Partisan Communication in Two-Stage Elections:

The Effect of Primaries on Intra-Campaign Positional Shifts in Congressional Elections

Mike Cowburn
European University Viadrina
cowburn@europa-uni.de

Marius Sältzer
Universität Oldenburg
marius.saeltzer@uni-oldenburg.de

The influence of congressional primary elections on candidate positioning remains disputed and poorly understood. We test whether candidates communicate artificially 'extreme' positions during the nomination, as revealed by moderation following a primary defeat. We apply a scaling method based on candidates' language on Twitter to estimate positions of 988 candidates in contested U.S. House of Representatives primaries in 2020 over time, demonstrating validity against NOMINATE (r > 0.93) where possible. Losing Democratic candidates moderated significantly after their primary defeat, indicating strategic position-taking for perceived electoral benefit, where the nomination contest induced artificially 'extreme' communication. We find no such effect among Republicans. These findings have implications for candidate strategy in two-stage elections and provide further evidence of elite partisan asymmetry.

Keywords: primary elections, partisan asymmetry, Congress, text-as-data, polarization

21st June 2023

1 Introduction

To become a member of Congress, most candidates must win two elections, with distinct incentives, actors, and electorates in each. Though positional differences between parties' primary and general electorates appear minimal (Abramowitz, 2008; Hirano and Snyder, 2019; Sides et al., 2020), policy demanders active in the party network play an important role during the nomination (Bawn et al., 2012; Cohen et al., 2008; Masket, 2009) and have distinct and 'extreme'[1] preferences (Hill and Huber, 2017; Kujala, 2019; Saunders and Abramowitz, 2004). Candidates must therefore appeal to non-centrist groups in the party network to become the nominee (Fiorina, Abrams and Pope, 2005) before attempting to garner wider support among a general electorate who prefer moderate candidates (Ansolabehere, Snyder and Stewart, 2001) and punish extremism (Canes-Wrone, Brady and Cogan, 2002). Accordingly, candidates are presented with a *strategic positioning dilemma* (Brady, Han and Pope, 2007) across the electoral cycle: which constituency should they appeal to?

Some research suggests that candidates move away from the center in primaries (Brady, Han and Pope, 2007; Burden, 2001), but a systematic study of candidate positions across a primary and general election cycle remains lacking, in part due to the limited availability of positional time series data of elected officials and losing candidates. Traditional ideal point estimates are only available for elected members of Congress (McCarty, Poole and Rosenthal, 2006) or aggregated across an entire election cycle (Bonica, 2014). To fill this gap, we measure changes in candidate positions both during and after the primary using an original dataset of dynamic social media-based positions. We use supervised machine learning (Goet, 2019; Green et al., 2020) to identify the liberal–conservative axis of 2,500,000 tweets by 988 candidates running for the U.S. House of Representatives in 2020. We validate our measure using NOMINATE scores of candidates in the sample who had ever served in Congress, our scores correlate at 0.93.

We use this measure to test candidate responses to the strategic positioning dilemma over the electoral cycle. Importantly for our design, our method enables us to continue positioning

[1] We use the term 'extreme' here in line with the established use in the primary election literature (e.g., Hall, 2015). 'Extremism' may result from positions far from the 'center', greater consistency, or some combination of these.

candidates after they lose a primary. Given that voters punish inconsistency (Canes-Wrone, Brady and Cogan, 2002), we expect that primary winners will maintain positions taken during the primary to prevent accusations of 'flip-flopping'. We argue instead that positional adaptation will only be observed among primary *losers* after their defeats, and use this movement to identify whether candidates took artificial positions during the nomination, comparing their communication during the primary campaign with their positions after they lose. In doing so, we test the adaptative rather than the selective effect of the nomination process—our interest is in the change in candidate behavior rather than election outcomes—and hypothesize that losing candidates will moderate after a primary defeat. In this paper we focus solely on the candidate side of the dilemma, we are explicitly not capturing voter responses to or reception of candidate positioning.

Among Democratic candidates, losing a primary was clearly associated with moderation following a defeat, suggesting the adoption of artificial or strategic positions during the nomination. This finding aligns with other scholarship about candidate behavior in two-stage elections (Brady, Han and Pope, 2007; Burden, 2001) and similar research on rhetorical position-shifting by presidential primary winners (Acree et al., 2020). We find no equivalent shift in the position of losing Republican candidates, indicating limited strategic position-taking and continued support for 'conservative' sentiment even when electoral incentives were absent. The party-level differences are likely explained by the asymmetric nature of the Republican and Democratic parties (Grossmann and Hopkins, 2016; Hacker and Pierson, 2006; Theriault, 2013). Our findings are significant at both the party and candidate levels, and when we restrict our analyses to tweets that explicitly contain policy content.

We proceed as follows: First, we review the literature on strategic positioning in campaign communication. Second, we consider the ability of existing measures to fully answer our question, introducing our scaling technique based on Twitter text. Next, we present our data and findings. Finally, we discuss explanations and implications of our results at both the party and candidate levels.

2 Candidate Incentives in Primaries

Before candidates can compete in a general election, they must first earn the party's nomination. To win the nomination, candidates must appease various party stakeholders or "policy demanders" (Bawn et al., 2012). Both theoretical expectations (May, 1973) and empirical evidence (Abramowitz, 2010; Converse, 1964) indicate that these groups—by virtue of being highly engaged and politically active—hold positions away from the center and prioritize candidates' positional congruence in their selection criteria.

Primary voters do not appear to share the distinct preferences of these policy demanders, with empirical studies of both presidential (Abramowitz, 2008; Norrander, 1989) and congressional primary electorates (Hirano and Snyder, 2019; Sides et al., 2020) finding little or no positional differences between primary and general election party voters. Despite these findings, primary electorates are frequently characterized as extreme by scholars (Burden, 2001; Fiorina, Abrams and Pope, 2005; Kamarck, 2014) and politicians (Schumer, 2014; Keisling, 2010) alike. Here, the *perceptions* of political actors are of particular importance given our focus on candidate behavior, where candidates might adopt artificial positions because they believe that primary voters hold non-centrist preferences with which they try to align. DeCrescenzo (2020) finds that elites behave as if primary voters want ideological candidates, despite limited evidence that these voters express any such preference.

Yet, winning a primary is not only dependent on positional alignment with voters. In presidential contests, Cohen et al. (2008) document the influence of party elites during the nomination. At the congressional level, Hassell (2018) similarly finds that actors in the party network play a key role in candidate selection. The UCLA school of parties (especially Bawn et al., 2012) highlights the importance of "policy demanders"—including donors, activists, interest groups, and even friendly partisan media—in determining candidate selection outcomes. In part because U.S. nominations are comparatively inclusive and decentralized (Hazan and Rahat, 2010; Cowburn and Kerr, 2023), formal party organizations have been "hollowed out" (Schlozman and Rosenfeld, 2019), transferring power from electability-focused formal structures toward comparatively non-centrist and policy-oriented "informal party organizations" (Masket, 2009). Alignment with these groups can help candidates secure the nomination in several ways.

Fundraising is a key indicator of a primary campaign's viability. Donors—and large donors in particular—hold more extreme and consistent positions than primary voters (Kujala, 2019), with distinct preferences and policy positions from non-donors (Gilens, 2009). In short, "Democratic contributors are more liberal than other Democrats and Republican contributors are more conservative than other Republicans" (Hill and Huber, 2017, 10) and donate to proximate candidates (Bonica, 2014). Consequently, non-centrist position-taking aligns with an increased ability to raise funds in both primary and general elections (Ensley, 2009).

Activists form an integral part of a wider network (Bawn et al., 2012) and are a vital resource during the nomination process (Masket, 2009) constituting primary campaigns on the ground. Like donors, these partisans are further from the political center than primary electorates (Hill and Huber, 2017; Saunders and Abramowitz, 2004). Interest groups can play a similar role, with evidence that candidates with interest group support have had increased success in congressional nominations in recent years (Manento, 2019). Both activists and interest groups hold distinct positions on the issues they care about and seek assurances that candidates are positionally aligned during the nomination. Providing assurances to multiple groups can pull candidates away from the center in a process of "conflict extension" (Layman et al., 2010), with evidence that primary candidates who receive more interest group support take positions further from the center (La Raja and Schaffner, 2015; Manento, 2019). The proliferation of partisan media may have further elevated ideological candidates through favorable coverage to an audience of party sympathizers (Heft et al., 2021).

Taken together, these factors help explain why candidates further from the center appear to be preferred even when primary electorates are moderate (Chen and Yang, 2002; Cooper and Munger, 2000). Consequently, there may be considerable benefit to candidates who can communicate non-centrist positions during the nomination.

2.1 Communication and Positional Change

Legislators signal preferences through roll-call voting (Canes-Wrone, Brady and Cogan, 2002) and other candidates need to make alternative credible claims of positions, such as by differentiating themselves through their policies, behavior, or language. Intra-party positioning may include drawing support from aligned allies, attacking a primary opponent on ideological grounds, or associating with an ideological faction (Blum, 2020). These types of differentiation are difficult to change during an election cycle. Perceptions of candidates' positions may also be based on information obtained prior to the election, giving campaigns limited ability to shift over time. Candidates may also perceive strategic disadvantages of moving positions, such as being labeled as inconsistent or of 'flip-flopping', which voters are liable to punish (DeBacker, 2008). Under the assumptions of the strategic positioning dilemma, we expect candidates to adopt non-median positions during the primary, with limited moderation of nominees in general election campaigns due to the electoral penalties attached to moving position. Because we do not expect primary winners to adapt their positions, we focus on losing candidates' positional adaptation after primary defeats to empirically identify artificial positioning during the primary.

Political communication—including press statements, interviews, and social media activity—allows more flexibility, enabling candidates not only to alter their policy positions but also to change emphasis (Meyer and Wagner, 2019). Candidates can reposition not only by changing their stances on issues but also by changing the issues that they talk about (Budge and Farlie, 1983). Candidates who present themselves away from the center in their policy positions are also non-centrist in their communication, demonstrated here by the alignment of positions derived from voting behavior and social media communication for candidates in our data who ever served in Congress.

Most losing candidates in our sample did not run for alternative public office following their defeat. Though most—not all—remained active partisans, relatively few faced continued deliberation or public votes on their positions. Some candidates ran for or continued to hold local public office, but the vast majority did not. We consider losers' social media communication after the primary as the best available approximation of 'sincere' preferences. We recognize that even this communication does not take place in a vacuum, as unsuccessful candidates likely wish to remain in good standing with their party, either to run for public office again or to hold an appointed position. Yet, social media posts likely play a minimal role in fulfilling these goals, and, though we acknowledge that candidates will not want to communicate anything that causes reputational damage, they are likely less strategic than contributions in party meetings or other formal venues. We also recognize that the dominant linguistic frames used by party leaders and other elites likely influence candidate communication but minimize the extent of such effects by comparing candidates' positions against themselves across a relatively short period. Empirically, we also expect that these strategic considerations likely decrease rather than accentuate positional movement compared to (unobservable) communication absent *any* external incentives.

3 Measuring Elite Positions

To determine whether candidates communicate artificial positions in primaries, we require positions *over time*. Common measures of positional estimation based on roll-call votes (Poole and Rosenthal, 1985) or campaign donations (Bonica, 2014) are either not available for all candidates or fail to provide the required temporal granularity. We therefore use an alternative measure placing candidates and officeholders on the same dimension by scaling social media communication. Social media allow political elites to communicate directly with potential voters in public. Twitter in particular has developed into an important campaign tool for parties and politicians that has gained substantial scholarly attention (Barberá et al., 2019; Cowburn and Oswald, 2020; Cowburn and Knüpfer, 2023; Russell, 2018). Tweets have become part of the news cycle and Twitter is now a rich source of information about the thematic emphases of politicians and their positions. In line with established literature on the subject (see e.g., Boireau, 2014.; Ceron, 2016; Sältzer, 2020), we analyze Twitter text to position candidates over time. Unsupervised text classification methods include Wordfish, which enables comparisons of election manifestos (Slapin and Proksch, 2008) and political speeches (Lauderdale and Herzog, 2016). One challenge of these approaches is a lack of agreement that the extracted dimensions relate to political ideology. Supervised text analysis ensures a correct understanding of the underlying dimension but requires 'training data' to teach algorithms which text aligns with different positions. Since ideology is continuous rather than categorical, methods such as Wordscores (Laver, Benoit and Garry, 2003) use scaling, but set fixed endpoints using anchor documents. Similar approaches have also been applied to newspapers (Gentzkow and Shapiro, 2010) and television channels (Martin and McCrain, 2019). To identify the dimension of partisan conflict, Goet (2019) and Green et al. (2020) use supervised learning on party labels to identify positions. We follow this approach here.

3.1 Data

We collected the timelines of social media accounts of candidates running as a Republican or Democrat in a contested primary for the U.S. House of Representatives in 2020. In line with the established literature (Boatright, 2013, 2014), we consider primaries as contested when two same-party candidates feature on a ballot. Twitter accounts were collected based on a search list created by sourcing ballotpedia.com. We restricted our sample to candidates in contested primaries with identifiable Twitter accounts who tweeted regularly enough for us to position them both before and after their primary election date. We include positional data from 988 of the total of 1,772 candidates that stood in a contested primary as a Democrat or Republican for the U.S. House of Representatives in the 2020 election cycle. Our sample is heavily skewed towards candidates with a realistic chance of winning the nomination, where a large proportion of excluded candidates did not raise money or actively campaign and received single-digit vote shares. Unsurprisingly, higher-performing candidates were more likely to have an active social media presence.[2] Our data include candidates from forty-nine states, as Louisiana does not hold congressional primaries.[3]

Accounts were cross-referenced with manually-collected candidate data (Cowburn, 2022), compiled throughout the 2020 primary cycle using certified data from state's websites. Tweets were collected using the Twitter API implementation rtweet (Kearney, 2018) for all candidates with Twitter accounts in June 2020. Having gathered the list of accounts in June, we constructed our

[2] Other studies of congressional primaries restrict inclusion based on vote share thresholds (Boatright, 2013, 2014) or advocate for financial measures (Thomsen, 2021). Restricting based on social media presence is analogous and excludes many of the same long-shot candidates.

[3] Given only eight districts in California or Washington featured same-party (Democratic) general elections we include these states. We repeat our main analysis without these districts in the supplementary material.

dataset between June 2020 and March 2021. To prepare the data, we removed all URLs, lower-cased, and cleaned for HTML code (such as emojis). We removed names, punctuation, numbers, and Quanteda's (Benoit et al., 2018) default English stopword lists to reduce computational requirements. We remove all hashtags and mentions in our main analysis after comparing validity across specifications (see supplementary materials).

3.2 Positions from Twitter Text

Following Goet (2019) and Green et al. (2020) we use a supervised machine learning model to estimate candidates' positions in Euclidean space (Laver, Benoit and Garry, 2003; Slapin and Proksch, 2008). We classify each candidate based on their party identification using a Naïve Bayes classifier. Our model uses a bag-of-words approach to predict the party membership of each candidate. Each word in the dataset is assigned a partisan value which can then be applied to any document to score how "partisan" it is. Traditional classifiers use binary classification to estimate the outcome, but, because we want a continuous measure, we use the (normalized) relative log-likelihood, giving a score that a document has a certain partisan "identity". In the case of individual positions (as in the validation) this 'document' is all tweets by a candidate in a given period.

Uncertainty: One disadvantage of this approach is the absence of confidence intervals. As the model estimates the likelihood of a text's partisanship, there is no natural interpretation of uncertainty. We can quantify how dependent the results are on specific cases and features, for example, if a candidate uses specific terminology in a manner distinct from their colleagues and changes the meaning. To account for this possibility, we compute *bootstrapped* positions. Instead of computing a single Naïve Bayes model, we resample all data by drawing ninety percent of them 400 times, rerunning the model, and storing the term weights. When predicting the positions of documents, we again predict 400 positions, computing the standard deviation to get an approximation of error. The results are normally-distributed positions around a mean, allowing us to quantify potential uncertainty.

To apply our data to our research question we compute candidate positions at different time points, before and after their respective primaries. We use a three-step process: training the Naïve Bayes model, computing positions of members of congress, validating these positions, and aggregating the data at different levels. We predict the party membership of a validation set of thirty percent of candidates using the other seventy percent as training data. We achieve an accuracy of 0.946, precision of 0.955, recall of 0.926, and F1 score of 0.940, indicating that the model is very good at predicting candidates' partisan affiliation.[4] Having trained the model at the individual level, we then apply the weights of these terms to tweets aggregated at the candidate level, the candidate level before and after the primary, and the party level over time (weeks). In other words, we train the model on partisan difference and then estimate the degree of partisanship.

Challenges of this approach include variation in the quantity of candidate-level data, with some candidates rarely tweeting and others so active that their tweets are capped by the API rate limitations Twitter imposes (3,200 tweets). Perhaps most importantly, our dataset includes a combination of political tweets mixed with apolitical tweets that do not indicate position. This mix of content has the potential to produce problems when scaling positions, where higher rates of non-political tweets could result in candidates being interpreted as moving toward the center (Grimmer and Stewart, 2013). We deal with this problem explicitly by also applying our model to policy-related tweets only.

Our approach has several advantages. We use the simplest possible model, driven by our desire to avoid overfitting, as a model that was too tuned to classify partisanship might neglect intra-party differences. A second advantage is the computational requirements where, because of the speed of Naïve Bayes, large bootstraps can still run on a single computer. This type of model also does not require stop criteria or a loss metric as it is solved on the document feature matrix (DFM), meaning it does not need to converge in the way that a deep learning model would.

External Validity: Introducing a new measurement for a latent construct requires external validation, we demonstrate our scores' predictive validity against other known estimates of congressional candidates. Given that one motivation for this study is the absence of such measures for all candidates, we compare our results with a subset of our data. The most widely used measure

[4]We also include the results of ten-fold cross-validation in the supplementary materials.

is NOMINATE (Poole and Rosenthal, 1985), based on members' roll-call voting in Congress. Of course, this measure is only available for members who have ever served in Congress. If these members are positioned in a meaningful way that captures the underlying dimension, other candidates placed on the same dimension should also align. In total, we validate our measure using over 2,000,000 Tweets by 518 members of Congress.

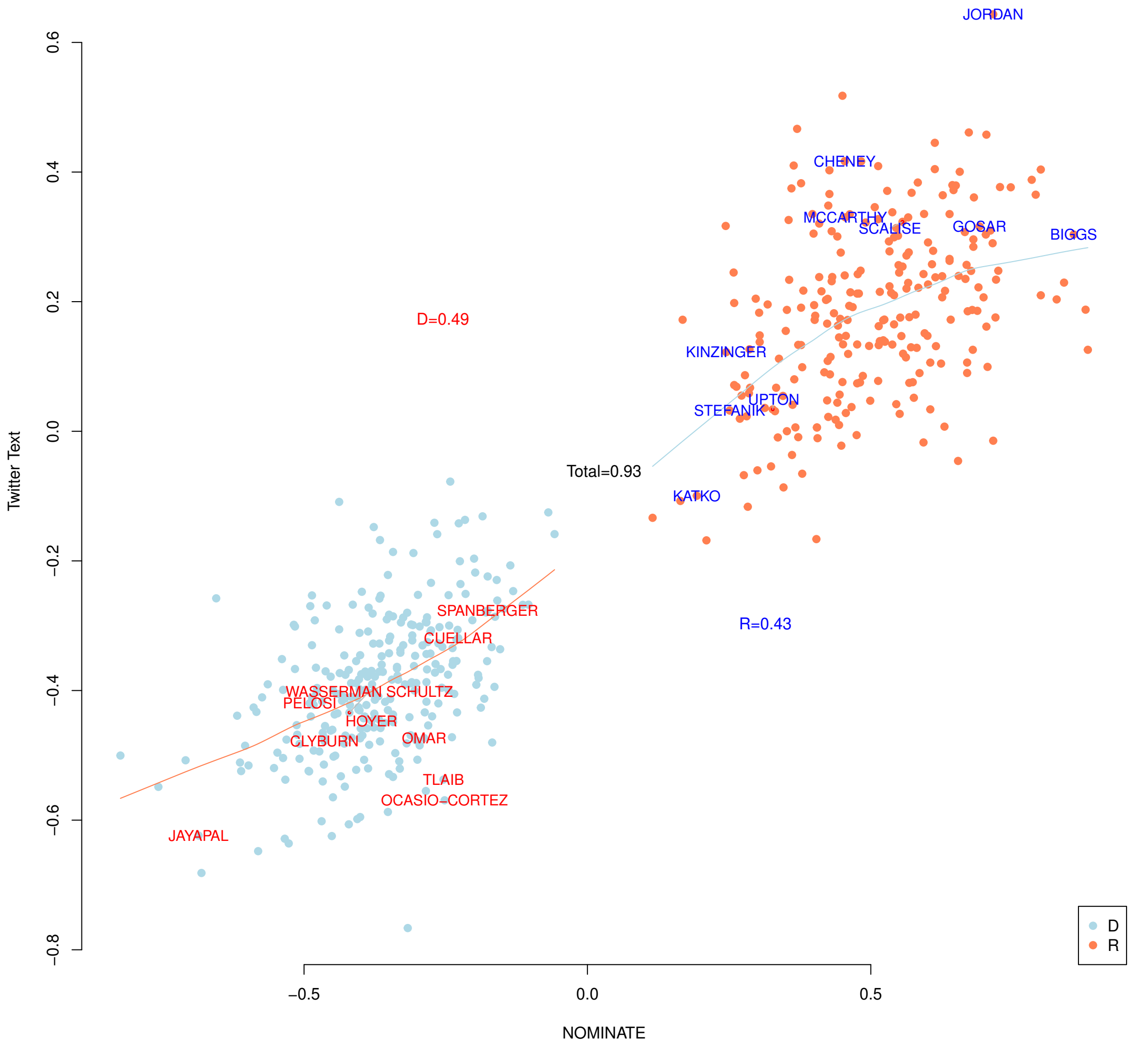

Figure 1: Validation Against NOMINATE for Members of Congress

Figure 1 shows this validation, with NOMINATE scores on the x-axis. The y-axis shows the average positions predicted by Twitter communication over the entire electoral cycle. To increase the number of data points against which to validate, and to give our model a hard test, we also include U.S. senators and incumbent representatives who retired in 2020 in this plot. Our model was not trained on these members' tweets, providing an ideal independent corpus against which to validate.[5]

The overall correlation is 0.93, with higher intra-party correlations than alternative recognized scaling measures such as follower network scores (Barberá, 2015) or CFscores (see Barber, 2022). We also demonstrate semantic validity by labeling some notable representatives' positions. In both parties, representatives who are commonly perceived as 'moderates'—including Abigail Spanberger, Henry Cuellar, John Katko, and Fred Upton—are also moderate by our measure. Similarly, representatives such as Pramila Jayapal and Jim Jordan, viewed as highly liberal and conservative respectively, are away from the center on our scale. In addition, Democratic representatives such as Alexandria Ocasio-Cortez and Rashida Tlaib, who are incorrectly positioned as moderates by NOMINATE due to their opposition to some Democratic bills,[6] are positioned as more liberal under our measure. These correlations give confidence that our measure is aligned with the liberal–conservative dimension structuring roll-call voting behavior, and suggest that in some cases where they differ, our measure may even serve as a more accurate proxy for ideology than NOMINATE.

Semantic Validity: Though we obtain predictive validity by comparing the positions generated with roll-call votes, we need to qualify our analysis by understanding the language that identifies our dimension. To do so, we interpret influential words that produce scores further from the center. Our measure can be said to have semantic validity if these scores are associated with parties' positions, campaign rhetoric, or policy issues.

Figure 2 shows the terms for each end of the dimension surrounding the positions estimated in Figure 1 that occur more than 1000 times in the entire corpus of tweets. Positions from Figure 1 are shown in the center of Figure 2. The lower (higher) the position of a word on the y-axis, the more indicative it is for the Democratic (Republican) Party and contributes to a score further to the left (right). Accordingly, representatives that tweet a lot about "illegals" and "rioters" receive scores further to the right than those who tweet about more moderate identifying terms such as "manufacturers" or "regulations". The positions of words on the x-axis are for presentation

[5] Senators' data are only used for validation and do not feature in our main analyses.
[6] See Lewis (2022) for details

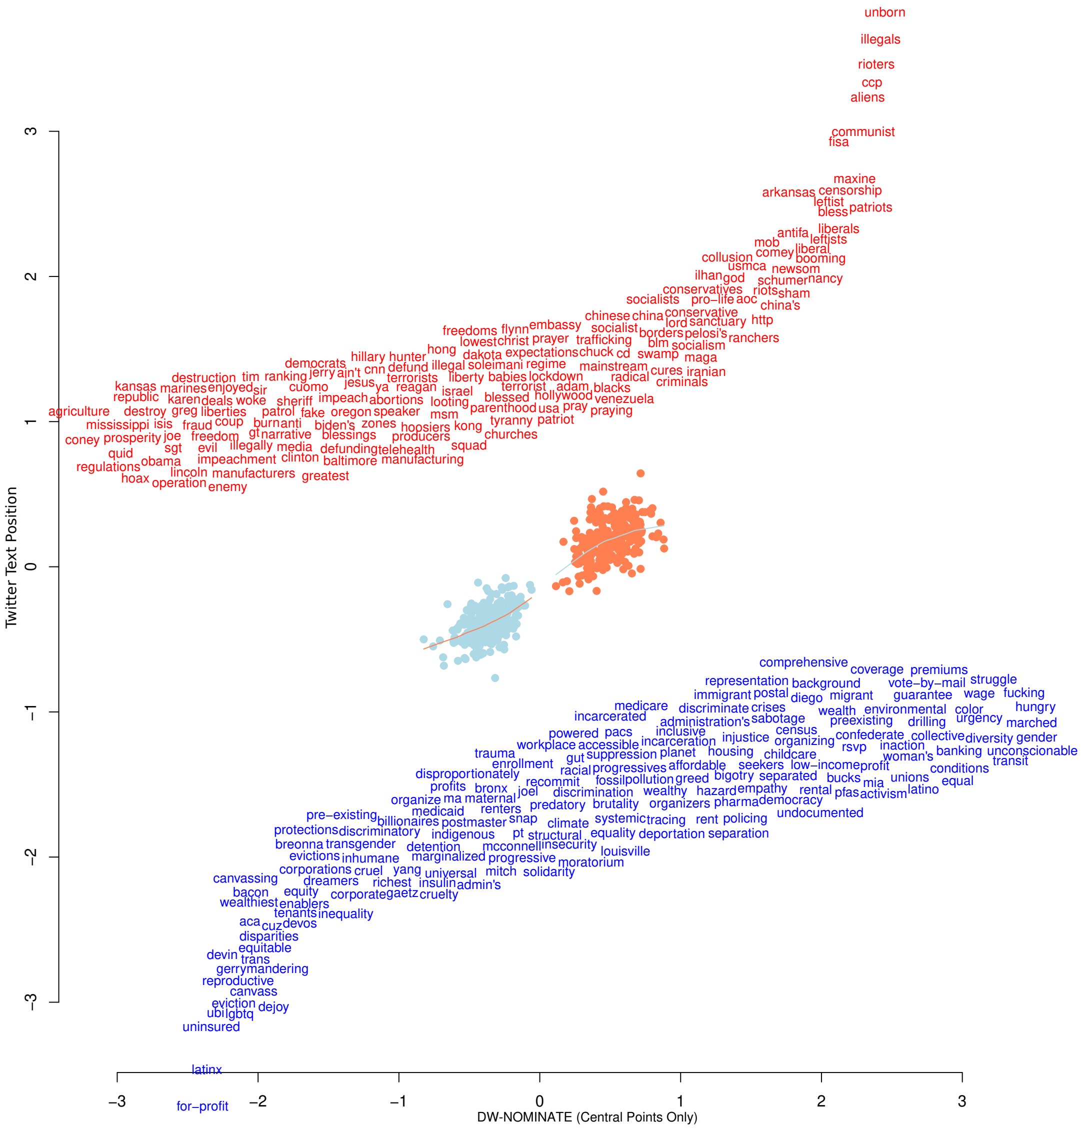

Figure 2: Validation with Terms

purposes only and have no substantive meaning. Figure 2 demonstrates that the terms that score highly in either a liberal or conservative direction are in line with partisan expectations, where terms at the bottom would be words expected to be used by Democrats and terms at the top of the figure expected to be used by Republicans. In other words, Figure 2 indicates that our approach has semantic validity.

These terms can broadly be grouped into three categories: policy-related, own-party rhetoric, and negative terms. Policy-related terms to the right included "illegals", "censorship" and "unborn". Republican own-party rhetorical terms included "patriots" and "conservatives". The terms "rioters", "communist", and "leftist" were used by Republican candidates to talk negatively about the Democratic Party and their supporters and were similarly scored to the right. Liberal policy-related terms included "uninsured", "ubi", and "for-profit". Democratic own-party rhetorical terms included "canvass" and "progressive", and terms such as "lgbtq" and "trans" referred to demographic groups who favor the party. The terms "inhumane" and "cruelty" were negative liberal identifiers. Given that the terms at each end of our scale can be broadly understood as having a partisan valence, we can say that our approach has semantic validity.

4 Findings

Following validation, we trust the model to infer positions. In our first analysis, we produce a model at the party level and focus on dynamics over time. To test the effect of primaries, we are not interested in the date, but the *relative* time to or since candidates' respective primaries. Because states hold nomination contests on different dates, we center the time around each intra-party election, using a time-to-primary variable for each tweet as weeks before or after the primary. We then aggregate at the following levels: party, whether the candidate won their primary, and time-to-primary (weeks). Each observation is the aggregate of terms used by members of a party who won or lost the nomination at the same relative time before or after their primary.[7]

4.1 Shifting after the Primary: The Party Perspective

Figure 3 shows the positions of winning and losing candidates in both parties as groups aggregated by week to or from their respective primary. As the figure indicates, Democratic candidates who do not become the nominee shift their position *towards the center* directly after their primary. Republican losers do not moderate following primary defeats.

To test the statistical significance of this effect, we run a comparative interrupted time series analysis (ITS) with the below specification (see also Linden, 2015). Our data are repeated observations of candidates' communication positions and we expect positions to change following the 'intervention'; the primary election date. We use a (comparative) ITS model given the obvious differences between many candidates who win and lose primary elections. Many candidates

[7] As a placebo test, we also randomized this date. See supplementary material for details.

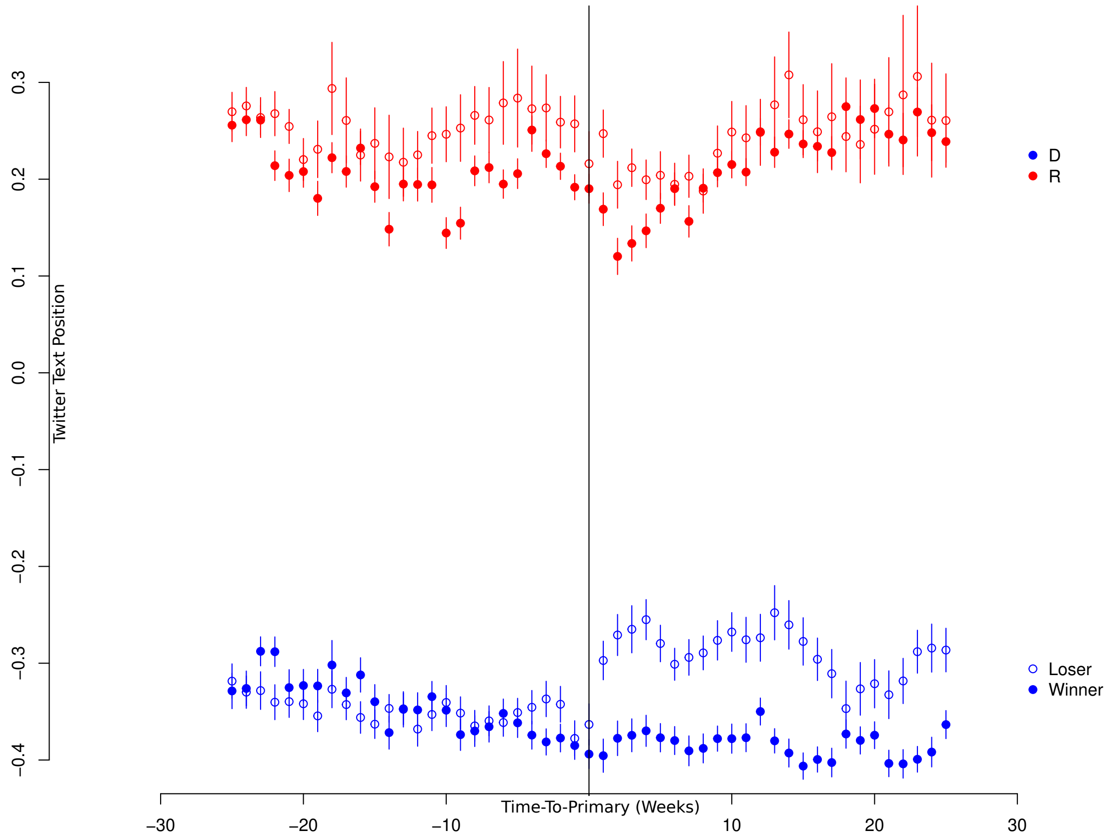

Figure 3: Party Level Positions Over Time

who win primary elections are either incumbent members of Congress or highly experienced and well-financed challengers. In contrast, many primary losers receive little to no support from party elites, have little financial support, and may be relatively unknown. Put simply, we conceive that there are too many differences between winning and losing primary candidates to control for, even using approaches such as matching, synthetic controls, or propensity score weighting. Instead, we use an ITS which allows us to compare groups and compare candidates' positions to themselves prior to the intervention. We do not expect primary winners to moderate immediately after the primary in this design. Conversely, we expect that losing candidates will be more moderate after the primary than they were during the nomination. Using an ITS rather than a two-way fixed effects model also allows us to include group characteristics that change gradually during the election cycle. Given that our data-generating process is independent for each time period, we do not include lagged variables in our models (see also Warner, 2019).[8] One drawback of this design is

[8]Empirically our data are independent at each time point, where the communication for a given week is not the result of communication beforehand. Yet, theoretical and empirical literature indicates that candidates benefit from positional consistency. Though our dependent variable

that the differences—both between winners and losers, and losers versus themselves in the previous period—mean our results are associational, and we cannot infer that the presence of the primary is what *caused* candidates to adopt artificial positions. We run separate models by party, with the following specification for our first models:

$$\mathrm{Y_{it}} = \beta_0 + \beta_1\mathrm{T_t} + \beta_2\mathrm{X_t} + \beta_3\mathrm{X_tT_t} + \beta_4\mathrm{Z_i} + \beta_5\mathrm{Z_iT_t} + \beta_6\mathrm{Z_iX_t} + \beta_7\mathrm{Z_iX_tT_t} + \varepsilon_t$$

Where $\mathrm{Y_{it}}$ is candidate position Y given membership of group[9] i measured at week t. $\mathrm{T_t}$ is the time in weeks to or since the primary. $\mathrm{X_t}$ is a dummy variable representing the primary election, where pre-primary observations take the value zero and post-primary observations the value one. $\mathrm{X_tT_t}$ is the interaction term between post-primary and time, meaning β_2 is the immediate change following the primary and β_3 gives the ongoing movement among all observations. $\mathrm{Z_i}$ is the group we expect to moderate, which takes the value one if a candidate lost and zero if a candidate won their primary. Coefficients β_4 to β_7 are the same as β_0 to β_3 interacted with losing ($\mathrm{Z_i}$), meaning β_6 gives the immediate change among losing candidates immediately after the primary and β_7 gives the ongoing movement following the primary. We expect moderation from losing candidates immediately after they lose their primary, meaning β_6 ($\mathrm{Z_iX_t}$) is our main object of interest for the first models.[10]

Given that our goal is not the causal identification of differences between winners and losers, we also include a second set of models that are restricted to losing candidates only. These models take the same form as the above specification with the removal of the loser variable $\mathrm{Z_i}$ and subsequent interactions, meaning $\mathrm{X_t}$ is the object of interest in these models. Our first models indicate how losing candidates were positioned relative to winners in the same week, whereas the second set of models identify how candidates moved relative to themselves in the previous period.

One potential issue with cross-sectional time series data is non-stationarity, where conditional means are dependent on the time period and where a variable has a unit root. To demonstrate that our models have $I(0)$ balance (Pickup and Kellstedt, 2022) and to understand the order of

does not depend linearly on its own previous values, we expect these values to be correlated. We therefore demonstrate the robustness of our findings by including a lagged version of candidate positions in the supplementary material.

[9] Democratic winners, Democratic losers, Republican winners, Republican losers.

[10] We use Newey-West standard errors to account for potential heteroskedasticity and serial autocorrelation.

integration we perform (augmented) Dickey-Fuller (Dickey and Fuller, 1979) tests on each of the four groups' dependent variables, with results reported in the supplementary material. In each case, our tests return significant values, indicating no unit root on the left-hand side of our models. We also account for variation in the trend stationary dependent variable by including T_t in our specification. Of our independent variables, both the primary (X_t) and winning or losing (Z_i) do not contain a stochastic component. The only term on the right-hand side of our equation that is stochastic is the error term; we demonstrate that the estimated errors (residuals) are indeed white noise in a further series of Dickey-Fuller tests, with the results reported in the supplementary material. These tests indicate that our equation is $I(0)$ balanced.

Table 1: ITS Results: Party Level

	All Candidates		Losers Only	
	Democrats	Republicans	Democrats	Republicans
Time (T_t)	-0.003***	–0.001***	–0.001**	0.000
	(0.000)	(0.001)	(0.000)	(0.001)
Post-Primary (X_t)	0.009	–0.044***	0.093***	–0.053***
	(0.009)	(0.014)	(0.010)	(0.013)
Post-Primary : Time (X_tT_t)	0.003***	0.007***	–0.000	0.003***
	(0.001)	(0.001)	(0.001)	(0.001)
Loser (Z_i)	0.025***	0.064***		
	(0.009)	(0.013)		
Loser : Time (Z_iT_t)	0.002***	0.001		
	(0.001)	(0.001)		
Loser : Post-Primary (Z_iX_t)	0.084***	–0.009		
	(0.013)	(0.020)		
Loser : Post-Primary : Time ($Z_iX_tT_t$)	–0.003***	–0.003*		
	(0.001)	(0.001)		
Intercept	–0.387***	0.190***	–0.362***	0.254***
	(0.006)	(0.010)	(0.007)	(0.009)
N	102	102	51	51
R^2	0.846	0.608	0.765	0.368
Adjusted R^2	0.834	0.579	0.750	0.328
Residual Std. Error	0.016	0.025	0.018	0.024
	(df = 94)	(df = 94)	(df = 47)	(df = 47)
F Statistic	73.747***	20.851***	50.927***	9.138***
	(df = 7; 94)	(df = 7; 94)	(df = 3; 47)	(df = 3; 47)

Newey-West Standard Errors Shown in Parentheses
*p < 0.1; **p < 0.05; ***p < 0.01

In line with the visual trend depicted in Figure 3, our first model in Table 1 shows that Democratic losers became significantly more moderate than winners immediately after the primary (Z_iX_t). In contrast to the weak time trend, the effect is almost five percent of the total range of the variable, this is the strongest identifier of position other than partisanship. In other words, losers shift their position after their primary relative to winners, and this shift is more than twenty times greater than the average weekly positional change (T_t). Losing Democratic candidates were more moderate than winners prior to the primary (Z_i) yet moved much further rightward following the primary (Z_iX_t). All other Democratic coefficients in this first model are substantively close to zero.

For Republican losers, Table 1 indicates no significant moderation following primary defeats relative to primary winners (Z_iX_t). It appears that Republican winners moderate slightly after the primary (X_t) then quickly move back towards their pre-primary positions in subsequent weeks (X_tT_t), also seen in Figure 3. Across the whole period, losing Republican primary candidates are consistently further to the right than winners (Z_i). All other coefficients in this first model are substantively close to zero.

In the second set of models, we consider the position of losers after the primary compared to their positions during the primary, indicated by the post-primary coefficient (X_t). Among Democratic losers, our finding is virtually unchanged, with Democratic candidates again positioned significantly further to the right immediately after the primary compared to their previous positions. Among Republicans, we also see evidence of moderation of losers in the immediate post-primary period as compared to their position during the primary. As depicted visually in Figure 3, it appears that all Republicans moderated immediately after the primary and then returned to their original positions over time. This movement is substantively far smaller than among losing Democrats.

Unsurprisingly, partisanship—shown here in the form of the intercept—is the strongest predictor of position for candidates in both parties. At the party level, we find a clear moderating effect among losing Democratic candidates.

4.2 Robustness to the Changing Salience of Non-Political Tweets

One identifiable problem of ideal point estimation over time is the changing salience of features that contribute to the dimension (Grimmer and Stewart, 2013). The appearance of moderation may stem from movement toward more centrist content—ideological moderation—or a reduction of political or policy-related content. Accordingly, it might be that candidates are merely tweeting less about politics and turning their account into a private platform after they lose a primary rather than continuing to discuss politics.

To ensure the robustness of our approach to this problem, we apply our method to a subset of explicitly policy-related tweets. To do so, we hand-coded a random set of 1,200 tweets using three categories; political (y/n), policy-related (y/n), and policy area (using policy fields established in the Comparative Agendas Project). Though the sample was too small to analyze policy areas individually, roughly half of the tweets in the sample were policy-related. We then trained a classifier for these tweets, using an English-language Bidirectional Encoder Representations from Transformers (BERT) (Devlin et al., 2019) model, which achieves a satisfactory F1 score of 0.8. We use this model to predict whether all 2,500,000 tweets in our original sample were policy-related (again, roughly half were) and estimate positions.[11] We then re-ran our analyses on this subset.

The results are shown in Table 2 and align with our main finding, with substantively significant moderation among Democratic losers after the primary, either compared to Democratic winners in the same period or to themselves during the primary. Movement immediately after the primary is again more than twenty times the size of the average weekly movement and is the strongest indicator of position other than partisanship. Our finding that Republican losers were more moderate after than during the primary is no longer significant when restricted to policy tweets, suggesting that this finding was at least partly the result of a shift in focus. This additional analysis gives confidence that our main result for Democrats is not an artifact of the changing saliency of policy-related tweets after primary defeats and is instead evidence of positional adaptation by losing candidates.[12]

[11] We repeated this process with political (y/n). Because roughly ninety percent of tweets were coded as political, this variable had limited analytical application.

[12] We again demonstrate stationarity and $I(0)$ balance by conducting Dickey-Fuller tests on our dependent variables and residuals in this subset, see supplementary material.

Table 2: ITS Results: Policy Tweets Only

	All Candidates		Losers Only	
	Democrats	Republicans	Democrats	Republicans
Time (T_t)	−0.003***	−0.001	−0.000	−0.000
	(0.001)	(0.001)	(0.001)	(0.001)
Post-Primary (X_t)	0.020*	−0.040***	0.059***	−0.024
	(0.011)	(0.015)	(0.011)	(0.020)
Post-Primary : Time (X_tT_t)	0.001	0.006***	−0.002**	0.002
	(0.001)	(0.002)	(0.001)	(0.001)
Loser (Z_i)	-0.003	0.033***		
	(0.009)	(0.019)		
Loser : Time (Z_iT_t)	0.002***	0.000		
	(0.001)	(0.001)		
Loser : Post-Primary (Z_iX_t)	0.041***	0.001		
	(0.014)	(0.028)		
Loser : Post-Primary : Time ($Z_iX_tT_t$)	−0.002***	–0.002		
	(0.001)	(0.002)		
Intercept	−0.342***	0.211***	−0.456***	0.152***
	(0.026)	(0.038)	(0.007)	(0.013)
N	102	102	51	51
R^2	0.851	0.610	0.534	0.055
Adjusted R^2	0.839	0.576	0.504	−0.005
Residual Std. Error	0.016	0.025	0.019	0.035
	(df = 93)	(df = 93)	(df = 47)	(df = 47)
F Statistic	66.578***	18.157***	17.937***	0.910
	(df = 8; 93)	(df = 8; 93)	(df = 3; 47)	(df = 3; 47)

Newey-West Standard Errors Shown in Parentheses
*p < 0.1; **p < 0.05; ***p < 0.01

4.3 Individual Level Robustness

To avoid the ecological fallacy, we also analyze the individual level. We do not have enough tweets at the individual level to reliably compute positions in the same density as at the party level[13] meaning we instead aggregate candidates' positions before and after their primary to enable the direct comparison of candidate-level movement. In this model, we control for incumbency given that incumbents may face additional pressures and incentives to maintain their positions because they have political records to uphold which can be held accountable by voters and opposition candidates. Given that district partisanship influences positional incentives in both primary and general elections, we control using *The Cook Political Report*'s (2017) partisan voting index (PVI), rescaled to a +/− Republican lean.[14]

[13]The number of candidates positioned is also reduced from 988 to 886.
[14]We repeat this analysis without controls in the supplementary material, our results are unchanged.

Figure 4 shows the individual-level results. These models use the difference (movement) in candidates' positions before and after their primary as the dependent variable, where positive coefficients indicate rightward movement and negative coefficients indicate leftward movement. We test using two dependent variables: absolute movement, and a variable of *significant* movement. This variable takes the value 1 if a candidate moves rightward three standard error confidence intervals and the value −1 if a candidate moves left to the same degree.

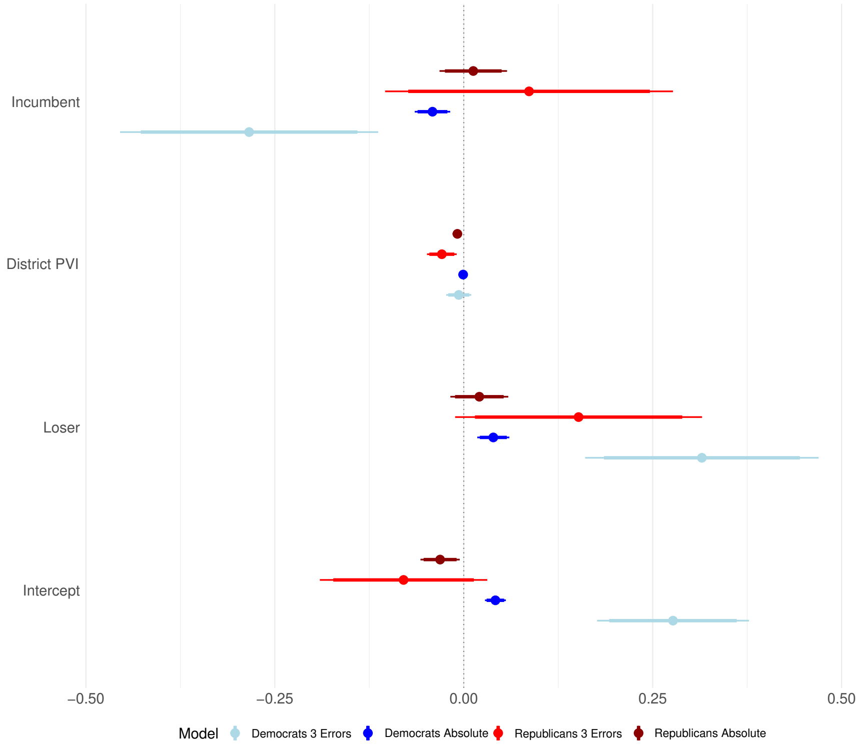

Figure 4: Individual Level Movement

In line with our party-level findings, Democratic losers took more moderate positions after the primary in both individual-level models, giving further confidence in our party-level findings. Republican losers also move slightly to the right, but the effect is not statistically significant. As in the party-level model, partisanship—the intercept—indicates moderation among all candidates at the individual level following the primary. Democratic incumbents moved slightly to the left at the individual level, with no significant effect among Republicans. District partisanship had no relationship to Democratic positioning and a small but significant association for Republicans, who took less conservative positions in districts that were less favored for the party.

5 Discussion

Our results indicate that primaries are associated with artificial position-taking among Democratic candidates only. We interpret these findings as support for the strategic positioning dilemma among Democratic candidates, who adopted artificial positions during the primary which they did not continue to hold once absent the (perceived) incentives to do so. Among Republican candidates, we find minimal evidence of artificial positioning, suggesting that communication during the primary was done out of conviction rather than for perceived advantage. Absent electoral incentives, losing Republican primary candidates continued to communicate highly conservative positions.

The moderation of losing Democratic candidates after the primary indicates our theorized effect that intra-party competition is associated with artificial extremism during the nomination. Grossmann and Hopkins (2016) suggest that the Democratic Party is a diverse coalition of group-oriented actors. Rather than being defined by ideological conflict, candidates advocate for different groups which are understood primarily in terms of demographics and identity. Consequently, Democratic candidates are less frequently ideological purists and so may be more comfortable adapting their positions. Because ideology is not a central binding force in the party, candidates are able to be more flexible and change positions than their Republican counterparts. If candidates perceive that important policy demanders are to their left, they may have additional incentives to adopt artificial positions during the nomination. The Democratic Party might therefore recruit more strategic candidates or be more selective in recruitment by actively seeking out candidates who can adapt positions. The ability to be flexible and strategically appeal to many of the diverse interest groups that make up the Democratic Party appears one important characteristic sought out by party elites and policy demanders in the party network who play a central role in candidate recruitment (Cohen et al., 2008; Hassell, 2018). These groups prefer candidates with a broad appeal during the nomination process (Masket, 2020), in part out of necessity because the party needs to carry some swing or even marginally Republican-favored districts in general elections to control the House. In short, recruitment strategies matter and are likely asymmetric (Maestas and Stewart, 2012). Intra-party power struggles likely provide further incentives for Democrats to moderate after a primary. Though progressives have made recent gains, the Democratic Party remains dominated by 'establishment' center-left moderates, meaning losing candidates who want to continue a career in the party are wise to moderate to appeal to like-minded individuals.

For Republicans, our results align with scholarship that positions candidates for Congress as more extreme, or at least more ideologically consistent, than other groups and voters in their party (Bafumi and Herron, 2010; Barber, 2016). These results run counter to the expectations of the strategic positioning dilemma. Candidates in the Republican Party take non-centrist positions out of conviction both during and after the primary, where losing a primary was not associated with moderation. That losing Republicans largely continue to communicate non-centrist positions likely reflects a reality where the only candidates running are located so firmly on the right of the political spectrum that they perceive little concern over strategic positioning during the nomination. This explanation aligns with scholarship indicating that the Republican Party has moved sharply rightward in recent years (Hacker and Pierson, 2006; Mann and Ornstein, 2012; Theriault, 2013), meaning losing primary candidates have less incentive to moderate to help their future career in the party. Republican partisans are also less tolerant of elite positional heterogeneity (Dunn, 2021), meaning party elites and other actors in the formal party organization may be more disposed to recruit loyal (or sincere) believers who hold consistent positions away from the political center. Given the (perceived) position of primary voters and policy demanders in the party, moderate Republicans may simply decide that running for Congress is not worthwhile (Thomsen, 2017). Institutional biases in general elections—including aggressive Republican gerrymandering in the previous redistricting cycle and the electorally inefficient clustering of Democratic voters in urban districts—may also have furthered a perception among Republican policy demanders and primary voters that candidates on the right of the political spectrum are electorally viable.

Given that our analysis is conducted over a single electoral cycle, we must also consider the relative effect of 2020 electoral conditions on the two parties. Boatright and Moscardelli (2018) demonstrate that congressional primaries have a "presidential pulse." In 2020, the Democratic Party was favored to win the presidency and expected a strong down-ballot performance, with higher numbers of primary candidates as a result. Higher levels of primary competition may have served as a further incentive to induce Democratic candidates to adopt artificial positions.

The party-level differences may also relate to demographic and ideological differences between Twitter and non-Twitter users. Twitter users are Democratic-leaning and disproportionately come from demographic groups which favor the party, such as young college-educated Whites with higher incomes (Wojcik and Hughes, 2019). Even among Democratic partisans, those on Twitter tend to hold more progressive or left-leaning positions (Cohn and Quealy, 2019), with fewer moderates active on social media (Hawkins et al., 2018). Democratic primary candidates may therefore have communicated positions on Twitter to appeal to a section of the electorate that they—correctly—perceived as non-centrist. In contrast, Republican candidates may perceive that fewer of their primary voters are on Twitter and so use the platform to communicate to journalists and media outlets, other candidates, or party figures.

Asymmetries in the parties' financial structures may further explain our findings. Basedau and Kollner show that "centripetal tendencies are better avoided when the channels of party finance are controlled by the party leadership" (2005, 19), and recent literature highlights clear partisan differences in this regard. Boatright and Albert (2021) show that independent expenditures were not particularly prevalent in financing primary challengers to Democratic incumbents in 2018. Assuming a similar pattern in 2020, the tighter financial control of the formal institutions of the Democratic Party may have incentivized losing candidates to moderate to retain favor with party leadership and advance their political careers. The asymmetric structure of media ecosystems, with greater pressure from the right and far-right of the ideological spectrum (Heft et al., 2021), may also have induced Republican candidates to maintain conservative positions. Pierson and Schickler (2020) find that meso-institutional structures pull Republicans away from the center more than Democrats. One interpretation of our findings is that these structures continue to affect Republicans' positions following primary defeats.

For general election voters, these results are not encouraging when considered in terms of spatial models of voting. Given that we find limited evidence of moderation among primary winners in either party,[15] voters in November appear to have been presented with polarized *choices*—as theorized by Fiorina, Abrams and Pope (2005)—albeit for contrasting partisan reasons, with

[15] This result aligns with the expectations and findings in Brady, Han and Pope (2007).

Democratic candidates having strategically adopted artificial positioning during the nomination and Republicans sincerely holding non-centrist positions out of conviction. Non-moderation of Democratic primary winners may indicate a perception among candidates that they must continue to hew to the preferences of policy demanders beyond the primary or reflect candidates' beliefs about the electoral risks associated with moving positions between a primary and general election. Among Republicans, our data suggest limited adaptation, and positions appear more deeply ingrained in the preferences of candidates.

6 Conclusion

We find that losing Democratic candidates moderate after the primary. We argue that this is evidence that candidates communicated artificial positions during the nomination to try and align with key policy demanders and the perceived positions of their primary voters during the nomination. Losing Republican candidates did not moderate following their primary defeats. These results align with scholarship indicating asymmetry in the ideological positions (Hacker and Pierson, 2006; Theriault, 2013) and identities (Grossmann and Hopkins, 2016) of the two major parties and the policy demanders active during the nomination process within each. These differences provide distinct partisan constraints and incentives to candidates both during and after primary elections.

The debate over whether primaries contribute to polarization in Congress is ongoing (Abramowitz, 2010; Fiorina, Abrams and Pope, 2005; Sides et al., 2020), yet, many studies only consider this question in terms of a selective effect from primary voters. We demonstrate a further way in which contested nominations may exacerbate partisan conflict in Congress: the adaptation of candidate positions during the nomination phase of the election cycle. If many candidates perceive that communicating artificial positions is beneficial during the primary and then feel compelled to maintain those positions during the general election, voters in November will be presented with more polarized choices as a result of the nomination process.

We find little movement among nominees in either party once they are selected, a potentially positive normative finding in terms of representation. Regardless of whether candidates adopt sincere or strategic positions, primary winners communicate positions in general election campaigns that are consonant with their positions during the nomination. How candidates communicate in a primary is at least consistent with what they advocate when they become the nominee—and, potentially, indicative of the policies they will support in Congress. This finding contrasts with the image of politicians as pandering to different groups for their own benefit (Lippmann, 1955; Jacobs and Shapiro, 2000).

References

Abramowitz, Alan I. 2008. "Don't Blame Primary Voters for Polarization." *The Forum* 5(4):1–11.

Abramowitz, Alan I. 2010. *The Disappearing Center: Engaged Citizens, Polarization, and American Democracy.* Yale University Press.

Acree, Brice D.L., Justin H. Gross, Noah A. Smith, Yanchuan Sim and Amber E. Boydstun. 2020. "Etch-a-Sketching: Evaluating the Post-Primary Rhetorical Moderation Hypothesis." *American Politics Research* 48(1):99–131.

Ansolabehere, Stephen, James M. Snyder and Charles Stewart. 2001. "Candidate Positioning in U.S. House Elections." *American Journal of Political Science* 45(1):136–159.

Bafumi, Joseph and Michael C. Herron. 2010. "Leapfrog Representation and Extremism: A Study of American Voters and Their Members in Congress." *The American Political Science Review* 104(3):519–542.

Barber, Michael. 2016. "Representing the Preferences of Donors, Partisans, and Voters in the US Senate." *Public Opinion Quarterly* 80:225–249.

Barber, Michael. 2022. "Comparing Campaign Finance and Vote-Based Measures of Ideology." *The Journal of Politics* 84(1):613–619.

Barberá, Pablo. 2015. "Birds of the Same Feather Tweet Together: Bayesian Ideal Point Estimation Using Twitter Data." *Political Analysis* 23(01):76–91.

Barberá, Pablo, Andrue Casas, Jonathan Nagler, Patrick J. Egan, Richard Bonneau, John T. Jost and Joshua A. Tucker. 2019. "Who Leads? Who Follows? Measuring Issue Attention and Agenda Setting by Legislators and the Mass Public Using Social Media Data." *American Political Science Review* 113(4):883–901.

Basedau, Matthias and Patrick Köllner. 2005. "Factionalism in Political Parties: An Analytical Framework for Comparative Studies." *SSRN Electronic Journal* .

Bawn, Kathleen, Martin Cohen, David Karol, Seth Masket, Hans Noel and John Zaller. 2012. "A Theory of Political Parties: Groups, Policy Demands and Nominations in American Politics." *Perspectives on Politics* 10(3):571–597.

Benoit, Kenneth, Paul Nulty, Stefan Müller, Adam Obeng, Kohei Watanabe and Akitaka Matsuo. 2018. "Quanteda: An R package for the Quantitative Analysis of Textual Data." *Journal of Open Source Software* 30(3):774–.

Blum, Rachel M. 2020. *How the Tea Party Captured the GOP: Insurgent Factions in American Politics.* Chicago: University of Chicago Press.

Boatright, Robert G. 2013. *Getting Primaried: The Changing Politics of Congressional Primary Challenges.* University of Michigan Press.

Boatright, Robert G. 2014. *Congressional Primary Elections.* Routledge.

Boatright, Robert G. and Vincent G. Moscardelli. 2018. Is There a Link Between Primary Competition and General Election Results. In *Routledge Handbook of Primary Elections*, ed. Robert G. Boatright. New York: Routledge and Taylor and Francis pp. 188–212.

Boatright, Robert G. and Zachary Albert. 2021. "Factional Conflict and Independent Expenditures in the 2018 Democratic House Primaries." *Congress & the Presidency* 48(1):50–77.

Boireau, Michaël, ed. 2014. *Determining Political Stances from Twitter Timelines: The Belgian Parliament Case*.

Bonica, Adam. 2014. "Mapping the Ideological Marketplace." *American Journal of Political Science* 58(2):367–386.

Brady, David W., Hahrie Han and Jeremy C. Pope. 2007. "Primary Elections and Candidate Ideology: Out of Step with the Primary Electorate?" *Legislative Studies Quarterly* 32(1):79–105.

Budge, Ian and Dennis J. Farlie. 1983. *Explaining and Predicting elections: Issue Effects and Party Strategies in 23 Democracies*. London: Allen and Unwin.

Burden, Barry C. 2001. The Polarizing Effects of Congressional Primaries. In *Congressional Primaries and the Politics of Representation*, ed. Peter F. Galderisi, Marni Ezra and Michael Lyons. Rowman & Littlefield.

Canes-Wrone, Brandice, David W. Brady and John F. Cogan. 2002. "Out of Step, Out of Office: Electoral Accountability and House Members' Voting." *The American Political Science Review* 96(1):127–140.

Ceron, Andrea. 2016. "Intra-Party Politics in 140 Characters." *Party Politics* 23(1):7–17.

Chen, Kong-Pin and Sheng-Zhang Yang. 2002. "Strategic Voting in Open Primaries." *Public Choice* 112(1/2):1–30.

Cohen, Marty, David Karol, Hans Noel and John Zaller. 2008. *The Party Decides: Presidential Nominations Before and After Reform*. Chicago: University of Chicago Press.

Cohn, Nate and Kevin Quealy. 2019. "The Democratic Electorate on Twitter Is Not the Actual Democratic Electorate." *The New York Times* .
URL: *www.nytimes.com/interactive/2019/04/08/upshot/democratic-electorate-twitter-real-life.html*

Converse, Philip E. 1964. "The Nature of Belief Systems in Mass Publics." *Critical Review* 18(1-3):1–74.

Cooper, Alexandra and Michael C. Munger. 2000. "The (Un)Predictability of Primaries with Many Candidates: Simulation Evidence." *Public Choice* 103(3/4):337–355.

Cowburn, Mike. 2022. "Partisan Polarization in Congressional Nominations: How Ideological & Factional Primaries Influence Candidate Positions." Doctoral Thesis Freie Universität Berlin.

Cowburn, Mike and Curd B. Knüpfer. 2023. "The Emerging Fault Line of Alternative News: Intra-Party Division in Republican Representatives' Media Engagement." *Party Politics* 0(0).

Cowburn, Mike and Michael Oswald. 2020. "Legislator Adoption of the Fake News Label: Ideological Differences in Republican Representative Use on Twitter." *The Forum* 18(3).

Cowburn, Mike and Rebecca Kerr. 2023. "Inclusivity and Decentralisation of Candidate Selectorates: Factional Consequences for Centre-Left Parties in England, Germany, and the United States." *Political Research Quarterly* 76(1):292–307.

DeBacker, Jason Matthew. 2008. "Flip-Flopping: Ideological Adjustment Costs in the United States Senate.".

DeCrescenzo, Michael G. 2020. "Do Primaries Work? Constituent Ideology and Congressional Nominations." Doctoral Thesis University of Wisconson-Madison.

Devlin, Jacob, Ming-Wei Chang, Kenton Lee and Kristina Toutanova. 2019. "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding.". arXiv:1810.04805.

Dickey, David A. and Wayne A. Fuller. 1979. "Distribution of the Estimators for Autoregressive Time Series with a Unit Root." *Journal of the American Statistical Association* 74(366a):427–431.

Dunn, Amina. 2021. "Two-Thirds of Republicans want Trump to Retain Major Political Role; 44% Want him to Run Again in 2024.".
URL: *www.pewresearch.org/fact-tank/2021/10/06/two-thirds-of-republicans-want-trump-to-retain-major-political-role-44-want-him-to-run-again-in-2024/*

Ensley, Michael J. 2009. "Individual Campaign Contributions and Candidate Ideology." *Public Choice* 138(1/2):221–238.

Fiorina, Morris P., Samuel J. Abrams and Jeremy C. Pope. 2005. *Culture War? The Myth of a Polarized America*. Longman.

Gentzkow, Matthew and Jesse M. Shapiro. 2010. "What Drives Media Slant? Evidence From U.S. Daily Newspapers." *Econometrica* 78(1):35–71.

Gilens, Martin. 2009. "Preference Gaps and Inequality in Representation." *Political Science &*

Politics 42(2):335–341.

Goet, Niels D. 2019. "Measuring Polarization with Text Analysis: Evidence from the UK House of Commons, 1811–2015." *Political Analysis* 27(4):518–539.

Green, Jon, Jared Edgerton, Daniel Naftel, Kelsey Shoub and Skyler J. Cranmer. 2020. "Elusive consensus: Polarization in Elite Communication on the COVID-19 Pandemic." *Science Advances* 6(28).

Greenacre, Michael J. 2007. *Correspondence Analysis in Practice*. Interdisciplinary Statistics Series 2nd ed. Chapman & Hall/CRC.

Grimmer, Justin and Brandon M. Stewart. 2013. "Text as Data: The Promise and Pitfalls of Automatic Content Analysis Methods for Political Texts." *Political Analysis* 21(3):267–297.

Grossmann, Matthew and David A. Hopkins. 2016. *Asymmetric Politics: Ideological Republicans and Group Interest Democrats*. Oxford University Press.

Hacker, Jacob S. and Paul Pierson. 2006. *Off Center: The Republican Revolution and the Erosion of American Democracy*. Yale University Press.

Hall, Andrew B. 2015. "What Happens When Extremists Win Primaries?" *American Political Science Review* 109(1):18–42.

Hassell, Hans J. G. 2018. *The Party's Primary: Control of Congressional Nominations*. Cambridge University Press.

Hawkins, Stephen, Daniel Yudkin, Míriam Juan-Torres and Tim Dixon. 2018. "Hidden Tribes: A Study of America's Polarized Landscape.".
URL: *www.hiddentribes.us/media/qfpekz4g/hidden$_t$ribes$_r$eport.pdf*

Hazan, Reuven Y. and Gideon Rahat. 2010. *Democracy within Parties: Candidate Selection Methods and Their Political Consequences*. Oxford University Press.

Heft, Annett, Curd B. Knüpfer, Susanne Reinhardt and Eva Mayerhöffer. 2021. "Toward a Transnational Information Ecology on the Right? Hyperlink Networking among Right-Wing Digital News Sites in Europe and the United States." *The International Journal of Press/Politics* 26(2):484–504.

Hill, Seth J. and Gregory A. Huber. 2017. "Representativeness and Motivations of the Contemporary Donorate: Results from Merged Survey and Administrative Records." *Political Behavior* 39(1):3–29.

Hirano, Shigeo and James M. Snyder. 2019. *Primary Elections in the United States*. Cambridge University Press.

Hopkins, Daniel J. and Hans Noel. 2021. "Trump and the Shifting Meaning of "Conservative": Using Activists' Pairwise Comparisons to Measure Politicians' Perceived Ideologies." *American Political Science Review* pp. 1–8.

Jacobs, Lawrence R. and Robert Y. Shapiro. 2000. *Politicians Don't Pander: Political Manipulation and the Loss of Democratic Responsiveness*. Chicago, IL: University of Chicago Press.

Kamarck, Elaine C. 2014. "Increasing Turnout in Congressional Primaries." *Brookings* .
URL: *www.brookings.edu/wp-content/uploads/2016/06/KamarckIncreasing-Turnout-in-Congressional-Primaries72614.pdf*

Kearney, Michael W. 2018. "rtweet: An Implementation of Calls Designed to Collect and Organize Twitter Data via Twitter's REST and Stream Application Program Interfaces.".

Keisling, Phil. 2010. "To Reduce Partisanship, Get Rid of Partisans." *The New York Times* .
URL: *www.nytimes.com/2010/03/22/opinion/22keisling.html*

Kujala, Jordan. 2019. "Donors, Primary Elections, and Polarization in the United States." *American Journal of Political Science* 00(0):1–16.

La Raja, Raymond J. and Brian F. Schaffner. 2015. *Campaign Finance and Political Polarization: When Purists Prevail*. University of Michigan Press.

Lauderdale, Benjamin E. and Alexander Herzog. 2016. "Measuring Political Positions from Legislative Speech." *Political Analysis* 24(3):374–394.

Laver, Michael, Kenneth Benoit and John Garry. 2003. "Extracting Policy Positions from Political Texts Using Words as Data." *American Political Science Review* 97(02).

Layman, Geoffrey C., Thomas M. Carsey, John C. Green, Richard Herrera and Rosalyn Cooperman. 2010. "Activists and Conflict Extension in American Party Politics." *The American Political Science Review* 104(2):324–346.

Lewis, Jeff. 2022. "Why are Ocasio-Cortez, Omar, Pressley, and Talib estimated to be moderates

by NOMINATE?".
URL: *https://voteview.com/articles/Ocasio-Cortez$_O$mar$_P$ressley$_T$laib*

Linden, Ariel. 2015. "Conducting Interrupted Time-series Analysis for Single- and Multiple-group Comparisons." *The Stata Journal* 15(2):480–500.

Lippmann, Walter. 1955. *The Public Philosophy*. New Brunswick: Routledge.

Maestas, Cherie D. and Melissa Stewart. 2012. Recruitment and Candidacy. In *New Directions in Congressional Politics*, ed. Jamie L. Carson. New York: Routledge pp. 25–44.

Manento, Cory. 2019. "Party Crashers: Interest Groups as a Latent Threat to Party Networks in Congressional Primaries." *Party Politics* 27(1):1–12.

Mann, Thomas E. and Norman J. Ornstein. 2012. *It's Even Worse Than It Looks: How the American Constitutional System Collided with the New Politics of Extremism*. New York, NY: Basic Books.

Martin, Gregory J. and Joshua McCrain. 2019. "Local News and National Politics." *American Political Science Review* 113(2):372–384.

Masket, Seth. 2009. *No Middle Ground: How Informal Party Organizations Control Nominations and Polarize Legislatures*. Ann Arbor: University of Michigan Press.

Masket, Seth. 2020. *Learning from Loss: The Democrats, 2016–2020*. New York: Cambridge University Press.

May, John D. 1973. "Opinion Structure of Political Parties: The Special Law of Curvilinear Disparity." *Political Studies* 21(2):135–151.

McCarty, Nolan, Keith T. Poole and Howard Rosenthal. 2006. *Polarized America: The Dance of Ideology and Unequal Riches*. MIT Press.

Meyer, Thomas M. and Markus Wagner. 2019. "It Sounds Like They are Moving: Understanding and Modeling Emphasis-Based Policy Change." *Political Science Research and Methods* 7(4):757–774.

Norrander, Barbara. 1989. "Ideological Representativeness of Presidential Primary Voters." *American Journal of Political Science* 33(3):570–587.

Pickup, Mark and Paul M. Kellstedt. 2022. "Balance as a Pre-Estimation Test for Time Series Analysis." *Political Analysis* pp. 1–10.

Pierson, Paul and Eric Schickler. 2020. "Madison's Constitution Under Stress: A Developmental Analysis of Political Polarization." *Annual Review of Political Science* 23(1):37–58.

Poole, Keith T. and Howard Rosenthal. 1985. "A Spatial Model for Legislative Roll Call Analysis." *American Journal of Political Science* 29(2):357.

Report, Cook Political. 2017. "Cook PVI.".
URL: *www.cookpolitical.com/pvi-0*

Russell, Annelise. 2018. "U.S. Senators on Twitter: Asymmetric Party Rhetoric in 140 Characters." *American Politics Research* 46(4):695–723.

Sältzer, Marius. 2020. "Finding the Bird's Wings: Dimensions of Factional Conflict on Twitter." *Party Politics* .

Saunders, Kyle L. and Alan I. Abramowitz. 2004. "Ideological Realignment and Active Partisans in the American Electorate." *American Politics Research* 32(3):285–309.

Schlozman, Daniel and Sam Rosenfeld. 2019. The Hollow Parties. In *Can America Govern Itself?*, ed. Frances E. Lee and Nolan McCarty. SSRC Anxieties of Democracy Cambridge: Cambridge University Press pp. 120–152.

Schumer, Charles. 2014. "End Partisan Primaries, Save America." *The New York Times* .
URL: *www.nytimes.com/2014/07/22/opinion/charles-schumer-adopt-the-open-primary.html*

Sides, John, Chris Tausanovitch, Lynn Vavreck and Christopher Warshaw. 2020. "On the Representativeness of Primary Electorates." *British Journal of Political Science* 50(2):1–9.

Slapin, Jonathan B. and Sven-Oliver Proksch. 2008. "A Scaling Model for Estimating Time-Series Party Positions from Texts." *American Journal of Political Science* 52(3):705–722.

Theriault, Sean M. 2013. *The Gingrich Senators: The Roots of Partisan Warfare in Congress*. New York: Oxford University Press.

Thomsen, Danielle M. 2017. *Opting Out of Congress: Partisan Polarization and the Decline of Moderate Candidates*. Cambridge: Cambridge University Press.

Thomsen, Danielle M. 2021. Competition in Congressional Primaries. In *Annual Conference of the Midwest Political Science Association (MPSA)*. Chicago (held virtually): .

Warner, Zach. 2019. "Conditional Relationships in Dynamic Models.".
URL: *zachwarner.net/download/Warner-Conditional-Relationships.pdf*
Wojcik, Stefan and Adam Hughes. 2019. "Sizing Up Twitter Users.".
URL: *www.pewresearch.org/internet/2019/04/24/sizing-up-twitter-users*

7 Supplementary Information

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

Table 3: Descriptive Statistics

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Movement	886	−0.055	0.883	−1	−1	1	1
NOMINATE	283	0.016	0.454	−0.747	−0.396	0.460	0.883
Position Before	886	−0.083	0.320	−0.862	−0.363	0.200	0.761
Position After	886	−0.102	0.305	−0.864	−0.366	0.158	1.395
Candidates (Contested)	1772	3.481	2.350	2.000	2.000	5.000	19.000

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	0.946	0.969	0.911	0.939
2	0.933	0.955	0.900	0.926
3	0.926	0.948	0.912	0.929
4	0.919	0.897	0.935	0.915
5	0.909	0.886	0.907	0.897
6	0.946	0.941	0.941	0.941
7	0.896	0.889	0.907	0.898
8	0.936	0.938	0.931	0.934
9	0.936	0.932	0.925	0.929
10	0.909	0.890	0.921	0.905

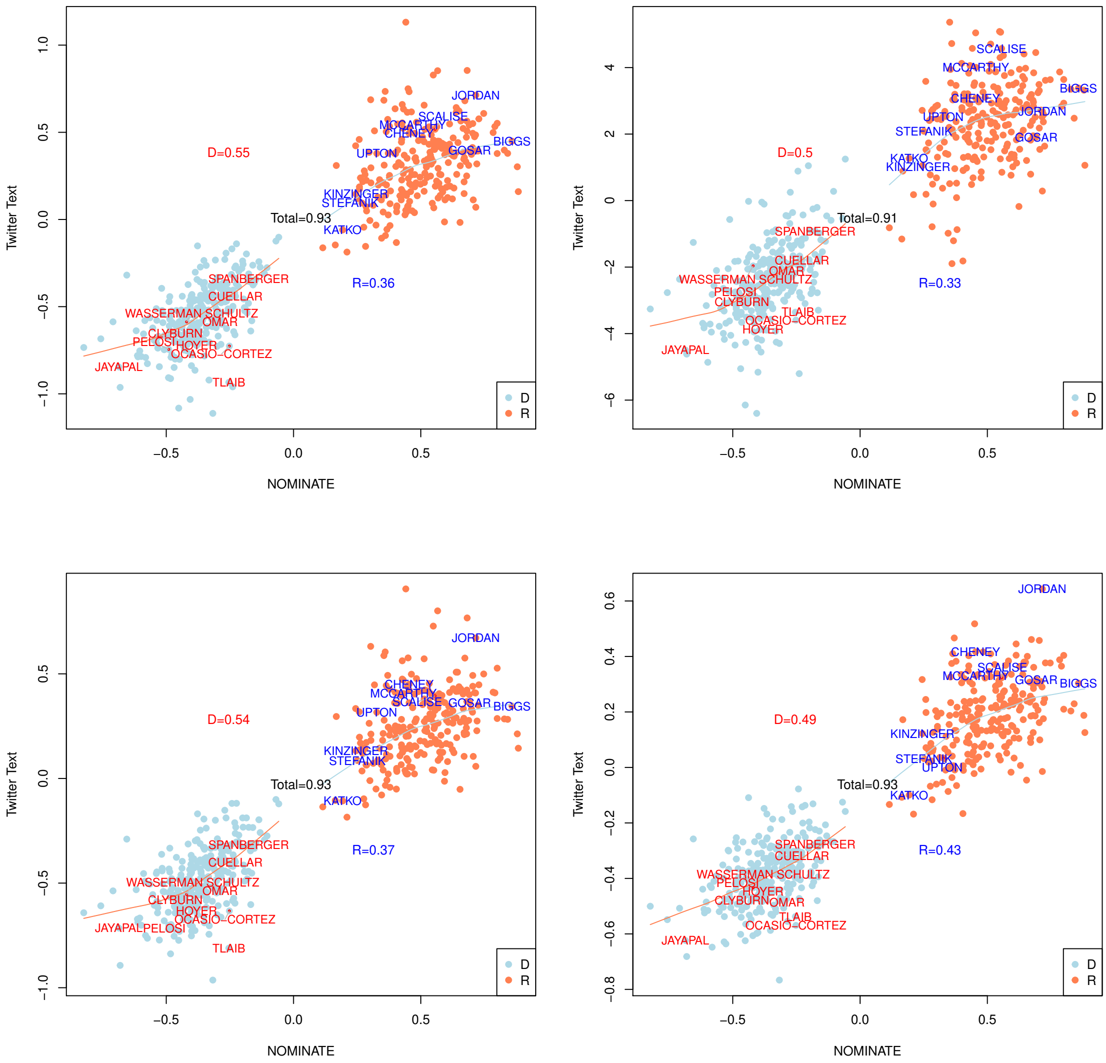

Figure 5: Comparative Validity of Alternative Measures

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 places 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.

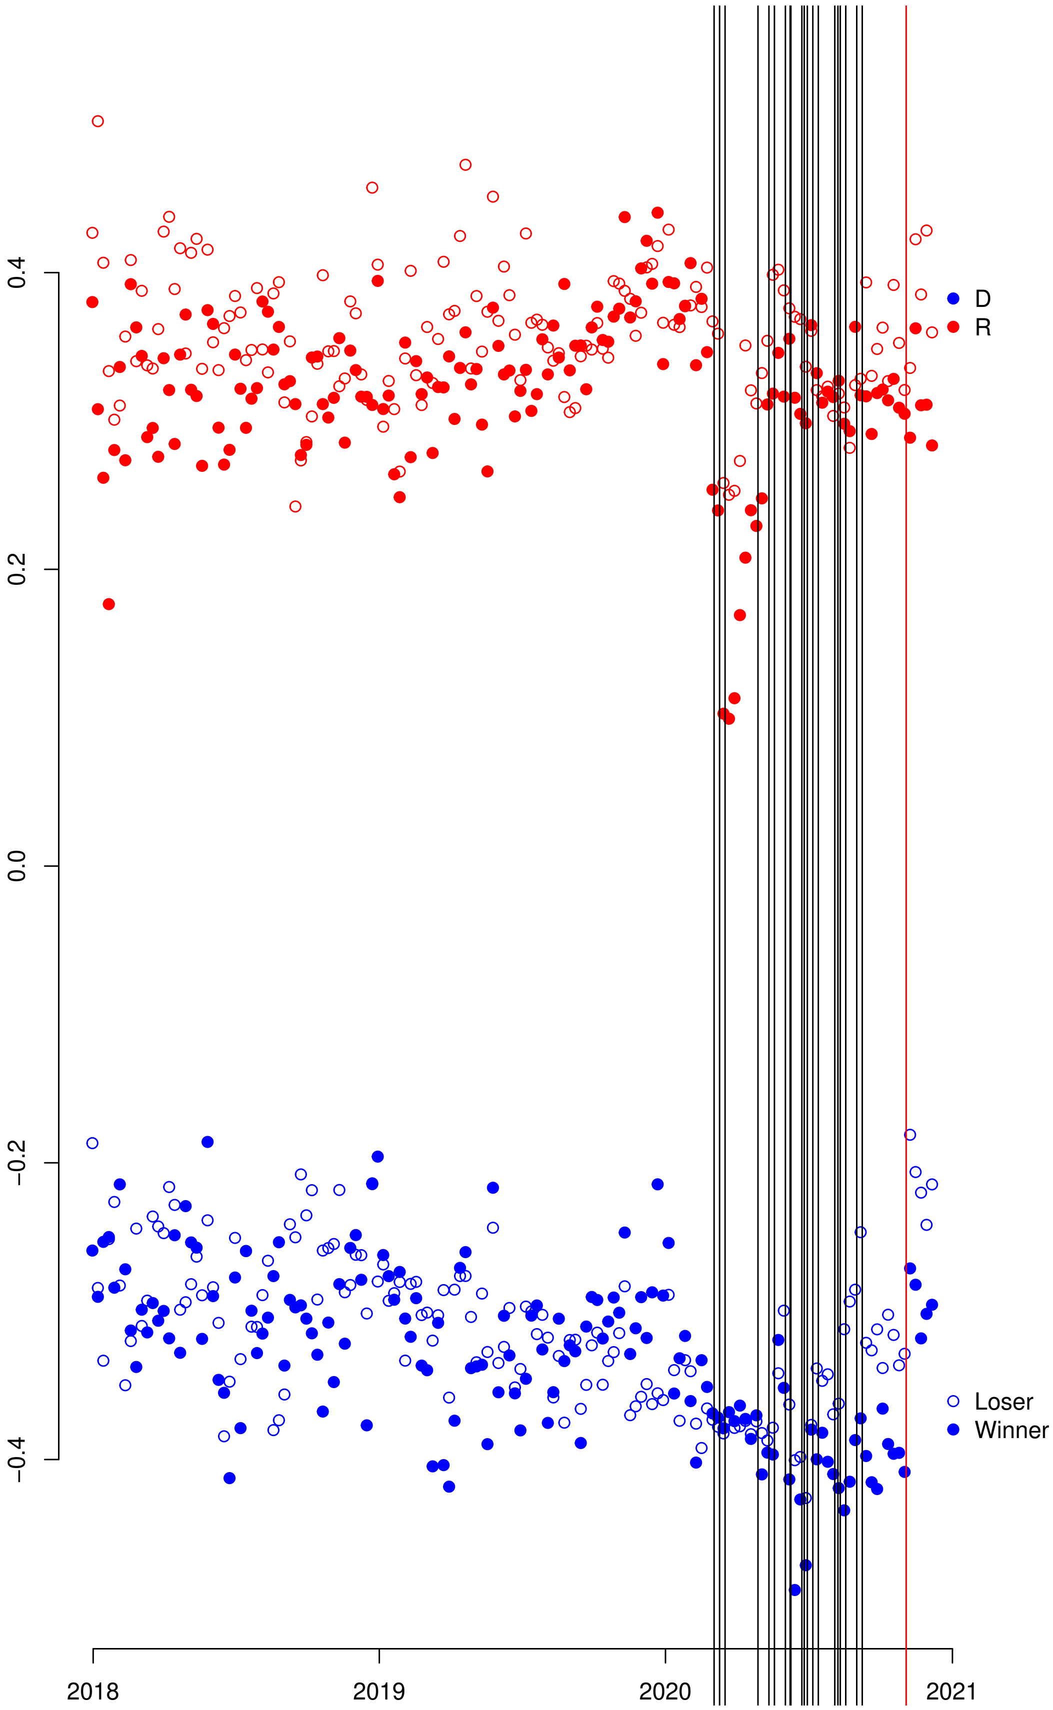

Figure 6: Natural Time

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.

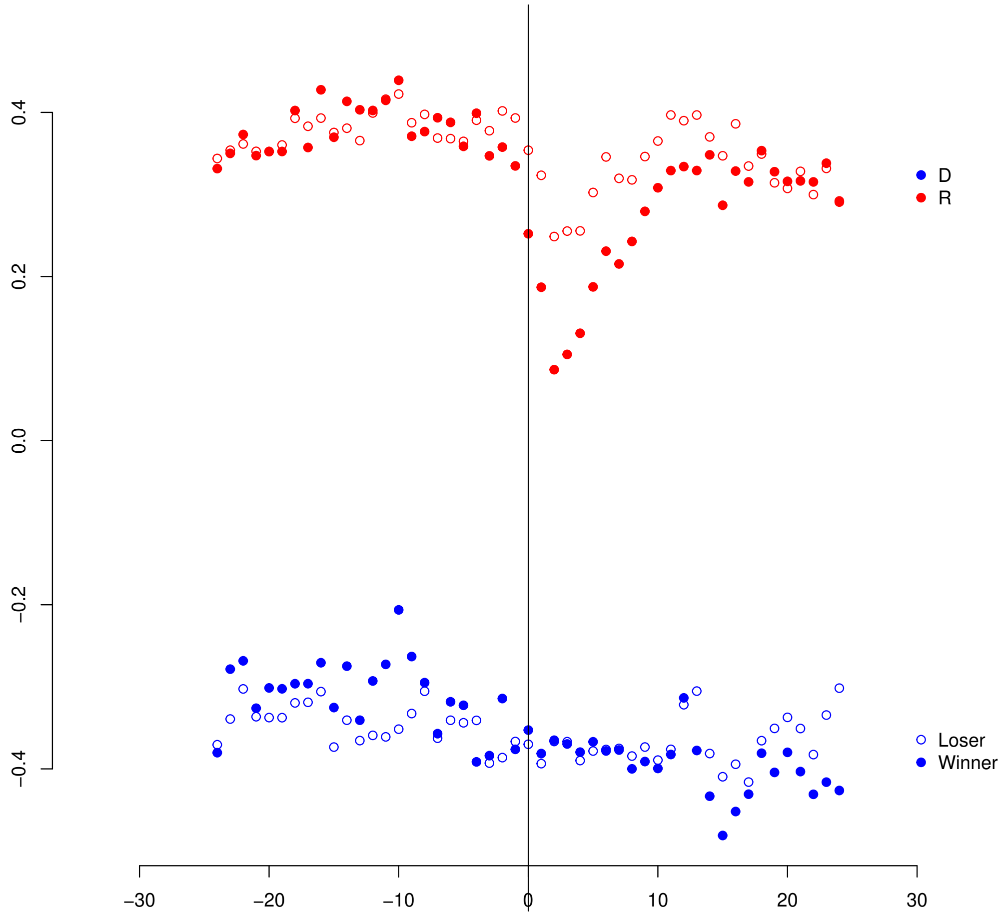

Figure 7: Placebo Test

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 (T_t)	-0.002^{**}
	(0.001)
Post-Pseudo Primary (X_t)	-0.029
	(0.018)
Post- Pseudo Primary : Time (X_tT_t)	-0.0003
	(0.001)
Loser (Z)	-0.027
	(0.018)
Loser : Time (ZT_t)	0.001
	(0.001)
Loser : Post-Pseudo Primary (ZX_t)	0.006
	(0.026)
Loser : Post-Pseudo Primary : Time (ZX_tT_t)	0.003^{*}
	(0.002)
Intercept	-0.337^{***}
	(0.012)
N	98
R^2	0.549
Adjusted R^2	0.514
Residual Std. Error	0.032 (df = 90)
F Statistic	15.628^{***} (df = 7; 90)

Newey-West Standard Errors Shown in Parentheses
$^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$

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.

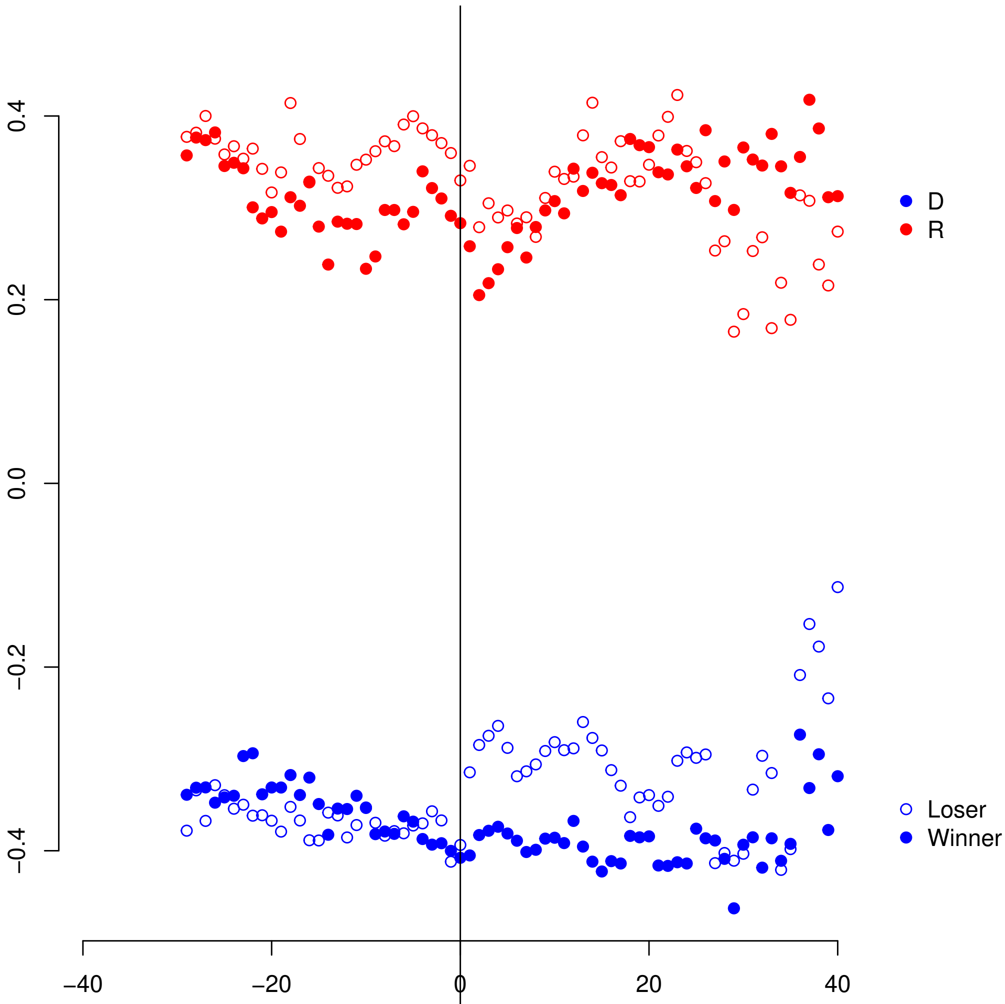

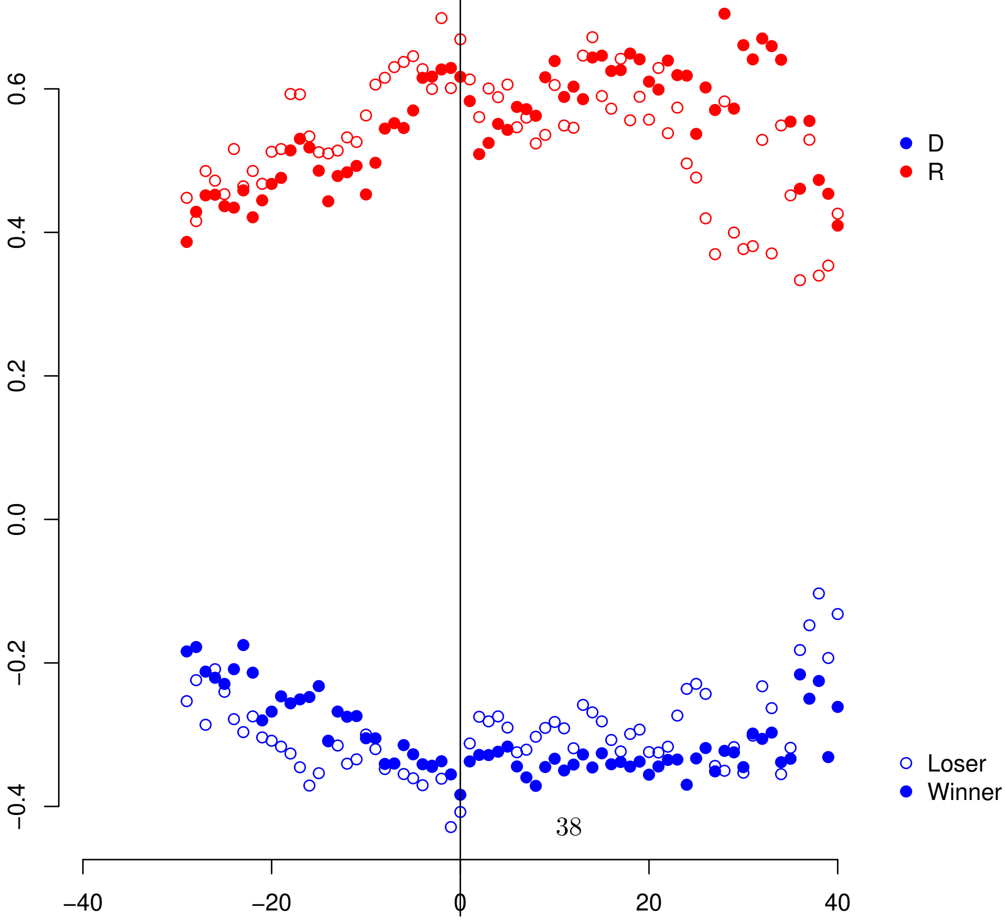

Figure 8: Naïve Bayes Approach & Canonical Correspondence Analysis Comparison

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.

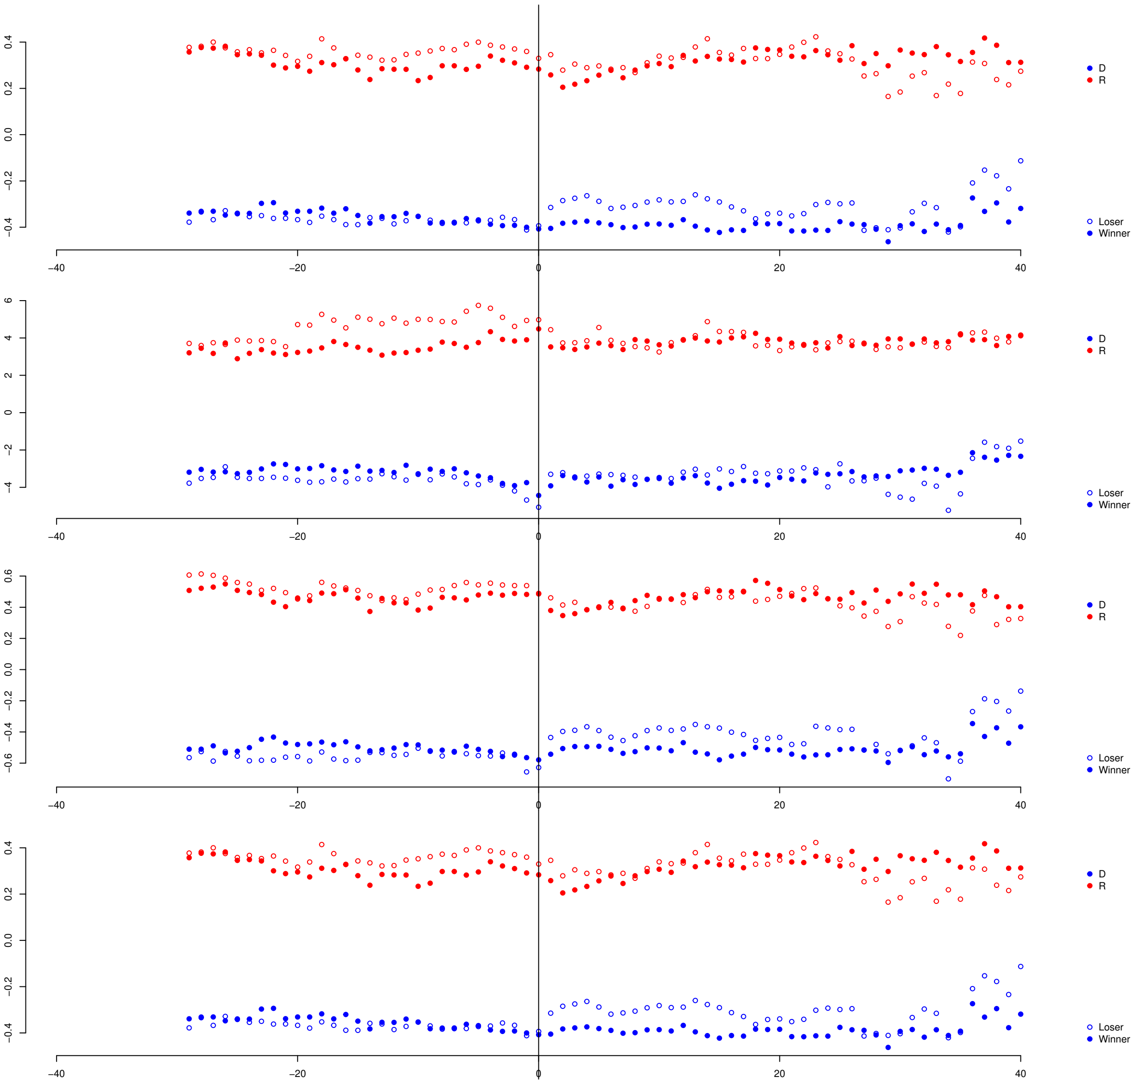

Figure 9: Main Analysis with Different Terms

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.

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	−0.043**	−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	0.148 (df = 882)	0.148 (df = 867)
F Statistic	16.088*** (df = 3; 882)	15.821*** (df = 3; 867)

*p < 0.1; **p < 0.05; ***p < 0.01

In Table 7 we demonstrate the robustness of our main individual results to three standard errors of movement.

Table 7: Individual Level Results (as Coefficient Plot)

	Democrats Absolute	Democrats Three Errors	Republicans Absolute	Republicans Three Errors
	(1)	(2)	(3)	(4)
Loser	−0.001	−0.007	−0.008***	−0.029***
	(0.001)	(0.008)	(0.002)	(0.010)
District PVI	−0.041***	−0.284***	0.013	0.086
	(0.012)	(0.087)	(0.023)	(0.097)
Incumbent	0.042***	0.277***	−0.031**	−0.080
	(0.007)	(0.051)	(0.013)	(0.057)
Intercept	0.039***	0.315***	0.021	0.152*
	(0.011)	(0.079)	(0.020)	(0.083)
N	472	472	414	414
R^2	0.094	0.099	0.033	0.025
Adjusted R^2	0.089	0.094	0.026	0.018
Residual Std. Error	0.102 (df = 468)	0.742 (df = 468)	0.184 (df = 410)	0.784 (df = 410)
F Statistic	16.272*** (df = 3; 468)	17.209*** (df = 3; 468)	4.623*** (df = 3; 410)	3.523** (df = 3; 410)

*p < 0.1; **p < 0.05; ***p < 0.01

In Table 8 we demonstrate the robustness of our individual results with the removal of the additional controls in our main analysis.

Table 8: Individual Robustness without Controls

	Absolute Movement (No Controls)	Three Errors (No Controls)	Absolute Movement (w/Controls)	Three Errors (w/Controls)
	(1)	(2)	(3)	(4)
Loser	0.057***	0.438***	0.043***	0.371***
	(0.014)	(0.071)	(0.015)	(0.076)
Republican	−0.079***	−0.359***	−0.073***	−0.339***
	(0.014)	(0.075)	(0.014)	(0.075)
District PVI			−0.003***	−0.010*
			(0.001)	(0.006)
Incumbent			−0.028**	−0.145**
			(0.012)	(0.061)
Loser : Republican	−0.043**	−0.324***	−0.037*	−0.298***
	(0.020)	(0.104)	(0.020)	(0.104)
Intercept	0.038***	0.252***	0.040***	0.263***
	(0.010)	(0.053)	(0.010)	(0.053)
N	886	886	886	886
R^2	0.052	0.056	0.066	0.064
Adjusted R^2	0.049	0.053	0.061	0.059
Residual Std. Error	0.148 (df = 882)	0.769 (df = 882)	0.147 (df = 880)	0.766 (df = 880)
F Statistic	16.088*** (df = 3; 882)	17.349*** (df = 3; 882)	12.407*** (df = 5; 880)	12.104*** (df = 5; 880)

*p < 0.1; **p < 0.05; ***p < 0.01

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 (T_t)	-0.003***	–0.001
	(0.001)	(0.001)
Post-Primary (X_t)	0.020*	–0.040***
	(0.011)	(0.015)
Post-Primary : Time ($X_t T_t$)	0.000*	-0.000
	(0.001)	(0.000)
Loser (Z_i)	0.024***	0.057***
	(0.009)	(0.018)
Loser : Time ($Z_i T_t$)	0.002***	0.000
	(0.001)	(0.001)
Loser : Post-Primary ($Z_i X_t$)	0.055***	–0.018
	(0.020)	(0.025)
Loser : Post-Primary : Time ($Z_i X_t T_t$)	0.002***	0.001
	(0.001)	(0.001)
Policy Tweets	–0.003***	–0.003
	(0.001)	(0.001)
Constant	–0.342***	0.211***
	(0.026)	(0.038)
N	102	102
R^2	0.851	0.610
Adjusted R^2	0.839	0.576
Residual Std. Error	0.016	0.025
F Statistic	66.578***	18.157***

Newey-West Standard Errors Shown in Parentheses
*p < .1; **p < .05; ***p < .01

As a further robustness check, we also validate our measure against Hopkins and Noel (2021)'s pair-wise 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.

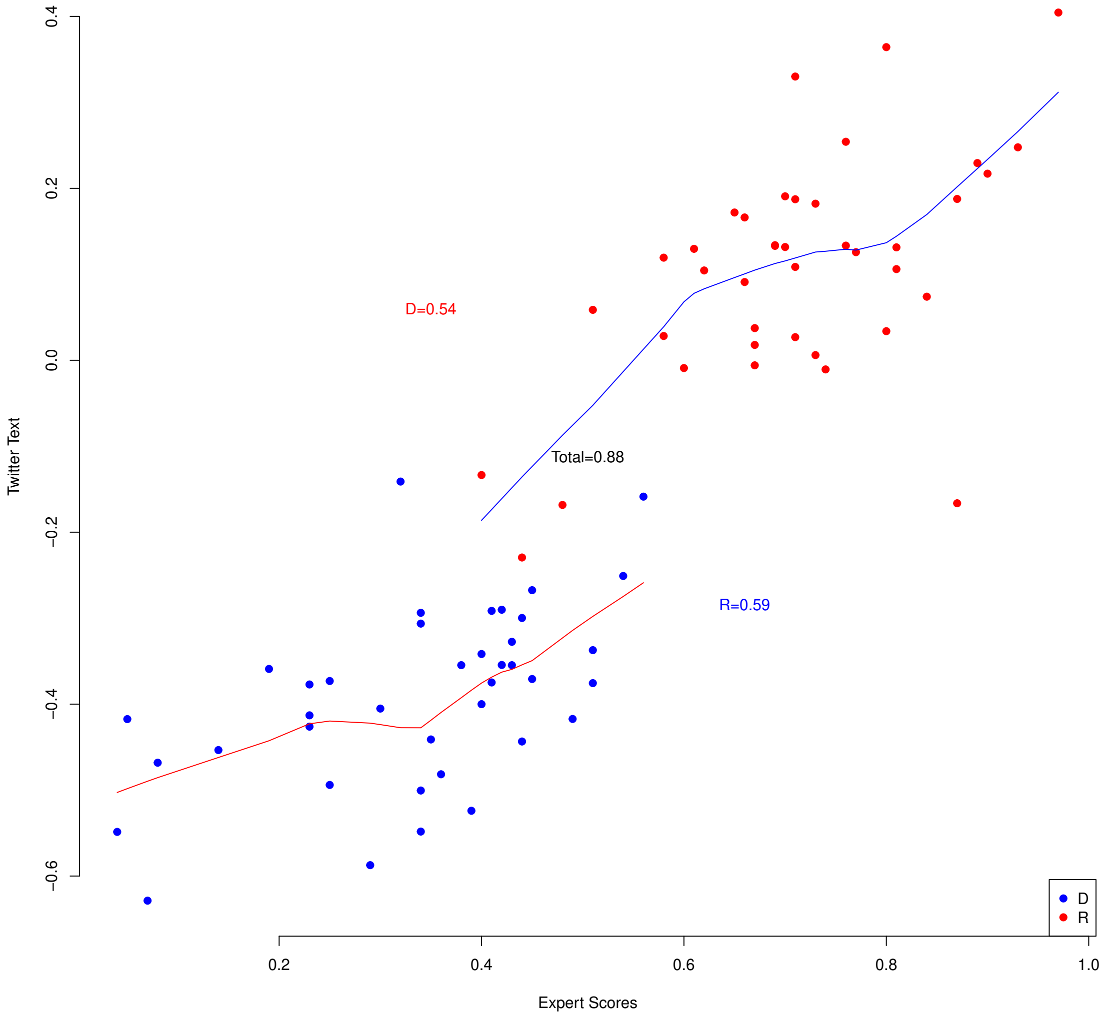

Figure 10: Validation Against Hopkins & Noel Pairwise Activist Scores

We also validate our model against Barberá's 2015 Follower Network in Figure 11.

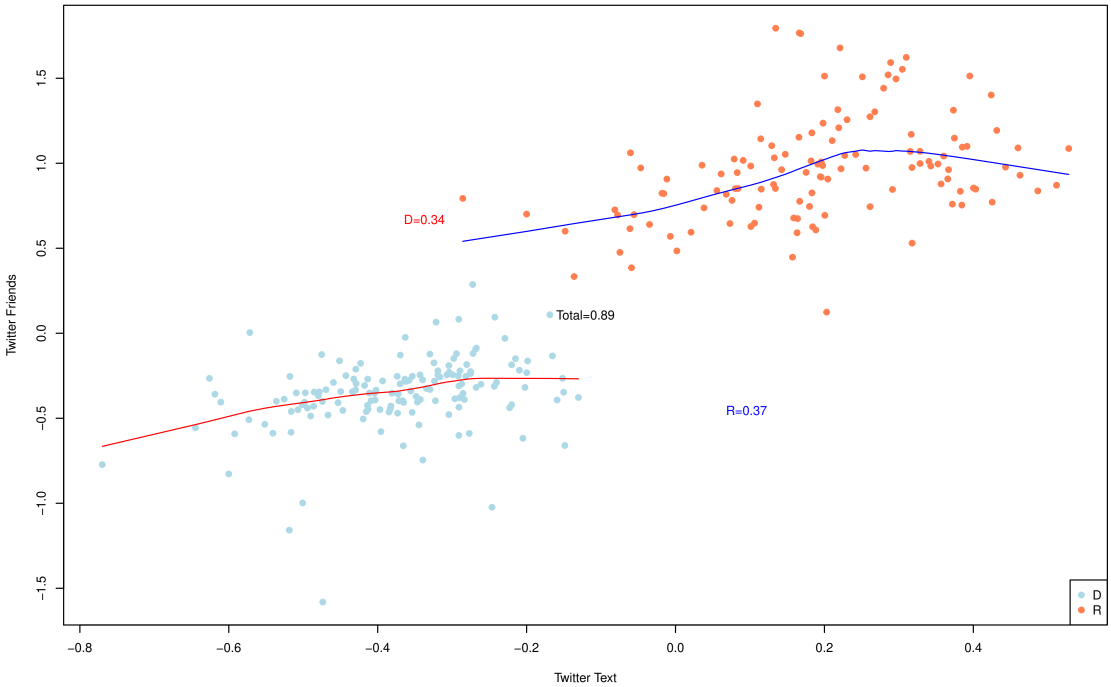

Figure 11: Validation Against Barberá's Follower Network

In Table 10 we report the results of our Dickey-Fuller tests for each of our four dependent variables in our main analysis. 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 10: Dickey-Fuller Tests for Unit Root in Dependent Variable

Group	Republican Losers	Republican Winners	Democratic Losers	Democratic Winners
Test Statistic	-6.137	-4.372	-3.762	-5.270
N	176	176	176	176
1% Critical Value	-4.015	-4.015	-4.015	-4.015
5% Critical Value	-3.440	-3.440	-3.440	-3.440
10% Critical Value	-3.140	-3.140	-3.140	-3.140
p-value	0.0000	0.0024	0.0186	0.0001

In Table 11 we repeat our Dickey-Fuller tests for each of our four dependent variables in our policy tweets subset of data. As in our full dataset, all four groups are stationary.

In addition, we test that our estimated errors (residuals) are white noise by running Dickey-Fuller tests on the residuals.

Table 11: Policy Tweets: Dickey-Fuller Tests for Unit Root in Dependent Variable

Group	Republican Losers	Republican Winners	Democratic Losers	Democratic Winners
Test Statistic	-6.957	-4.817	-4.821	-5.757
N	176	176	176	176
1% Critical Value	-4.015	-4.015	-4.015	-4.015
5% Critical Value	-3.440	-3.440	-3.440	-3.440
10% Critical Value	-3.140	-3.140	-3.140	-3.140
p-value	0.0000	0.0004	0.0004	0.0000

Table 12: Dickey-Fuller Tests for Unit Root in Residuals

Group	Republican Losers	Republican Winners	Democratic Losers	Democratic Winners
Test Statistic	-7.021	-4.739	-5.242	-6.316
N	176	176	176	176
1% Critical Value	-4.015	-4.015	-4.015	-4.015
5% Critical Value	-3.440	-3.440	-3.440	-3.440
10% Critical Value	-3.140	-3.140	-3.140	-3.140
p-value	0.0000	0.0006	0.0001	0.0000

Table 13: Policy Tweets: Dickey-Fuller Tests for Unit Root in Residuals

Group	Republican Losers	Republican Winners	Democratic Losers	Democratic Winners
Test Statistic	-7.949	-4.997	-5.549	-6.314
N	176	176	176	176
1% Critical Value	-4.015	-4.015	-4.015	-4.015
5% Critical Value	-3.440	-3.440	-3.440	-3.440
10% Critical Value	-3.140	-3.140	-3.140	-3.140
p-value	0.0000	0.0002	0.0000	0.0000

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*.

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 14: Party Level Fixed Effects

	Position	
	Democrats	Republicans
	(1)	(2)
Loser	−0.003	0.047***
	(0.006)	(0.009)
Post-Primary	−0.039***	0.009
	(0.006)	(0.009)
Loser : Post-Primary	0.098***	−0.019
	(0.009)	(0.013)
Intercept	−0.345***	0.206***
	(0.004)	(0.007)
N	102	102
R^2	0.709	0.255
Adjusted R^2	0.701	0.232
Residual Std. Error (df = 98)	0.022	0.034
F Statistic (df = 3; 98)	79.777***	11.171***

*p < 0.1; **p < 0.05; ***p < 0.01

Table 15: Lagged DV As Additional Control: All Tweets

	All Tweets		Policy Tweets Only	
	Democrats	Republicans	Democrats	Republicans
Time (T_t)	−0.002***	−0.000	−0.002***	0.001
	(0.001)	(0.001)	(0.001)	(0.001)
Post-Primary (X_t)	0.009	−0.024*	0.015	−0.017
	(0.008)	(0.013)	(0.009)	(0.017)
Post-Primary : Time (X_tT_t)	0.002***	0.003***	0.001	0.001
	(0.001)	(0.001)	(0.001)	(0.001)
Loser (Z_i)	0.019**	0.030**	−0.0002	0.011
	(0.008)	(0.014)	(0.009)	(0.017)
Loser : Time (Z_iT_t)	0.002***	0.000	0.002**	−0.000
	(0.001)	(0.001)	(0.001)	(0.001)
Loser : Post-Primary (Z_iX_t)	0.049***	−0.001	0.030**	0.008
	(0.013)	(0.018)	(0.013)	(0.024)
Loser : Post-Primary : Time ($Z_iX_tT_t$)	−0.003***	−0.001	−0.002**	−0.001
	(0.001)	(0.001)	(0.001)	(0.002)
Lagged Position	0.420***	0.472***	0.323***	0.458***
	(0.082)	(0.093)	(0.089)	(0.092)
Intercept	−0.229***	0.102***	−0.310***	0.070***
	(0.031)	(0.020)	(0.041)	(0.016)
N	100	100	100	100
R^2	0.887	0.699	0.675	0.405
Adjusted R^2	0.877	0.673	0.647	0.352
Residual Std. Error (df = 91)	0.014	0.022	0.016	0.030
F Statistic (df = 8; 91)	88.850***	26.420***	23.664***	7.732***

*p < 0.1; **p < 0.05; ***p < 0.01