INSTITUTE for REPLICATION

No. 36
I4R DISCUSSION PAPER SERIES

Flight to Safety: COVID-Induced Changes in the Intensity of Status Quo Preference and Voting Behavior: A Comment on Bisbee and Honig

Alice Malmberg

Daniel Scates

June 2023

INSTITUTE for REPLICATION

I4R DISCUSSION PAPER SERIES

I4R DP No. 36

Flight to Safety: COVID-Induced Changes in the Intensity of Status Quo Preference and Voting Behavior: A Comment on Bisbee and Honig

Alice Malmberg[1], Daniel Scates[1]

[1]UC Davis, CA./USA

JUNE 2023

E-Mail: joerg.peters@rwi-essen.de
RWI – Leibniz Institute for Economic Research
Hohenzollernstraße 1-3
45128 Essen/Germany
www.i4replication.org

ISSN: 2752-1931

Flight to Safety: COVID-Induced Changes in the Intensity of Status Quo Preference and Voting Behavior: A Comment on Bisbee and Honig

Bisbee and Honig (2022) *APSR*, I4R-2022-80

Alice Malmberg and Daniel Scates, UC Davis

Revised May 2023

Abstract

Bisbee and Honig (2022) examine the effect of the COVID-19 pandemic on voting for Bernie Sanders in the 2020 Democratic Party primary using a difference-in-differences design, finding evidence that exposure to COVID-19 resulted in a 7-15 percentage point increase in voting for Biden. The study also uses a regression design with district-level fixed effects to estimate the effect of the COVID-19 pandemic on voting for anti-establishment candidates during the US 2020 House primaries. It finds evidence that an increase in COVID cases was associated with a decline in voting for anti-establishment candidates in general, and for those endorsed by the Tea Party.

We re-run the code for all tests in this paper, successfully reproducing its results in a preliminary replication. We then use the De Chaisemartin and D'Haultfœuille difference-in-differences estimator to replicate their main results, finding that though the coefficient remains negative, the results are not statistically significant.

We also replicate their tests regarding US House primary candidates using a different measure of anti-establishment candidates. Here, we find that the interaction term between anti-establishment candidates and COVID-19 remain statistically significant, with the same sign. Finally, we employ an expanded dataset that includes Congressional primary candidates that were omitted in the initial dataset, as well as a re-coded extremism variable that also includes candidates endorsed by Donald Trump. These updated findings corroborate the paper's initial results. However, due to a restrictive number of observations that interfered with our application of the De Chaisemartin and D'Haultfoeuille estimator, we believe that the expanded U.S. House primary results constitute the more robust half of our replication.

1 Introduction

Recent efforts in political science seek to ensure that research is conducted transparently, and that findings are robust and replicable. As part of the Institute for Replication's efforts to test the replicability of findings in the top political science journals, we replicate Bisbee and Honig's 2022 *APSR* article, "Flight to Safety: COVID-Induced Changes in the Intensity of Status Quo Preference and Voting Behavior" (I4R-2022-80)[1].[2]

Bisbee and Honig (2022) hypothesize that anxiety produced by a crisis, in this case the COVID-19 pandemic, produces a "flight to safety" where voters gravitated towards establishment candidates. The paper tests this theory using data from the 2020 presidential and Congressional primary elections in the United States and municipal elections in France.

The paper's main analysis, which tests this theory using the 2020 Democratic presidential primary results, finds evidence that COVID provided an electoral benefit for Biden, at Bernie Sanders' expense. Bisbee and Honig (2022) also conduct a survey experiment, finding that an anxiety-inducing prompt caused respondents to gravitate towards a less disruptive candidate. Finally, the authors test their hypotheses in the context of US House primary elections, and 2020 French municipal elections, finding further evidence in support of this theory.

We first reproduce Bisbee and Honig (2022)'s findings, figures, and tables using the code provided via the Harvard Dataverse.[3] We find that all code clearly reproduces the paper's results. We then conduct two robustness replications: one on the main results of the paper using a new difference-in-difference (DiD) estimator (De Chaisemartin, D'Haultfœuille, and Guyonvarch 2019) and another that stratifies the analysis of US House primaries by partisanship. Finally, we conduct a direct replication by expanding upon the original data set of US House candidates.

Our use of De Chaisemartin and D'Haultfœuille's DiD estimator to test the robustness of the paper's main findings did not produce statistically significant results, though the sign of the coefficient for COVID exposure remains the same as in the paper's original model. However, due to a restrictive number of observations that interfered with our application of the De Chaisemartin and D'Haultfœuille estimator, we believe that the expanded U.S. House primary results constitute the more robust half of our replication. To this end, our expansion of the US House primary candidate section produced results comparable to the author's original findings that similarly support their theory. Consequently, we find that Bisbee and Honig (2022)'s findings are replicable overall and robust

[1] Code and data for this replication and extension is available at https://github.com/Dmscates/Bisbee-and-Honig-2022-Flight-to-Safety-Replication.

[2] Acknowledgements: We would like to thank Dr. Lauren Young for her encouragement and guidance with this project. We would also like to thank the authors of the original piece, Drs. James Bisbee and Dan Honig, for their time and assistance in understanding their data set and code. We would also like to thank Drs. Chris Hare and Ben Highton for their advice related to this undertaking.

[3] https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/S5YMS7

to our changes.

2 2020 Democratic Presidential Primary Outcomes

2.1 Replication of Presidential Primary Outcomes

Bisbee and Honig (2022) use a difference-in-difference design, in which they compared the difference between the vote share for Sanders in exposed and insulated counties before and after the outbreak of COVID-19. The original paper finds evidence that counties that were more exposed to COVID were, *a priori*, more likely to vote for Sanders than Biden in the 2020 Democratic Party primary. They also find evidence that, in counties that had COVID cases as of March 17th were less supportive of Sanders than similar counties that voted prior to their eventual exposure. This provides evidence in support of their hypothesis that anxiety over COVID depressed vote share for Sanders.

In their second analysis, Bisbee and Honig (2022) use a two-way fixed effects (TWFE) model. This model incorporates fixed effects for the date and the designated market area (DMA) of a county. The DMAs in this model correspond to the media markets where counties are located. Because campaigns purchase television advertisement space on a per-media market basis, this affects how much information voters in a given DMA were exposed to; in theory, those residing within the same media market receive the same information. Since DMAs cut across state lines, Bisbee and Honig (2022) are able to compare counties that received similar information, but had different presidential primary election dates.

Bisbee and Honig (2022)'s presidential primary outcomes replicate perfectly using the R code provided on the Harvard Dataverse repository. Below is a recreation of the paper's Table 1, showing the effects of COVID-19 exposure on primary vote share for Sanders, as well as of their Figure 4, showing the coefficient plots for their DiD analysis.

Table 1: Sanders Two-Way Vote Share on Exposure

	Dependent variable: Sanders 2020 two-way					
	Full Sample			March and April 7th		
	Basic	Match	Weight	Basic	Match	Weight
	(1)	(2)	(3)	(4)	(5)	(6)
Exposure Dummy	−1.321*	−1.066*	−1.348***	−0.885*	−0.574**	−0.885**
	(0.552)	(0.478)	(0.380)	(0.382)	(0.204)	(0.325)
	[0.0179]	[0.0273]	[0.001]	[0.022]	[0.006]	[0.007]
Turnout 2020	0.175***	0.189***	0.137***	0.197***	0.219***	0.169***
	(0.026)	(0.052)	(0.038)	(0.026)	(0.052)	(0.041)
	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
Sanders 2016	0.027***	0.031***	0.032***	0.026***	0.033***	0.034***
	(0.004)	(0.007)	(0.006)	(0.004)	(0.006)	(0.006)
	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
Caucus switch	0.628	0.434	0.874	0.205	0.735	0.764
	(0.959)	(0.394)	(0.557)	(0.771)	(0.587)	(0.465)
	[0.513]	[0.273]	[0.119]	[0.790]	[0.213]	[0.102]
Caucus dummy	1.227***	0.023	0.764†			
	(0.221)	(0.433)	(0.417)			
	[0.000]	[0.958]	[0.069]			
Observations	1,882	666	1,882	1,710	666	1,710
R^2	0.830	0.895	0.883	0.860	0.903	0.898

Note: *p<0.1; **p<0.05; ***p<0.01
Standard errors in parentheses; P-values in brackets

2.2 Robustness Checks of Presidential Primary Outcomes

We find that the study's research design provides a convincing method of estimating the effect of exposure to COVID-19 at the county level. However, there are two notes we make that inspire our robustness check. First, much of the paper's theory is based on the DMA level. Indeed, the measure for whether a county was exposed to the pandemic was whether there was a confirmed COVID case in the DMA at a given time.

Second, recent work highlights a potential flaw with the TWFE DiD model, in which results may be biased if negative weights arise (De Chaisemartin and D'Haultfœuille 2020). This occurs in cases where the average treatment effects are heterogeneous across time or across groups.

To summarize De Chaisemartin and D'Haultfœuille's argument, let us assume that all observations in the same group and time cell *(g, t)* have the same treatment, and that treatment is binary. For our purposes, this assumption holds because all observations in the same DMA are assigned the same value of COVID-19 cases.

Here, let us consider the regression of Y_{igt}, the outcome of unit i in group g at period t with group and time fixed effects, and a treatment of D_{it}. Let $\hat{\beta}_{fe}$ represent the coefficient of D_{gt}, and β_{fe} denote its expectation. Under the parallel trends assumption, De Chaisemartin and D'Haultfœuille (2020, 2964-2965) shows that $\hat{\beta}_{fe} = E(\sum_{gt,D_{gt}=1} W_{gt}\Delta_{gt})$, where Δ_{gt} is the is the average treatment effect (ATE) for group g at time t, and W_{gt} are the model's weights, which can be any real number, but must sum to 1. If treatment effects are heterogeneous, it is possible for $\hat{\beta}_{fe}$ to be negative, despite each observation's ATE being positive. This happens as a result of the weights being negative.

In order to address this issue, we use De Chaisemartin, D'Haultfœuille, and Guyonvarch (2019)'s DiD estimator, which is robust to heterogeneous treatment effects across time and groups. This estimator compares the outcome changes of switchers – groups that flip from being untreated to being treated across two different dates – to those of groups that remain untreated across both dates. Standard errors are then calculated using bootstrapping methods.

We also carry out our analysis at the DMA level, rather than the county level. This has some significant drawbacks, but lines up closely with the idea that media coverage of new cases primarily drove COVID-19-related anxiety early in the pandemic (Bisbee and Honig 2022). We use De Chaisemartin and D'Haultfœuille's *did_multiplegt* Stata package[4] to estimate the effect of COVID-19 exposure on the change in vote share for Sanders in 2020 between DMAs that spanned across state borders and thus had Democratic primaries that took place across multiple dates. We control for the vote share that Sanders received in the 2016 Democratic presidential primaries.

Table 2 and Figure 1 on the following page display the results of this test. Table 2 shows that though the coefficient is in the correct direction to substantiate Bisbee and Honig's results, the standard error is too large for the results

[4]https://ideas.repec.org/c/boc/bocode/s458643.html

to be statistically significant. This helps us glean some information about the underlying data; of 40 DMAs in the data set, we can see that just 18 had confirmed COVID cases as of the date of the second Democratic primary held in that media market.

Relatedly, use of the new DiD approach has some significant drawbacks, but lines up closely with the idea that media coverage of new cases primarily drove COVID-19-related anxiety early in the pandemic (Bisbee and Honig 2022). The biggest issue with this decision is that there are actually only 40 unique DMAs in the data set. This small N leads to large standard errors, and means that the method we use cannot control adequately for more than one variable. This is because the bootstrapping method used results in more controls than observations in many cases.

Table 2: Effect of COVID Exposure on Sanders Vote Share

	Estimate
Exposure	-.0698
	(0.603)
	[0.454]
Observations	40
Switchers	18

Standard errors in parentheses; P-values in brackets

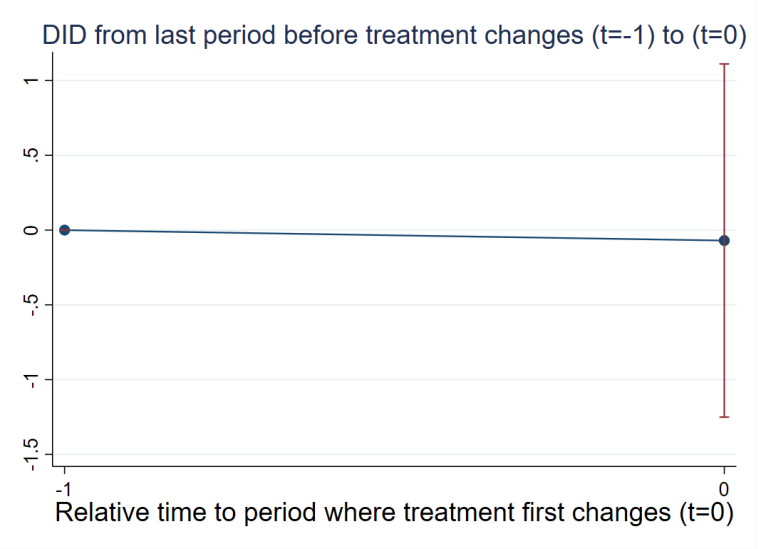

Figure 1: Coefficient Plot Showing Effect of COVID Exposure on Primary Vote

3 U.S. House Primary Data

3.1 Replication of U.S. House Primary Data

Bisbee and Honig (2022)'s U.S. House Primary Data replicates perfectly using the provided R code from the Harvard Dataverse. Below is a recreation of the article's Table 3: Antiestablishmment Vote Share as a Function of COVID-19 Exposure:

Table 3: Anti-establishment Vote Share as a Function of COVID-19 Exposure

	Dependent variable: Vote Share		
	Extreme (Any)	Justice Dems	Tea Party
Anti-Est.	0.342^{***}	0.134^{**}	0.364^{***}
	(0.043)	(0.050)	(0.043)
	$[2.1e^{-14}]$	[0.0082]	$[7.09e^{-16}]$
Cases (ln)	−0.044	−0.014	−0.031
	(0.150)	(0.151)	(0.151)
	[0.769]	[0.925]	[0.835]
Anti-Est. X Cases (ln)	-0.016^{**}	0.014	-0.020^{**}
	(0.006)	(0.012)	(0.006)
	[0.0085]	[0.256]	[0.0015]
District FE	Y	Y	Y
Observations	2,019	2,019	2,019
Anti-Est. Candidates	70	8	62
R^2	0.516	0.506	0.515

Note: †p<0.1; *p<0.05; **p<0.01; ***p<0.001
Standard errors in parentheses; P-values in brackets

3.2 Robustness Check of U.S. House Primary Data

As a robustness check on the authors' original findings we subset the original data by party and run two additional models: each estimates primary vote share for Justice Democrats and Tea Party endorsees as a function of logged COVID-19 cases given the total sample of Democratic and Republican Congressional candidates, respectively.

In doing so, we endeavor to check that the paper's aforementioned findings from Table 2 above hold up when stratifying the total sample of 2020 Congressional primary candidates by partisanship. Results of these models are in Table 4 on the following page.

Table 4: Anti-establishment Vote Share as a Function of COVID-19 Exposure Stratified by Candidate Party

	Dependent variable: Vote Share	
	Dem Subset	GOP Subset
Anti-Est.	0.387**	0.496***
	(0.194)	(0.052)
	[0.047]	[$< 2e^{-16}$]
Cases (ln)	0.047	−0.071
	(0.195)	(0.231)
	[0.809]	[0.758]
Anti-Est. X Cases (ln)	0.005	−0.019**
	(0.025)	(0.007)
	[0.847]	[0.0047]
District FE	Y	Y
Observations	921	1,033
R^2	0.696	0.799

Note: *p<0.1; **p<0.05; ***p<0.01
Standard errors in parentheses; P-values in brackets

We compare the coefficients of the "Justice Dems" and "Tea Party" columns of the original results in Table 2 to their respective partisan subsets in Table 4. Stratifying these data by partisanship retains the the significance of the anti-establishment coefficient for both parties as well as that of the interaction between anti-establishment and COVID cases within the GOP data set, though there is a modest increase in the magnitude of both the Democratic and and GOP anti-establishment coefficients. Interestingly, there is a sign change for the COVID case coefficient for Democrats; when subsetting by partisanship, the directionality of this variable changes from negative to positive. However, as in the original model, this coefficient also fails to achieve statistical significance.

Overall, we find that estimating these models after stratifying the sample of 2020 Congressional primary candidates by partisanship does not substantively change the paper's findings; thus, this robustness check substantiates its conclusions.

4 Expansion of U.S. House Primary Data

4.1 Imputing Missing Observations

In order to expand the U.S. House Primary Data section of the article, we focused on providing a more thorough data set of the candidates coded as extremists.

To do so, we first cross-referenced those candidates listed as anti establishment left-wing and right-wing candidates in 2020 with official endorsement lists accessed from the Tea Party Patriots[5] and Justice Democrats[6] official websites and the post history for the Justice Democrats official Facebook page.[7] We identified eight additional Democratic candidates who received Justice Democrats endorsements[8] and two additional Republican candidates who received Tea Party Network endorsements.[9]

We also find that many House candidates competing in primaries in Louisiana and North Carolina are missing from the original data set. To remedy this, we use Ballotpedia listings to compile the names of candidates who competed in these states' primaries and data from the *New York Times* COVID Data Tracker Github repository[10] to round out the original data set to include observations from these states. Doing so adds 33 additional candidates from Louisiana and 52 from North Carolina, increasing the total number of observations from 2,019 to 2,105.

Table 5 shows the original and updated candidate counts:

Table 5: Updated Observations

	Original	Updated Original
Extreme (Any)	70	80
Justice Dems	8	16
Tea Party	62	64
Total N	2,019	2,105

Table 6 on the following page re-estimates the paper's original model from Table 3 to include the expanded sample of candidates receiving Justice Democrats and Tea Party endorsements, as well as the additional candidates identified from Louisiana and North Carolina. Doing so demonstrates that the coefficients for the the Tea Party and bipartisan extremist models remain largely consistent overall, though there is a modest decrease in the magnitudes of both variables

[5]https://teapartypatriots.ning.com/
[6]https://justicedemocrats.com/
[7]https://www.facebook.com/justicedemocrats
[8]Cori Bush (MO-1), Marie Newman (IL-3), Andom Ghebreghiorgis (NY-16), Rashida Tlaib (MI-13), Morgan Harper (OH-3), Jessica Cisneros (TX-28), Alexandria Ocasio-Cortez (NY-14), and Ayanna Pressley (MA-7)
[9]Clay Higgins (LA-3) and Richard Hudson (NC-8)
[10]https://github.com/nytimes/covid-19-data

compared to the initial model. Interestingly, though the magnitude of the Justice Democrats variable is comparable to that in the initial model (0.132 versus 0.134), the significance in the new model goes away entirely.

The most interesting finding from the re-estimated model is that the sign of the COVID case variable flips from negative to positive across all three samples of extremist candidates. However, for both the bipartisan extremist and Tea Party models, this variable nevertheless still fails to attain significance at the 0.05 level. Though the COVID case variable is significant in the Justice Democrats model, both the anti-establishment candidate variable and the interaction between the anti-establishment and COVID case variables are not statistically significant. Therefore, re-estimating Bisbee and Honig's initial models after expanding the data to include the missing Justice Democrat and Tea Party-endorsed primary candidates and the observations from North Carolina and Louisiana also does not substantively change the paper's findings.

Table 6: Re-Estimated Model of Anti-establishment Vote Share as a Function of COVID-19 Exposure with Missing Candidates Included

	Dependent variable: Vote Share		
	Extreme (Any)	Justice Dems	Tea Party
Anti-Est.	0.336^{***}	0.132	0.334^{***}
	(0.045)	(0.117)	(0.046)
	[$5.34e^{-13}$]	[0.263]	[$2.91e^{-12}$]
Cases (ln)	0.018^{*}	0.029^{***}	0.019^{**}
	(0.010)	(0.009)	(0.009)
	[0.067]	[0.0014]	[0.046]
Anti-Est. X Cases (ln)	-0.014^{**}	0.003	-0.015^{**}
	(0.006)	(0.014)	(0.007)
	[0.028]	[0.814]	[0.0296]
District FE	Y	Y	Y
Observations	2,089	2,089	2,089
R^2	0.514	0.503	0.512

Note: *p<0.1; **p<0.05; ***p<0.01
Standard errors in parentheses; P-values in brackets

4.2 Adding a New Measure of Right-Wing Extremism

Additionally, we theorize that in 2020, an endorsement by then-president Donald Trump may have served as a more salient indicator of right wing extremism than affiliation with the Tea Party Movement. Indeed, many political

pundits have observed that the Tea Party eventually morphed into support for Trump in the 2016 and 2020 elections (Pew 2019; Elving 2022). Consequently, we expand the data set of anti-establishment rightist candidates to include those who received President Trump's endorsement in 2020, hypothesizing that COVID outbreaks should have a comparable effect on support for Trump-endorsed Congressional primary candidates to that for Tea Party endorsees.

Finally, to reflect the fact that we are using two different sources of rightest candidates as we are still including Tea Party endorsees, we create two additional new categories: Tea and Trump encapsulates those were endorsed by both the Tea Party and Trump in 2020 and Tea or Trump reflects those who earned either right wing endorsement.

Table 7 builds off of the previous Table 5 to show the newly expanded counts of extremist candidates:

Table 7: Updated and Expanded Observations

	Original	Updated Original	Expanded
Extreme (Any)	70	80	183
Justice Dems	8	16	16
Tea Party	62	64	64
Trump	NA	NA	143
Tea and Trump	NA	NA	40
Tea or Trump	NA	NA	167
Total N	2,019	2,105	2,105

After compiling these updated data, we utilize Bisbee and Honig's original House Primary code to analyze these new observations, regressing the total number of extremist candidates in each category on the logged total number of COVID cases per Congressional district at the time of each state's 2020 House primaries. Also in line with the paper's original methodology, we interact these variables. Table 8 on the following page displays the results of this analysis.

The coefficients of the original three models of extremist candidates (overall extremism, Justice Democrats, and Tea Party), remain consistent with those in the earlier updated models shown in Table 6. There is also a significance change in the interaction of the COVID case count and vote share variables; for the overall extremism and Tea Party models, this interaction now achieves significance at the 0.05 level.

We now turn to the new models (Trump endorsees, Tea Party or Trump, and Tea Party and Trump). For the bipartisan extremist, Tea Party, and Tea or Trump models, the coefficient of the interaction term between anti-establishment candidates and COVID case count is negative and statistically significant. The direction and magnitude of these findings are in line with those in the original Extremist and Tea Party models in Table 3, as well as the re-estimated versions of both models in Table 6. It is important, however, to note

that the magnitude of all coefficients is extremely modest.

Both the bipartisan extremist model as well as two out of the three right-wing extremist models including Tea Party endorsed candidates (the base Tea Party model and the Tea Party or Trump model) substantiate Bisbee and Honig's original findings Congressional districts with higher COVID case numbers had lower primary vote shares for extremist right wing primary candidates.

Table 8: Expansion of Anti-establishment Vote Share as a Function of COVID-19 Exposure

	Dependent variable: Vote Share					
	Extreme	Justice Dems	Tea	Trump	Tea and Trump	Tea or Trump
Anti-Est.	0.336^{***}	0.132	0.334^{***}	0.321^{***}	0.323^{***}	0.326^{***}
	(0.045)	(0.117)	(0.046)	(0.037)	(0.055)	(0.033)
	$[5.34e^{-13}]$	[0.263]	$[2.91e^{-12}]$	$[< 2e^{-16}]$	$[9.61e^{-9}]$	$[< 2e^{-16}]$
Cases (ln)	0.018^{*}	0.029^{**}	0.019^{**}	0.008	0.019^{**}	0.009
	(0.010)	(0.009)	(0.009)	(0.008)	(0.009)	(0.008)
	[0.067]	[0.0014]	[0.046]	[0.311]	[0.046]	[0.299]
Anti-Est. X Cases (ln)	-0.014^{**}	0.003	-0.015^{**}	-0.009^{*}	-0.014^{*}	-0.010^{**}
	(0.006)	(0.014)	(0.007)	(0.005)	(0.008)	(0.005)
	[0.028]	[0.814]	[0.029]	[0.068]	[0.072]	[0.030]
District FE	Y	Y	Y	Y		
Observations	2,089	2,089	2,089	2,089	2,089	2,089
R^2	0.514	0.503	0.512	0.527	0.508	0.531

Note: *p<0.1; **p<0.05; ***p<0.01
Standard errors in parentheses; P-values in brackets

Of the interactions in the added models of Trump endorsees and the model with candidates endorsed by both right-wing groups (Tea and Trump), both coefficients are in the same direction and similar magnitude to the Extremist and Tea Party models; however, neither interaction attains significance at the 0.05 level. Similarly, as in the original paper, the interaction term for the updated Justice Democrats model is positive and not statistically significant. Consequently, there does not seem to be any clear, significant relationship for Trump-endorsed right-wing candidates, or for Justice Democrats-endorsed left-wing candidates.

5 Conclusion

Our replication and extension of Bisbee and Honig (2022) produce mixed results; though ones that overall support the paper's original conclusions. Using an application of De Chaisemartin and D'Haultfœuille's new DiD method, we first test whether the findings that rapid spread of the COVID-19 pandemic caused 2020 Democratic presidential primary voters to gravitate toward establishment candidates, thereby decreasing vote share for Sanders.

Using the updated method, we find that though the effect of Covid continues to have a negative effect, the results are no longer statistically significant. That said, this method suffers from a low number of DMAs, with the small N resulting from our aggregation likely driving this result.

We also replicated the paper's test of whether a similar desire for candidate moderation was evident in 2020 US House primaries, extending it to include candidates that were omitted from the original analysis and adding several new variables to capture support for Trump-endorsed Congressional candidates. The original paper's results largely remained robust after incorporating these changes. Overall, the original findings of Bisbee and Honig (2022) largely remained robust to our replications and extensions.

6 References

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.

De Chaisemartin, Clément, and Xavier D'Haultfoeuille (2022). "Two-Way Fixed Effects and Differences-in-Differences with Heterogeneous Treatment Effects: A Survey." *National Bureau of Economic Research.* https://www.nber.org/papers/w29691 (April 23, 2023).

De Chaisemartin, Clément, and Xavier D'Haultfœuille (2020). "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects." *American Economic Review.* 110(9): 2964–96.

De Chaisemartin, Clément, Xavier D'Haultfoeuille, and Yannick Guyonvarch (2019). "DIDMULTIPLEGT: Stata module to estimate sharp Difference-in-Difference designs with multiple groups and periods," Statistical Software Components S458643, Boston College Department of Economics, revised 04 May 2023

Elving, Ron (2022). "Trump's MAGA is marching down a trail blazed by the Tea Party." *NPR.* https://www.npr.org/2022/05/21/1100386445/trumps-maga-is-marching-down-a-trail-blazed-by-the-tea-party.

"Trump's Staunch GOP Supporters Have Roots in the Tea Party." *Pew* (2019). https://www.pewresearch.org/politics/2019/05/16/trumps-staunch-gop-supporters-have-roots-in-the-tea-party/.