7 Supplementary Information

We present the descriptive statistics of our data in Table 3.

Table 3: Descriptive Statistics

Statistic	N	Mean	St. Dev.	Min	Max
NOMINATE	283	0.016	0.454	−0.747	0.883
Movement	886	−0.018	0.152	−0.830	0.670
Position before	886	−0.055	0.321	−0.852	0.784
Position after	886	−0.073	0.307	−0.844	1.434
Uncontested	886	0.296	0.457	0	1

In Table 4 we present the results of a ten-fold cross-validation of our machine learning approach. Cross-validation is a resampling procedure used to evaluate machine learning models on a limited data sample. Cross-validation estimates the skill of a machine learning model on unseen data. That is, using a limited sample to estimate how the model is expected to perform in general when used to make predictions on data not used during the training of the model.

Table 4: 10-Fold Cross-Validation

	Accuracy	Precision	Recall	F1
1	0.926	0.935	0.891	0.913
2	0.933	0.956	0.902	0.928
3	0.926	0.933	0.906	0.920
4	0.902	0.897	0.903	0.900
5	0.906	0.920	0.881	0.900
6	0.889	0.906	0.863	0.884
7	0.949	0.950	0.943	0.947
8	0.919	0.895	0.944	0.919
9	0.906	0.924	0.901	0.912
10	0.923	0.938	0.906	0.922

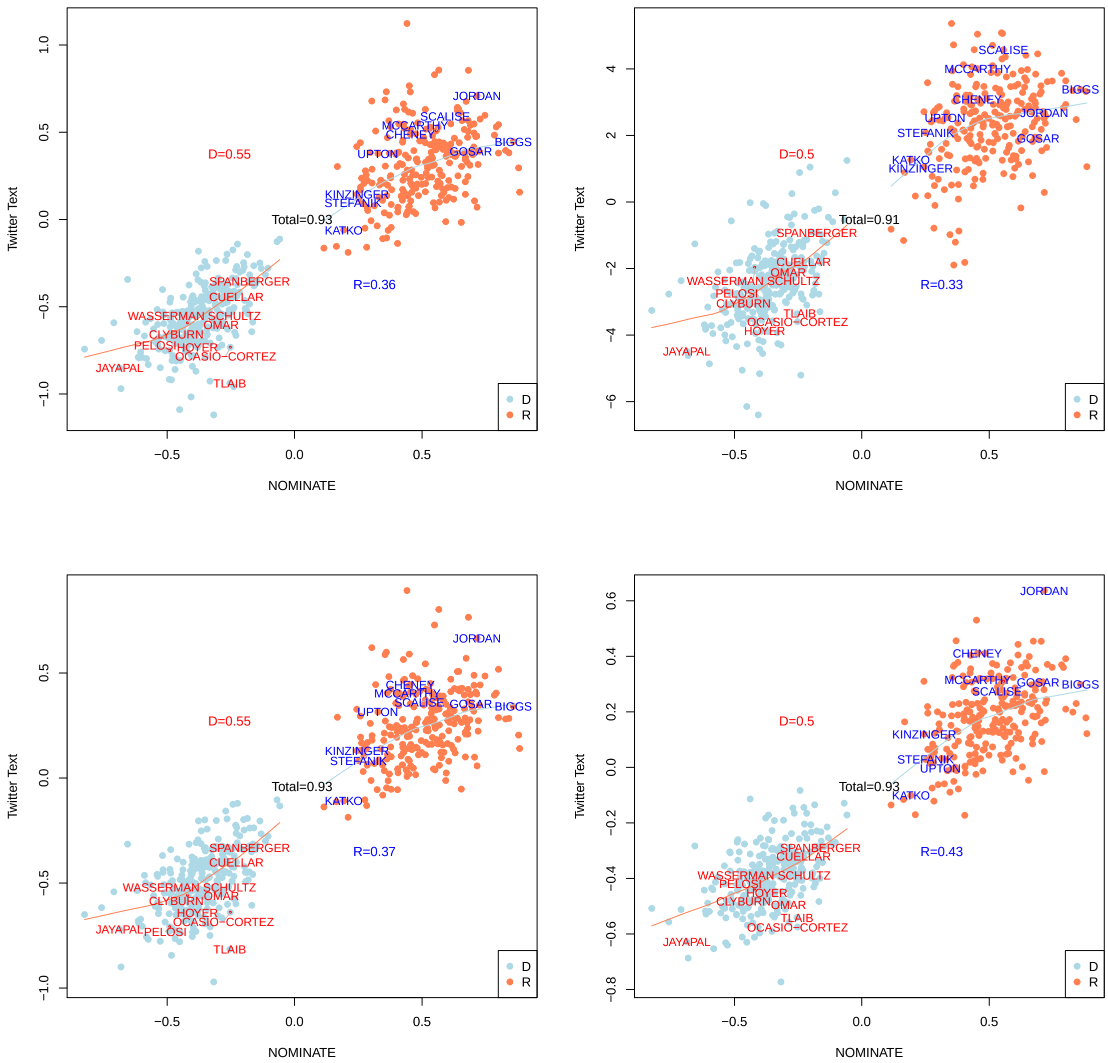

Figure 5: Comparative Validity of Alternative Measures

Positional scaling often depends on the exact choice of specification. We chose to remove all Twitter-specific references, hashtags, and @-mentions from the data. Figure 5 shows the correlation with NOMINATE for all terms (top-left), only @-mentions (top-right), hashtags but no @-mentions (bottom-left), and only plain text (bottom-right). We use only plain terms in our main analysis (see Figure 1) as they are most balanced between Republicans and Democrats in terms of intra-party correlations and have the most semantic validation in terms of the positions of individual representatives.

We recognize that the unusual political climate in the summer of 2020 may impact the generalizability of our findings. In Figure 6 we plot the main figure using the true calendar date rather than the 'time-to-primary' variable we use elsewhere. This figure shows that the murder of George Floyd (25th May

2020) and the subsequent national protests, which were at their height between 26th May and 9th June, do not appear to have impacted the positioning of candidates in either party in real time. Given that ten of the forty-nine states' primaries took place prior to 25th May and we see no difference in the behavior of candidates in these contests compared to the twelve states which had their primaries shortly after this date, or compared to the twenty-seven states who held their primaries later in the summer, we are confident that our findings are not impacted by these events.

Figure 6 does indicate an influence of the outbreak of the COVID-19 pandemic on Republican positioning, where both winners and losers took more 'moderate' positions in March 2020. This moderation was driven by an increased focus on healthcare policy, a domain traditionally considered a Democratic issue. After a short period of speaking about this issue in a way similar to Democrats, Republicans quickly found their own language to talk about COVID and our models place them in similar positions as in February 2020. Given our non-finding for the Republican Party and the fact that this movement occurs prior to most (though not all) primary elections, we believe it does not adversely affect our findings, though it may contribute to the wider Republican confidence intervals in our main analysis.

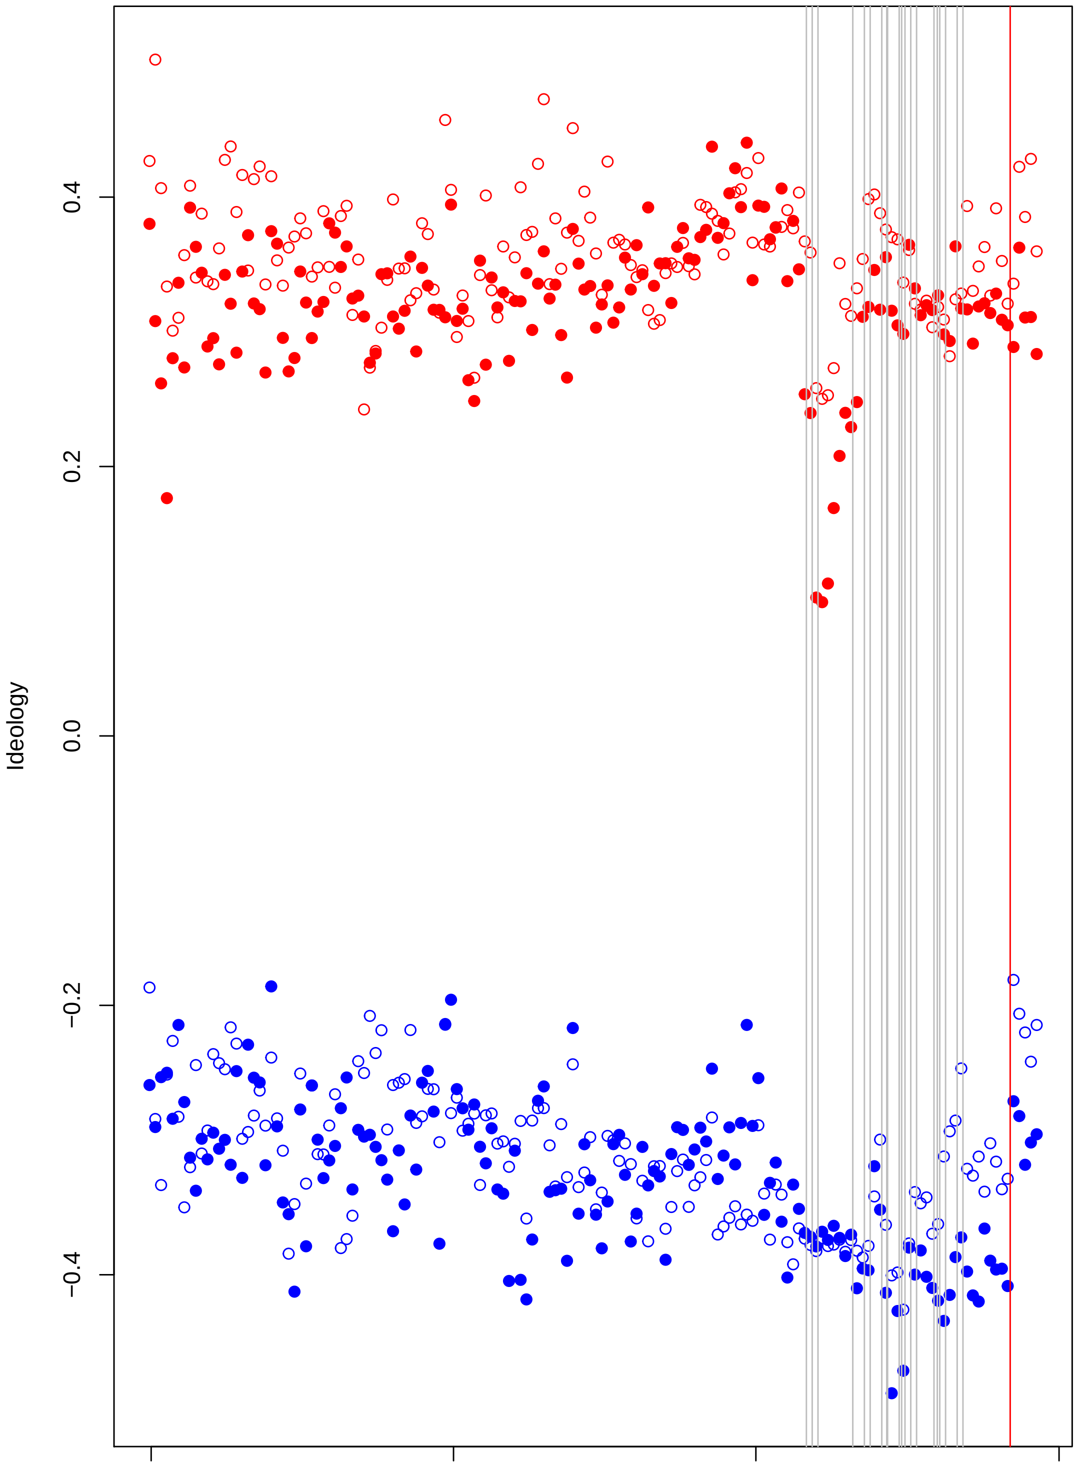

Figure 6: Natural Time

As an additional check of our analysis, we run a placebo test. Figure 7 shows the positions of winning and losing primary candidates by party over time. Instead of using the real time-to-primary variable, we randomized the primary date for each individual from all real primary dates and aggregated the positions over week to these fictitious primaries. If there was a confounder correlated with the primary date, it would still systematically affect the dependent variable over time rather than at the date of the primary. Since we only randomize across nineteen weeks in total (weeks that had primary elections), there is still a relevant time trend in the data.

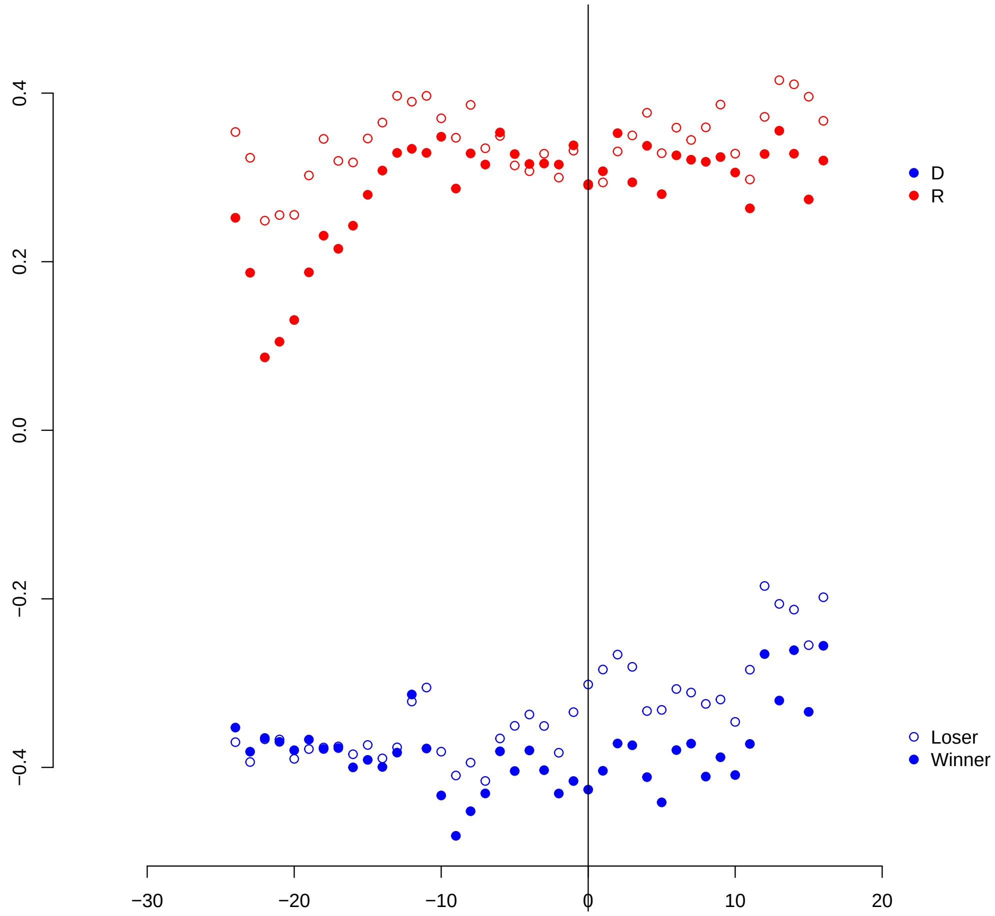

Figure 7: Placebo Test

As the plot demonstrates, the main effect for Democratic candidates is no longer present and only emerges once all primaries have concluded. We do not see this for the Republicans. In other words, if we set new primary dates for each Republican candidate, we would still observe the same overall behavior. This is **not** the case for the Democrats.

We also run our ITS model for Democrats using a randomized primary date at the candidate level. We present our results in Table 5. As expected, when we randomize the primary date there is no significant effect (ZX_t).

Table 5: ITS Results: Placebo Dates

	Position Democrats
Time to Pseudo Primary (T_t)	-0.008^{***}
	(0.001)
After Pseudo Primary (X_t)	0.012
	(0.022)
Loser (Z_i)	0.028
	(0.021)
TTP:After Pseudo Primary ($X_t T_t$)	0.012^{***}
	(0.002)
After Pseudo Primary: Loser ($Z_i T_t$)	-0.027
	(0.031)
TTP: Loser	0.004^{**}
	(0.002)
TTP:After Pseudo Primary: Loser ($Z_i X_t T_t$)	-0.002
	(0.002)
Intercept	-0.401^{***}
	(0.015)
N	98
R^2	0.582
Adjusted R^2	0.550
Residual Std. Error	0.038 (df = 90)
F Statistic	17.926^{***} (df = 7; 90)

$^{*}p < .1$; $^{**}p < .05$; $^{***}p < .01$

So what drives the result in Figure 7 for Republicans? Time series analysis of political positions has numerous challenges, the most severe of which is the effect of changing saliency that might introduce exogenous shocks into the data. Because many candidates use Twitter to respond to events and current developments, convergence may result from the whole 'system' (all candidates) moving and tweeting about the same issues. As we measure the relative emphasis of specific terms, systemic movement can be problematic, with issues varying in prevalence over time. As an example, healthcare is more commonly emphasized by Democratic candidates, but, as discussed above, the COVID-19 pandemic also led to Republicans emphasizing this traditionally 'Democratic' issue.

To tackle this problem we 'detrend' the data, using canonical correspondence analysis to control for time effects. The common use of correspondence or factor analysis is to extract values for the main dimension, controlling for additional variables and implicitly computing positions of third variables extracted from word weights. By using time as an explanatory variable, we only observe differences in emphasis. If the saliency of an issue rises collectively, we put less weight on it. This process of 'detrending' provides more consistent positions and removes time trends from the data, where the model subtracts the time-based component from the word weight (Greenacre, 2007). Figure 8 compares approaches, where the upper plot shows the Naïve Bayes approach used in our main analysis, and the lower uses the Canonical Correspondence Analysis discussed here. These effects are substantively the same, with the additional caveat that the detrending produces stronger time effects for the Republicans.

In combination with the placebo test, we conclude that the COVID-19 pandemic affected the political positions of the Republicans, as healthcare, typically a Democratic issue, made the agenda. Before Republicans formulated their own framing, they used similar language to Democrats. This effect leads to a strong time-based overlay in the data that cannot be eliminated at this point, but which requires additional data from future elections.

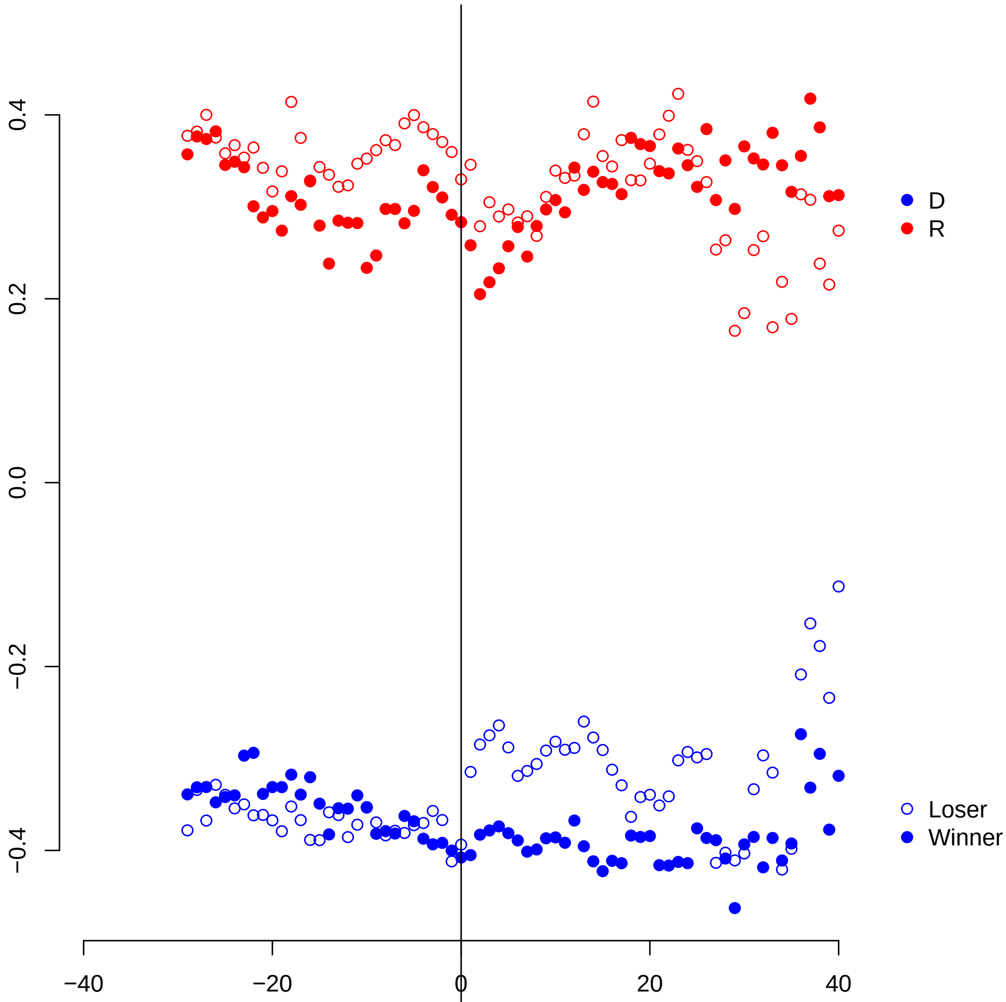

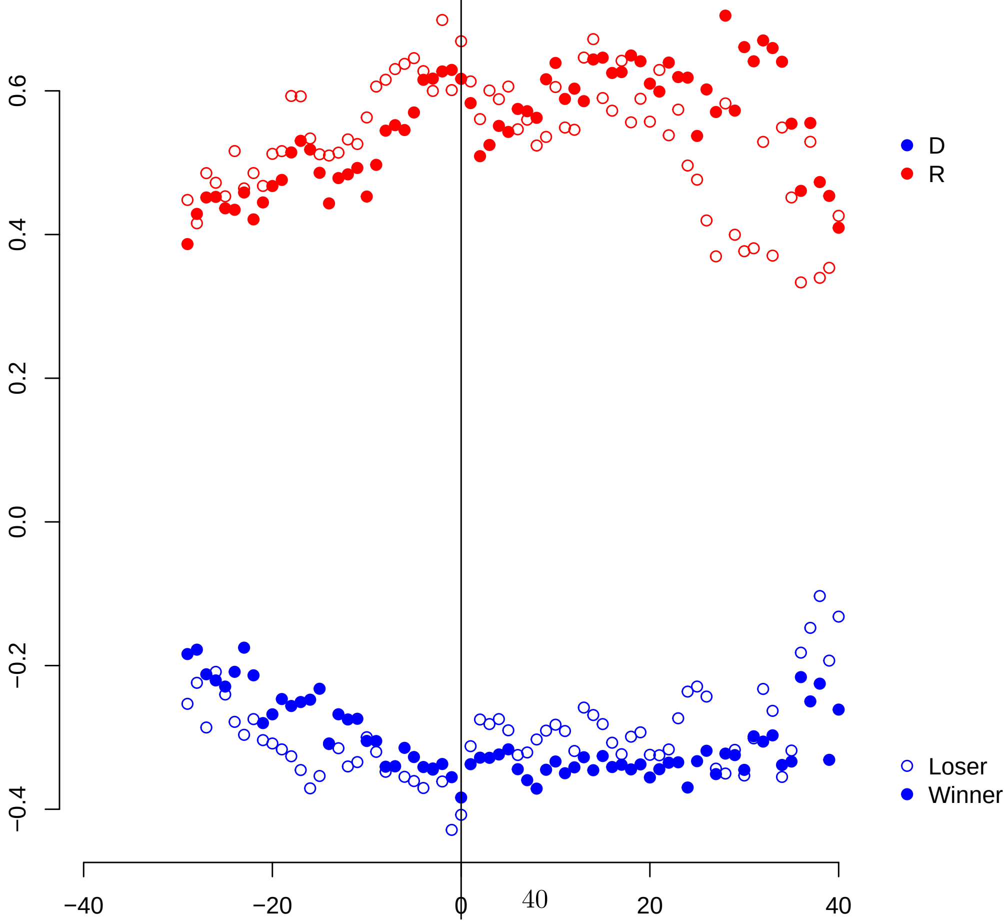

Figure 8: Naïve Bayes Approach & Canonical Correspondence Analysis Comparison

We further demonstrate the robustness to over-time trends by showing that our approach is not affected by the choice of which terms to include in the analysis in Figure 9. When we include all terms (first plot), hashtags (third plot), and hashtags and @-mentions (fourth plot), our results remain present. Only if we restrict our data only to @-mentions (second plot) is our effect no longer present.

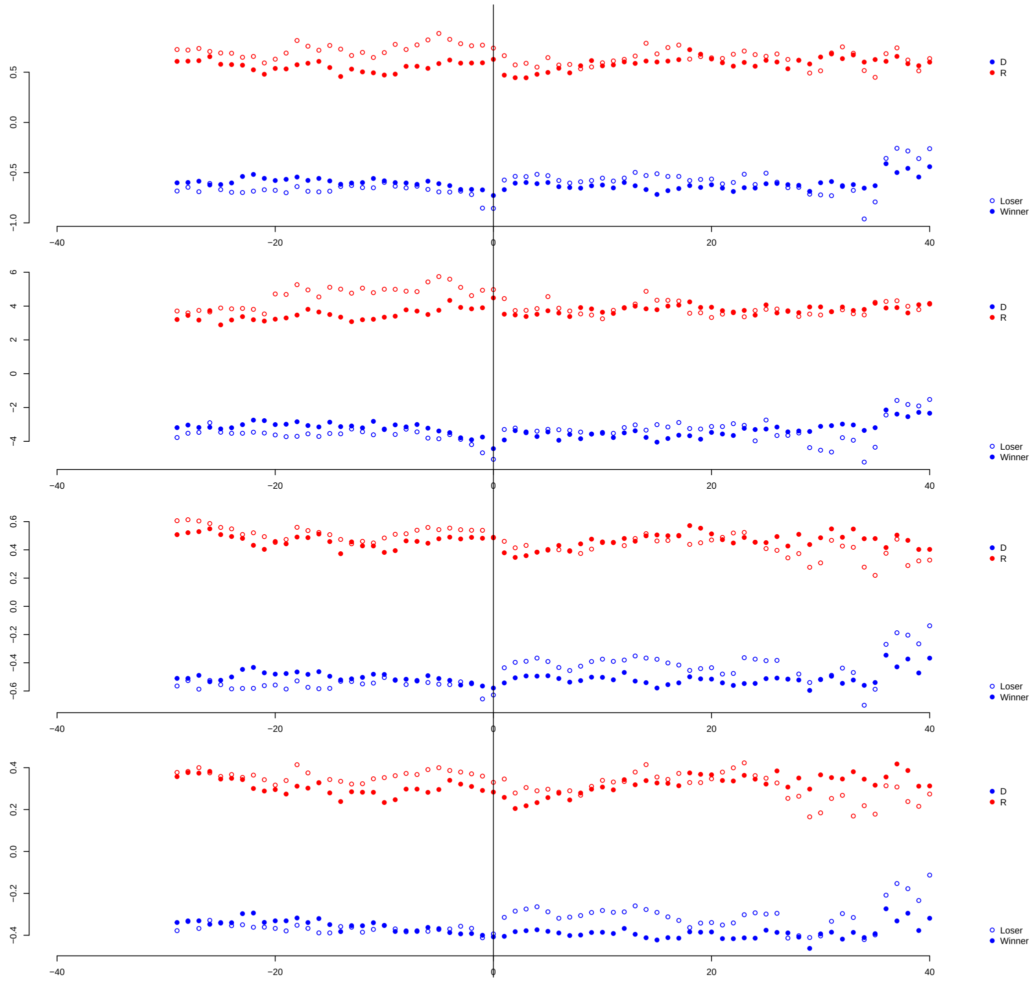

Figure 9: Main Analysis with Different Terms

In Table 6 we repeat our individual-level analysis with the removal of candidates in the eight districts that saw same-party (all Democrat vs Democrat) general elections as a result of California or Washington's top-two primary system. Those districts were CA-12, CA-18, CA-29, CA-34, CA-38, CA-44, CA-53, and WA-10. Our results are substantively unchanged with the removal of these districts. In Table 7 we demonstrate the robustness of our main individual results to three standard errors of movement. In Table 8 we demonstrate the robustness of our individual results with the removal of the additional controls in our main analysis.

Table 6: Original Analysis and Removal of Same Party

	Movement Right	
	Original Analysis	Without Same Party
	(1)	(2)
Loser	0.057***	0.056***
	(0.014)	(0.014)
Republican	−0.079***	−0.079***
	(0.014)	(0.015)
Republican Loser	−0.044**	−0.041**
	(0.020)	(0.020)
Constant	0.038***	0.039***
	(0.010)	(0.010)
N	886	871
R^2	0.052	0.052
Adjusted R^2	0.049	0.049
Residual Std. Error	0.149 (df = 882)	0.149 (df = 867)
F Statistic	16.117*** (df = 3; 882)	15.845*** (df = 3; 867)

*p < .1; **p < .05; ***p < .01

Table 7: Individual Level Results (as Coefficient Plot)

	Movement Right Democrats Absolut	Movement 3 Errors Democrats 3 Errors	Movement Right Republicans Absolut	Movement 3 Errors Republicans 3 Errors
	(1)	(2)	(3)	(4)
Intercept	0.039^{***}	0.307^{***}	0.021	0.151^{*}
	(0.011)	(0.080)	(0.020)	(0.084)
Loser	-0.001	-0.004	-0.008^{***}	-0.026^{**}
	(0.001)	(0.009)	(0.002)	(0.010)
Lean Conservative	-0.042^{***}	-0.272^{***}	0.013	0.072
	(0.012)	(0.088)	(0.023)	(0.098)
Incumbent	0.042^{***}	0.281^{***}	-0.032^{**}	-0.068
	(0.007)	(0.052)	(0.013)	(0.057)
N	472	472	414	414
R^2	0.094	0.092	0.033	0.021
Adjusted R^2	0.088	0.086	0.026	0.014
Residual Std. Error	0.102 (df = 468)	0.750 (df = 468)	0.185 (df = 410)	0.790 (df = 410)
F Statistic	16.223^{***} (df = 3; 468)	15.740^{***} (df = 3; 468)	4.639^{***} (df = 3; 410)	2.928^{**} (df = 3; 410)

$^{*}p < .1$; $^{**}p < .05$; $^{***}p < .01$

Table 8: Individual Robustness without Controls

	Movement Right Absolut	Movement 3 Errors 3 Errors	Movement Right Absolut + Controls	Movement 3 Errors 3 Errors + Controls
	(1)	(2)	(3)	(4)
Loser	0.057***	0.424***	0.043***	0.358***
	(0.014)	(0.072)	(0.015)	(0.076)
Party Republican	−0.079***	−0.349***	−0.073***	−0.332***
	(0.014)	(0.076)	(0.015)	(0.076)
PVI			−0.003***	−0.008
			(0.001)	(0.006)
Incumbent			−0.029**	−0.143**
			(0.012)	(0.062)
Loser*Party Republican	−0.044**	−0.304***	−0.037*	−0.281***
	(0.020)	(0.104)	(0.020)	(0.105)
Constant	0.038***	0.257***	0.040***	0.268***
	(0.010)	(0.053)	(0.010)	(0.053)
N	886	886	886	886
R^2	0.052	0.053	0.066	0.060
Adjusted R^2	0.049	0.050	0.061	0.055
Residual Std. Error	0.149 (df = 882)	0.774 (df = 882)	0.148 (df = 880)	0.772 (df = 880)
F Statistic	16.117*** (df = 3; 882)	16.366*** (df = 3; 882)	12.448*** (df = 5; 880)	11.231*** (df = 5; 880)

*p < .1; **p < .05; ***p < .01

As an additional check on our approach of running our analysis on the subset of policy-related tweets, we also run a separate analysis on the entire corpus with a control for policy-related tweets. We present our results in Table 9. As with our other robustness checks, our main finding that Democratic losers moderate remains substantively significant.

Table 9: ITS Results: Policy Tweets Control

	Position	
	Democrats	Republicans
	(1)	(2)
Time to Primary (T_t)	-0.003^{***}	-0.001
	(0.001)	(0.001)
After Primary (X_t)	0.020^{*}	-0.040^{**}
	(0.011)	(0.015)
Loser (Z_i)	0.023^{***}	0.057^{***}
	(0.009)	(0.018)
No. Policy Tweets	-0.00001^{*}	-0.00001
	(0.00000)	(0.00001)
TTP:After Primary ($X_t T_t$)	0.001	0.006^{***}
	(0.001)	(0.002)
After Primary: Loser ($Z_i T_t$)	0.054^{***}	-0.019
	(0.020)	(0.025)
TTP: Loser	0.002^{***}	0.001
	(0.001)	(0.001)
ttp:After Primary: Loser ($Z_i X_t T_t$)	-0.003^{***}	-0.003
	(0.001)	(0.002)
Intercept	-0.313^{***}	0.240^{***}
	(0.026)	(0.038)
N	102	102
R^2	0.848	0.606
Adjusted R^2	0.835	0.573
Residual Std. Error (df = 93)	0.016	0.025
F Statistic (df = 8; 93)	65.023^{***}	17.912^{***}

$^{*}p < .1$; $^{**}p < .05$; $^{***}p < .01$

As a further robustness check, we also validate our measure against Hopkins and Noel (2021)'s pairwise activist scores for senators in Figure 10. As noted in the main text, we do not train our model on senators' tweets, making these tweets an excellent independent corpus against which to validate our approach. We also validate our model against Barberá's 2015 Follower Network in Figure 11.

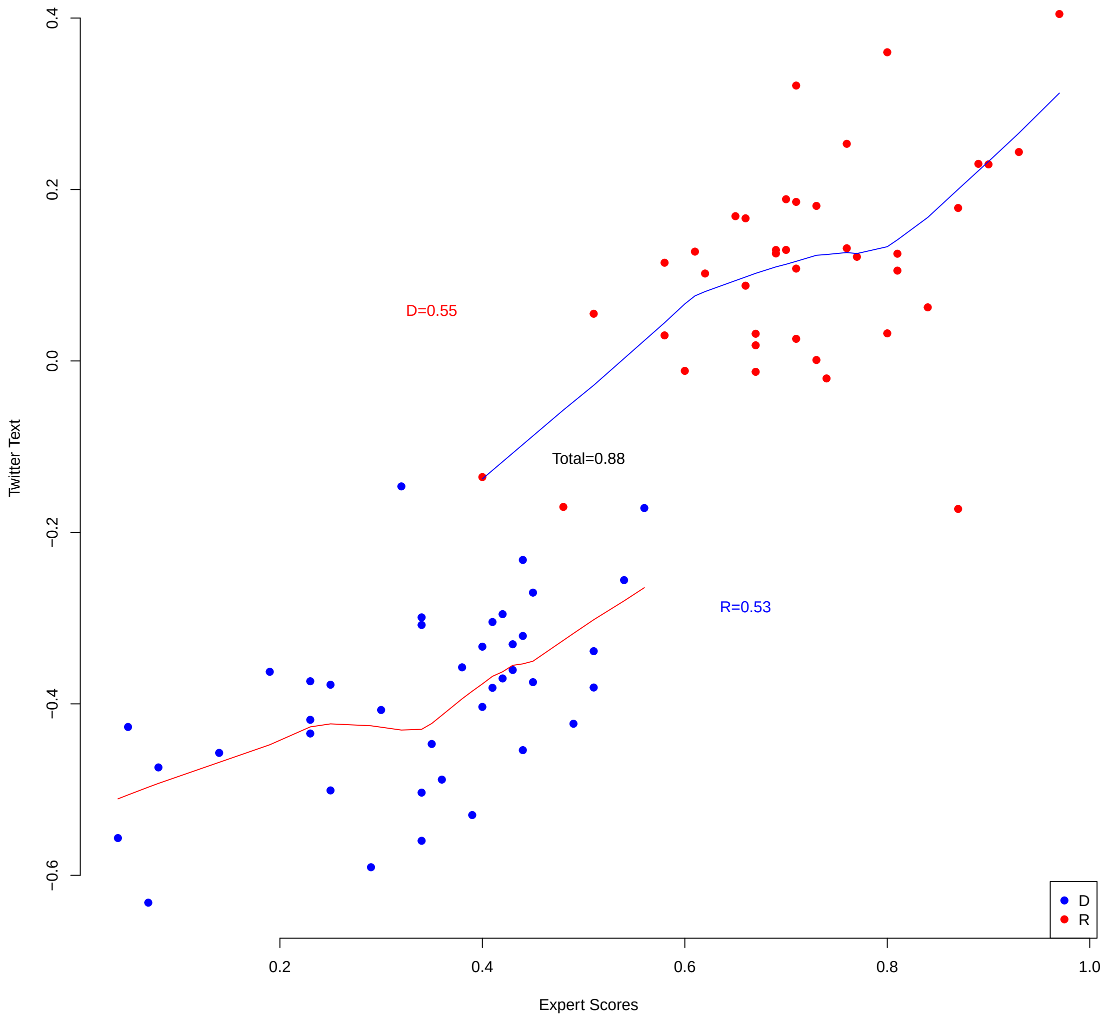

Figure 10: Validation Against Hopkins & Noel Pairwise Activist Scores

Table 10: Dickey-Fuller Tests for Unit Root in Dependent Variable (All Tweets)

Republican_Loser	Republican_Winner	Democratic_Loser	Democratic_Winner
-6.291072	-4.530455	-3.727189	-5.429676
5.26e-07	.0013454	.0206281	.0000317
-4.014667	-4.014667	-4.014667	-4.014667
-3.439867	-3.439867	-3.439867	-3.439867
-3.139867	-3.139867	-3.139867	-3.139867

Table 11: Dickey-Fuller Tests for Unit Root in Dependent Variable (Policy Only)

Republican_Loser	Republican_Winner	Democratic_Loser	Democratic_Winner
-6.952026	-5.063808	-4.841324	-5.957281
1.81e-08	.0001572	.0003962	2.70e-06
-4.014667	-4.014667	-4.014667	-4.014667
-3.439867	-3.439867	-3.439867	-3.439867
-3.139867	-3.139867	-3.139867	-3.139867

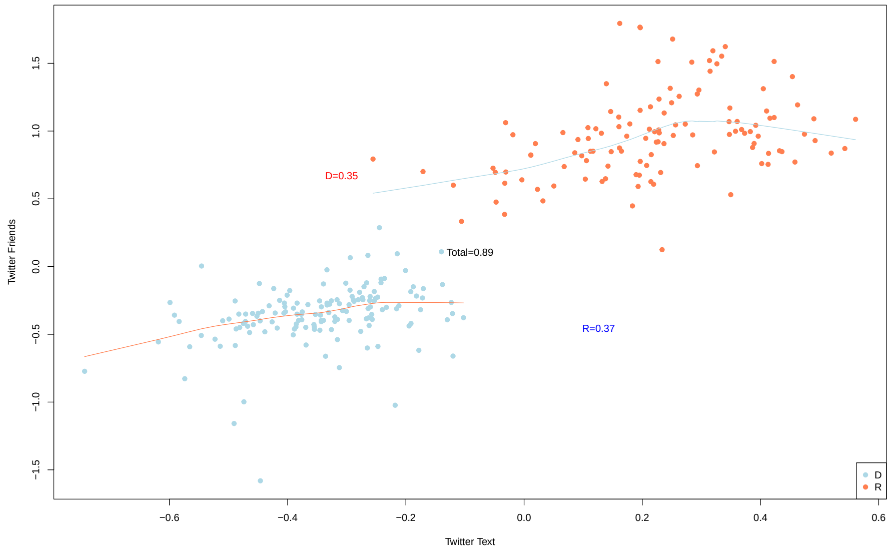

Figure 11: Validation Against Barberá's Follower Network

In Table 10 (all tweets) and Table 11 (policy only) we report the results of our Dickey-Fuller tests for each of our four dependent variables in our main analysis. In addition, we test that our estimated errors (residuals) are white noise by running Dickey-Fuller tests on the residuals, presenting the results in Table 12 (all tweets) and Table 13 (policy only). Each Dickey-Fuller test tests the null hypothesis that a unit root is present, meaning stationarity is the alternative hypothesis. In all four cases, our p-values are below 0.001, indicating stationarity in these variables.

Table 12: Dickey-Fuller Tests for Unit Root in Residuals (All Tweets)

Republican_Loser	Republican_Winner	Democratic_Loser	Democratic_Winner
-7.006754	-4.939502	-5.135386	-6.409418
1.36e-08	.0002648	.0001158	2.91e-07
-4.014667	-4.014667	-4.014667	-4.014667
-3.439867	-3.439867	-3.439867	-3.439867
-3.139867	-3.139867	-3.139867	-3.139867

Table 13: Dickey-Fuller Tests for Unit Root in Residuals (Policy Only)

Republican_Loser	Republican_Winner	Democratic_Loser	Democratic_Winner
-7.823222	-5.253872	-5.545008	-6.414199
1.83e-10	.0000692	.0000188	2.84e-07
-4.014667	-4.014667	-4.014667	-4.014667
-3.439867	-3.439867	-3.439867	-3.439867
-3.139867	-3.139867	-3.139867	-3.139867

To test the robustness of our main finding to an alternative specification, we also run a two-way fixed effects regression for each party with the results shown in Table 14. Here, we use a two-way fixed effects regression as it allows us to adjust for unobserved unit-specific and time-specific confounders at the same time. In this model we instead treat time as a dichotomous indicator with the value one after the primary; the 'intervention'. Given that we expect moderation from losing candidates, our panel variable takes the value one for those candidates who do not win and zero for winning candidates. We are unable to demonstrate the necessary assumptions for a difference-in-differences (DiD) design with our data, most obviously the likely violation of stable unit treatment value assumption (SUTVA) given the clear differences between many winning and losing candidates. As a result, we are unable to say that the presence of the primary is what caused candidates to adopt artificial positions, though a clear trend of post-primary moderation among Democratic losers is observed. This estimator takes the following specification:

$$Y_{it} = \alpha_i + \gamma_t + \beta_1 X_{it} + \varepsilon_{it}$$

Where Y_{it} is candidate position Y given membership of group i measured at time t. α_i is the difference between primary winners and losers over the entire period. γ_t is the difference between winning candidates' positions before and after the primary election. $\beta_1 X_{it}$ is our main object of interest and is the interaction term between time and losing, and ε_{it} is the error term.

We present our findings in Table 14, with positive coefficients indicating rightward positioning and negative coefficients indicating leftward movement. In line with the visual trend depicted in Figure

3, Table 14 shows that Democratic losers became significantly more moderate immediately after they lose the primary election (Loser : Post-Primary). Though Democratic losers are no different from winners during the primary campaign (Loser), once the primary finishes these candidates move rightward. Interestingly, winning Democratic candidates do not moderate after the primary, and are, on average somewhat further to the left than during the primary campaign (Post-Primary).

Among Republicans, Table 14 indicates no significant moderation following primary defeats (Loser : Post-Primary). On average, Republican winners are slightly further to the right than winners across the entire time period (Loser), with no movement among winners following a primary (Post-Primary). As in our main analysis, partisanship is the strongest predictor of position among candidates in both parties. As in the model included in our main analysis, we find a clear moderating effect among losing Democratic candidates *only*.

Table 14: Two-Way-Fixed-Effect Version

	Position	
	Democrats	Republicans
	(1)	(2)
Lagged Position	-0.004	0.046^{***}
	(0.006)	(0.009)
Time to Primary (T_t)	-0.039^{***}	0.009
	(0.006)	(0.009)
After Primary (X_t)	0.098^{***}	-0.018
	(0.009)	(0.013)
Loser (Z_i)	-0.318^{***}	0.236^{***}
	(0.004)	(0.007)
N	102	102
R^2	0.703	0.252
Adjusted R^2	0.694	0.229
Residual Std. Error (df = 98)	0.022	0.034
F Statistic (df = 3; 98)	77.432^{***}	10.985^{***}

$^{*}p < .1$; $^{**}p < .05$; $^{***}p < .01$

Though our observations of candidate positions are not linearly related to their positions in other time periods, we note that there is an extensive literature indicating that voters reward positional consistency among candidates. Accordingly, we demonstrate the robustness of our main finding to the inclusion of a lagged version of the dependent variable. We present the results in Table 15.

Table 15: Lagged DV As Additional Control

	Position		Policy Only Position	
	Democrats	Republicans	Democrats	Republicans
	(1)	(2)	(3)	(4)
Lagged Position	0.420***	0.472***		
	(0.080)	(0.093)		
Lagged Position Policy Only			0.324***	0.456***
			(0.089)	(0.092)
Time to Primary (T_t)	−0.002***	−0.0004	−0.002***	0.001
	(0.0005)	(0.001)	(0.0005)	(0.001)
After Primary (X_t)	0.009	−0.024*	0.015	−0.017
	(0.008)	(0.013)	(0.009)	(0.018)
Loser (Z_i)	0.019**	0.030**	−0.001	0.011
	(0.008)	(0.014)	(0.009)	(0.017)
TTP:After Primary ($X_t T_t$)	0.002***	0.003***	0.001	0.001
	(0.001)	(0.001)	(0.001)	(0.001)
After Primary: Loser ($Z_i T_t$)	0.049***	−0.001	0.029**	0.008
	(0.013)	(0.018)	(0.013)	(0.025)
TTP: Loser	0.002***	0.0004	0.002**	−0.0004
	(0.001)	(0.001)	(0.001)	(0.001)
ttp:After Primary: Loser ($Z_i X_t T_t$)	−0.003***	−0.001	−0.002**	−0.001
	(0.001)	(0.001)	(0.001)	(0.002)
Intercept	−0.213***	0.118***	−0.291***	0.086***
	(0.029)	(0.022)	(0.038)	(0.018)
N	100	100	100	100
R^2	0.884	0.696	0.673	0.401
Adjusted R^2	0.874	0.670	0.644	0.348
Residual Std. Error (df = 91)	0.014	0.022	0.017	0.031
F Statistic (df = 8; 91)	86.786***	26.070***	23.427***	7.606***

*p < .1; **p < .05; ***p < .01