260.9	79,194	0	4.1
18	% LTHS	14.9	13.5	0	4.7

Table 1: Balance table comparing counties that we drop (0) and that we keep (1) with DMA fixed effects.

D Considering Alternative Explanations

D.1 Evidence of Anxiety: Robustness

Our theorized mechanism rests on the assumption that differences in the exposure to the pandemic cause differences in anxiety, which then generate differences in observed vote shares for Sanders. In the body of our paper, we demonstrate that there is daily evidence of cross-sectional relationships between exposure to the virus and Google searches for the term “coronavirus”, which we interpret as a proxy for anxiety.

Our second measure of anxiety is a weekly measure of county-level mobility, derived from GPS-enabled cell phones. We posit that a reduction in movement reflects increased concern about the health risks associated with the virus and is thus a proxy for anxiety that varies over space and time in the first weeks of March. We obtain data on county-level mobility by week from Cuebiq (2020). Cuebiq partners with 86 smartphone apps to collect individual-level location data from opted-in users. These data are aggregated from individual cell phone GPS data which measures a box around all locations observed for users on a given day and calculates mobility as the (logged) distance between opposite corners of this box.¹ We measure these county-level values every week from January 20th (when the first US-based cases were reported in Washington state) to April 13th. Figure 2 plots a series of regressions predicting changes in county-level mobility as a function of DMA-level cases of Covid-19. As illustrated, there is no meaningful relationship between this behavior and cases until after Super Tuesday, at which point the coefficients become significantly negative, suggesting that anxiety over the health risks of the pandemic grew over the first few weeks of March and, importantly, that they were correlated with geographic variation in exposure during these weeks.

These estimates are likely confounded with changes in federal and state policies on social distancing, particularly if these policies are correlated with geographic variation in the outbreak. However, we emphasize that these estimates are based on DMA-level cases, many of which cut across state borders. Since the vast majority of the policies aimed at reducing movement were enacted at the state level, a pure story of responding to policy initiatives is unable to fully explain the variation in DMA-level cases that we use to identify these coefficients.² In sum, we argue that these negative coefficients suggest that the geographic

¹These individual daily data are aggregated up to the county-week and placed on an index ranging from 1 to 5 based on the median mobility of all users in a county. Values on this index correspond to approximately:

- 1: 10 meters
- 2: 100 meters
- 3: 1 kilometer
- 4: 10 kilometers
- 5: 100 kilometers

²All states except ND, SD, NE, IA, WY, and AR had stay at home orders; all states but SD, NE, WY,

variation in exposure that we use to predict vote choice is also correlated with reduced mobility in ways that aren't purely reflecting government policies, which we interpret as evidence of a meaningful association between exposure and anxiety.




Figure 2: Weekly (x-axis) coefficients (y-axis) on the relationship between DMA-level cases and county-level mobility, including state random effects. Vertical dashed bars indicate primary dates.

D.2 “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

UT, and AR closed non-essential businesses; all states but ND banned large gatherings; every state declared a state of emergency.

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 and compare the diff-in-diff results prior to, and following Super Tuesday. 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 3, our results are all null, regardless of whether we are comparing the pre-party decides voting behavior to Super Tuesday, March 10th, or March 17th.




Figure 3: Permutation test results for different choices of the pre- and post-outbreak cutoff, estimated using diff-in-diff. Densities represent 100 bootstrapped estimates of the impact of exposure on the change in support for Sanders when the DMA-level Covid-19 cases are randomly reshuffled. Consistent null results suggest that our main findings are not spuriously conflating the timing of the outbreak with the Democratic Party rallying around Biden.

The main diff-in-diff results use outbreak dates to separate treated and control elections as per Table 2, 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 4 summarizes these results for every specification at our disposal. The Democratic party consolidated support behind Biden ahead of Super Tuesday. As Figure 4 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.

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 2: Treatment and control elections by outbreak “start date”. February (Feb) primaries include IA, NH, NV, and SC. Super Tuesday (ST) primaries include AL, AR, CA, CO, ME, MA, MN, NC, OK, TN, TX, UT, VA and VT. 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.)

These results also serve as simple placebo tests by treating the later election as the control data, and the earlier as the treated data. As illustrated, these cases reveal a positive estimate, suggesting that Sanders did better in those areas that would be exposed on March 17th, but were not yet. Similar results hold if we look instead at the difference-in-differences specification election-by-election. As illustrated in Figure 5, the penalty against Sanders in Covid-exposed counties did not begin until after Super Tuesday.

D.3 Turnout

An alternative explanation for the results summarized in our manuscript is 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, it would suggest 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, working against our main results.

Nevertheless, we predict variation in primary turnout by exposure across age-groups,

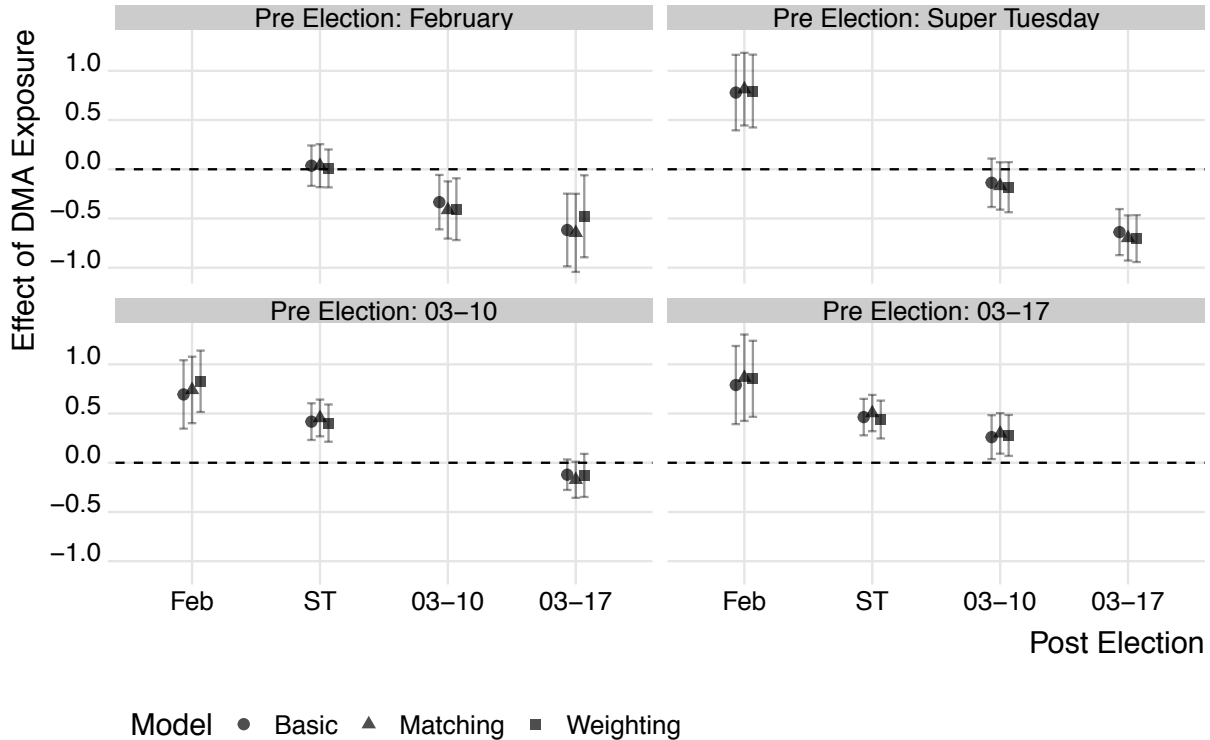


Figure 4: Pairwise election comparisons by definition of pre-outbreak election (panel titles) and post-outbreak comparison elections (x-axes).

estimating an interacted specification of the following form:

$$\text{turnout}_{ct} = \beta_0 + \beta_1 \text{Covid}_{ct} + \beta_2 \text{Age}_c + \beta_3 \text{Covid}_{ct} \times \text{Age}_c + \gamma \mathbf{X}'_c + \varepsilon_{ct} \quad (1)$$

where Age_c is the share of the county’s population that is older than 60 or younger than 30. We are interested in the β_3 coefficient which captures the interacted relationship between turnout and exposure to Covid-19 by age. As illustrated in Figure 6, there is little evidence to suggest that such an age-based turnout dynamic is active.

An alternative turnout story is that the Sanders campaign was effectively finished after Biden’s convincing victory in South Carolina on February 29th. In this scenario, a number of would-be Sanders voters were planning on casting what were effectively protest votes, and those in COVID-exposed areas didn’t bother since the “costs” of doing so were higher. If this were the case, we should expect to see a decline in turnout in more pro-Sanders counties following the South Carolina primary. As illustrated in Figures 7 and 8, there is little descriptive evidence that pro-Sanders voters stayed home after the South Carolina primaries.

To confirm this visual intuition, we run a set of regressions date-by-date, in each case predicting county-level turnout in 2020 with the same county’s support for Sanders in 2016, controlling for the share of the population over 60, the share with a college degree, the unemployment rate, and the share that is white. As illustrated in Figure 9, there is

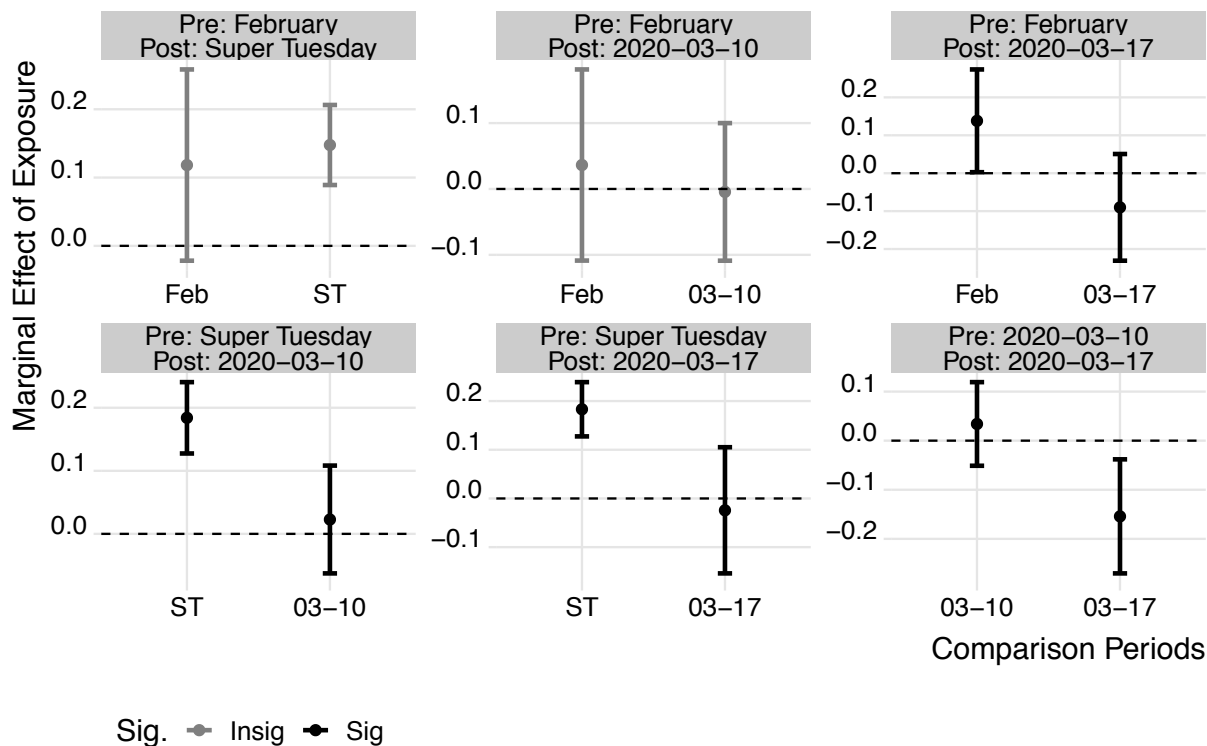


Figure 5: Pairwise election comparisons by definition of pre- and post-outbreak election.

no evidence suggesting that 2016 Bernie supporters stayed home after the South Carolina primary. To the contrary, there is a significant but small positive association between pro-Sanders counties in 2016 and 2020 turnout on Super Tuesday.

D.4 Selection Effects

Our results would be spurious if the outbreak disproportionately affected parts of the country that were already anti-Bernie to begin with. Although our matching and weighting strategies are one solution to minimizing this risk, we can also evaluate the identification challenge directly. We replace our main outcome variable with Sanders’ 2016 voteshare, testing whether 2020 exposure rates also predict 2016 Sanders support. We find, if anything, a source conservative bias as shown in Figure 10. Specifically, the counties that were more exposed to the outbreak in 2020 were, if anything, *more* supportive of Sanders in 2016, revealing a selection effect that works against our main results.

D.5 Economic Policy Preferences

The Covid-19 pandemic influenced more than just individuals’ anxiety over their health. It also precipitated a painful economic contraction in which the stock market lost

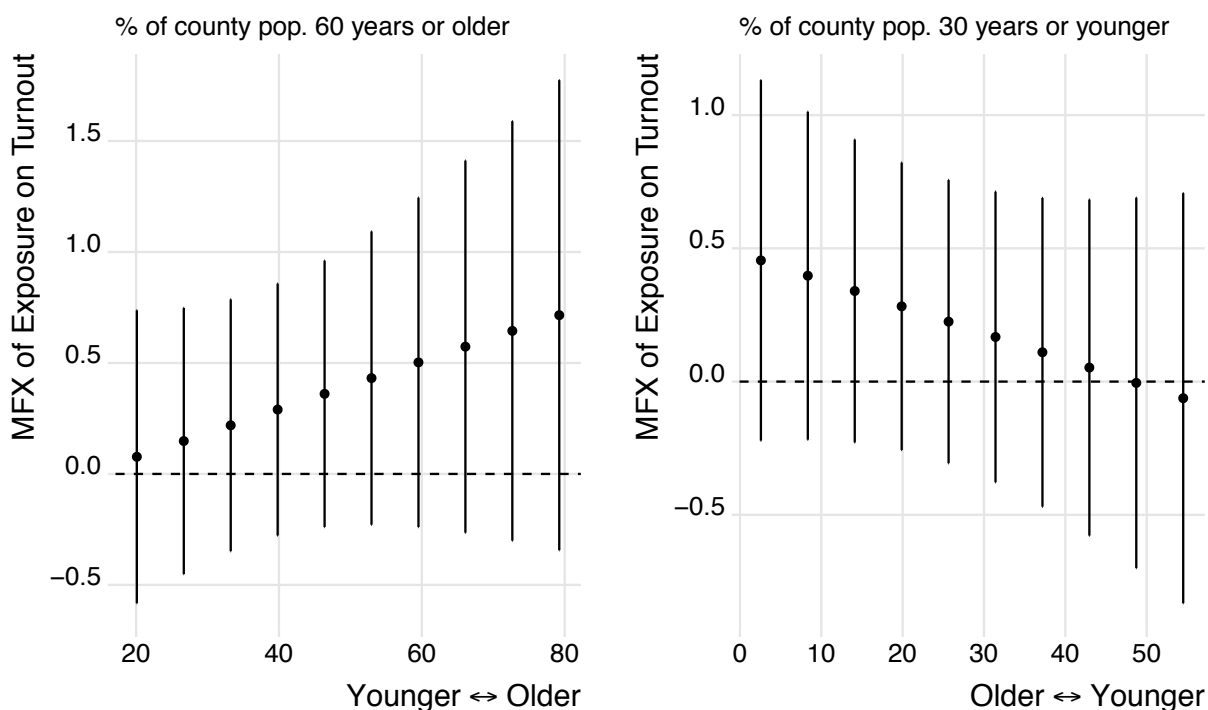


Figure 6: Marginal effects of exposure on logged turnout (y-axes) across counties with smaller and larger proportions of their population older than 60 years of age (left panel) and younger than 30 years of age (right panel).

over 20% of its value in the span of a week. An alternative story explaining the shift toward Biden focuses on the stock market contraction and the increased value of his relatively Wall Street-friendly policy position.

To examine whether the decline in support for Sanders was due not to a political flight to safety but rather due to his economic policy position, we examine how evaluations of the economy evolved between insulated and exposed areas in the early days of March. Our main results exploit variation in exposure and voting behavior across both geography and time. For the economy story to be true, it would require that places more exposed to the virus grew concerned about the economy earlier than those relatively insulated.

Survey results from the Nationscape survey (Tausanovitch et al., 2019) illustrate a clear pessimistic shift in March of 2020 in both insulated and exposed congressional districts (see Figure 11). While the exposed districts were more pessimistic in March, they were also more pessimistic in the months prior to the outbreak. A difference-in-differences regression confirms that the pre- and March differences between exposed and insulated areas are not significantly different from each other, as illustrated in Figure 12.

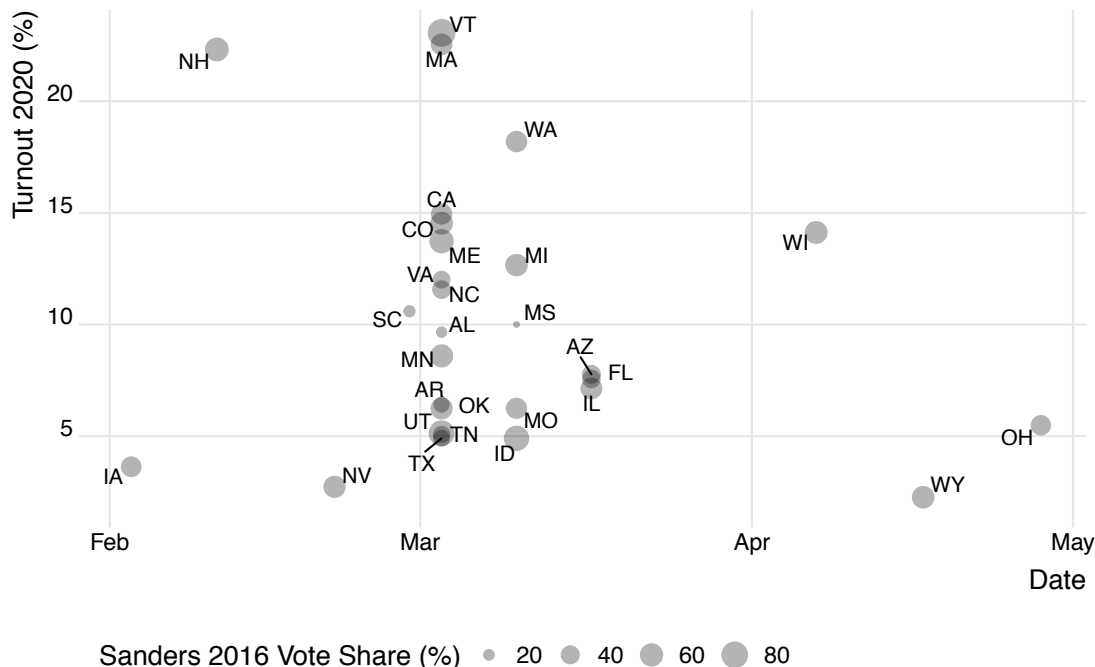


Figure 7: Average state-level turnout in the 2020 primary election (y-axis) by primary date (x-axis). Each point represents a state, sized by the overall average support for Sanders in the 2016 primary election. While there appears to be a secular decline in turnout over time, there is little evidence that this decline was starker in pro-Bernie states following Biden’s victory in South Carolina on February 29th.

E Balance and Weighting Robustness

We achieve good balance on both the matching and weighting strategies employed in the body of our paper. Figure 13 plots the improvements to balance on observables between treated and control units generated by our choice of nearest-neighbor matching using minimized Mahalanobis distance. And Table 3 summarizes the differences in treated and control covariates prior to, and following the `cbps` 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 `CBPS` method of Blackwell et al. (2009) with optimal weights (Zubizarreta, 2015), 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 in Table 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

Table 3: Balance Results for CBPS

	Covs	Diff_Unm	Bal_Test_Unm	Diff_Match	Bal_Test_Match
1	Prop.score			0.150	
2	Lths	-0.600	Not Balanced, >0.05	0	Balanced, <0.05
3	Collup	0.590	Not Balanced, >0.05	0.040	Balanced, <0.05
4	Caucus switch	0.060	Not Balanced, >0.05	0	Balanced, <0.05
5	Lt30yo	-0.190	Not Balanced, >0.05	-0.060	Not Balanced, >0.05
6	60up	0.160	Not Balanced, >0.05	0.070	Not Balanced, >0.05
7	Below poverty level	-0.570	Not Balanced, >0.05	-0.020	Balanced, <0.05
8	Female HH no husband	-0.510	Not Balanced, >0.05	-0.010	Balanced, <0.05
9	Unem rate	-0.240	Not Balanced, >0.05	0.010	Balanced, <0.05
10	LFPR	0.360	Not Balanced, >0.05	-0.050	Balanced, <0.05
11	Manufacturing	-0.210	Not Balanced, >0.05	-0.040	Balanced, <0.05
12	Md inc HH	0.620	Not Balanced, >0.05	0.040	Balanced, <0.05
13	Rural	-0.560	Not Balanced, >0.05	-0.050	Balanced, <0.05
14	Speak only english	-0.310	Not Balanced, >0.05	-0.050	Not Balanced, >0.05
15	White	0.330	Not Balanced, >0.05	-0.050	Balanced, <0.05
16	Black or african american	-1.170	Not Balanced, >0.05	0.020	Balanced, <0.05
17	Ln tot pop	0.510	Not Balanced, >0.05	0.050	Not Balanced, >0.05
18	Turnout pct 20	0.420	Not Balanced, >0.05	0.030	Balanced, <0.05

Table 4: Balance results for optweight

	Covs	Diff_Unm	Bal_Test_Unm	Diff_Match	Bal_Test_Match
1	Lths	-0.530	Not Balanced, >0.05	0	Balanced, <0.05
2	Collup	0.540	Not Balanced, >0.05	0	Balanced, <0.05
3	Lt30yo	-0.170	Not Balanced, >0.05	0	Balanced, <0.05
4	60up	0.140	Not Balanced, >0.05	0	Balanced, <0.05
5	Below poverty level	-0.490	Not Balanced, >0.05	0	Balanced, <0.05
6	Female HH no husband	-0.460	Not Balanced, >0.05	0	Balanced, <0.05
7	Unem rate	-0.180	Not Balanced, >0.05	0	Balanced, <0.05
8	LFPR	0.290	Not Balanced, >0.05	0	Balanced, <0.05
9	Manufacturing	-0.260	Not Balanced, >0.05	0	Balanced, <0.05
10	Md inc HH	0.590	Not Balanced, >0.05	0	Balanced, <0.05
11	Rural	-0.550	Not Balanced, >0.05	0	Balanced, <0.05
12	Speak only english	-0.330	Not Balanced, >0.05	0	Balanced, <0.05
13	White	0.300	Not Balanced, >0.05	0	Balanced, <0.05
14	Black or african american	-1.160	Not Balanced, >0.05	0	Balanced, <0.05
15	Ln tot pop	0.520	Not Balanced, >0.05	0	Balanced, <0.05
16	Turnout pct 20	0.450	Not Balanced, >0.05	0	Balanced, <0.05

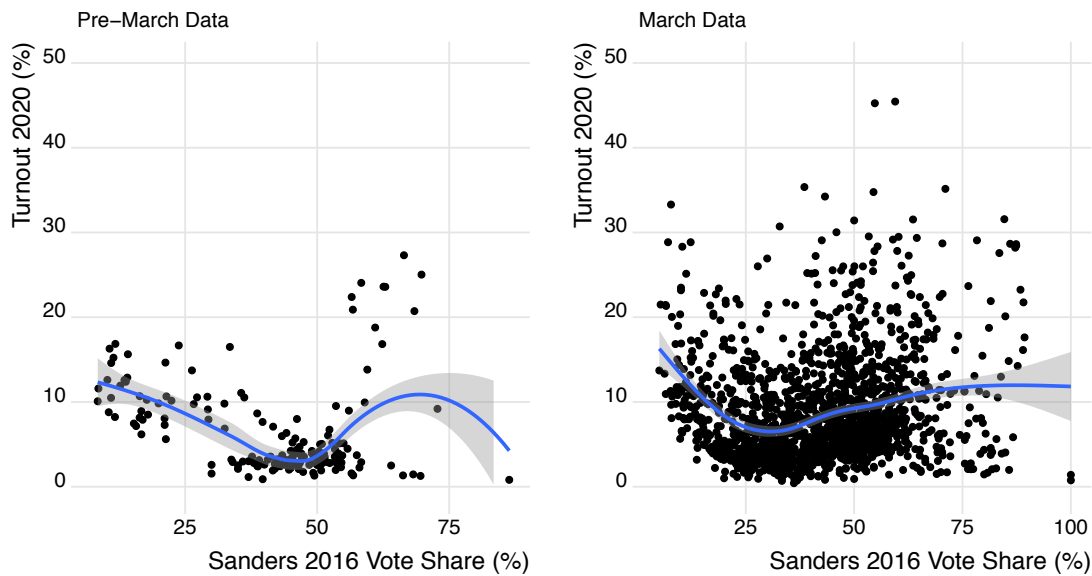


Figure 8: Scatter plots of Sanders’ 2016 vote share (x-axes) against turnout in the 2020 primary (y-axes), measured prior to the March elections, and during the March elections. While the most anti-Sanders counties had higher turnout in both periods, there is little evidence suggesting that Bernie supporters disproportionately stayed home after the Biden’s victory in South Carolina on February 29th.

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:

- % 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 5. The balance test results are visualized in Figure 14.

Table 5: CEM Robustness

	<i>Dependent variable:</i>					
	(1)	(2)	(3)	(4)	(5)	(6)
Treatbin	-1.053*** (0.338)	-0.835** (0.353)	-1.141*** (0.337)	-0.915*** (0.345)	-1.083** (0.453)	-0.950*** (0.358)
LTHS	-0.033 (0.049)	-0.083 (0.125)	-0.044 (0.100)	-0.051 (0.052)	-0.008 (0.158)	0.019 (0.108)
Coll. Up	0.064 (0.055)	0.300 (0.196)	0.145* (0.079)	0.051 (0.045)	0.187 (0.146)	0.071 (0.081)
Lt30yo	-0.065 (0.053)	-0.095 (0.267)	-0.067 (0.073)	-0.071 (0.043)	-0.245 (0.171)	-0.035 (0.073)
60up	-0.235*** (0.047)	-0.422* (0.215)	-0.312*** (0.076)	-0.279*** (0.039)	-0.584*** (0.160)	-0.345*** (0.072)
Below poverty level	0.054 (0.034)	0.076 (0.097)	-0.025 (0.051)	0.052 (0.034)	0.120 (0.146)	0.001 (0.055)
Female HH no hus	-0.050 (0.038)	0.063 (0.171)	-0.002 (0.060)	-0.062* (0.037)	-0.225 (0.204)	-0.061 (0.050)
Unem rate	-0.034 (0.027)	0.059 (0.095)	-0.049 (0.042)	-0.041 (0.030)	0.070 (0.136)	-0.037 (0.046)
LFPR	0.072* (0.038)	-0.238** (0.096)	-0.009 (0.050)	0.064 (0.040)	-0.101 (0.112)	-0.024 (0.056)
Manufacturing	-0.065** (0.032)	0.107 (0.066)	-0.010 (0.040)	-0.075*** (0.022)	0.026 (0.072)	-0.043 (0.039)
Med HH Inc	-0.215*** (0.053)	-0.096 (0.138)	-0.285*** (0.067)	-0.256*** (0.054)	-0.208 (0.170)	-0.228*** (0.069)
Rural	-0.041 (0.027)	-0.027 (0.108)	-0.018 (0.044)	-0.020 (0.024)	-0.033 (0.150)	0.020 (0.043)
Speak only english	-0.103*** (0.037)	-0.248 (0.151)	-0.067 (0.069)	-0.120*** (0.040)	-0.188 (0.180)	-0.002 (0.068)
White	0.029 (0.051)	-0.024 (0.227)	0.067 (0.060)	0.034 (0.054)	0.120 (0.347)	0.052 (0.068)
Black	-0.297*** (0.055)	-0.514** (0.220)	-0.236** (0.090)	-0.334*** (0.065)	-0.328 (0.386)	-0.247** (0.104)
Tot pop	0.015 (0.025)	0.031 (0.073)	-0.001 (0.029)	0.030 (0.028)	-0.008 (0.094)	0.025 (0.030)
Turnout pct 20	0.187*** (0.037)	0.204 (0.178)	0.153** (0.063)	0.246*** (0.025)	0.387*** (0.133)	0.283*** (0.047)
Observations	1,882	456	1,882	1,710	370	1,710
R ²	0.790	0.774	0.876	0.834	0.873	0.900

Note: DMA-cluster robust standard errors in parentheses. *p<0.1; **p<0.05; ***p<0.01

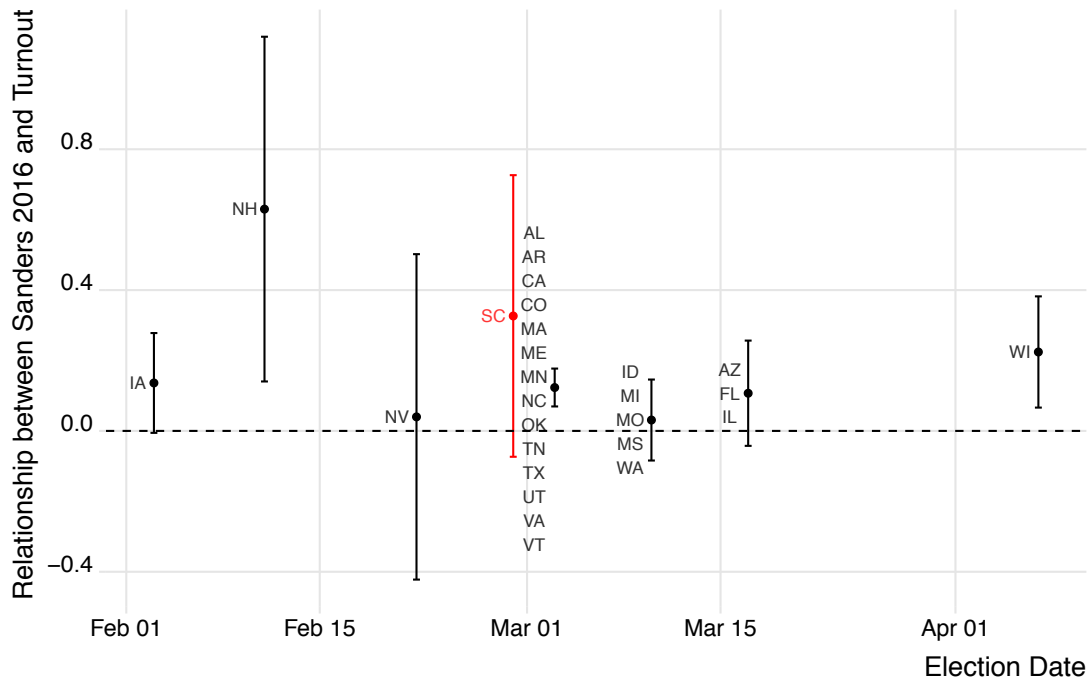


Figure 9: Coefficients (y-axis) estimate the relationship between a county’s 2020 turnout and their support for Sanders in 2016, estimated date-by-date (x-axis). South Carolina primary indicated in red. Ensuing primaries show no evidence of a negative relationship between 2020 turnout and 2016 support for Sanders.

F Survey Experiment

We fielded an online survey experiment between May 11th and May 20th, 2020. We used a convenience sample of 650 Amazon Mechanical Turk workers (“Turkers”). While the Turker population is not representative of Americans writ large (Berinsky, Huber and Lenz, 2012), researchers have been successful in replicating lab experiments on the platform (Berinsky, Huber and Lenz, 2012; Crump, McDonnell and Gureckis, 2013; Clifford, Jewell and Waggoner, 2015). We conducted power analysis prior to fielding the survey using the `DeclareDesign` package for R to identify this sample size as necessary for identifying effects commensurate to half what we observe in our analysis of the real-world primary vote share. IRB approval was obtained prior to fielding the survey.

Our survey consisted of four parts and was designed to be completed in approximately five minutes. The first section following the consent page asked the respondents to read a paragraph describing the COVID-19 pandemic and projecting the severity of the outbreak over the next year (i.e., into summer of 2021). After reading this description, respondents were presented with descriptions of two hypothetical challengers running for executive office, and asked to indicate which candidate they preferred on a 4-item Likert scale. The third section of the survey asked respondents to provide basic demographic information (age, sex, race, party affiliation) and also to indicate whether they personally had been infected by




Figure 10: Coefficient estimates connecting Sanders’ 2016 vote share with 2020 exposure to Covid-19, as of their election date. X-axis indicates different choices for fixed effects.

the virus and whether they knew personally anyone who had. The last section presented respondents with the alternative description of the pandemic, ensuring that all participants were given equal information.

Our core quantity of interest is the relationship between support for mainstream candidates in response to heightened anxiety. To operationalize this, we randomly varied the description of the future of the COVID-19 pandemic to be either reassuring or pessimistic. We then described the two hypothetical challengers as either mainstream or anti-establishment. The treatment text and the candidate descriptions are summarized below.

Treatment Text

Optimistic: Some experts believe the worst of the COVID-19 pandemic is behind us. Multiple states, and perhaps the country as a whole, are likely beyond their initial peak of COVID infections.³ The death toll predicted by leading models has been lowered, suggesting the effect of the virus will not be as terrible as initially feared.⁴ There are more than 60

³New York Times. “Coronavirus in the US: Latest Map and Case Count.”

⁴Raymond, Adam. “Key Coronavirus Model Now Predicts Many Fewer US Deaths.” New York Magazine; Shaw, Adam. “Top Coronavirus Model Significantly Lowers Total Estimates of US Deaths in New

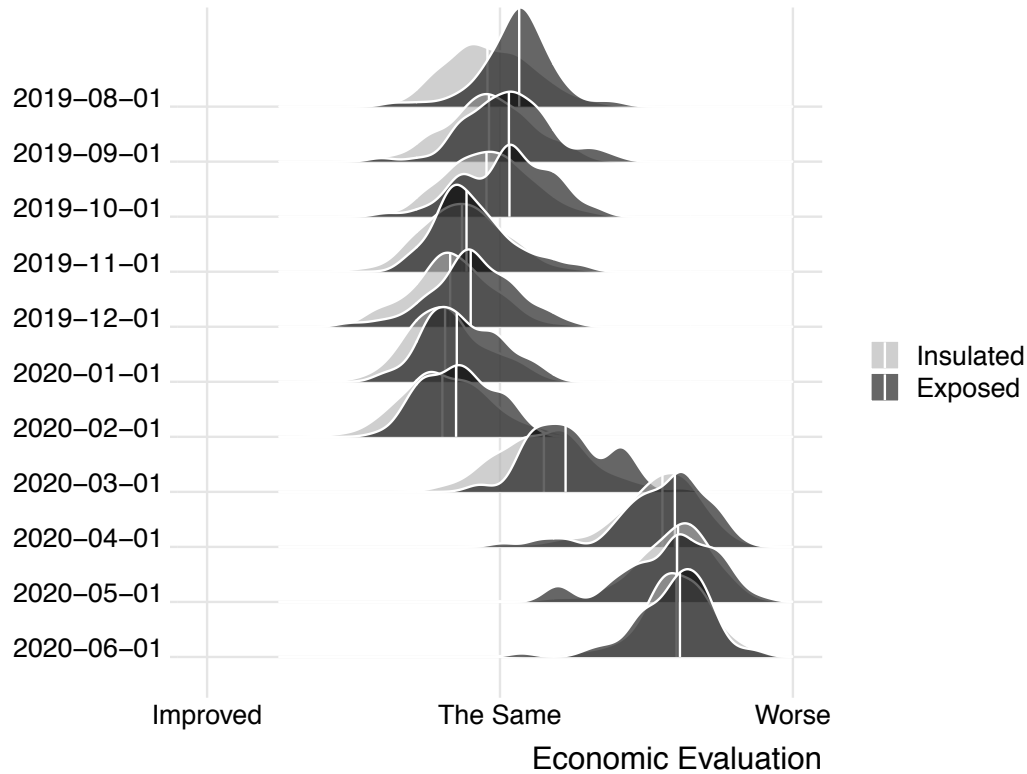


Figure 11: Economic evaluations by month between insulated (light gray) and exposed (dark gray) Congressional Districts, measured in March.

candidate vaccines now in development worldwide, and several have entered early clinical trials in human volunteers.⁵ There are also promising signs that existing anti-retroviral drugs may be effective in substantially reducing the severity and lethality of COVID-19 infection.⁶

Pessimistic: Some experts believe that COVID pandemic will continue to rage for many months to come. A well-respected group of pandemic experts believe COVID is “likely to keep spreading for at least another 18 months to two years—until 60% to 70% of the population has been infected.”⁷ Many experts also believe that loosening COVID-induced restrictions on mobility will lead to a 2nd wave of infections even worse than the initial wave.⁸ In addition, current figures may be markedly underestimating the death toll to date⁹ – suggesting that not only will the situation worsen in the months to come, but will do so by further declining from a status quo that is already worse than many realize.

Projection.” Fox News.

⁵Lanese, Nicoletta. “When Will a COVID-19 Vaccine be ready?” Livescience

⁶Feuerstein Herper. “Early peek at data on Gilead coronavirus drug suggests patients are responding to treatment.” Statnews.

⁷Fox, Maggie. “Expert Report Predicts Up to Two More Years of Pandemic Misery.” CNN.

⁸Weiss, Elizabeth. “When Will a Second Wave of the Coronavirus Hit, and What Will it Look Like?” USA Today.

⁹Wu, Jin and Allison McCann. “28,000 Missing Deaths: Tracking the True Toll of the Coronavirus Crisis.” New York Times.

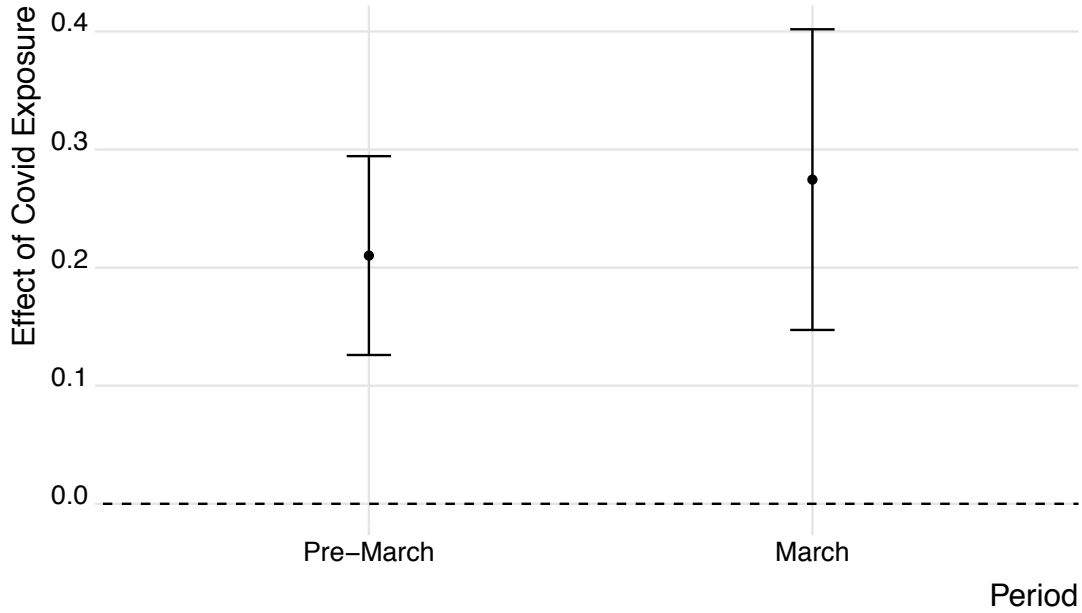


Figure 12: Difference-in-differences results comparing economic evaluations in insulated and exposed districts prior to March and in March. Positive values indicate more negative views of the economy.

Candidate Descriptions

The following descriptions were provided in which candidate A was always the anti-establishment candidate (or “disruptive”) and candidate B was always the mainstream candidate (or “safe”). The remainder of the sentences were randomly assigned to either candidate A or candidate B.

Disruptive: Candidate A seeks fundamental transformation of the economic, social, and political order. He believes that the system is broken, and the time for radical change is now. [SENTENCE 2]. [SENTENCE 3]. [SENTENCE 4].

Safe: Candidate B believes in strengthening existing economic, social, and political institutions. He believes that we must come together and re-invest in our system, strengthening its foundations to support generations to come. [SENTENCE 2]. [SENTENCE 3]. [SENTENCE 4].

Sentence 2

He is 48 years old, and was born and brought up in your area, before going to university to study physics.

He is 45 years old; he lives in the district and studied business at university.

Sentence 3:




Figure 13: Balance of treated and control covariates before (red) and after (blue) matching. 45 degree line indicates perfect match.

After university he trained as an accountant, and set up a company ten years ago; it now employs seven people.

He is a lawyer and runs a busy local practice.

Sentence 4 - Policy Platform:

NONE (Roughly half of respondents received no 4th sentence)




Figure 14: Balance performance across covariates using CEM (Blackwell et al., 2009).

He is passionate about achieving universal access to high-quality health care.

He is passionate about achieving universal access to high-quality education.

F.1 Ethical Considerations

To allow full assessment of the experiment, and in line with APSA Principles Guidance, we outline here the basic procedures we follow. This research, conducted with IRB approval, occurred entirely on the online platform M-Turk. The IRB ruled this research exempt, determining there was no harm to participants – in part because we did not collect any personally identifying information. Participants were restricted to US residents of eligible voting age (above 18; thus no children were in the subject pool), and were compensated at what would be substantially more than the US minimum wage (75 cents for an activity taking less than 5 minutes – the average time to completion was 3.67 minutes). While the experimental design allows us only limited understanding of the subject pool’s diversity, of 664 respondents the sample was 60% male, 39% female, and 1% other; the sample was 10% Hispanic, 80% white, 9% African American, and 11% identifying as another race. We do not believe that participation likely differentially benefited or harmed any participating groups.

F.2 Survey Analysis

Our main results assume that the spread of the Covid-19 virus across the United States was as-good-as-random. However, it may be that the virus hit anti-Sanders locations earlier than others. If this were to be true, it would mean that our findings of an anti-Sanders effect of exposure is a spurious byproduct of selection bias. To test the possibility of this selection effect, we rely on weekly survey data from Tausanovitch et al. (2019) which measures Sanders favorability by Congressional District. These data cover the period between July 2019 and April 2020, allowing us to examine whether the pandemic hit areas predisposed to vote again Sanders earlier in its spread.

We look for descriptive evidence of the pandemic’s spread penalizing Sanders in two ways. First, we plot the distribution of Sanders’ support by districts that were insulated (gray) and exposed (black) by month in Figure 15. We define districts as “insulated” or “exposed” based on their March and April number of cases. Specifically, if these districts were in the top quartile of Covid-19 cases during these months, they are classified as exposed. Otherwise they are insulated. Figure 15 illustrates that, if anything, the pandemic hit more pro-Sanders districts earlier, meaning that the selection bias should work against our findings that exposed voters were less likely to vote for Sanders.

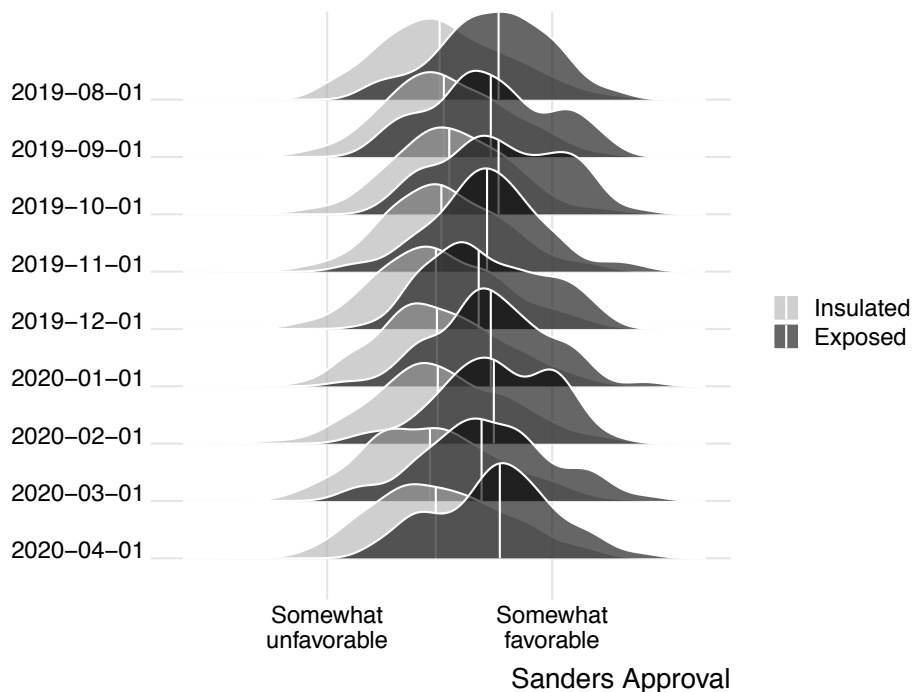


Figure 15: Histograms of favorability (x-axis) toward Bernie Sanders by month (y-axis) in congressional districts that were in the top quartile (black) or lower quartiles (gray) of Covid-19 cases in the months of March and April, 2020.

Second, we plot the week-by-week relationship between the number of Covid-19 cases and Sanders approval in Figure 16. Again, we find evidence suggesting that areas with more

cases were also those more favorable toward Bernie Sanders, regardless of what week we examine. Taken together, these results reaffirm that, to the extent that the virus did not spread randomly with respect to politics, it did so in a way to bias against our findings.




Figure 16: Week-by-week scatterplots of logged Covid-19 cases (x-axis) and average Sanders favorability (y-axis).

G Sanders v Biden Ideological Placement

We interpret the choice facing Democratic primary voters as one between the safety of the familiar (Biden) versus the uncertainty of the extreme (Sanders). We list a number of pieces of evidence in support of this interpretation in our paper, including exit poll data from CNN, the candidates’ own self-image, and evidence from *Azevedo, Jost and Rothmund (2017)* suggesting that Sanders supporters were less likely to support the economic, gender, or general systems of societal organization. We also used survey data from Pew Research Center’s American News Pathways data tool to summarize how Democrats themselves placed Sanders and Biden in terms of their ideology between February and March of 2020.¹⁰ As illustrated in Figure 17, there is convincing evidence that Democratic voters saw Sanders as the more ideologically extreme candidate, and Biden as the more moderate candidate. Insofar

¹⁰https://www.pewresearch.org/pathways-2020/DEM20IDEO_a/political_party_ideology/us_adults

as these ideological placements map on to feelings of safety (for Biden) versus uncertainty (for Sanders), we argue that these patterns further bolster our interpretation of the primary choice as between the safety of the familiar versus the uncertainty of the unknown.

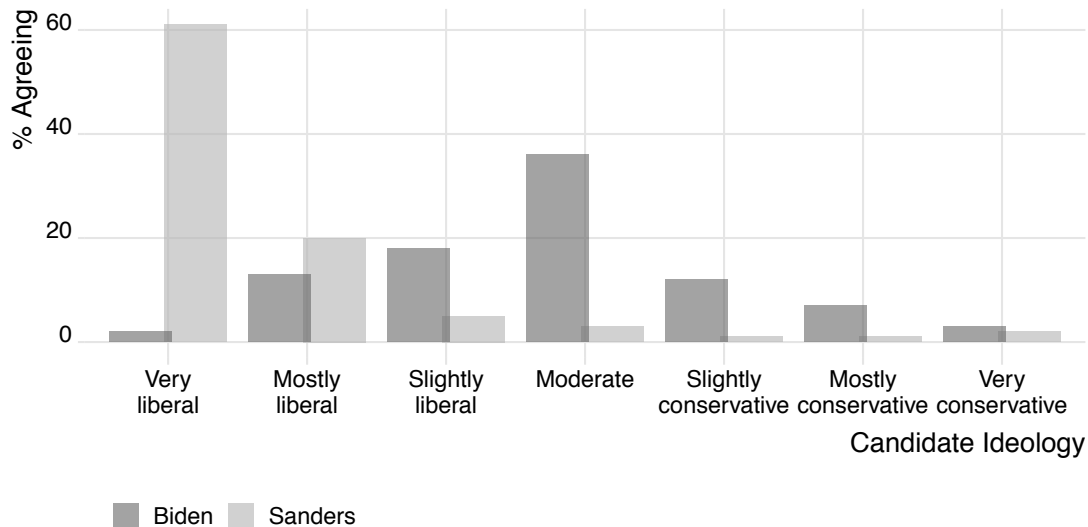


Figure 17: Average placement of Biden (dark gray) and Sanders (light gray) on a left-right ideological spectrum (x-axis). Data from Pew Research Center’s American News Pathways data tool.

H Generalizability

H.1 Democratic Party House Primary

Simple descriptive plots summarized in Figure 18 suggest that the electoral fortunes of the anti-establishment candidates for the house declined both over the course of the 2020 primary season, as well as over the geographic variation in Covid-19 exposure.

H.2 France Diff-in-Diff

The French municipal elections were held on March 15th, 2020 and on June 28th, 2020. Voting data at the Department level is obtained from <https://www.data.gouv.fr/fr/datasets/elections-municipales-2020-resultats-1er-tour/> for the first wave, and <https://www.data.gouv.fr/fr/datasets/municipales-2020-resultats-2nd-tour/> for the second wave of voting. We mapped these voting results to the number of Covid-19 cases in each area using data from <https://www.data.gouv.fr/fr/datasets/donnees-de-laboratoires-infra-departementales-durant-lepidemie-covid-19/>.




Figure 18: Extreme (black) versus mainstream (gray) vote shares (y-axes) by date (top panel) and logged cases (bottom panel).

We estimate a diff-in-diff model predicting the vote share for mainstream and non-mainstream parties for Department d with the following specification:

$$VS_d = \lambda_d + \beta_1 Main_d + \beta_2 Post + \beta_3 Main_d \times Post + \varepsilon_d \quad (2)$$

where λ_d are Department-fixed effects.

This specification only focuses on the temporal source of variation, and compares the electoral fortunes of non-mainstream parties between when voters went to the polls in March and June. The inclusion of Department fixed effects makes this a comparison of how these parties fared *in the same area*. However, it is possible that all non-mainstream parties experienced a secular decline in support for some other reason, such as the strategic decision of “first you choose, then you discard”.¹¹ Under the two-round system, strategic voting can result in a penalty against less popular candidates and parties in the second round, providing an alternative explanation for the penalty to non-mainstream parties we observe in the data (Dolez, Laurent and Blais, 2017).

We implement an alternative specification that predicts non-mainstream vote share

¹¹<https://blogs.lse.ac.uk/europpblog/2017/05/12/understanding-the-campaign-dynamics-of-the-french-presidential-election/>

as a function of geographic variation in exposure to Covid-19, measured as logged deaths. Formally, we estimate:

$$VS_{d,t} = \lambda_d + \delta_t + \beta_1 Main_d + \beta_2 Deaths_{d,t} + \beta_3 Main_d \times Deaths_{d,t} + \varepsilon_{d,t} \tag{3}$$

Here, the subscript t refers to the election date – either March 15th or June 28th. By implementing both Department (λ_d) and date (δ_t) fixed effects, we isolate variation due to the change in Covid-19 cases between the two elections across Departments. In both specifications, we find substantively similar results, depicted as marginal effect plots in Figure 19. The decline in support for the non-mainstream parties is strongly correlated with exposure to Covid-19, consistent with our theorized mechanism of a political flight to safety.

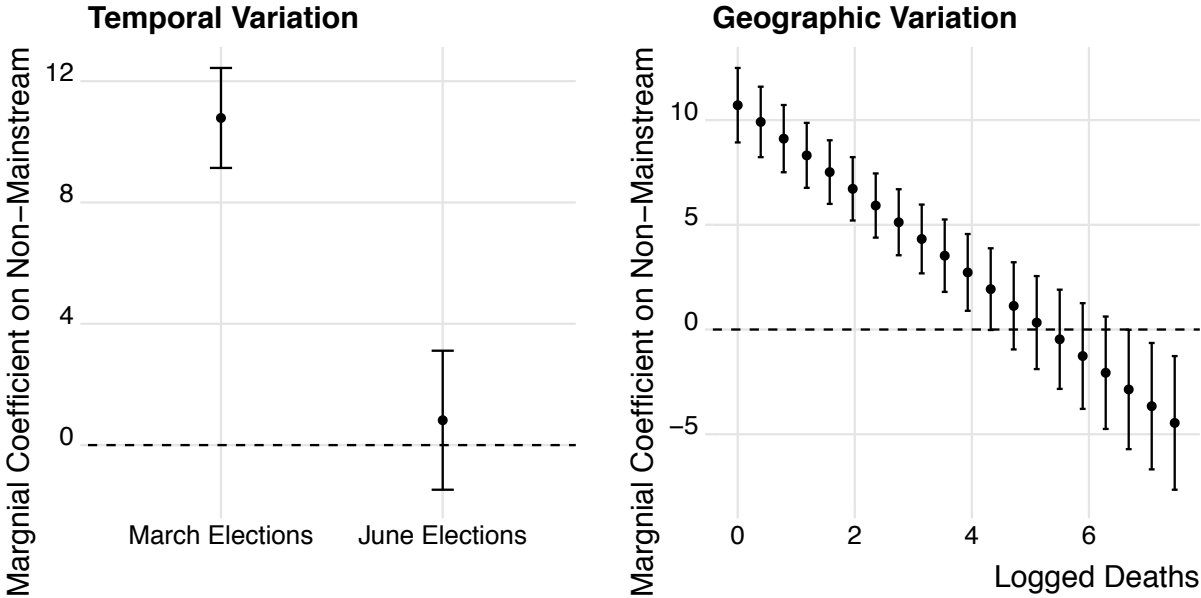


Figure 19: Diff-in-diff results for support for non-mainstream parties in the French municipal elections of 2020. Left panel summarizes Equation 2 in which the difference between mainstream and non-mainstream vote shares is compared across the March and June rounds. Right panel summarizes Equation 3 in which the difference between mainstream and non-mainstream vote shares is compared across different levels of Covid-19 exposure, measured with logged deaths.

References

- Azevedo, Flávio, John T Jost and Tobias Rothmund. 2017. ““Making America great again”: System justification in the US presidential election of 2016.” *Translational Issues in Psychological Science* 3(3):231.
- Berinsky, Adam J, Gregory A Huber and Gabriel S Lenz. 2012. “Evaluating online labor markets for experimental research: Amazon. com’s Mechanical Turk.” *Political analysis* 20(3):351–368.
- Blackwell, Matthew, Stefano Iacus, Gary King and Giuseppe Porro. 2009. “cem: Coarsened exact matching in Stata.” *The Stata Journal* 9(4):524–546.
- Clifford, Scott, Ryan M Jewell and Philip D Waggoner. 2015. “Are samples drawn from Mechanical Turk valid for research on political ideology?” *Research & Politics* 2(4):2053168015622072.
- Crump, Matthew JC, John V McDonnell and Todd M Gureckis. 2013. “Evaluating Amazon’s Mechanical Turk as a tool for experimental behavioral research.” *PloS one* 8(3).
- Cuebiq, Analysis. 2020. “Mobility Insights.”
URL: <https://www.cuebiq.com/visitation-insights-covid19/>
- Dolez, Bernard, Annie Laurent and André Blais. 2017. “Strategic voting in the second round of a two-round system: The 2014 French municipal elections.” *French Politics* 15(1):27–42.
- Hazlett, Chad and Yiqing Xu. 2018. “Trajectory balancing: A general reweighting approach to causal inference with time-series cross-sectional data.” *Available at SSRN 3214231* .
- Tausanovitch, Chris, Lynn Vavreck, Tyler Reny, Alex Rossell Hayes and Aaron Rudkin. 2019. “Democracy fund+ UCLA nationscape methodology and representativeness assessment.”
- 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.