The Rigidity of Legislative Positions: Incumbents Do Not Respond to Extreme Challengers

Kisoo Kim*

June 27, 2026

Abstract

Models of electoral accountability predict that legislators should respond strategically to electoral threats. I ask whether U.S. House incumbents change their legislative behavior when an ideologically extreme challenger wins the opposing party's nomination, becoming their general-election opponent. Using a regression discontinuity design in close opposing-party primaries and constructing pre-primary ideological scores, I compare incumbents who narrowly face extreme rather than more moderate general-election opponents. I find no substantial post-primary response in party unity, ideological extremity, votes with party leaders, roll-call participation, or bill sponsorship. The estimates are consistently small across specifications and outcome measures, in line with stable legislative commitments rather than strategic adjustment in either ideological direction. Combined with existing evidence that incumbents respond to same-party primary threats, the results raise the possibility that electoral accountability operates selectively: legislative behavior responds to some forms of electoral pressure but not to variation in opposing-party challenger ideology.

Keywords: Electoral accountability, electoral competition, candidate ideology, position-taking, U.S. House of Representatives

*Email: kisoo@virginia.edu

1 Introduction

The 2010 Tea Party wave reshaped Republican primaries across the country, producing nominees markedly more conservative than their predecessors. In dozens of congressional districts, Democratic incumbents learned they would face general election opponents positioned far to the right of the typical Republican challenger. Did these incumbents adjust their legislative voting in response? More broadly, when faced with a more or less ideologically extreme opponent, does the incumbent change their legislative behavior?

Electoral accountability rests on a simple premise: elections discipline incumbents through electoral competition. Yet which specific features of competition shape incumbent behavior remains an open question. If incumbents respond strategically to changes in electoral competition—the ideological character of their opponent, for example—then the emergence of an extreme challenger would trigger behavioral adjustments. If legislative positions are sticky, driven by personal conviction, party discipline, or constituency preferences rather than strategic calculation, then challenger extremity should have no effect on incumbent behavior. Whether incumbents respond to extreme challengers informs our understanding of the specific mechanisms through which electoral competition shapes legislative behavior, as well as the sources of political polarization, which many attribute in part to the rise of extreme candidates.

In this article, I investigate one specific channel of electoral competition: whether facing an ideologically extreme challenger from the opposing party causes incumbent legislators to change their legislative behavior. Theoretical predictions about how incumbents respond to extreme challengers are genuinely ambiguous. On the one hand, legislators may hold genuine ideological commitments that resist short-term strategic adjustment. Repositioning is also costly—it risks alienating existing supporters and invites accusations of flip-flopping (Lee et al., 2004; Fowler and Hall, 2016). These factors suggest that legislative positions may be sticky regardless of challenger characteristics, a pattern I term the inertia hypothesis.

On the other hand, theories of strategic response predict behavioral change, though they disagree on direction. The electoral slack hypothesis, grounded in policy-motivated spatial voting (Calvert, 1985; Wittman, 1983), predicts that incumbents facing extreme challengers gain greater electoral leeway to pursue their own ideological preferences. The base mobilization hypothesis, grounded in Hall and Thompson's (2018) analysis of turnout dynamics, offers a similar prediction through a different mechanism: an extreme opponent activates the incumbent's own base, and an externally activated base both permits and rewards more partisan positioning. Both predict a positive treatment effect on incumbents' ideological positioning (movement toward their party's pole). By contrast, under pure office motivation, a vote-share maximizing model (Downs, 1957) predicts the opposite: facing a more extreme opponent, the incumbent should move toward the opponent's position rather than away from it, capturing voters in the ideological space between them. This Downsian moderation hypothesis predicts a negative treatment effect on the same outcome.

To distinguish among these predictions, I identify the causal effect of challenger extremity using a regression discontinuity design in opposing-party primary elections. When an ideologically extreme candidate narrowly wins versus narrowly loses the opposing party's primary, the incumbent faces quasi-randomly assigned variation in whether their general election opponent is extreme or moderate—variation generated by an election in which the incumbent does not participate.

I measure incumbent response using four primary outcomes—party unity voting, ideological extremity (movement toward the own-party pole), votes with party leadership, and roll-call participation—supplemented by bill sponsorship counts as a secondary measure of legislative activity. Critically, I observe behavior in the post-primary, pre-general election period within the same Congress, so the sample is not conditioned on which incumbents survive the general election—a post-treatment outcome that challenger extremity itself plausibly affects.

I find no meaningful changes in incumbent voting behavior when facing extreme challengers. Across all five behavioral outcomes the estimated effects are small and statistically indistinguishable from zero, and the result is uniform: it persists across bandwidth choices, donut-hole exclusions, incumbent-clustered standard errors, and a restriction to the most ideologically extreme challengers. The estimates are precise enough to rule out large strategic shifts in either direction—the extremization that slack and turnout-based theories predict, and large Downsian moderation alike—even as the design retains less leverage over small, few-point effects. The dominant pattern is rigidity—incumbent legislative behavior does not track the ideology of the general-election opponent.

My study makes three contributions. First, I provide causal evidence on whether incumbents adjust their *legislative voting behavior* in response to opposing-party challenger extremity. A contemporaneous working paper by Schoenenberger (2024), using a closely related RDD design, reports that incumbents moderate in response to extremist challengers; my null finding adds evidence from a methodologically distinct approach—pre-primary DIME-based ideology scores and direct outcome measures across five behavioral channels rather than an indirect agreement-weighted single outcome—to an empirical question that remains unsettled. My results also complement Di Tella et al.'s (2025) analysis of *campaign language*: whatever adjustments incumbents make to campaign rhetoric, I do not observe corresponding changes in legislative behavior.

Second, combining my null finding with the positive findings of Meyer (2022) and Cowburn and Theriault (2025) for same-party primary challengers suggests a behavioral asymmetry: incumbent roll-call records appear to respond to threat existence but not, in this specification, to ideological variation within an already-present opposing-party threat. Because the pattern emerges from comparison across studies with different institutions, treatments, and designs, I advance this asymmetry as a hypothesis the present paper raises and partially probes rather than one it establishes; adjudicating it—and whether it reflects an intrinsic difference between threat types or the differential bindingness of the two electoral channels in the typical district—requires a unified design that observes both threat types for the same incumbents.

Third, I measure challenger ideology using pre-primary DIME scores—a measurement approach that captures candidate positioning before primary outcomes are determined, avoiding contamination by post-primary fundraising dynamics. This enables precise measurement of challenger extremity at the moment incumbents learn who their opponent will be. As scholarly attention to primary elections grows, this timing discipline—dating a candidate's ideology to the moment of nomination, before the post-primary broadening of the donor base can shift donor-based scores—offers a transferable measurement strategy that future studies of candidate positioning and primary-to-general-election dynamics can readily adopt.

These findings connect to the broader literature on political polarization. Hall (2015) demonstrates that extreme nominees face electoral penalties in general elections, and Fowler et al. (2023) show that voters prefer genuinely moderate candidates. These dynamics create conditions where incumbents might be expected to respond strategically. My null finding suggests that incumbents do not exploit this electoral latitude—legislative positions are sticky, and challenger characteristics alone do not move them. Polarization dynamics therefore operate more through candidate entry and selection (Thomsen, 2014) than through incumbent adjustment.

2 Theory and Hypotheses

Electoral accountability represents a cornerstone of democratic theory. Citizens hold representatives accountable through the threat of electoral defeat, and this threat should induce responsiveness to constituent preferences (Canes-Wrone et al., 2002). When legislators deviate from district preferences, they face electoral punishment. The logic suggests that incumbents should adjust their positions in response to electoral threats, including the identity and ideology of their challengers.

I examine how incumbents respond to a specific type of electoral threat: facing an ideologically extreme challenger from the opposing party. Extreme challengers are those whose policy positions lie far from the center of the ideological spectrum, in the direction of their party's pole. A liberal Democrat or a conservative Republican who wins their party's nomination creates a distinctive electoral context for the incumbent. Theory offers competing predictions spanning all three possible directions of response: two mechanisms predict movement toward the incumbent's own ideological pole, one predicts movement toward the opponent, and one predicts no movement at all.

2.1 The Electoral Slack Hypothesis

Classical spatial voting theory suggests that incumbents facing extreme challengers should have greater latitude to pursue their own ideological preferences. In Downsian models, voters choose the candidate closest to their ideal point. When an extreme candidate wins the opposing party's nomination, the incumbent gains electoral slack: they can move toward their own party's ideological pole while still remaining closer to the median voter than their extreme opponent—provided that the incumbent's ideal position is more extreme than their current position. Far from inducing moderation, an extreme challenger provides the incumbent with room to become more partisan.

This logic finds support in extensions of the basic Downsian framework. Calvert (1985) and Wittman (1983) show that policy-motivated candidates diverge from the median when electoral competition permits. An extreme challenger increases the space for such divergence. Groseclose (2001) demonstrates that candidates with valence advantages can afford more extreme positions; analogously, an incumbent whose opponent is ideologically distant enjoys a positional advantage that permits ideological indulgence.

I expect that if the electoral slack hypothesis holds, incumbents facing extreme challengers should increase their party unity voting and move toward more extreme ideological positions. The treatment effect on party unity should be positive as incumbents exploit their increased electoral latitude.

2.2 The Base Mobilization Hypothesis

A turnout-based perspective predicts extremization through a different mechanism. Hall and Thompson (2018) show that extreme nominees face electoral penalties primarily through differential turnout: while an extremist may energize their own party's base, they invigorate the opposing party's base even more strongly. The asymmetry follows from a concave spatial utility function—because voters dislike ideologically distant candidates at an increasing rate, the opposing base, already farthest from the extremist, responds most sharply. When an extreme challenger wins the opposing party's primary, the incumbent's own base is therefore activated from outside: mobilized not by anything the incumbent does, but by opposition to the extremist.

External activation can translate into more partisan incumbent behavior through either of two variants of the mechanism. In the freedom variant, turnout that ordinarily must be earned through positioning is now secured by the opponent's extremity; released from tending the base, the incumbent can indulge more partisan positions without risking demobilization. In the pressure variant, an activated base pays closer attention and demands loyalty, pulling the incumbent toward the base's preferred positions. The two variants attribute different motivations to the incumbent—one permissive, one coercive—but they push behavior in the same direction, toward the incumbent's own party pole. My design cannot separate them, and I treat base mobilization as a single hypothesis defined by its turnout logic rather than by which variant carries it.

Under the base mobilization hypothesis, the treatment effect on party unity should be positive. The prediction coincides in direction with electoral slack—both anticipate extremization—but the mechanisms differ: slack operates through the spatial logic of positioning relative to the median voter, base mobilization through the turnout calculus of who shows up. They are observationally equivalent in the sign of the predicted effect and cannot be distinguished on sign alone.

2.3 The Downsian Moderation Hypothesis

Office-motivated spatial competition points in the opposite direction. In a vote-maximizing framework (Downs, 1957), a candidate's optimal position tracks the electorate to be assembled, not the party's pole. When the opposing party nominates an extremist, the ideological space between the incumbent and the opponent widens, and the voters in that space become winnable: by moving toward the opponent's position, the incumbent captures moderates who would otherwise divide between the candidates. The empirical record on general elections reinforces the incentive. Extreme nominees systematically underperform at the polls (Hall, 2015), and voters reward genuinely moderate candidates (Fowler et al., 2023), so an office-motivated incumbent handed an extremist opponent has a clear play available—moderate, and let the contrast do the campaigning.

The moderation hypothesis shares with electoral slack the premise that an extreme opponent relaxes the incumbent's electoral constraint, but the two differ in what the incumbent is assumed to maximize. Electoral slack describes a policy-motivated incumbent who spends the windfall on ideological indulgence; Downsian moderation describes an office-motivated incumbent who converts the same windfall into votes. Under the moderation hypothesis, the treatment effect on ideological extremity should be negative—incumbents move away from their own party's pole—and party unity should, if anything, fall as incumbents break with their party to occupy the vacated center.

2.4 The Inertia Hypothesis

A final perspective, grounded in Lee et al.'s (2004) analysis of policy divergence, predicts no movement in either direction: legislative positions are largely sticky. Representatives' voting behavior reflects personal ideology, party discipline, and constituency preferences that are largely fixed in the short term. Strategic repositioning is costly: it alienates existing supporters, invites charges of inconsistency, and may conflict with genuine policy commitments.

The empirical record supports this stickiness. Ansolabehere et al. (2001) show that House candidates position themselves with their party far more than with their district, leaving little scope for opponent-specific tailoring. Lee et al. find that barely-winning Democrats and Republicans do not converge in their voting behavior, suggesting limited electoral responsiveness even in highly competitive districts, and Fowler and Hall (2016) confirm the elusiveness of convergence across multiple empirical settings. If legislative positions are driven by factors other than strategic calculation vis-a-vis specific opponents, challenger characteristics should elicit no response. Nor is within-session adjustment easy even for a willing incumbent: roll-call votes are cast in a party-organized legislature with strong floor agenda control, committee structures, and leadership cues that constrain individual members' choices. Short-term deviation from established voting patterns requires not only willingness but also opportunity—and party-line votes, by definition, leave little room for individual repositioning.

Under the inertia hypothesis, I expect incumbents to maintain their positions regardless of whether they face an extreme or moderate challenger. The treatment effect on party unity should be close to zero as incumbents exhibit position rigidity.

2.5 Distinguishing Among Hypotheses

My empirical design distinguishes among three directional predictions: movement toward the incumbent's own pole (positive treatment effects, consistent with electoral slack and base mobilization), movement toward the opponent (negative effects, consistent with Downsian moderation), and no movement (null effects, consistent with inertia). The ideological extremity outcome, constructed symmetrically as movement toward the incumbent's own party pole, is the most direct adjudicator among the three: its sign separates extremization from moderation, and its magnitude separates both from inertia. Electoral slack and base mobilization cannot be distinguished from each other on sign alone; I maintain them as separate hypotheses because they derive from distinct theoretical traditions (spatial models versus turnout models) and would generate different predictions in designs capable of measuring turnout composition or swing-voter behavior directly. One asymmetry in the test deserves emphasis in advance: a positive or negative estimate can affirmatively support extremization or moderation, but a null supports inertia over moderation only insofar as the confidence interval is tight enough to exclude meaningful negative effects. I return to this standard when interpreting the results.

Theory is genuinely ambiguous about which prediction should dominate, and this ambiguity is precisely what makes the research question important. The electoral slack, base mobilization, and Downsian moderation logics all assume that incumbents respond to shifts in electoral competition—through policy-motivated indulgence, turnout-driven freedom, or office-motivated vote-seeking, respectively. The inertia hypothesis, grounded in the extensive literature on legislative stability (Poole and Rosenthal, 2007; Lee et al., 2004; Fowler and Hall, 2016), suggests that none of these adjustments occurs in practice because legislative positions reflect stable commitments that are costly to change.

2.6 What the Null Identifies (and What It Does Not)

The treatment my design isolates is specific: variation in an opposing-party general-election challenger's ideology, conditional on a contested primary having produced one of two close winners. This is variation in challenger *ideology*, not in challenger *presence*. The incumbent already faces an opposing-party challenger of some ideological type; my RDD asks whether the challenger being one step more extreme alters subsequent roll-call behavior.

A related boundary concerns timing. Strategic-politician accounts hold that electoral competition does much of its work before candidates are even nominated: prospective challengers enter selectively, and incumbents position themselves in part to deter threatening entry (Jacobson and Kernell, 1983). If incumbents adjust their records in anticipation of the challenger they might face, a post-primary window observes behavior after the relevant adjustment has occurred, and a post-primary null could coexist with substantial earlier responsiveness. The design bounds this concern in two ways. Anticipatory positioning common to both sides of the threshold—adjustment to the mere possibility of an extremist opponent—is differenced out by the close-election comparison and lies outside the estimand; the question here is whether the realized identity of the opponent moves behavior once the primary resolves it. Anticipation that instead tracks the likely primary winner would leave a footprint the design can detect: incumbents just above and just below the threshold would enter the post-primary period with different records. The lagged-outcome placebo tests in Appendix Table A4 find no pre-primary discontinuity in any of the five outcomes, so within the close-primary sample, incumbents about to face an extremist were not already adjusting in advance. What remains genuinely open is responsiveness concentrated entirely at the entry and deterrence stage—a channel this design cannot reach and a natural target for future work.

The closest comparison is a contemporaneous working paper by Schoenenberger (2024), which applies the same close-opposing-party-primary RDD to roughly the same question and reports incumbents moderating in response to extremist challengers—the response the Downsian moderation hypothesis predicts—with effects concentrated in marginal districts. Two methodological differences separate his design from the present one and offer the most traceable explanations for the divergence between his positive finding and this paper's null. First, candidate ideology measurement: Schoenenberger uses Hall-Snyder donation-weighted scores, whereas the present paper uses DIME-based scores restricted to pre-primary contributions (see Research Design). Both approaches recover candidate ideology from donation data, but they construct and date the measure differently, so the two designs need not classify the same near-threshold primaries as contests between an extremist and a moderate. Second, outcome measure: his primary outcome is an agreement-rate-weighted estimate of post-primary roll-call extremism that imputes each incumbent's position from peers in the same Congress, whereas the present paper uses five direct measures of the incumbent's own legislative behavior. His effective sample for the marginal-district heterogeneity finding is restricted to 1996 onward by data availability, overlapping substantially with the 2000–2018 sample used here, so the divergence is unlikely to be a period effect.

Meyer (2022) and Cowburn and Theriault (2025) offer a separate kind of comparison: they study same-party primary threats, where the treatment is the existence and credibility of an intraparty challenger rather than the ideology of an already-present opposing-party challenger. Their positive findings combined with the present null have a more straightforward interpretation than the Schoenenberger contrast: incumbent roll-call behavior responds to threat existence regardless of ideological content, while the divergence with Schoenenberger is methodologically traceable rather than channel-based.

A second alternative interpretation attributes the cross-study pattern to which threat is binding rather than to anything intrinsic about threat type. Most U.S. House general elections are not closely contested; an incumbent facing an extreme rather than moderate opposing-party challenger may rationally ignore the variation when the general-election outcome is dominated by district partisanship. Credible same-party primary challenges, by contrast, typically emerge in safely partisan districts where the primary is the genuine contest. Under this account, the observed cross-study asymmetry reflects the differential bindingness of the two thresholds in observed samples, not a structural property of inter- versus intraparty competition. Schoenenberger's finding that opposing-party effects concentrate in marginal districts is consistent with this bindingness logic: where the general election is genuinely at stake, incumbents respond to opposing-party challenger ideology; where it is not, they do not.

The district-competitiveness analysis I report below provides a direct empirical challenge to the marginal-district heterogeneity reading. Splitting the 172 baseline races at the median district presidential margin, I find no detectable effect of challenger extremity in either the more-competitive or the safer subsample (gap of 1.5 percentage points, statistically indistinguishable from zero). Under the marginal-district bindingness reading, the more-competitive subsample should produce a positive response, since the general election is genuinely at stake; the absence of any response in that subsample weighs against the bindingness account as the full explanation. The check is partial in two respects. First, the split threshold is a 7-point presidential margin, so the more-competitive half is not composed exclusively of toss-ups—though its own median margin of 3.9 points places it much closer to genuine swing-district territory than the threshold suggests, which strengthens the force of the null within it. Second, even in competitive general elections, the marginal leverage of challenger ideology specifically—relative to challenger resources, name recognition, or campaign quality—on the incumbent's vote share may be small, consistent with the modest electoral penalty for extreme nominees that Hall (2015) documents.

The empirical question Schoenenberger and the present paper both ask—does opposing-party challenger ideology move incumbent legislative behavior?—therefore remains unsettled across contemporaneous working-paper evidence. The present paper contributes a null finding that is methodologically distinct from his positive result in ways that matter: challenger ideology dated to the pre-primary period, direct outcomes across five behavioral channels rather than an indirect agreement-weighted single outcome, and robustness across competitiveness subsamples that directly challenges his heterogeneity reading. Whether the divergence ultimately reflects measurement-method sensitivity, outcome granularity, the specific way bindingness interacts with district characteristics, or other specification differences is itself a substantive question for replication-oriented work. What can be said now is that the broadest theoretical predictions of incumbent strategic adjustment to opposing-party challenger ideology—predictions drawn from spatial-positioning and base-mobilization models—are not consistently supported across closely related empirical designs, and the channel they posit may operate under narrower conditions than the theory itself acknowledges.

3 Research Design

3.1 Identification Strategy

I identify the causal effect of facing an extreme challenger using a regression discontinuity design in opposing-party primary elections. The key insight is that in close primary elections, whether the extreme candidate wins or loses approaches random assignment. Incumbents whose opposing-party primaries are decided by narrow margins face quasi-randomly assigned variation in challenger extremity.

The running variable in my design is the vote margin by which the most extreme candidate won or lost the opposing party's primary among the two finalists. I operationalize it as the extremist candidate's vote share among the top two finishers minus 0.5, which is equivalent to half the finalist margin. Positive values indicate the extreme candidate won (incumbent faces an extreme challenger); negative values indicate the extreme candidate lost (incumbent faces the other finalist, who is relatively more moderate). The threshold is zero, where the two finalists tied.

This design addresses the Marshall (2024) critique of politician characteristic regression discontinuity designs. Marshall demonstrates that RDD estimates of politician characteristics suffer from posttreatment bias because characteristics that help candidates win are confounded with compensating differentials that keep elections close. My design sidesteps this critique by using the opposing non-incumbent party's primary as the source of variation. The incumbent does not participate in this election, so compensating differentials that affect the incumbent's characteristics do not apply.

3.2 Sample Construction

My analysis draws on U.S. House elections from 2000 to 2018, focusing on incumbents seeking reelection who face contested opposing-party primaries. The final scored analysis file contains 344 incumbent-race observations with an identifiable top-two opposing-party primary, measured ideology for both top-two candidates using DIME contribution data, and observed post-primary incumbent outcomes. My baseline specification then restricts attention to contests with above-median ideological contrast between the top-two primary candidates ($absdist \geq 0.154$), yielding the 172-race analytic sample used in the main text. Appendix Table A9 summarizes this analysis-sample flow.

The effective sample for RDD estimation is smaller: the MSE-optimal bandwidth selection (Calonico et al., 2014) yields approximately 109 observations near the threshold ($N_{\text{eff}} = 109$). This is the number of observations within the optimal bandwidth on either side of the cutoff, not the full analytic sample. The distinction matters for interpreting statistical power: my standard errors and minimum detectable effects are functions of N_{eff}, not N.

In multi-candidate primaries, I restrict attention to the top two finishers and designate the more ideologically extreme of those two finalists as the "extreme" candidate (most liberal for Democrats, most conservative for Republicans). The running variable is then computed from that finalist pair.

3.3 Measuring Challenger Ideology: Dynamic Pre-Primary Campaign Finance Score

A key measurement approach in this paper is my construction of dynamic pre-primary DIME scores to measure challenger ideology. The Database on Ideology, Money in Politics, and Elections (DIME; Bonica, 2014) estimates candidate ideology from the mix of campaign contributions they receive. Bonica's dynamic CFscore methodology works as follows: (1) run correspondence analysis on the full dataset to estimate static contributor ideal points and static recipient ideal points simultaneously; (2) fix the static contributor ideal points; (3) for each recipient-period (election cycle), calculate a period-specific recipient CFscore as the weighted average of contributors' static CFscores, weighted by donation amount; (4) require at least 25 unique contributors per cycle for inclusion.

I extend this dynamic CFscore approach by modifying which donations enter the period-specific calculation. Standard dynamic CFscores incorporate all contributions a candidate receives throughout an election cycle. I instead include only contributions from two sources: the pre-primary period of the current cycle, and an optimal number of earlier cycles.

The pre-primary restriction ensures that my ideology measure captures the candidate's ideological profile at the moment the incumbent learns who their opponent will be, rather than being contaminated by post-primary fundraising dynamics. Post-primary contributions may reflect strategic giving by donors responding to nomination outcomes rather than initial ideological positioning.

The optimal inclusion window addresses the trade-off between using more information (longer windows with more contribution history) and capturing current ideological positioning (shorter windows that reflect recent donor coalitions). I systematically test different lookback windows (n = 0, 1, 2, 3, 5, 10, and all available prior cycles) and select the window that maximizes the correlation with career DW-NOMINATE scores, following Bonica (2014). The optimal lookback is n = 3 prior cycles: the correlation with career DW-NOMINATE is r = 0.849 (N = 8,015), while correlations are essentially flat for $n \geq 1$, indicating that additional lookback beyond 3 cycles adds negligible information.

I validate my pre-primary scores against three independent benchmarks; the correlations reported here are pooled across parties. First, the correlation with career DW-NOMINATE scores is r = 0.849 (N = 8,015). Second, the correlation with DIME dynamic CFscores is r = 0.869 (N = 11,714). Third, the correlation with Nokken-Poole scores (Nokken and Poole, 2004) is r = 0.791 for all winners (N = 6,357) and r = 0.733 for first-term members only (N = 103). The first-term correlation provides the most demanding test: it asks whether pre-primary contribution patterns predict the future voting behavior of candidates who have never served, and is therefore free of circularity concerns.

3.4 The Post-Primary Observation Window

I observe incumbent behavior in the post-primary, pre-general election period within the same Congress. Congressional sessions run from January through the end of the year, while primary elections are spread from March through September, most concentrated in the spring and early summer; the post-primary period therefore typically encompasses a substantial portion of the legislative session. It is also the period when electoral concerns are most heightened: the incumbent knows the identity and ideological profile of the general-election challenger and faces an imminent election, so if strategic repositioning occurs, this is when it should be observable. Measuring outcomes before the general election has the additional virtue that the sample is not conditioned on which incumbents survive it—a post-treatment outcome that challenger extremity itself plausibly affects. Primary timing does generate variation in window length: states with early primaries (March–June) produce longer post-primary windows than states with late primaries (August–September), a variation that bears on the measurement of bill sponsorship in particular and that Appendix Table A10 examines directly.

3.5 Outcome Variables

I examine five outcome measures—four roll-call-based outcomes and one exploratory measure of legislative activity—capturing different dimensions of legislative behavior in the post-primary, pre-general election period:

1. **Party Unity Change**: The change in the proportion of party-line votes—those in which a majority of one party opposes a majority of the other—where the incumbent votes with their party majority. This is my primary outcome.
2. **Ideological Extremity**: Change in voting in the direction of the incumbent's own party pole. Constructed as the change in conservative voting score multiplied by party direction (+1 for Republican incumbents, -1 for Democratic incumbents), so that a positive value denotes movement *toward* the incumbent's own ideological pole regardless of party.
3. **Votes with Party Leader**: Change in the proportion of votes aligned with party leadership positions.
4. **Roll-Call Participation**: Change in overall roll-call participation rate.
5. **Bill Sponsorship**: Change in bill sponsorship activity, measured as a pre-computed change score. I treat this outcome as exploratory: earlier primaries mechanically produce longer post-primary windows and therefore more opportunity for sponsorship, and although the change score accounts for session-level variation, the count-based measure warrants more caution than the four proportion-based roll-call outcomes.

All outcomes are measured as changes from the pre-primary to post-primary period within the same Congress. This differencing removes time-invariant incumbent characteristics and focuses on within-session behavioral changes. For roll-call voting outcomes, the proportion-based measurement is naturally scale-free.

3.6 Estimation

I implement the bias-corrected RD estimator with MSE-optimal bandwidth and a triangular kernel (Calonico et al., 2014; De Magalhães et al., 2025).

The estimand is the average treatment effect at the threshold: the difference in incumbent behavior when the extreme candidate barely wins versus barely loses the opposing primary. Identification rests on continuity: expected potential outcomes evolve smoothly through the threshold, so incumbents just above and just below it are comparable in expectation and the discontinuity carries a causal interpretation.

3.7 Validity Tests

The McCrary (2008) density test yields $t = -1.03$ ($p = 0.303$), providing no evidence of manipulation at the threshold (Appendix Figure A1).

Pre-treatment incumbent and district characteristics show no discontinuity at the threshold; Table A1 reports the full balance results. The one exception is ideological distance between incumbent and challenger, where the imbalance is marginal ($p = 0.052$) but expected: by construction, extreme challengers are more ideologically distant from incumbents.

Three considerations address the ideological-distance imbalance. First, and most importantly, ideological distance is partly a post-treatment characteristic: it mechanically changes when the treatment (extreme vs. moderate challenger) changes, making it an invalid covariate for balance testing (Eggers et al., 2015). Second, controlling for ideological distance in the RDD specification does not materially change the estimated treatment effect (Section 5). Third, donut-hole RDD estimates excluding near-threshold observations show that the imbalance attenuates (Appendix Table A3), suggesting it is concentrated among the closest races.

Placebo cutoffs at margins of -10, -5, $+5$, and $+10$ percentage points yield no significant discontinuities for the primary outcome (Appendix Table A2), confirming that the party-unity result is concentrated at the true threshold.

To rule out anticipation effects, I estimate the RDD using pre-treatment outcomes as dependent variables. Incumbent DW-NOMINATE scores and incumbent ideological extremity show no discontinuity at the threshold in the full baseline sample (Appendix Table A4). In the subset of 38 baseline incumbents observed in consecutive Congresses, prior-Congress party unity, votes with party leader, ideological extremity, and participation likewise show no discontinuity ($p \geq 0.12$). These checks are noisier given the small effective sample but provide no evidence of pre-existing behavioral trends.

4 Results

4.1 Main Finding: No Detectable Legislative Response to Extreme Challengers

Figure 1 presents the RD plot for my primary outcome, party unity change, showing binned averages and local polynomial fits on each side of the threshold. There is no visible discontinuity at the cutoff.

Incumbents do not change their voting behavior when facing ideologically extreme challengers from the opposing party. The estimated effect on party unity is 1.7 percentage points ($p = 0.236$). The 95% confidence interval ranges from −1.1 to 4.6 percentage points, comfortably including zero. This point estimate is substantively small relative to baseline party unity rates of 85–90%: even the upper bound of the confidence interval would represent a modest shift in partisan voting.

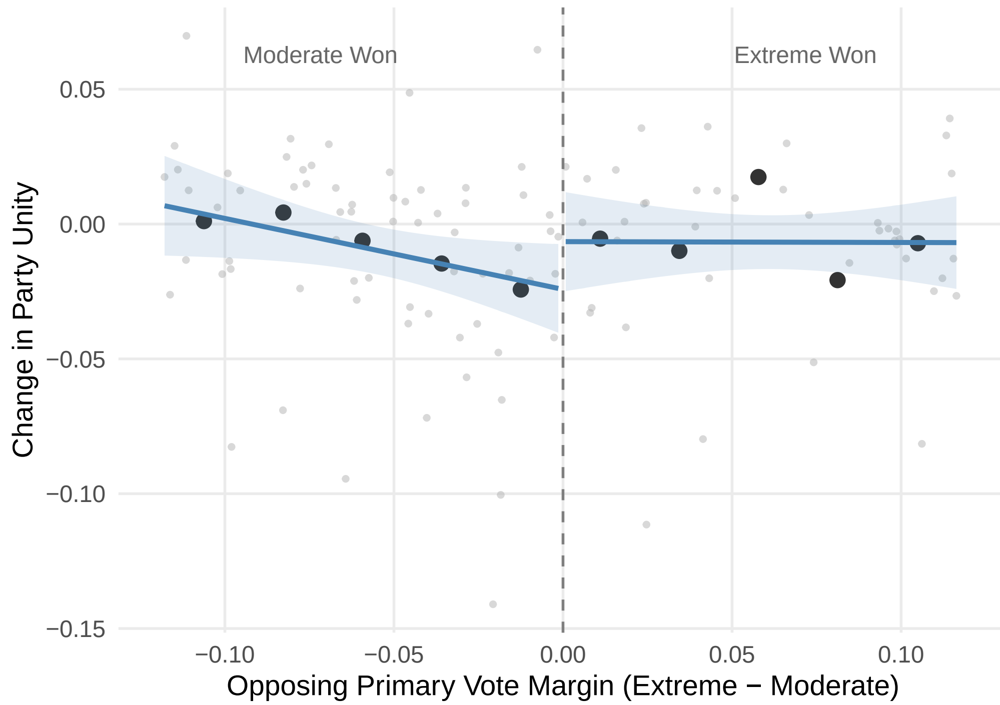

Figure 1: RD Plot for Party Unity Change

Note: Binned means and local polynomial fits are shown on each side of the opposing-party primary threshold. The running variable is the extreme candidate's top-two vote-share margin.

4.2 Statistical Power and Minimum Detectable Effect

At the MSE-optimal bandwidth ($N_{\text{eff}} = 109$), my design is powered to detect effects of at least 4.1 percentage points at 80% power (α = 0.05, two-tailed). The 95% confidence interval of [-1.1, 4.6] rules out effects larger than about 4.6 percentage points. This is an informative null: the confidence interval rules out the large positive shifts that slack and turnout theories predict, and its lower bound excludes all but modest declines in party unity of the kind Downsian moderation implies. Power varies across outcomes, because each is estimated at its own MSE-optimal bandwidth: the minimum detectable effect at 80% power is 3.7 percentage points for votes with party leader (effective sample 94), 4.9 for ideological extremity (57), 5.3 for roll-call participation (79), and 15.5 for bill sponsorship (76; Appendix Table A11). The sponsorship figure is an outlier that reinforces treating that outcome as exploratory—the design is barely powered to detect anything short of an implausibly large sponsorship shift. What the design cannot do is reject small effects, and a formal equivalence test makes this concrete. Anchoring the equivalence bound to the roughly 3-percentage-point same-party response reported by Meyer (2022)—a magnitude in the range this literature typically reports, which I treat as approximate rather than source-pinned—two one-sided tests fail to establish equivalence to zero at a 3-percentage-point bound for any of the five outcomes (party unity p = 0.191). Even at a 4-point bound, party unity narrowly misses the equivalence bar (p = 0.060), and only the two most precisely estimated outcomes, votes with party leader and participation, clear it (Appendix Table A12). A same-party-sized effect therefore cannot be statistically rejected: the null is informative about large strategic shifts, not about effects of a few percentage points.

Table 1: Main Results—Effect of Facing Extreme Challenger on Incumbent Behavior

	(1) Party Unity	(2) Extremity	(3) Party Leader	(4) Participation	(5) Sponsorship
Panel A: Baseline sample—above-median ideological contrast ($N = 172$)					
Extreme Challenger	0.017	−0.017	0.014	0.001	−0.002
	(0.015)	(0.018)	(0.013)	(0.019)	(0.055)
	[0.236]	[0.328]	[0.308]	[0.969]	[0.978]
95% CI Lower	−0.011	−0.052	−0.013	−0.036	−0.110
95% CI Upper	0.046	0.017	0.040	0.038	0.107
Effective N	109	57	94	79	76
Panel B: Full scored sample—no contrast restriction ($N = 344$)					
Extreme Challenger	0.019	0.007	0.014	0.045	0.041
	(0.012)	(0.012)	(0.012)	(0.048)	(0.049)
	[0.107]	[0.574]	[0.234]	[0.344]	[0.404]
95% CI Lower	−0.004	−0.017	−0.009	−0.048	−0.056
95% CI Upper	0.042	0.031	0.036	0.139	0.138
Effective N	203	161	188	165	146

Note: Local linear regression with a triangular kernel and MSE-optimal bandwidth, estimated separately in each panel. Entries are robust bias-corrected point estimates, standard errors (in parentheses), and p-values (in brackets), on the outcome's proportion scale ($0.017 = 1.7$ percentage points). Panel A restricts to races with above-median ideological contrast between the two primary finalists (the baseline analytic sample used throughout the main text); Panel B uses the full scored sample without that restriction. All outcomes are measured as changes from the pre-primary to the post-primary period.

4.3 Treatment Contrast at the Threshold

A key consideration for interpreting my null finding is the treatment contrast in actual challenger ideology at the threshold. Appendix Table A5 and Figure A2 report the first-stage discontinuity: when the extreme candidate barely wins versus barely loses the opposing primary, the winning nominee's DIME score shifts by 0.135 points, an estimate statistically indistinguishable from zero. This contrast is modest relative to the full-sample variation in challenger ideology, reflecting the fact that the two finalists in close primaries are often not dramatically different in their ideological positioning.

This modest first stage does not invalidate my design—the treatment in a sharp RDD is whether the extreme candidate wins, not the magnitude of the ideology gap, and the reduced-form estimate directly captures the causal effect of this binary treatment. However, it does bear on interpretation. Under a weak-treatment reading, the "extreme" and "moderate" challengers produced by close primaries may not differ enough to elicit a behavioral response. I think the weight of evidence favors the inertia interpretation for two reasons. First, the null holds even at the Hall-style top-quartile restriction (3.9 pp, statistically indistinguishable from zero; Section 4.7), which retains only races with large ideological contrasts between primary candidates and in which the first-stage discontinuity strengthens materially (Appendix Table A5). Second, the null extends uniformly across five distinct outcome measures—a pattern more consistent with genuine non-responsiveness than with a treatment that simply failed to register.

4.4 Robustness Across Outcomes

The null finding extends to all five outcome variables, as shown in Table 1. Votes with party leader shows an effect of 1.4 percentage points, roll-call participation 0.1 percentage points, and bill sponsorship -0.2 percentage points; all are statistically indistinguishable from zero. The ideological extremity outcome carries the three-way adjudication among extremization, moderation, and inertia, and deserves a closer reading. Its point estimate is -1.7 percentage points ($p = 0.328$)—negative, the direction Downsian moderation predicts—with a 95% confidence interval running from -5.2 to +1.7 percentage points. The upper bound rules out all but trivial movement toward the incumbent's own pole, clear evidence against the electoral slack and base mobilization hypotheses. The lower bound excludes large moderation, but the interval cannot separate a small moderation effect from exactly zero; this outcome is also the least powered in the design (effective sample of 57, barely half of party unity's 109), so the inertia reading rests on the estimate's magnitude and on the uniformity of the pattern across outcomes rather than on a decisive rejection of moderation. Across the four roll-call outcomes and the exploratory sponsorship measure alike, the estimates are substantively small—the pattern inertia predicts and neither directional alternative does.

4.5 Robustness to Specification Choices

I assess the stability of my findings across alternative specifications. Table 2 presents estimates for the primary outcome (party unity) across different bandwidth choices, polynomial orders, kernel specifications, and with covariates.

Table 2: Robustness of Party Unity Estimates Across Specifications

Specification	Estimate	SE	p-value	Effective N
Panel A: Bandwidth Variation				
$0.5\times$ Optimal	0.004	0.022	0.851	57
$0.75\times$ Optimal	0.009	0.018	0.617	79
$1.0\times$ Optimal	0.011	0.016	0.470	109
$1.25\times$ Optimal	0.015	0.015	0.318	119
$1.5\times$ Optimal	0.016	0.014	0.276	129
$2.0\times$ Optimal	0.018	0.014	0.203	153
Panel B: Polynomial Order				
Local linear (baseline)	0.017	0.015	0.236	109
Quadratic	0.007	0.018	0.682	95
Panel C: Kernel Choice				
Triangular (baseline)	0.017	0.015	0.236	109
Uniform	0.021	0.015	0.170	99
Epanechnikov	0.020	0.015	0.172	101
Panel D: Covariate Adjustment				
No added covariates	0.017	0.015	0.236	109
+ Ideological Distance	0.019	0.015	0.199	110

Note: All entries are robust bias-corrected `rdrobust` estimates for party unity change.
Panel A fixes the bandwidth at multiples of the main optimal choice; Panels B–D use the Table 1 baseline as their reference row.
The reproducible control specification currently available in the live pipeline adds incumbent-challenger ideological distance.

Results are stable across bandwidth variation, polynomial order, and kernel choice. Across fixed bandwidths from $0.5\times$ to $2.0\times$ the party-unity estimate ranges from 0.4 to 1.8 percentage points and never approaches conventional statistical significance. Quadratic and alternative-kernel specifications remain similarly small and null. Adding ideological distance as a control—addressing the covariate imbalance concern—yields an estimate of 1.9 percentage points, very close to the main estimate, confirming that the imbalance does not confound the substantive conclusion. Dependence from repeated incumbents is also limited: the 172-race baseline sample contains 158 unique incumbents, and only 13 incumbents appear more than once. Re-estimating the main specification with incumbent-clustered standard errors yields a party-unity estimate of 1.8 percentage points, with similarly negligible changes across the other outcomes (Appendix Table A6). All specifications remain non-significant, confirming the robustness of the null finding.

4.6 Exploratory Heterogeneity by District Competitiveness

As a direct district-baseline competitiveness check, I merge recent presidential vote share by congressional district and convert it into incumbent-party support in the district. I define district competitiveness as the absolute distance between that incumbent-party presidential vote share and 50%, so lower values indicate electorally closer districts. The measure covers all baseline races: for Congresses 108-115 I use Daily Kos Elections' precinct-aggregated district-level returns; for the 1990s-redistricting Congresses 106-107 and a handful of districts affected by mid-decade redistricting that Daily Kos does not cover, I aggregate Dave Leip county-level returns to congressional districts using the Ferrara et al. (2022) population-based crosswalks for each Congress's specific boundary file. Splitting at the median competitiveness distance (0.070) yields 86 more competitive and 86 safer districts.

Using the pooled-sample MSE-optimal bandwidth for both subgroup estimates, the party-unity estimate is 0.4 percentage points in more competitive districts and 1.9 percentage points in safer districts; the pooled estimate for this restricted sample is 1.7 percentage points. The corresponding estimates for ideological extremity, votes with party leader, roll-call participation, and sponsorship are likewise small, with no consistent competitive-safer pattern.

The competitive-safer gap on the primary outcome is 1.5 percentage points, well within sampling variation, and no other outcome shows a competitive-safer difference outside the baseline confidence interval. A supplemental split by lagged incumbent general-election vote (38 races; see Appendix) likewise shows no consistent subgroup pattern. I therefore find no evidence that incumbents respond more strongly to extreme challengers in electorally competitive districts.

4.7 Sample Restriction by Ideological Distance

Following Hall (2015), I progressively tighten the sample inclusion criterion to retain only races with a stark ideological contrast between primary candidates. The top-quartile restriction (absdist ≥ 0.297, N = 86) yields a point estimate of 3.9 percentage points for party unity, still statistically indistinguishable from zero, and the top-decile restriction (absdist ≥ 0.441, N = 35) produces an estimate of exactly 0.0 percentage points. These restrictions do not merely shrink the sample; they sharpen the treatment. The first-stage discontinuity in winner ideology rises from 0.14 DIME units at baseline to 0.33 at the top quartile and 0.58 at the top decile, where it is conventionally significant ($p = 0.022$, $F = 5.27$; Appendix Table A5). The reduced-form null therefore cannot be attributed to an absent ideological contrast: where the design delivers its sharpest separation between extreme and moderate challengers, the party-unity response is zero.

The baseline restriction is not driving the null from the other direction either. Estimated on the full scored sample of 344 races, without the above-median-contrast gate, the party-unity effect is 1.9 percentage points ($p = 0.107$; see Table 1, Panel B) with tighter precision than the baseline, and a continuous sweep of the contrast threshold across 46 cuts finds no threshold at which the effect is conventionally significant (the sweep peaks at 3.8 percentage points, $p \approx 0.10$; Appendix Figure A4). The median cut used in the main text neither manufactures nor suppresses a result.

Bill sponsorship shows a marginally significant positive effect at the top quartile (0.147) and a nominally significant effect at the top decile (0.257). These results do not survive multiple testing correction. Within this section's 15 tests (5 outcomes $\times$ 3 sample restrictions), the Bonferroni threshold is $p < 0.003$, and Benjamini-Hochberg correction yields an adjusted p-value of 0.108 for the top-decile result. A uniform count across the paper's full specification grid makes the point more forcefully: of the 101 tests of the challenger-extremity estimand reported anywhere in this paper—across outcomes, bandwidths, donut holes, sample restrictions, and subgroup splits—exactly three are nominally significant at the 5% level (the top-decile sponsorship estimate; the 2-percentage-point donut extremity estimate; and the pooled extremity estimate in the sparse 38-race lagged-general-election-vote subsample, $p = 0.004$), against 5.05 expected by chance under a global null. None of the three survives Benjamini-Hochberg or Bonferroni correction applied uniformly to the grid; the observed nominal-positive rate falls below the chance rate (Appendix Table A13). The top-decile sponsorship estimates additionally rest on very small effective samples (36 and 18 observations), and a timing check indicates they are not driven by longer post-primary windows in early-primary states: in the 54-race subsample with observed primary dates, splitting at the median primary month yields null sponsorship estimates both for May-or-earlier (-0.044) and June-or-later primaries (-0.032; Appendix Table A10). I therefore treat the bill sponsorship results as exploratory and do not interpret them as confirmatory evidence of an effect.

5 Conclusion

Incumbents do not adjust their legislative behavior when they face ideologically extreme challengers from the opposing party. This null finding is robust across four roll-call outcomes and an exploratory sponsorship measure, ten specifications, donut-hole exclusions, and subsample restrictions. Opposing-party challenger extremity does not move roll-call voting.

The asymmetry across studies is the more forward-looking contribution. Combining my null finding with the positive findings of Meyer (2022) and Cowburn and Theriault (2025) for same-party primary challengers suggests that accountability operates selectively: legislators respond to threats that imperil renomination, but perhaps not to threats that merely adjust the general-election margin. Cross-study comparison cannot isolate this asymmetry from differences in institution, treatment, and design, so I advance it as a hypothesis the present evidence motivates rather than establishes—and as the natural target for a unified within-study test.

Of the three directional predictions the theory section develops, the two extremization hypotheses fare worst: electoral slack and base mobilization both predict positive treatment effects, and the extremity outcome's confidence interval excludes all but trivial movement toward the incumbent's own pole. Downsian moderation survives only in attenuated form: the extremity point estimate is negative, as moderation predicts, but it is small and statistically indistinguishable from zero, and the interval excludes moderating shifts larger than about five percentage points while remaining unable to separate a small moderation effect from exactly zero. The pattern most consistent with the full set of estimates is inertia, and it aligns with Lee et al.'s (2004) finding that legislative positions are sticky: incumbents' roll-call records reflect stable commitments that do not bend to shifts in electoral competition. The treatment contrast at the threshold is modest at baseline (Appendix Table A5), but the possibility that a sharper contrast would elicit a response has been tested directly rather than left open: the first-stage separation strengthens roughly fourfold under the Hall-style restrictions, and the null persists exactly where the ideological contrast is sharpest. The result extends across all five outcome measures. The most parsimonious reading is that opposing-party challenger identity is simply not a lever that moves legislative behavior.

These findings bear on the polarization debate, within limits the design makes explicit. Polarization in Congress over the past two decades has been documented through multiple channels: ideological replacement of moderates by extremists (Thomsen, 2014); sorting of partisan donors and activists (Bonica, 2014); changes in party leadership control over the floor agenda. A further channel—incumbent adaptation to a polarized opposition—has often been assumed rather than tested. My findings show that one specific version of this channel does not operate: when incumbents face an unusually extreme opposing-party challenger rather than a mainstream one, they do not shift their roll-call records in response. The results do not license the broader conclusion that adaptation plays no role in polarization—incumbents demonstrably respond to same-party threats (Meyer, 2022; Cowburn and Theriault, 2025), and adaptation at the candidate-entry stage or outside the roll-call record is untested here. What the results support is narrower: for the specific mechanism of sitting incumbents extremizing in response to a polarizing opposition, the evidence points toward selection—who enters and stays in Congress—rather than in-office adjustment as the operative source of legislator-level polarization.

The findings also speak to electoral accountability more broadly. Standard accountability models hold that incumbents respond to electoral threats by adjusting their behavior toward constituent or median-voter preferences (Canes-Wrone et al., 2002). Combining the present null with positive findings from same-party primary contexts (Meyer, 2022; Cowburn and Theriault, 2025) yields a narrower implication than "all electoral threats discipline behavior": incumbent roll-call behavior responds to the existence of credible threats but does not respond to ideological variation within an already-present opposing-party threat. Whether this distinction reflects a structural property of the two threat types or, instead, the differential bindingness of primary versus general-election competition in the typical district is not resolved by cross-study comparison alone. What the evidence already supports is the narrower claim: that the channels through which electoral accountability operates are more specific than the broadest accountability claims imply, and that strategic-positioning models predicting incumbent adjustment in response to opposing-party challenger ideology are not borne out by within-Congress roll-call evidence.

Several scope conditions bound these conclusions. My design rules out effects larger than 4.6 percentage points but not smaller ones; I establish that opposing-party challengers do not produce large behavioral shifts, not that the effect is exactly zero. My outcomes are legislative and positional—incumbents may respond instead through their issue agendas, which demonstrably absorb challengers' campaign themes (Sulkin, 2005), or through campaign rhetoric (Di Tella et al., 2025) and fundraising, without changing roll-call positions. Baseline party unity of 85–90% constrains upward movement, though the null extends to less ceiling-constrained outcomes. My sample spans 2000–2018; exploratory temporal splits (Appendix Table A7) are imprecise. And my results speak to the U.S. House; the Senate, with higher-salience elections, may differ.

The most productive next step is a unified design that tests same-party and opposing-party threats for the same incumbents, directly measuring the asymmetry I identify across studies. Extending the analysis to non-legislative channels—campaign messaging, fundraising, constituency service—would clarify whether the rigidity I document in roll-call behavior extends to other dimensions of representation. And replicating the design in the Senate, where elections are more salient and states more heterogeneous, would test whether legislative inertia is a general feature of congressional behavior or specific to the House.

References

Ansolabehere, S., Snyder, James M., J., and Stewart, Charles, I. (2001). Candidate positioning in U.S. house elections. *American Journal of Political Science*, 45(1):136–159.

Bonica, A. (2014). Mapping the ideological marketplace. *American Journal of Political Science*, 58(2):367–386.

Calonico, S., Cattaneo, M. D., and Titiunik, R. (2014). Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica*, 82(6):2295–2326.

Calvert, R. L. (1985). Robustness of the multidimensional voting model: Candidate motivations, uncertainty, and convergence. *American Journal of Political Science*, 29(1):69–95.

Canes-Wrone, B., Brady, D. W., and Cogan, J. F. (2002). Out of step, out of office: Electoral accountability and house members' voting. *American Political Science Review*, 96(1):127–140.

Cowburn, M. and Theriault, S. M. (2025). Preventative polarization: Republican senators' positional adaptation in the Tea Party era. *American Politics Research*, 53(2):125–139.

De Magalhães, L., Hangartner, D., Hirvonen, S., Meriläinen, J., Ruiz, N. A., and Tukiainen, J. (2025). When can we trust regression discontinuity design estimates from close elections? Evidence from experimental benchmarks. *Political Analysis*, 33(3):258–265.

Di Tella, R., Kotti, R., Le Pennec, C., and Pons, V. (2025). Keep your enemies closer: Strategic platform adjustments during US and French elections. *American Economic Review*, 115(8):2488–2528.

Downs, A. (1957). *An Economic Theory of Democracy*. Harper and Row, New York.

Eggers, A. C., Fowler, A., Hainmueller, J., Hall, A. B., and Snyder, James M., J. (2015). On the validity of the regression discontinuity design for estimating electoral effects: New evidence from over 40,000 close races. *American Journal of Political Science*, 59(1):259–274.

Ferrara, A., Testa, P. A., and Zhou, L. (2022). New area- and population-based geographic crosswalks for U.S. counties and congressional districts, 1790–2020. *Working Paper*. Available at: `https://sites.google.com/view/andreas-ferrara/data-and-code`.

Fowler, A. and Hall, A. B. (2016). The elusive quest for convergence. *Quarterly Journal of Political Science*, 11(1):131–149.

Fowler, A., Hill, S. J., Lewis, J. B., Tausanovitch, C., Vavreck, L., and Warshaw, C. (2023). Moderates. *American Political Science Review*, 117(2):643–660.

Groseclose, T. (2001). A model of candidate location when one candidate has a valence advantage. *American Journal of Political Science*, 45(4):862–886.

Hall, A. B. (2015). What happens when extremists win primaries? *American Political Science Review*, 109(1):18–42.

Hall, A. B. and Thompson, D. M. (2018). Who punishes extremist nominees? Candidate ideology and turning out the base in US elections. *American Political Science Review*, 112(3):509–524.

Jacobson, G. C. and Kernell, S. (1983). *Strategy and Choice in Congressional Elections*. Yale University Press, New Haven, CT, 2nd edition.

Lee, D. S., Moretti, E., and Butler, M. J. (2004). Do voters affect or elect policies? Evidence from the U.S. house. *Quarterly Journal of Economics*, 119(3):807–859.

Marshall, J. (2024). Can close election regression discontinuity designs identify effects of winning politician characteristics? *American Journal of Political Science*, 68(2):494–510.

McCrary, J. (2008). Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of Econometrics*, 142(2):698–714.

Meyer, C. B. (2022). Getting 'Primaried' in the Senate: Primary challengers and the roll-call voting behavior of sitting senators. *Congress & the Presidency*, 49(2):230–251.

Nokken, T. P. and Poole, K. T. (2004). Congressional party defection in American history. *Legislative Studies Quarterly*, 29(4):545–568.

Poole, K. T. and Rosenthal, H. (2007). *Ideology and Congress: A Political Economic History of Roll Call Voting*. Transaction Publishers, New Brunswick, NJ, 2nd edition.

Schoenenberger, F. (2024). Strategic policy responsiveness to challenger platforms: Evidence from U.S. house incumbents. Working paper, Università della Svizzera italiana.

Sulkin, T. (2005). *Issue Politics in Congress*. Cambridge University Press, New York.

Thomsen, D. M. (2014). Ideological moderates won't run: How party fit matters for partisan polarization in congress. *Journal of Politics*, 76(3):786–797.

Wittman, D. (1983). Candidate motivation: A synthesis of alternative theories. *American Political Science Review*, 77(1):142–157.

6 Appendix

6.1 Figure A1: McCrary Density Test

The density test for manipulation of the running variable yields a t-statistic of -1.03 (p = 0.303). Figure A1 shows no evidence of bunching at the threshold.

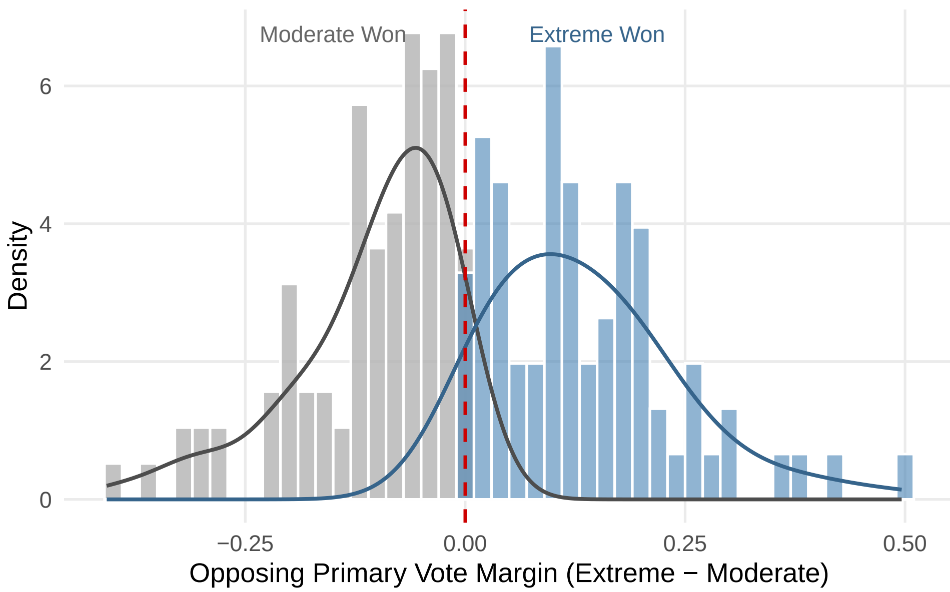

Figure A1: McCrary Density Test for the Running Variable

Note: The running variable is the top-two vote-share margin of the most extreme candidate in the opposing-party primary. The estimated density is smooth through the threshold.

6.2 Table A1: Covariate Balance at the Threshold

Table A1: Covariate Balance at the Threshold

Covariate	RDD Estimate	SE	p-value	N_{eff}	Balance
Incumbent DW-NOMINATE	0.006	0.192	0.973	95	Pass
Incumbent Extremity (|DW-NOM|)	−0.055	0.070	0.433	80	Pass
Incumbent is Democrat	−0.150	0.226	0.507	115	Pass
Prior General Election Vote Share	0.023	0.067	0.729	32	Pass
Recent District Presidential Dem Share	−0.097	0.076	0.204	34	Pass
Ideological Distance (mechanical)	0.201	0.104	0.052	101	Mechanical

Note: RDD estimates use local linear regression with a triangular kernel.
District electoral covariates are observed for 54 of the 172 baseline races and should be read as coverage-limited balance checks.
Ideological distance is partly mechanical (post-treatment) and is better interpreted as a first-stage diagnostic than as a balance violation.

6.3 Table A2: Placebo Cutoff Tests

Table A2: Placebo Cutoff Tests for Party Unity Change

Placebo Cutoff	Estimate	SE	p-value	N_{eff}
−10 pp	−0.025	0.022	0.269	77
−5 pp	−0.009	0.016	0.571	85
+5 pp	0.005	0.023	0.834	81
+10 pp	−0.014	0.016	0.387	40

Note: Each row re-estimates the main party-unity specification at a false cutoff. None of the placebo discontinuities is statistically distinguishable from zero.

6.4 Table A3: Donut-Hole RDD Estimates

To address concerns that observations very close to the threshold may be affected by manipulation or measurement error, I estimate the RDD excluding races within 1 and 2 percentage points of the cutoff. Excluding $|\text{margin}| < 1$ pp (12 observations removed, N = 160) yields an estimate of 3.5 percentage points for party unity. Excluding $|\text{margin}| < 2$ pp (23 observations removed, N = 149) yields 7.1 percentage points. Neither estimate is statistically distinguishable from zero, and under the 1 pp exclusion the same holds for every outcome (Table A3).

While the point estimates increase with the exclusion zone, the standard errors generally grow even faster, reflecting the rapid loss of precision as observations nearest the threshold are removed. This pattern is consistent with noise amplification in shrinking samples rather than a genuine effect gradient. The one exception to uniform non-significance is ideological extremity under the 2 pp exclusion, which flips sign relative to the baseline estimate (-1.7 pp in Table 1) and reaches nominal significance (8.2 pp, $p = 0.008$). I read this cell as part of the same small-sample instability rather than as evidence of an effect: it is a single sign-reversed estimate among the ten outcome estimates in Table A3, it emerges only after the sample shrinks to N = 149, and it does not survive a Bonferroni correction across those ten tests (threshold $p < 0.005$). The covariate imbalance on ideological distance also attenuates in the donut samples, suggesting it is concentrated among the closest races and does not affect the broader pattern.

Table A3: Donut-Hole RDD Estimates

	Exclude $\lvert$margin$\rvert < 1$ pp			Exclude $\lvert$margin$\rvert < 2$ pp		
Outcome	Estimate	SE	p-value	Estimate	SE	p-value
Party Unity	0.035	0.032	0.265	0.071	0.097	0.463
Extremity	0.030	0.032	0.342	0.082	0.031	0.008
Party Leader	0.030	0.030	0.308	0.072	0.096	0.458
Participation	−0.025	0.043	0.563	−0.016	0.078	0.836
Sponsorship	0.059	0.090	0.510	−0.121	0.102	0.236
Ideo Distance (balance)	−0.234	0.184	0.204	−0.175	0.264	0.507

Note: Donut-hole RDD excluding observations closest to the threshold. $N = 160$ for the 1 pp exclusion and $N = 149$ for the 2 pp exclusion.

6.5 Table A4: Lagged-Outcome Placebo Tests

Table A4: Lagged-Outcome Placebo Tests

Pre-Treatment Outcome	RDD Estimate	SE	p-value	N_{eff}
Incumbent DW-NOMINATE	0.006	0.192	0.973	95
Incumbent Extremity	−0.055	0.070	0.433	80
Prior-Congress Party Unity Change	−0.026	0.046	0.579	17
Prior-Congress Votes with Leader	−0.058	0.038	0.121	15
Prior-Congress Ideological Extremity	−0.030	0.030	0.320	25
Prior-Congress Participation	−0.041	0.042	0.329	15

Note: RDD estimates use pre-treatment outcomes as dependent variables. Prior-Congress outcomes are available for 38 baseline races and should be read as limited but consistent placebo checks.

6.6 Table A5: First-Stage RDD — Challenger Ideology at the Threshold

Table A5: First-Stage RDD: Winning Challenger Ideology Across Sample Restrictions

Sample	Estimate	SE	p-value	F-stat	N_{eff}	95% CI
Full sample ($N = 344$)	0.070	0.088	0.425	0.64	168	[−0.102, 0.243]
Baseline, above-median contrast ($N = 172$)	0.135	0.147	0.355	0.85	78	[−0.152, 0.423]
Top quartile, Hall ($N = 86$)	0.331	0.260	0.203	1.62	40	[−0.178, 0.840]
Top decile, Hall ($N = 35$)	0.576	0.339	0.090	2.88	22	[−0.089, 1.241]
Extremist score (placebo, baseline)	−0.060	0.122	0.620	—	95	[−0.299, 0.178]

Note: Each row reports the robust bias-corrected discontinuity in the winning nominee's DIME score at the vote-margin threshold, estimated on progressively tighter above-median-contrast restrictions of the scored sample. The first stage strengthens monotonically as the ideological contrast between the two finalists sharpens (0.07 $\rightarrow$ 0.14 $\rightarrow$ 0.33 $\rightarrow$ 0.58 across the four samples). The top-decile robust estimate is only marginal ($p = 0.090$); the conventional (non-bias-corrected) estimate reaches conventional significance (0.587, $p = 0.022$, $F = 5.27$). The F-statistic stays below 6 throughout, so the design is read as reduced-form / intention-to-treat rather than a rescaled fuzzy-RD LATE. The placebo row estimates the (spurious) discontinuity in the losing extremist candidate's own DIME score on the baseline sample.

6.7 Figure A2: First-Stage RD Plot — Winner Ideology

Figure A2 plots the winning nominee's DIME score against the extreme candidate's vote-share margin. The modest discontinuity at the threshold is visible: the jump in winner ideology when the extreme candidate barely wins is small relative to the overall variation.

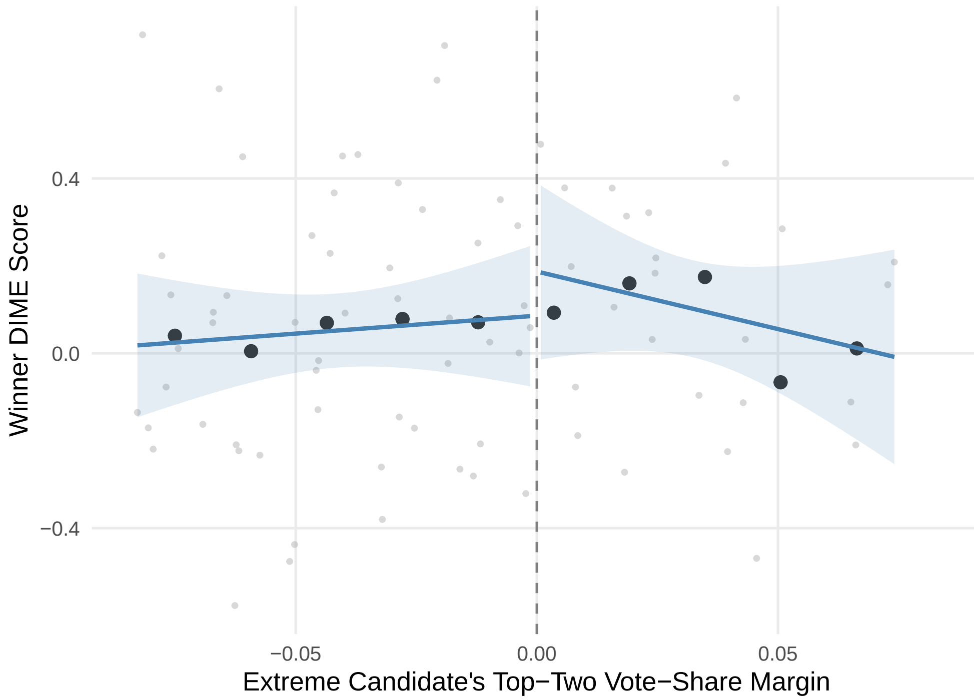

Figure A2: First-Stage RD Plot: Winner DIME Score

Note: Binned means and local polynomial fits for the winning nominee's DIME score on each side of the opposing-party primary threshold. The modest discontinuity reflects the limited ideological contrast between the extreme and moderate finalists in close primaries.

6.8 Table A6: Incumbent-Clustered Standard Errors

Table A6: Incumbent-Clustered Standard Errors

Outcome	Clust. Est.	Clust. SE	Clust. p	N_{eff}
Party Unity	0.018	0.015	0.219	110
Extremity	−0.017	0.018	0.355	57
Party Leader	0.014	0.013	0.308	94
Participation	0.001	0.019	0.973	79
Sponsorship	−0.002	0.056	0.979	77

Note: Estimates are re-computed with incumbent-level clustering and remain close to the baseline robust-SE results. The baseline sample contains 158 unique incumbents, 13 of whom appear more than once, accounting for 27 observations. There are 0 duplicated election rows.

6.9 Table A7: Temporal Heterogeneity (Pre/Post 2010)

Table A7: Temporal Heterogeneity: Pre- and Post-2010

	Pre-2010 ($N = 61$)			Post-2010 ($N = 111$)		
Outcome	Estimate	SE	p-value	Estimate	SE	p-value
Party Unity	0.035	0.039	0.373	0.011	0.017	0.535
Extremity	0.010	0.045	0.824	−0.008	0.017	0.654
Party Leader	0.027	0.045	0.553	0.012	0.015	0.416
Participation	−0.008	0.049	0.866	−0.000	0.018	0.985
Sponsorship	−0.050	0.122	0.680	0.006	0.060	0.916

Note: The sample is split at 2010. Estimates remain small and statistically indistinguishable from zero in both subperiods.

6.10 Table A8: Specification Summary for Party Unity

Table A8: Specification Summary for Party Unity

Specification Family	Values Tested	Estimate Range	Summary
Bandwidth variation	$0.5\times$, $0.75\times$, $1.0\times$, $1.25\times$, $1.5\times$, $2.0\times$ optimal bandwidth	0.004 to 0.018	All estimates positive; none significant at 5%
Polynomial order	Local linear, local quadratic	0.007 to 0.017	Small and null
Kernel choice	Triangular, uniform, Epanechnikov	0.017 to 0.021	Stable across kernels
Controlled specifications	No added covariates, + ideological distance	0.017 to 0.019	Near baseline

Note: Entries summarize the party-unity specifications reported in the main text and appendix. Across all reported variants, the estimated discontinuity remains substantively small.

6.11 Figure A3: Specification Curve for Party Unity

Figure A3 displays the point estimate and 95% confidence interval for party unity across all ten reported specifications, sorted by estimate magnitude. All estimates are positive, range from 0.4 to 2.1 percentage points, and all confidence intervals include zero.

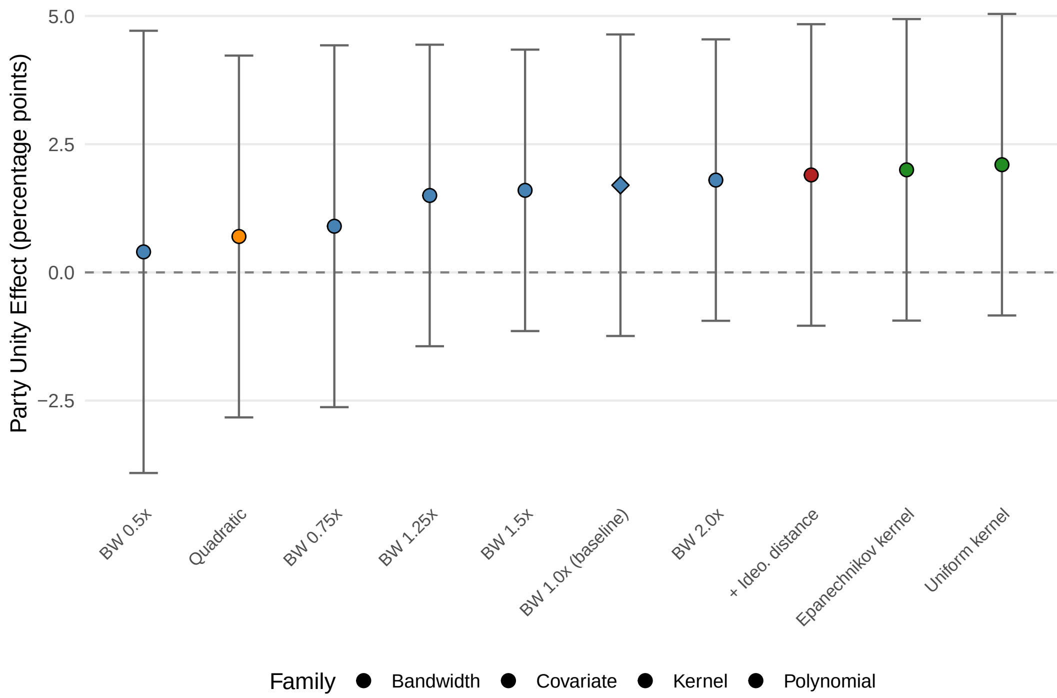

Figure A3: Specification Curve: Party Unity Estimates Across All Specifications

Note: Each point is a robust bias-corrected `rdrobust` estimate for party unity change. Error bars show 95% confidence intervals. The diamond marks the baseline specification. Specifications vary bandwidth, polynomial order, kernel choice, and covariate inclusion.

6.12 Table A9: Analytic Sample Construction

Table A9: Analytic Sample Construction

Sample Step	Observations
Scored top-two opposing-party primaries with observed party-unity outcome	344
Restrict to above-median ideological contrast ($absdist \geq 0.154$)	172
Treated races in baseline sample	76
Control races in baseline sample	96
Effective sample for main party-unity RDD (N_{eff})	109

Note: Counts refer to the scored analysis file used for estimation. The baseline restriction retains races where the ideological contrast between the top-two primary candidates is at or above the sample median.

6.13 Table A10: Heterogeneity by Primary Timing

Table A10: Heterogeneity by Primary Timing

	May or Earlier ($N = 28$)			June or Later ($N = 26$)		
Outcome	Estimate	SE	p-value	Estimate	SE	p-value
Party Unity	0.023	0.046	0.610	0.031	0.039	0.429
Extremity	0.018	0.052	0.728	0.023	0.026	0.385
Party Leader	0.015	0.037	0.678	0.006	0.089	0.944
Participation	−0.001	0.029	0.964	−0.399	0.341	0.242
Sponsorship	−0.044	0.228	0.847	−0.032	0.184	0.860

Note: Primary dates are observed for 54 races in the baseline sample.
The split is defined by the median observed primary month (5), yielding a balanced early/late comparison.
Sponsorship estimates remain null in both subsamples, though some late-primary effective samples are small.

6.14 Table A11: Per-Outcome Minimum Detectable Effects

Table A11: Per-Outcome Minimum Detectable Effect

Outcome	Robust SE	N_{eff}	MDE (80%)	MDE (90%)
Party Unity	0.015	109	0.041	0.047
Extremity	0.018	57	0.049	0.057
Party Leader	0.013	94	0.037	0.043
Participation	0.019	79	0.053	0.061
Sponsorship	0.055	76	0.155	0.180

Note: Each outcome is estimated at its own MSE-optimal bandwidth, so both its robust standard error and effective sample differ. The minimum detectable effect is $(z_{0.975} + z_{\text{power}}) \times$ SE at 80% and 90% power ($\alpha = 0.05$, two-tailed), on the proportion scale ($0.041 = 4.1$ percentage points). Effective sample and detection threshold rank the outcomes differently: ideological extremity is least powered by effective sample ($N_{\text{eff}} = 57$), whereas bill sponsorship is least powered by detectable magnitude, its MDE more than three times party unity's because the count-based measure is an order of magnitude noisier than the proportion-based outcomes.

6.15 Table A12: Equivalence (TOST) Tests

Table A12: Equivalence (Two One-Sided) Tests

			p_{TOST} at equivalence bound Δ		
Outcome	Estimate	SE	$\Delta = 0.02$	$\Delta = 0.03$	$\Delta = 0.04$
Party Unity	0.017	0.015	0.426	0.191	0.060
Extremity	−0.017	0.018	0.436	0.232	0.096
Party Leader	0.014	0.013	0.316	0.109	**0.024**
Participation	0.001	0.019	0.155	0.061	**0.019**
Sponsorship	−0.002	0.055	0.370	0.304	0.244

Note: Post-hoc two one-sided tests on the Table 1 estimates, using the same Normal approximation `rdrobust` uses for its reported p-values. $p_{\text{TOST}} = \max(p_{\text{lower}}, p_{\text{upper}})$; equivalence to zero is concluded at $\alpha = 0.05$ when $p_{\text{TOST}} < 0.05$ (**bold**). The equivalence bound Δ is on the proportion scale ($0.03 = 3$ percentage points), with $\Delta = 0.03$ anchored to the roughly three-point same-party response reported by Meyer (2022). At the primary $\Delta = 0.03$ bound no outcome is equivalent to zero, including party unity ($p = 0.191$); only at $\Delta = 0.04$ do votes with party leader and participation clear the bar.

6.16 Table A13: Uniform Multiple-Testing Summary

Table A13: Uniform Multiple-Testing Summary Across the Full Specification Grid

Quantity	Value
Total tests of the challenger-extremity estimand	101
Nominally significant ($p < 0.05$)	3
Expected under a global null (0.05×101)	5.05
Surviving Benjamini-Hochberg ($q = 0.05$)	0
Surviving Bonferroni ($\alpha = 0.05/101$)	0

The three nominally significant cells:

Outcome	Test family	Specification	Raw p	BH-adj. p
Ideological Extremity	Competitiveness subgroups	pooled_lagged_gen	0.004	0.392
Ideological Extremity	Donut-hole	donut_2pp	0.008	0.392
Bill Sponsorship	Hall sample restriction	hall_q90	0.012	0.402

Note: Every discrete test of the estimand reported anywhere in the paper—across outcomes, bandwidths, donut-hole exclusions, sample restrictions, subgroup and timing splits, and clustered standard errors—assembled into one grid (101 tests). Balance, placebo, first-stage, equivalence, and minimum-detectable-effect calculations belong to different hypothesis families and are excluded. The observed number of nominal positives (3) falls below the 5.05 expected by chance under a global null, and none of the three survives Benjamini-Hochberg or Bonferroni correction applied uniformly to the grid.

6.17 Figure A4: Sensitivity to the Ideological-Contrast Cutoff

Figure A4 traces the party-unity estimate as the ideological-contrast cutoff is swept continuously from the full scored sample to the most restrictive tail, cross-validating the three discrete sample-restriction anchors reported in the text.

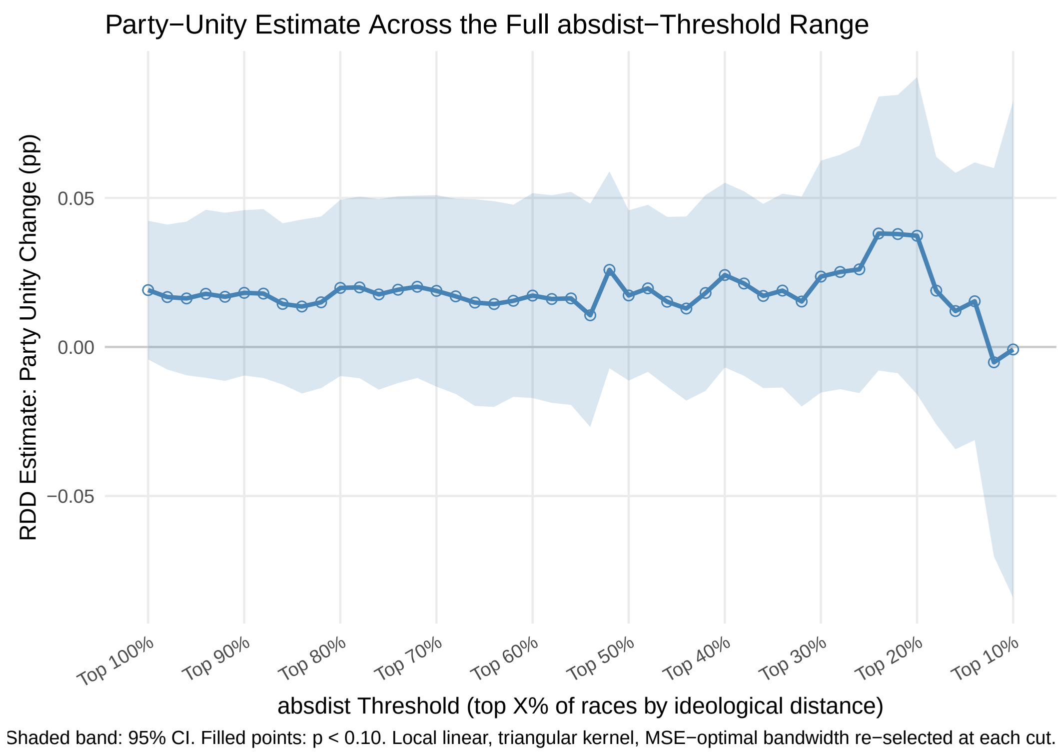

Figure A4: Party-Unity Estimate Across the Continuous Ideological-Contrast Cutoff

Note: Each point is the robust bias-corrected party-unity estimate at a given above-contrast cutoff, from x = 0 (full scored sample, N = 344) to x = 0.90 (top-decile restriction, N = 35), in steps of 0.02. The three discrete anchors reported in the text (full sample, top quartile, top decile) are reproduced by the corresponding points on the curve. No cutoff in the continuous range produces a conventionally significant positive effect.