Compromise Against Extremism: Strategic Responsiveness of Legislators to Challengers in the U.S. House

Preliminary Draft

Felix Schönenberger*

January 2022

Abstract

Theoretically, public policy in representative democracies is determined by strategic interactions between candidates prior to elections. Empirically, it has remained unclear whether and how politicians compete on policies. This paper sheds light on these fundamental questions by identifying challengers as an important and empirically disregarded driver of legislative behavior and public policy. Leveraging "as good as" random assignment of incumbent lawmakers in the U.S. House to relatively extreme challengers generated by close primary elections of the sitting legislator's opponent party, I use a regression discontinuity design to evaluate changes in incumbent's roll call voting behavior in response to an extremist challenger. I find robust evidence that incumbents *do react* differentially to extreme as opposed to moderate competitors by *moderating* their policy position. Results support theory predicting *(partial) policy convergence* and *strategic complementarity* of candidates' policy positions.

Keywords: electoral competition, electoral incentives, legislator behavior, policy convergence
JEL Classification: D72, D78, D82, P16

*Felix Schönenberger is a PhD student at the Institute of Economics, Università della Svizzera Italiana, Switzerland. E-mail: felix.schoenenberger@usi.ch.
I am indebted to my supervisor Patricia Funk for her insightful comments.

1 Introduction

Electoral competition is at the heart of representative democarcy. Yet, we know surprisingly little about how politicians compete against each other. At least since Hotelling (1929) and Downs (1957), political competition between candidates is at the core of political economy models: policies are the outcome of strategic interactions between candidates prior to elections. However, theoretical models of electoral competition sharply differ on one key aspect of strategic interaction, and, consequently, on their predicitons on legislative responsiveness of incumbent politicians to challengers.

Models in the tradition of Osborne and Slivinski (1996) and Besley and Coate (1997) assume that citizen-candidates cannot credibly commit to policy platforms. Instead, citizen-candidates of different ideology strategically decide whether or not to run for public office. Once elected, politicians implement their preferred policy throughout their term in office. The main prediciton of these models is *incumbent policy persistence*: Legislators do not adjust policymaking in any strategic way, and their voting records are irresponsive to challengers' policy position. On the other hand, models in the Downsian tradition assume that candidates compete on policy platforms to which they can credibly commit. Political competition leads to *(partial) policy convergence*, where re-election seeking legislators are predicted to strategically adjust policymaking in the direction of their challengers position, etiher because of purely opportunistic motives or because of policy motivation, i.e. to prevent election of an "extremist" candidate from the opposing party with a policy platform distant from the incumbent's ideal point.

This paper contributes to this debate by asking how legislators react to challengers. While theoretical work on strategic aspects of policymaking has abounded (see, e.g., Duggan and Martinelli (2017) for a recent review), direct empirical evidence on whether and how challengers affect legislator behavior is extremely scant.[1] Of course, credibly testing for incumbent responsiveness to challenger positions requires to overcome the obstacle that incumbent behavior and challenger position are jointly determined for two reasons: First, the very strategic nature of political competition implies that candidates' policy platforms are interdependent, and therefore subject to simultaneity bias. Second, incumbents and challengers compete for votes of the same electorate whose preferences and unobservable characteristics affect the political stands of both candidates, making any naive correlations prone to omitted variable bias.

This is, to the best of my knowledge, the first paper that overcomes both of these obstacles and provides a credibly causal effect of extremist challengers on incumbents' voting behavior. Following Hall (2015), I exploit the fact that in most U.S. sates, each of the two relevant parties' candidates are nominated in primary elections. Using a regression discontinuity design, I leverage plausibly exogenous variation from close primary elections in which an extremist narrowly wins

[1] The to date most comprehensive study of relative candidate positioning in U.S. House Elections (Ansolabehere et al., 2001) is correlational and essentially descriptive.

over a more moderate primary candidate, where relative extremism is inferred from pre-primary campaign donations. I thus examine the post-primary voting records of re-election seeking U.S. House incumbents in their last four months before the general elections, depending on whether their opposing party's primary electorate "as good as" randomly nominated a relative extremist as a challenger.

I find that the nomination of an extremist challenger causes a statistically sigificant and economically meaningful moderation in the incumbent's voting behavror. Based on the most conservative estimates, incumbents competing against a barely nominated extremist challenger are 5 percentage points less likely to vote in party-line compared to otherwise identical incumbents facing a relatively moderate challenger. Similarly, incumbents challenged by an extremist score 0.25 standard deviations higher on my DW-NOMINATE-based moderation measure, which – applying a 0.25 standard deviation change to DW-NOMINATE scores in the U.S. Senate (2019-2021) – approximately corresponds to a move from Mitch McConnell to Mitt Romney, or from Bernard Sanders to Chuck Schumer.

This paper is most closely related to the empirical study of partial policy convergence in the U.S. House, which has found mixed results. Perhaps the most influential work on this topic is that of Lee et al. (2004), who use a regression discontinuity design exploiting the fact that a narrow Democratic win in the election for the previous term, by virtue of the incumbency advantage, generates a sharp increase in the electoral strangth of the Democratic party in the election for the subsequent term. They find that changes in roll call behavior from the previous to the subsequent term are entirely explained by selection, concluding that voters do not "affect" policies in the sense that sitting legislators, once in office, are irresponsive to changes in electoral strength, as predicted by citizen-candidate models. In contrast, Jones and Walsh (2018) find that changes in the partisan composition of the electorate, as induced by redistricting, affect roll call voting behavior of U.S. legislators not only through selection, but also at the intensive margin: Increases in the Democratic vote share lead to a higher porbability of a Democratic win, and to sitting incumbents, if re-elected, voting more liberally.

While these divergent results may be due to different empirical approaches, both studies have in common that they evaluate legislator responsiveness to a revealed shift in voter preferences. Responsiveness to the electorate may or may not occur for strategic motives, [2] and tells little on whether the positions of candidates running in a given district diverge or converge.[3] The present paper adds to these studies by examining strategic responsiveness of legislators to challengers. This allows for a *direct* test of the hypothesis of partial policy convergence, for which we find

[2]Responsiveness to revealed shifts in voter preferences is also consistent with the view that well-intentioned legislators may have incomplete information and elections lead them to update their beliefs about their constituencies' interests (see, e.g., Matsusaka (1992))

[3]Conceivably, an increase in the Democratic (Republican) vote share induced by redistricting can cause a backlash among minorized Republicans (Democrats) who may nominate a more conservative (liberal) candidate, which leads to policy *divergence*.

strong empirical support: Legislators facing an extremist challenger moderate their voting behavior. Thus, incumbent's policy extremism is *decreasing* in challger extremism, suggesting that policy positions are *strategic complements* in the electoral game.

2 Empirical Strategy and Data

2.1 Setting and Identification Strategy

With a few exceptions, U.S. states require that the Republican and Democratic party hold separate primary elections, where the the plurality of voters nominates the party's candidate running for the district's seat. Primary elections take place several months before the general election. Thus, the primary electorate of the opponent party determines the challenger of re-election seeking incumbents who then have the opportunity to alter their roll call voting behavior, thereby committing to a strategically adjusted policy platform prior to Election Day.[4] For competitive primaries where at least two candidates run for nomination, we have precise knowledge of the assignment mechanism that determines whether the incumbent runs against a relatively extreme or moderate challenger, i.e. whether a Democratic (Republican) incumbent runs against a more or less conservative Republican (more or less liberal Democrat). The incumbent gets assigned to an extremist challenger if and only if the more extreme candidate gets a pluratiy of the vote in of the opponent party's primary election.

Using a sharp regression discontinuity design, I exploit that close primary elections of the incumbents' opponent party generate "as if " random assignment to an extremist instead of a more moderate challenger. Assuming that in close enough primaries, agents have at best "imprecise control" (Lee and Lemieux, 2010) over the nomination outcome, allows me to recover a local average treatment effect comparing post-primary roll call voting behavior of otherwise identical incumbents who only differ in whether the more extremist or more moderate potential challenger won nomination by a narrow margin. Notably, local randomization occurs at the district-level, which – given single-member districts – coincides with the incumbent-level. Thus, the design directly addresses the twofold identification challenge of i) simultaneity due to strategic candidates choosing their positions interdependently, and ii) omitted variable bias stemming from unobserved voter preferences that affect both candidates' policy stances. Of course, the design is only valid if the distribution of pre-determined incumbent-level and district-level characteristics is continuous around the cutoff. I provide evidence in support of this assumption in section 3.

Formally, I implement the design by estimating equations of the following form:

$$Y_{i(d)} = \alpha + \theta T_{i(d)} + \beta_1 X_{i(d)} + \beta_2 X_{i(d)} T_{i(d)} + [\beta_3 (X^2_{i(d)} + \beta_4 X^2_{i(d)} T_{i(d)}] + \epsilon_{i(d)} \quad (1)$$

[4]Re-election seeking incumbents almost never lose in primaries. In my data, only 6 out of 523 re-election seeking incumbents were defeated in primary election. In what follows, I do not consider incumbents who lost renomination. Since primaries of both parties are held on the same date, I do not worry about endogenous sample selection.

where the outcome $Y_{i(d)}$ denotes one of two alternative measures of district d's incumbent i's roll call voting moderation, as described in more detail in the data section 2.3. $T_{i(d)}$ is a dummy variable equal to one if incumbent i's oppenent party nominates the more extreme candidate of the top-two candidates as a challenger, and equal to zero if it nominates the more moderate candidate. Relative extremism of the oppoent party's candidates is inferred from campaign receipts, as discussed section 2.2. The assignment variable $X_{i(d)}$ is the extremist's top-two candidate primary vote share, normalized such that $T_{i(d)} = 1$ if $X_{i(d)} > 0$ and $T_{i(d)} = 0$ if $X_{i(d)} < 0$.

For estimation I follow Calonico et al. (2014) and Calonico et al. (2019), using a non-parametric approach with bias-adjusted confidence interval estimates and MSERD-optimal bandwidths, fitting local low-order polynomial regression (linear or quadratic) on both sides of the cutoff. Given my relatively small sample size of 517 obseravtions, my preferred specification includes quadratic polynomials of the running variabe to reduce finite-sample approximation error, although I show that all findings are robust to local-linear specifications.

2.2 Measuring Challenger Extremism

Empirical research on the U.S, Congress conventionally useses DW-NOMINATE scores (Poole and Rosenthal, 1997) to scale legislator's roll call voting behavior on a liberal-conservative dimension with support ranging from -1 (very lberal) to 1 (very conservative). Measures of the policy position of non-incumbent candidates, who never cast a vote in Congress, are more difficult to obtain. They need to be inferred indirectly, either from candidate surveys (Ansolabehere et al., 2001) or from campaign contributions (Bonica, 2014; Bonica, 2018). While candidate surveys provide very limited coverage, especially of candidates that lost nomination, the most popular donation-based scalings rely also on post-primary, hence potentially endogenous donations. I therefore follow the approach of Hall and Snyder (2015), and use pre-primary contributions from common donors to bridge roll call based scalings of incumbents to non-incumbent candidates. The underlying intution is straghtforward and in line with the proximity assumption behind spatial voting models. Similarly to proximity voters electing policies closest to their ideal point, spatial giving posits that donors prefer donating to ideologically proximate candidates.

Hall-Snyder donation-based candidate scores are then calculated in two steps. First, incumbent's DW-NOMINATE scores are mapped to their donors:

$$DonorScore_j = \frac{\sum_i Contribution_{ij} Nominate_i}{\sum_i Contribution_{ij}}$$

where donor j's score is the contribution-weighted average DW-NOMINATE of the incumbents i to whom donor j contributed.

In a second step, donor scores are mapped to non-incumbent candidates:

$$CandidateScore_k = \frac{\sum_k Contribution_{jk} DonorScore_k}{\sum_k Contribution_{jk}}$$

where candidate k's score is the contribution-weighted average donor score of all donors j that contributed to k.

I calculate these scores from transaction-level data of the Federal Election Commission (FEC) for U.S. House elections from 1980 to 2018, whereby I impose a number of restrictions. First, I exclude several types of transactions, including loans, transfer payments, and contributions against a candidate. Second, for each candidate k, I do not consider transactions made after the primary election date in both steps of calculating k's candidate score. Third, I attenuate measurement error by excluding donors (candidates) who contribute (receive) to (from) fewer than 5 distinct candidates (donors). This procedure leaves me with a sample of 517 competitive primaries with valid candidate scores for both top-two candidates running for nomination against a re-election seeking incumbent of the opposing party.

Given that contribution-based scores are weighted averages of DW-NOMINATE scores, contribution-based scores are also bounded between -1 and 1, where increasing values indivate more conservative candidates. It is then straghtforward to infer the more extreme of the top-two primary candidates. For Republican primaries, I define the relative extremist as the more conservative candidate with a score closer to 1, while for Democratic primaries I classify the more lineral candidate with a score close to -1 as the relative extremist.

FIGURE 1: DONATIONS-BASED SCORES AND DW-NOMINATE SCORES

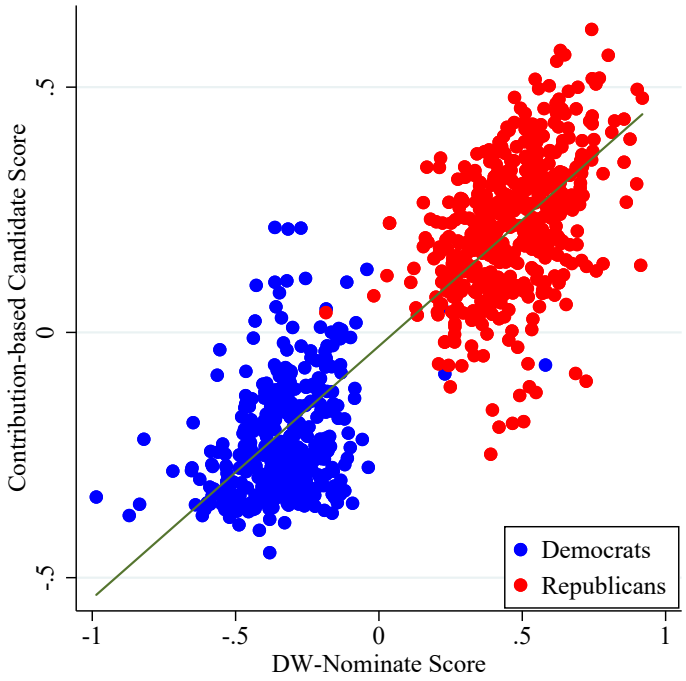

Notes: The figure plots incumbents' cotribution-based candidate scores against their roll call based DW-NOMINATE scores.

Donation-based scores can be validated by comparing them to roll call based DW-Nominate scores for candidates who eventually take office. As shown in Figure 1, donation-based scores and observed DW-NOMINATE scores are strongly, although not perfectly correlated ($r = 0.87$).[5]. Clearly, contribution-based scores measure candidates' postions with some error. It is worth noting, however, that I only use transactions made *before* the primary election. Thus, in my application of an RD design focusing on close primary elections, this results in classical measurement error that biases effects toward zero because I may missclassify moderate primary candidates as extremists. I will therefore interpret my estimates as lower bounds of the true effect. I probe robustness of this claim in section 2.2. where estimate equation 1 on different samples, successiviely exluding primaries where top-two candidates have relativeley close candidate scores.

2.3 Measuring Post-Primary Incumbent Moderation

We are interested in whether and how incumbents, prior to elections, strategically adjust their policy position on the liberal-conservative dimension in response to more or less extreme challenger, as revealed by the incumbent's post-primary roll call voting behavior. While it is natural to rely on DW-NOMINATE scores for measuring incumbent's roll call voting behavior on liberal-

[5]Within-party correaltions are weaker ($r = 0.36$ for Republicans, $r = 0.44$ for Democrats)

conservative scale, it is well known that DW-NOMINATE scores are based on a static model, thus being time-invariant over legislartor's career. Yet, we require a spatial measure specific for the post-primary time-period.

For the pusposes of this paper, I use a simple and transparent measure of incumbents' roll call moderation, which is based on DW-NOMINATE as well as it's underlying idea that legislators with a similar voting record should be placed close to each other on the liberal-conservative scale. Analogously to the measure of candidate extremism, I proceed in two steps. First, I calculate each incumbent i's *indirect DW-NOMINATE* as the agreement-rate weighted average of other incumbents $j \neq i$' s DW-NOMINATE:

$$IndirectNominate_i = \frac{\sum_j \alpha_{ij} Nominate_j}{\sum_j \alpha_{ij}}$$

where α_{ij} is the agreement rate between Representative i and j, i.e. the share of roll calls for which both i and j vote for the same side.[6] For this calculation, I only include post-primary roll calls held no later than 120 days before the general election, and I exclude uninformative lopsided votes on which more than 85% of all House members agree.

Following the logic of the previous subsection, I use the fact that DW-NOMINATE scores are bounded between -1 (very liberal) and 1 (very conservative), as must be the *indirect DW-NOMINATE*. I thus claculate roll call moderation as the distance of incumbents' *indirect DW-NOMINATE* from their partiy's theoretical extreme:

$$Moderation_{i(p)} = \begin{cases} 1 - IndirectNominate_{i(p)} & \text{if } p = \text{Republican} \\ |-1 - IndirectNominate_{i(p)}| & \text{if} p = \text{Democtrat} \end{cases}$$

This measure has a clear spatial interpretation. An increase in moderation means that a Democratic (Republican) incumbent's roll call voting behavior becomes more conservative (liberal), moving in the direction of the oppnent party. While qualitatively informative of directional changes in the incumbent's position taking, the magnitude has no direct quantitative interpretation. Although I address this by offering an indirect interpretation of the magnitudes by relating standard deviations of the moderation score to standard deviations in the U.S. Senates DW-NOMINATE scores, I also check robustness of my results to using a directly interpretable measure of roll call extremism, i.e. party-line voting on divisive issues on which the majority of Democrats disagrees with the majority of Republicans. Deviating from party-line in divisive votes is commonly regarded as a strong and costly signal of moderation to voters (see , e.g., Canes-Wrone et al., 2002; Carson et al., 2010). Substantively, I reach the same conclusion with both outcomes: Incumbents compromise on their policy position by moderating thier roll call voting behavior in response to

[6]Formally $\alpha_{ij} = \frac{1}{n} \sum_{r=1}^{n} I(v_{ir} = v_{jr})$, where I is a dummy variable = 1 if $v_{ir} = v_{jr}$, and v_{kr} is a dummy variable taking the value 1 if representative $k \in \{i, j\}$ votes "yea", and 0 if k votes "nay" in roll call r.

an extremist challanger.

3 Validity Checks on the Identification Assumption

Before presenting my results, I check the validity of my identification assumption that agents have imprecise control over close primary election outcomes by testing three of its implications. First, pre-determined incumbent-specific characteristics should be distributed smoothly around the cut-off. Crucially, estimating the strategic reaction of incumbents to extremists requires that incumbents' *pre-primary* policy position is unrelated to the nomination of an extremist challenger. I thus test for a discontinuity in incumbent characteristics and *pre-primary* voting behavior by estimating equation (1) with these pre-determined incumbent observables as outcome variables. Figure 2 plots the estimated discontinuities along with 95% confidence intervals. Point estimates of pre-treatment outcomes are close to zero and precisely estimated. Incumbents facing narrowly elected extremist exhibit *ex ante* equally moderate roll call voting behavior and are just as likely to vote in party-line on divisive issues as are incumbents running against a barely nominated moderate. Reassuringly, incumbents challenged by a closely nominated extremist neither differ in legislative experience, party affiliation, sex, age, nor birth cohort. I caveat that coefficient estimates on sex and party affiliation are somewhat imprecise such that I cannot exclude small discontinuities at the cutoff. However, I anticipate that the inclusion of incumbent characteristics as covariates does not affect my results.

FIGURE 2: CONTINUITY OF PREDETERMINED INCUMBENT-LEVEL COVARIATES AND PRE-PRIMARY OUTCOMES AROUND THE CUTOFF

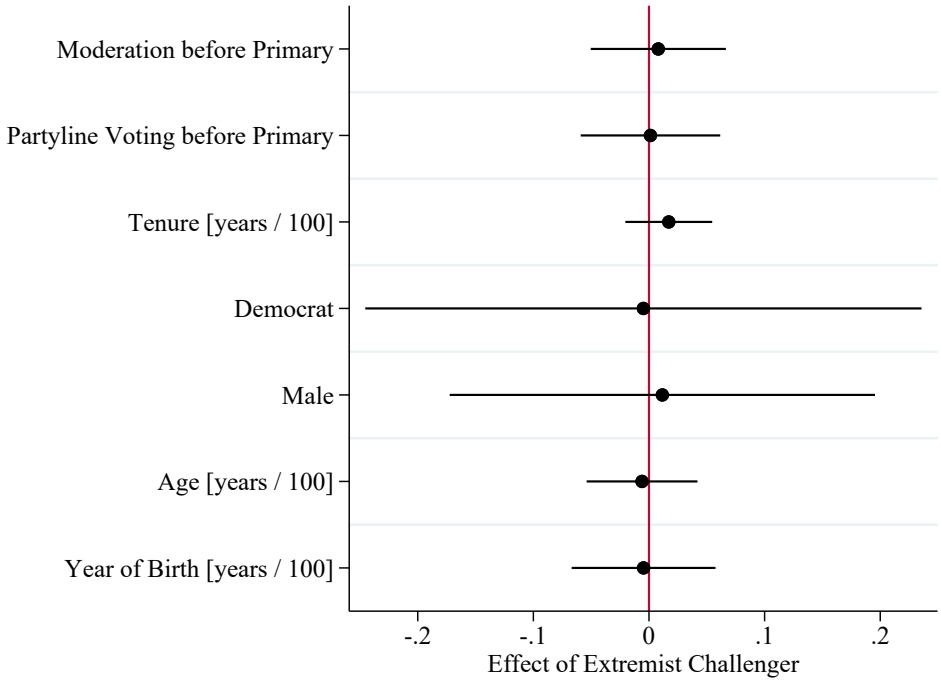

Notes: The figure plots discontinuities of incumbent-level observables at the cutoff along with 95% confidence intervals. Each estimate is obtained from a local-linear regression using a triangular kernel and heteroskedasticity-robust nearest neighbor variance estimation under the MSERD-optimal bandwidth.

Secondly, the identification assumtpion implies that districts whose non-incumbent party nominates an extremist challenger by a narrow margin, are comparable to districts whose non-incumbent party nominates a relatively moderate challenger. For identification of the incumbent's strategic response to an extremist challenger, it is crucial that the nomination of an extremist in close primary elections is unrelated to *pre-primary* voter preferences that shape both challenger and incumbent's position taking. As shown in Figure 3, a close primary win of an extremist challenger is unrelated to pre-determined voter preferences, as revealed by previous election results. I also fail to find a discontinutiy in other district attributes that likely predict policy preferences. Districts close to either side of the cutoff do not significantly differ in income, education, unemployment rate, racial composition, or population size.

FIGURE 3: CONTINUITY OF PREDETERMINED DISTRICT-LEVEL COVARIATES AROUND THE CUTOFF

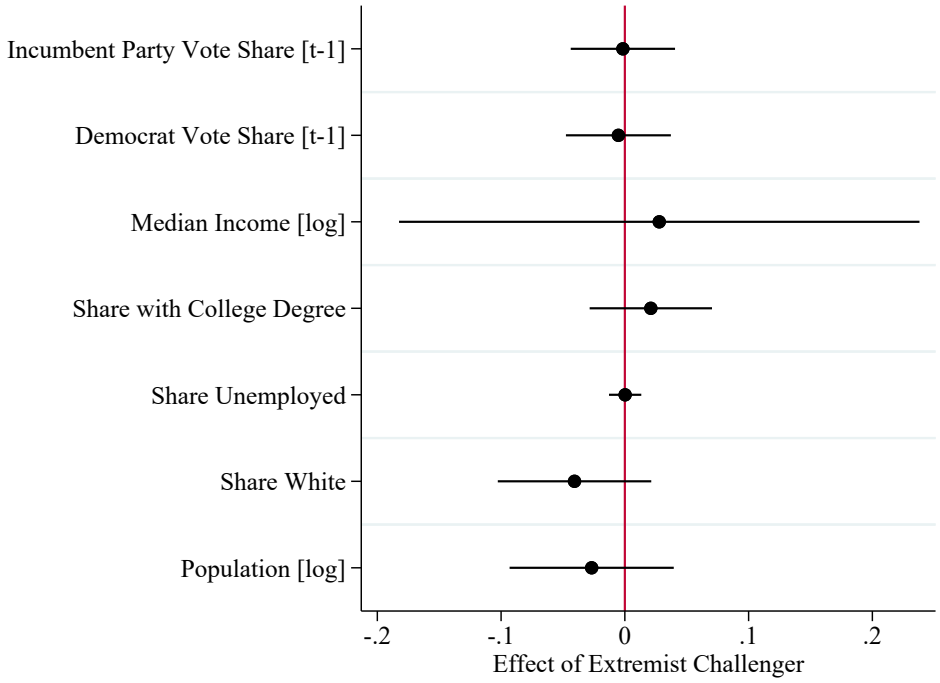

Notes: The figure plots discontinuities of district-level observables at the cutoff along with 95% confidence intervals. Each estimate is obtained from a local-linear regression using a triangular kernel and heteroskedasticity-robust nearest neighbor variance estimation under the MSERD-optimal bandwidth.

Finally, if relative extremists were differentially able to win close primary elections, we might expect "bunching" around the cutoff leading to a discontinuity in the distribution of the assignment variable. Figure 4, Panel A provides *prima facie* evidence agianst the concern of sorting close the cutoff, showing no discontinuous jump in the distribution of the extremist's vote share in my sample. A formal test devised by Cattaneo et al. (2020), in the spirit of McCrary (2008), fails to reject equal the density of extremist and moderate wins at the cutoff ($p = 0.786$). Figure 4, Panel B provides a graphical representation of the density test, showing estimated densities near to each other, and with overlapping 95% confidence intervals at the cutoff. Overall, I find no evidence for manipulation of the assignment variable and sorting around the cutoff that might invalidate the idetification assumption.

FIGURE 4: DISTRIBUTION AND ESTIMATED DENISTY OF EXTREMIST'S PRIMARY VOTE SHARE AROUND THE CUTOFF

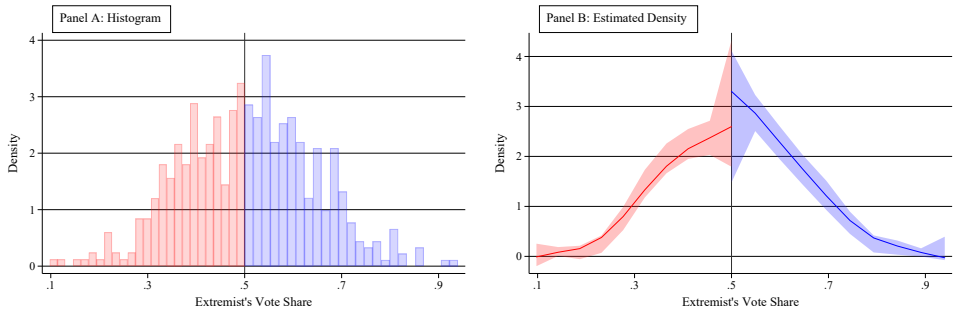

Notes: The figure plots the distribution of the extremists's primary vote share (Panel A), and estimated densities along with 95% confidence intervals (Panel B).

4 Results

4.1 Graphical Analysis

Before turning to formal discontinuity estimates, I provide suggestive graphical evidence examining the relationship between incumbent's roll call moderation and the opponent party's primary vote share of the extremist candidate in a narrow bandwidth around the cutoff.

To facilitate visualization, I reduce noise by using first-differenced outcomes.

Figure 5 plots binned averages of the outcome variable (vertical axis) along with 95% confidence intervals against the extremist's top-two candidate vote share 10 percentage points around the 0.5 cutoff (horizontal axis).[7] A separate quadratic polynomial spline is fitted on either side of the assignment variable's cutoff value. We observe a sharp discontinuity at the 0.5 threshold of the Extremist's primary vote share. An extremist challenger has a large postivie impact on the incumbent's moderation in roll call voting. An incumbent facing a narrowly nominated extremist scores 0.03 units (approximately a 1/4 standard deviation) higher on the moderation measure, compared to an incumbent whose opponent party's primary electorate nominated a more moderate challenger by a narrow margin. Similarly, incumbents running against a narrowly nominated extremist are 5 percentage points less likely to vote in party-line on divisive issues. While economically large, it is not immediate to infer whether these effects are statistically significant by graphical inspection. I thus proceed to fromal tests.

[7] The range of the extremist's primary vote share from 40% to 60% approximateely corresponds to the mserd-optimal bandwidth used for formal estimation, and includes more than 50% of the overall sample.

FIGURE 5: MODERATION AND PARTY-LINE VOTING OF INCUMBENTS WITH NARROWLY ELECTED CHALLENGERS

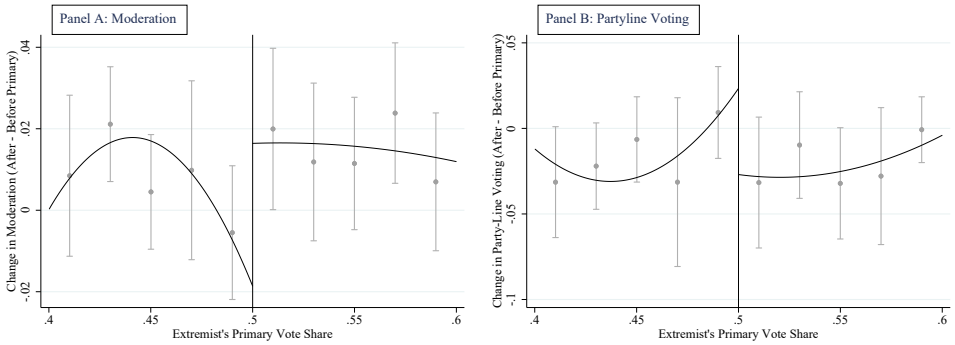

Notes: The figure plots binned averages of changes in incumbent moderation (Panel A) and incumbent's party line voting in divisive votes (Panel B) on the vertical axis against the extremist's top-two candidate vote share on horizontal axis. Averages are calculated within 0.2-percentage-point-wide bins. A separate quadratic polynomial spline is fitted on either side of the assignment variable's cutoff value of 0.5.

4.2 Regression Analysis

I next present formal estimates of equation (1). Given my limited sample size and non-negligible noise in the data, my preferred specification uses second-order polynomials and first-differenced outcomes, although I show robustness to a variety of specifications. Most importantly, to rule out concerns that discontinuities in first-differenced outcomes are driven by pre-existing differences in the dependent variable, I first estimate equation (1) separately for pre-primary and post-primary outcomes.

Table 1 presents the reults for the moderation measure of incumbent's roll call voting behavior, for both local-linear (Panel A) and local-quadratic specifications (Panel B). In column 1, I test for a discontinuities in the pre-primary outcome. The estimate is indistinguishable from zero, both statistically and economically. This holds irrespectively of the chosen polynomial, and of the inclusion of the full set of pre-determined incumbent-level and district-level covariates reported in section 3 (column 2). When considering post-primary moderation, the imprecisely estimated discontinuity jumps by several orders of magnitude (column 3). Importantly, the effect does not change substatively when controlling for the pre-primary outcome (column 4), as well as pre-determined incumbent- and district-level covariates (column 5), but is more precisely estimated and statistically different from zero ($p < 0.05$). I then directly test for a discontinuity in the difference between post-primary and pre-primary moderation using first-differenced outcomes (column 6). Reassuringly, estimates are virtually identical across specifications, and robust to conditioning on the base-level (column 7) and covariates (column 8). Overall, I find consistent

evidence that incumbents moderate their roll call behavior when confronted with a relatively extreme challenger. The most conservative estimates imply a sizeable increase in the moderation score by 0.03, that is a 0.25 standad deviation shift of the incumbent's roll call position towards the other party. To give an enconomically meaningflul interpretation of the effect size, I compare this to a 0.25 standard deviation shift in the Senators' DW-NOMINATE, which roughly correpsonds to a move from (the relatively conservative) Republican Senate Leader Mitch McConnell to (the relatively liberal) Mitt Romney, or from the (relatively liberal) Bernard Sanders to (the relatively conservative) Democratic Senate Leader Chuck Schumer.

Table 1: Effect of Extremist Challenger on Incumbent Moderation

	Before Primary		After Primary			Change (After - Before Primary)		
Panel A: Local Linear	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	0.010	0.001	0.042	0.031**	0.029**	0.032**	0.031**	0.029**
	(0.037)	(0.025)	(0.042)	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)
	[0.832]	[0.814]	[0.414]	[0.023]	[0.044]	[0.028]	[0.023]	[0.044]
MSERD-Optimal Bandwidth	0.068	0.068	0.068	0.068	0.068	0.068	0.068	0.068
Effective Observations	191	191	191	191	191	191	191	191
Panel B: Local Quadratic								
	0.011	-0.002	0.050	0.038**	0.037**	0.039**	0.038**	0.037**
	(0.046)	(0.032)	(0.053)	(0.016)	(0.016)	(0.017)	(0.016)	(0.016)
	[0.905]	[0.873]	[0.490]	[0.030]	[0.050]	[0.037]	[0.030]	[0.050]
MSERD-Optimal Bandwidth	0.101	0.101	0.101	0.101	0.101	0.101	0.101	0.101
Effective Observations	270	270	270	270	270	270	270	270
Observations	517	517	517	517	517	517	517	517
Outcome Before Primary	-	-	N	Y	Y	N	Y	Y
Covariates	N	Y	N	N	Y	N	N	Y

Notes: Each column reports results from two local polynomial regression using a triangular kernel and heteroskedasticity-robust nearest neighbor variance estimators, fitting separate polynomials of order 1 (Panel A) or order 2 (Panel B) on each side of the cutoff. The outcome variable is incumbent moderation prior to the opponent party's primary (columns 1 and 2), post-primary incumbent moderation in the 120 days before the general election (colums 3-5), or the difference in post- and pre-primary outcomes (columns 6-8). The independent variable of interest is a dummy equal to 1 if the more extreme of the top-two candidates in the opponent party's primary election gathered more than 50 percent of the top-two candidates vote total. Optimal bandwidths are derived under the MSERD procedure for the first-differenced outcome and standard RD estimation without covariates (column 6). Standard errors in parentheses. Robust p-values based on bias-adjusted estimates in brackets: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Next, I consider party-loyalty in divisive votes as an alternative measure of roll call extremism. Table 2 presents the results for party-line voting. Similarly to the previously discussed results for the moderation measure, I do not find any discontinuity in pre-primary party-line voting, but a sizeable post-primary decrease in party-line voting by statistically significant 5 percentage points ($p < 0.05$) when an incumbent faces an extremist as opposed to a moderate challenger. Both sets of results tell a qualitatively consistent story: Incumbents compromise on their roll call voting behavior in response to an extremist challenger, moving in the direction of their opponent party's position, which is consistent with (partial) policy convergence, and suggests that policy platforms are strategic complements.

TABLE 2: EFFECT OF EXTREMIST CHALLENGER ON INCUMBENT'S PARTY-LINE VOTING

	Before Primary		After Primary			Change (After - Before Primary)		
PANEL A: LOCAL LINEAR	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	0.006	0.008	-0.042	-0.050**	-0.043**	-0.049**	-0.050**	-0.043**
	(0.033)	(0.030)	(0.046)	(0.025)	(0.024)	(0.026)	(0.025)	(0.024)
	[0.925]	[0.972]	[0.251]	[0.010]	[0.025]	[0.015]	[0.010]	[0.025]
MSERD-Optimal Bandwidth	0.072	0.072	0.072	0.072	0.072	0.072	0.072	0.072
Effective Observations	203	203	203	203	203	203	203	203
PANEL B: LOCAL QUADRATIC								
	0.005	0.010	-0.056	-0.062**	-0.056**	-0.061**	-0.062**	-0.056**
	(0.042)	(0.037)	(0.056)	(0.030)	(0.029)	(0.031)	(0.030)	(0.029)
	[0.769]	[0.766]	[0.363]	[0.013]	[0.023]	[0.020]	[0.013]	[0.023]
MSERD-Optimal Bandwidth	0.106	0.106	0.106	0.106	0.106	0.106	0.106	0.106
Effective Observations	286	286	286	286	286	286	286	286
Observations	517	517	517	517	517	517	517	517
Outcome Before Primary	-	-	N	Y	Y	N	Y	Y
Covariates	N	Y	N	N	Y	N	N	Y

Notes: Each column reports results from two local polynomial regression using a triangular kernel and heteroskedasticity-robust nearest neighbor variance estimators, fitting separate polynomials of order 1 (Panel A) or order 2 (Panel B) on each side of the cutoff. The outcome variable is incumbent's share of votes cast in line with the majoirity of his own party and against the majority the opponent party in divisive roll calls prior to the opponent party's primary (columns 1 and 2), in all divisive post-primary roll calls held in the 120 days before the general election (colums 3-5), or the difference in post- and pre-primary outcomes (columns 6-8). The independent variable of interest is a dummy equal to 1 if the more extreme of the top-two candidates in the opponent party's primary election gathered more than 50 percent of the top-two candidates vote total. Optimal bandwidths are derived under the MSERD procedure for the first-differenced outcome and standard RD estimation without covariates (column 6). Standard errors in parentheses. Robust p-values based on bias-adjusted estimates in brackets: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5 Robustness Checks

5.1 Restults Depending on Bandwidth

My preferred specifications presented thus far rely on MSERD-optimal bandwidths which reduce asymptotic bias the most (Calonico et al., 2019). Nonetheless, I probe robustness to alternative bandwidth choices. I begin by replicating the whole set of results presented in Tables 1 and 2 with optimal bandwidths dervied und the IK-procedure (Imbens and Kalyanaraman, 2012). As showin in Appendix Tables A.1 and A.2, discontinuity estimates obtained with IK-optimal bandwidths are nearly identical to my main results presented above. For both outcome variables, results do not depend on the procedure of optimal bandwidth selection.

I next probe sensitivity of the discontinuity estimates for an extensive neighborhood around optimal bandwidths. More precisely, I re-estimate equation (1) on first-differenced outcomes across bandwidths from 0.02 up to twice the mserd-optimal bandwidth, for both local-linear and local-quadratic specifications. As shown in Figures 6 and 7, estimated discontinuity gaps are highly stable across aribitrarily chosen bandwidths. Reflecting the bias-variance tradeoff inherent to regression discontinuity designs, coefficients get somewhat smaller in magnitude but more precisely estimated with increasing bandwidth. Importantly, coefficients never change sign and significantly differ from zero across a wide range of bandwiths. I thus conclude that my main findings are not driven by bandwidth selection.

FIGURE 6: ROBUSTNESS TO BANDWIDTH CHOICE: MODERATION

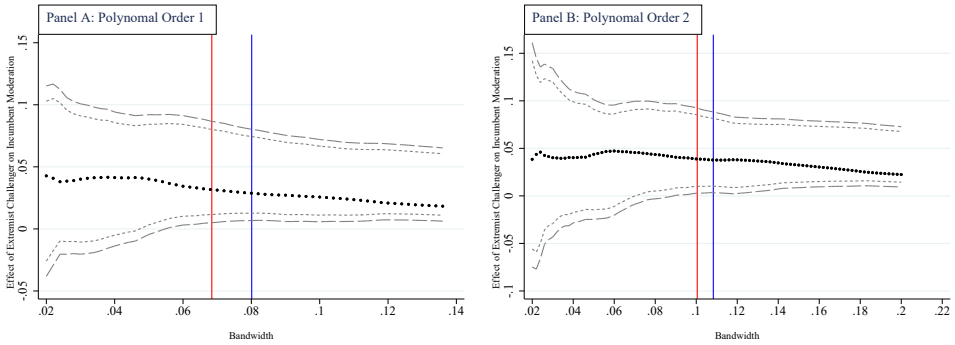

Notes: The figure plots estimated discontinuities in (first-differenced) incumbent moderation depending on bandwidth, for local-linear (Panel A) and local-quadratic soecifications (Panel B). Dots depict the estimated discontinuity with long-dashed and short-dashed lines representing 95% and 90% confidence intervals, respectively. The figure reports estimates for bandwidths between 0.02 and twice the mserd-optimal bandwidth, in increments of 0.002. The vertical red and blue lines indicate msered-optimal and IK-optimal bandwidths, respectively.

FIGURE 7: ROBUSTNESS TO BANDWIDTH CHOICE: PARTY-LINE VOTING

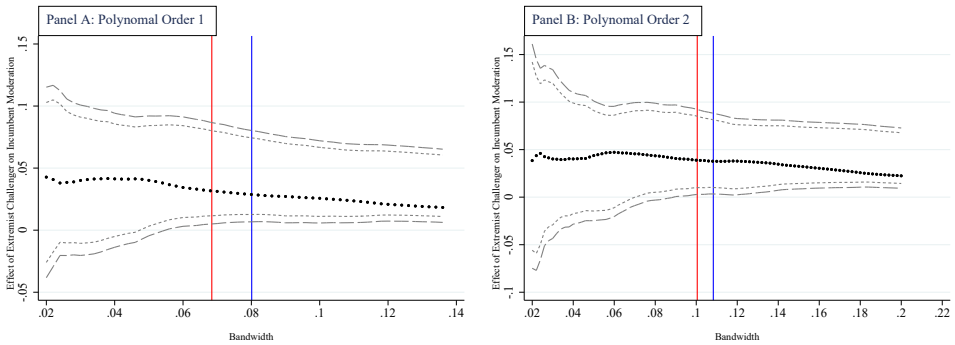

Notes: The figure plots estimated discontinuities in (first-differenced) incumbent party-line voring in divisive votes depending on bandwidth, for local-linear (Panel A) and local-quadratic soecifications (Panel B). Dots depict the estimated discontinuity with long-dashed and short-dashed lines representing 95% and 90% confidence intervals, respectively. The figure reports estimates for bandwidths between 0.02 and twice the mserd-optimal bandwidth, in increments of 0.002. The vertical red and blue lines indicate msered-optimal and IK-optimal bandwidths, respectively.

5.2 Restults Depending on Potential Challengers' Relative Extremism

Results presented thus far rely on the whole sample, including primaries where the top-two candidates have very similar donation-based scores. As noted in section 2.2, donation-based scores are measured with error, which may lead to missclassification of the more moderate candidate as the extremist and vice-versa. Although using only pre-primary donations for the calculation of the scores ensures that missclassification is not endogeneous, the resulting "classical" measurement error may be of concern not only because it leads to attenuation. As local regression discontinuity estimates are based on a finite sample in a narrow bandwidth around the cutoff, one may worry that missclassification of a few observations close to the threshold may drive the results. I address this concern by calculating primary candidates' relative extremism as the absolute distance of their donation-based scores, and successively exluding primaries where candidates have less distant donor scores, i.e. those observations where missclassification due to measurement error in the candidate scores is most likely to occur.

I thus re-estimate my preferred specification as in Table 1, Panel B, column 6 on different subsamples, starting with the whole sample from which I sequentially remove the lowest decile of the distribution of relative extremism. The results are summarized in Figure 8, which plots coefficient estimates along with 95% confidence intervals obtained from different subsamples. The leftmost estimate is based on the whole sample, and identical to the one reported in Table 1, Panel B, column 6. Moving rightwards on the x-axis, coefficient estimates are based on ever smaller subsamples with increasing relative extremism of the praimary candidates. Figure 9 replicates the exercise for the alternative outcome of party-line voting in divisive votes.

FIGURE 8: EFFECT OF EXTREMIST CHALLENGER ON INCUMBENT MODERATION DEPENDING ON RELATIVE EXTREMISM

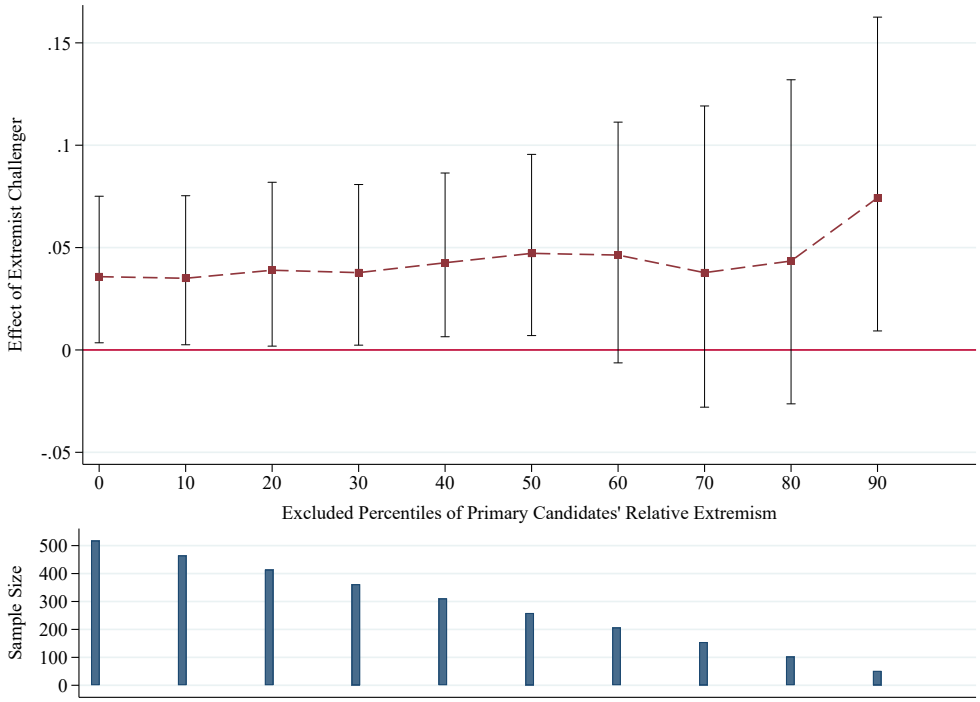

Notes: The top panel plots discontinuity estimates along with robust 95% confidence intervals (vertical axis) across different subsamples successively exluding observations in lowest decile of the distribution of relative extremism, i.e. the estimated ideological distance between opponent party's primary candidates (x-axis). The outcome variable is the difference between post- and pre-primary incumbent moderation. Estimates are obtained from local polynomial regressions using a triangular kernel and heteroskedasticity-robust nearest neighbor variance estimators, fitting separate polynomials of order 2 on either side of the cutoff. Each regression uses optimal bandwidths are derived under the MSERD procedure. The bottom panel indicates the size of the subsamples used for estimation.

FIGURE 9: EFFECT OF EXTREMIST CHALLENGER ON INCUMBENT'S PARTY-LINE VOTING DEPENDING ON RELATIVE EXTREMISM

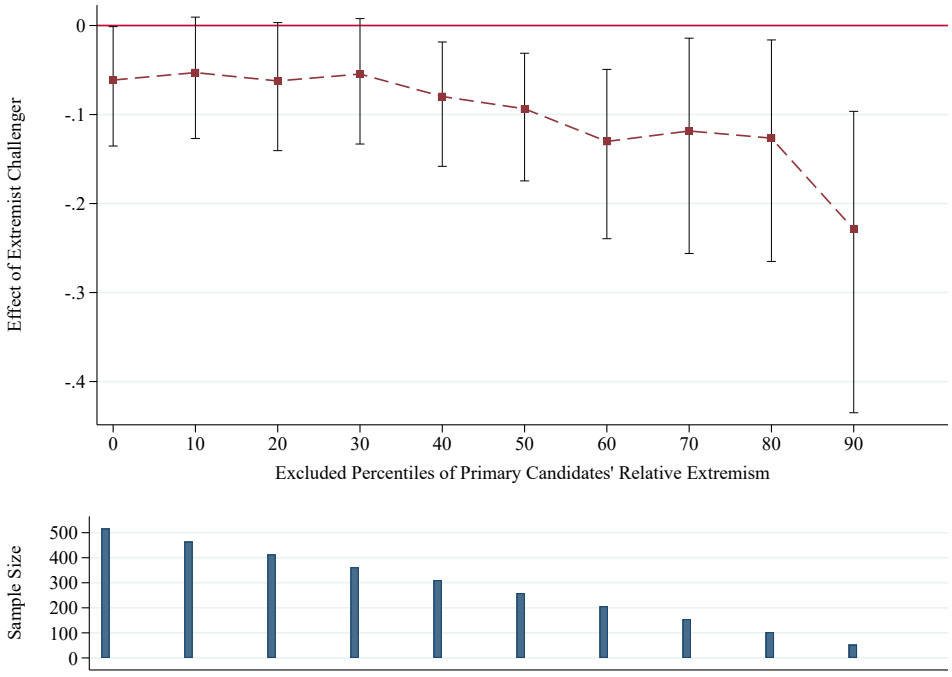

Notes: The top panel plots discontinuity estimates along with robust 95% confidence intervals (vertical axis) across different subsamples successively exluding observations in lowest decile of the distribution of relative extremism, i.e. the estimated ideological distance between opponent party's primary candidates (x-axis). The outcome variable is the difference between post- and pre-primary incumbent partyline voting in divisive votes. Estimates are obtained from local polynomial regressions using a triangular kernel and heteroskedasticity-robust nearest neighbor variance estimators, fitting separate polynomials of order 2 on either side of the cut-off. Each regression uses optimal bandwidths are derived under the MSERD procedure. The bottom panel indicates the size of the subsamples used for estimation.

Both sets of results reveal a consistent pattern. When excluding observations with low relative extremism of potential challengers, coefficients never change sign but get larger in magnitude. This is strong evidence against the concern that my baseline results are driven by missclassification of primary candidates. If anything, my results get stronger when excluding observations where donation-based scores predict primary candidates to be idelogogically close to each other. This pattern is consistent with attenuation bias due to classical measurement error stemming from missclassification of ideolgocially proximate candidates. But it is also consistent with the idea that larger differences in the ideological distance between prospective challengers reflect an increase in treatment intensity. However, it is not obvious how to discriminate between these two potential drivers. That said, I abide with a conservative interpretation of the baseline reults reported in the previous section 4.2 as lower bounds on the true effect of extremist challengers on incumbents'

roll call voting behavior.

6 Preliminary Conclusion and Further Questions

How do legislators react to challengers? This paper provides the first causal estimate of challengers' impact on incumbents' roll call behvior. Exploiting quasi-random assignment of incumbents to more or less extreme challgners of the opposite party, I find robust evidence that incumbents moderate their roll call voting behavior in response to an extremist challenger. This is an important result, both for the empirical and the theoretical literature on legislator behavior. The emprical literature has identified the influence of legislators' own ideology (Levitt, 1996) , their peers (Harmon et al., 2019) and daughters (Washington, 2008), their voters' preferences (Mian et al., 2010), media coverage (Snyder and Strömberg, 2010), on legislators' vote. This paper not only amends to the emirical literture by identifying challengers as another key driver of legislative decision making, but also directly speaks to the theoretical literature by providing direct evidence for partial policy convergence with policy positions as strategic complements.

Yet, the preliminary evidence presented thus far tells little about the mechanisms leading incumbents to comporomise vis-à-vis extremist challengers. There are at least two aspects to be considerd. First, it is not clear *to what* incumbents react. Is incumbent's roll-call moderation a strategic response to extremist's policy platform, or to extremist's valence characteristis? Second, *the motive why* incumbents react to extremists challengers by moderating their voting behavior remains obscure. Is it because of purely opportunistic office-seeking motives, or because of policy motivation? In future research, I intend to elaborate on these questions, both empirically and theoretically.

References

Ansolabehere, Stephen, James M. Snyder Jr., and Charles Stewart III, "Candidate Positioning in US House elections," *American Journal of Political Science*, 2001, *45* (1), 136–159.

Besley, Timothy and Stephen Coate, "An Economic Model of Representative Democracy," *The Quarterly Journal of Economics*, 1997, *112* (1), 85–114.

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

_ , "Inferring Roll-Call Scores From Campaign Contributions Using Supervised Machine Learning," *American Journal of Political Science*, 2018, *62* (4), 830–848.

Calonico, Sebastian, Matias D. Cattaneo, and Rocio Titiunik, "Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs," *Econometrica*, 2014, *82* (6), 2295–2326.

_ **,** _ **, Max H. Farrell, and Rocio Titiunik**, "Regression Discontinuity Designs Using Covariates," *Review of Economics and Statistics*, 2019, *101* (3), 442–451.

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

Carson, Jamie L., Gregory Koger, Matthew J. Lebo, and Everett Young, "The Electoral Costs of Party Loyalty in Congress," *American Journal of Political Science*, 2010, *54* (3), 598–616.

Cattaneo, Matias D., Michael Jansson, and Xinwei Ma, "Simple Local Polynomial Density Estimators," *Journal of the American Statistical Association*, 2020, *115* (531), 1449–1455.

Downs, Anthony, *An Economic Theory of Democracy*, Harper and Row, New York, 1957.

Duggan, John and César Martinelli, "The Political Eeconomy of Dynamic Elections: Accountability, Commitment, and Responsiveness," *Journal of Economic Literature*, 2017, *55* (3), 916–84.

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

_ **and James M. Snyder Jr.**, "Candidate Ideology and Electoral Success," 2015. Unpublished Manuscript.

Harmon, Nikolaj, Raymond Fisman, and Emir Kamenica, "Peer Effects in Legislative Voting," *American Economic Journal: Applied Economics*, 2019, *11* (4), 156–80.

Hotelling, Harold, "Stability in Competition," *The Economic Journal*, 1929, *39* (153), 41–57.

Imbens, Guido and Karthik Kalyanaraman, "Optimal Bandwidth Choice for the Regression Discontinuity Estimator," *The Review of Economic Studies*, 2012, *79* (3), 933–959.

Jones, Daniel B. and Randall Walsh, "How Do Voters Matter? Evidence from US Congressional Redistricting," *Journal of Public Economics*, 2018, *158*, 25–47.

Lee, David S. and Thomas Lemieux, "Regression Discontinuity Designs in Economics," *Journal of Economic Literature*, 2010, *48* (2), 281–355.

_ , Enrico Moretti, and Matthew J. Butler, "Do Voters Affect or Elect Policies? Evidence from the US House," *The Quarterly Journal of Economics*, 2004, *119* (3), 807–859.

Levitt, Steven D., "How Do Senators Vote? Disentangling the Role of Voter Preferences, Party Affiliation, and Senator Ideology," *The American Economic Review*, 1996, *86* (3), 425–441.

Matsusaka, John G., "Economics of Direct Legislation," *The Quarterly Journal of Economics*, 1992, *107* (2), 541–571.

McCrary, Justin, "Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test," *Journal of Econometrics*, 2008, *142* (2), 698–714.

Mian, Atif, Amir Sufi, and Francesco Trebbi, "The Political Economy of the US Mortgage Default Crisis," *American Economic Review*, 2010, *100* (5), 1967–98.

Osborne, Martin J. and Al Slivinski, "A Model of Political Competition with Citizen-Candidates," *The Quarterly Journal of Economics*, 1996, *111* (1), 65–96.

Poole, Keith T. and Howard Rosenthal, *Congress: A Political-Economic History of Roll Call Voting*, Oxford University Press, 1997.

Snyder, James M. Jr. and David Strömberg, "Press Coverage and Political Accountability," *Journal of Political Economy*, 2010, *118* (2), 355–408.

Washington, Ebonya L., "Female Socialization: How Daughters Affect Their Legislator Fathers," *American Economic Review*, 2008, *98* (1), 311–32.

Appendix

TABLE A.1: ROBUSTNESS: EFFECT OF EXTREMIST CHALLENGER ON INCUMBENT MODERATION WITH IK-OPTIMAL BANDWIDTHS

	Before Primary		After Primary			Change (After - Before Primary)		
PANEL A: LOCAL LINEAR	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	0.011	-0.000	0.040	0.028**	0.027**	0.029**	0.028**	0.027**
	(0.034)	(0.023)	(0.039)	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)
	[0.847]	[0.885]	[0.393]	[0.017]	[0.031]	[0.020]	[0.017]	[0.031]
IK-Optimal Bandwidth	0.080	0.080	0.080	0.080	0.080	0.080	0.080	0.080
Effective Observations	222	222	222	222	222	222	222	222
PANEL B: LOCAL QUADRATIC								
	0.016	0.001	0.054	0.037**	0.036**	0.038**	0.037**	0.036**
	(0.044)	(0.030)	(0.050)	(0.016)	(0.016)	(0.016)	(0.016)	(0.016)
	[0.952]	[0.972]	[0.511]	[0.028]	[0.043]	[0.034]	[0.028]	[0.043]
IK-Optimal Bandwidth	0.108	0.108	0.108	0.108	0.108	0.108	0.108	0.108
Effective Observations	288	288	288	288	288	288	288	288
Observations	517	517	517	517	517	517	517	517
Outcome Before Primary	-	-	N	Y	Y	N	Y	Y
Covariates	N	Y	N	N	Y	N	N	Y

Notes: Each column reports results from two local polynomial regression using a triangular kernel and heteroskedasticity-robust nearest neighbor variance estimators, fitting separate polynomials of order 1 (Panel A) or order 2 (Panel B) on each side of the cutoff. The outcome variable is incumbent moderation prior to the opponent party's primary (columns 1 and 2), post-primary incumbent moderation in the 120 days before the general election (colums 3-5), or the difference in post- and pre-primary outcomes (columns 6-8). The independent variable of interest is a dummy equal to 1 if the more extreme of the top-two candidates in the opponent party's primary election gathered more than 50 percent of the top-two candidates vote total. Bandwidths are optimal under the IK procedure for the first-differenced outcome and standard RD estimation without covariates (column 6). Standard errors in parentheses. Robust p-values based on bias-adjusted estimates in brackets: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

TABLE A.2: ROBUSTNESS: EFFECT OF EXTREMIST CHALLENGER ON INCUMBENT'S PARTY-LINE VOTING WITH IK-OPTIMAL BANDWIDTHS

	Before Primary		After Primary			Change (After - Before Primary)		
PANEL A: LOCAL LINEAR	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	0.006	0.011	-0.027	-0.034**	-0.031*	-0.033**	-0.034**	-0.031*
	(0.026)	(0.023)	(0.036)	(0.020)	(0.019)	(0.020)	(0.020)	(0.019)
	[0.837]	[0.906]	[0.207]	[0.045]	[0.059]	[0.044]	[0.045]	[0.059]
IK-Optimal Bandwidth	0.130	0.130	0.130	0.130	0.130	0.130	0.130	0.130
Effective Observations	332	332	332	332	332	332	332	332
PANEL B: LOCAL QUADRATIC								
	-0.003	0.001	-0.063	-0.059**	-0.054**	-0.059**	-0.059**	-0.054**
	(0.039)	(0.035)	(0.053)	(0.029)	(0.028)	(0.029)	(0.029)	(0.028)
	[0.690]	[0.643]	[0.437]	[0.020]	[0.034]	[0.032]	[0.020]	[0.034]
IK-Optimal Bandwidth	0.118	0.118	0.118	0.118	0.118	0.118	0.118	0.118
Effective Observations	313	313	313	313	313	313	313	313
Observations	517	517	517	517	517	517	517	517
Outcome Before Primary	-	-	N	Y	Y	N	Y	Y
Covariates	N	Y	N	N	Y	N	N	Y

Notes: Each column reports results from two local polynomial regression using a triangular kernel and heteroskedasticity-robust nearest neighbor variance estimators, fitting separate polynomials of order 1 (Panel A) or order 2 (Panel B) on each side of the cutoff. The outcome variable is incumbent's share of votes cast in line with the majoirity of his own party and against the majority the opponent party in divisive roll calls prior to the opponent party's primary (columns 1 and 2), in all divisive post-primary roll calls held in the 120 days before the general election (colums 3-5), or the difference in post- and pre-primary outcomes (columns 6-8). The independent variable of interest is a dummy equal to 1 if the more extreme of the top-two candidates in the opponent party's primary election gathered more than 50 percent of the top-two candidates vote total. Bandwidths are optimal under the IK procedure for the first-differenced outcome and standard RD estimation without covariates (columns 6). Standard errors in parentheses. Robust p-values based on bias-adjusted estimates in brackets: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.