Online Appendix:

Do Donors Punish Extremist Primary Nominees?

Evidence from Congress and American State Legislatures

American Political Science Review

Andrew C. W. Myers
Stanford University

Contents

A	Strategic Donating and Post-Treatment Bias	A.1
B	RD Balance Tests	A.8
C	Characteristics of Moderate and Extremist Bare-Winners	A.10
D	RD Estimates by Corporate Industry	A.11
E	Replicating Results Using State Legislative Roll-Call Voting Records	A.12
F	Additional RD Estimates for Individuals and Corporate PACs	A.13
G	RD Estimates Over Time	A.14
H	RD Estimates Using Only In-State Donors to Scale Candidates	A.15
I	RD Estimates Using Scalings that Adjust for Protest Voting	A.16
J	Donor Partisanship Over Time	A.17
K	Attenuation Bias Simulations	A.18
L	Roll-Call Classification Exercise	A.20

A Strategic Donating and Post-Treatment Bias

In this section, I outline two forms of endogeneity that complicate RD-based analyses of general-election fundraising totals when candidates are scaled on the basis of both their primary- and general-election contributions, and I document how my preferred scaling method addresses these concerns.

Post-Treatment Bias

The first concern with jointly scaling a candidate based on the contributions they receive both before and after the primary election is that the candidate's position in the scaling could be partially a function of their primary-election outcome. This possibility is problematic because it may cause bare primary winners and bare primary losers to appear systematically different, or even for their classification as relative moderates and extremists to be flipped.

Such a scenario would arise if the composition of a candidate's donorate changes after they secure their party's primary nomination. Using the FEC and NIMSP contribution data described in the main text, Figure A.1 illustrates two such compositional changes. The horizontal axis of this figure reports the number of election cycles until a given candidate wins their first primary nomination, with primary and general elections separated for the election cycle containing a candidate's first primary victory and pooled for all remaining election cycles. To ensure that I am capturing within-candidate changes in donor composition (rather than between-candidate differences), I restrict this analysis to candidates who win a primary election at some point in their career.

For each election cycle, the vertical axis of Panel A plots the share of a candidate's contributions that are from corporate PACs. The results are averaged across all candidates within each horizontal axis bin. The results indicate that winning a primary election causes a substantial increase in the share of contributions a candidate receives from corporate PACs.

To further illustrate these compositional effects, I introduce the concept of an "incumbent donor." For election cycle t and candidate i, I define an incumbent donor as a donor that has contributed to at least one incumbent by the time of election t that is not candidate i.[1] Incumbent donors are critical contributors because my method relies precisely on donors who contribute to both incumbents and non-incumbents to bridge roll-call voting scores from the former to the latter. I calculate the share of each candidate's donors that are "incumbent donors," weighted by contribution amounts, and again restrict the analysis to candidates who eventually win at least one primary election.[2] Panel B of Figure A.1 plots this share averaged across candidates in a given horizontal axis. Overall, I find that candidates' individual donorates become significantly more connected to other incumbents after they win their first primary.

Taken together, the results presented in panels A and B of Figure A.1 provide strong evidence that winning a primary election may alter candidates' relative ideological scaling if they are scaled in part based on their general-election receipts. In Figure A.2, I formally

[1] The restriction on t ensures that future donations do not affect prior donor classifications. The restriction on i prevents the incumbent donor share from mechanically becoming one after a candidate wins their first general election.

[2] The results, however, are highly similar without donation weights.

Figure A.1 – Effect of Winning Primary Election and Subsequent Legislative Experience on Donor Composition in Congress, 1980-22, and State Legislatures, 1996-22. This figure plots the share of a candidate's contributions that come from corporate PACs (vertical axis, Panel A) and incumbent donors (vertical axis, Panel B), averaged across all candidates with equal experience (horizontal axis). For election cycle t and candidate i, an incumbent donor is a donor that contributed to at least one incumbent by the time of election t that is not candidate i. The sample is restricted to candidates who win at least one primary election. Winning a primary election causes a large jump in contributions from corporate PACs, and subsequent legislative experience attracts better-connected donors.

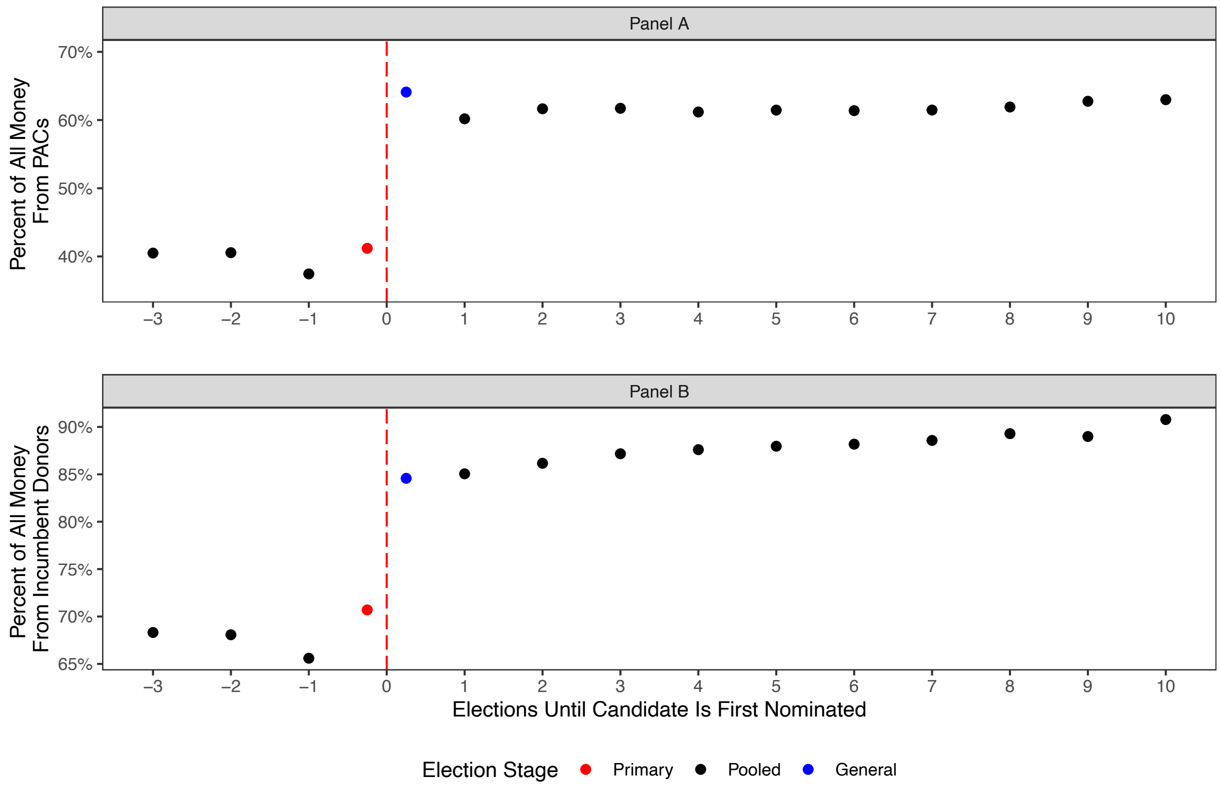

estimate the compositional effects identified in Figure A.1 using an RD.

Specifically, the unit of analysis in Figure A.2 is the candidate, and the running variable is the candidate's primary-election winning margin. This RD identifies the causal effect of winning the primary election on the composition of that candidate's contributions. In the first row of Figure A.2, I plot the RD estimate of the effect of winning a primary on a candidate's share of *all* contributions from incumbent donors or corporate PACs. As is apparent, winning a primary election substantially increases the candidate's share of *all* contributions from incumbent donors and corporate PACs.

To address concerns about post-treatment bias, in the main text I describe how I exclude all contributions made during the general election when scaling candidates. In the second row of Figure A.2, I reestimate the candidate-level RD after restricting the outcomes to

Figure A.2 – RD Estimate of Effect of Winning Primary Election on Donor Composition in Congress, 1980-22, and State Legislatures, 1996-22. Winning the primary election increases a candidate's share of *all* contributions and the share of *all* contributions from incumbent donors (first row), but this effect disappears when the sample is restricted to *primary-election* contributions (second row).

(a) Share of **All** Contributions
From Incumbent Donors

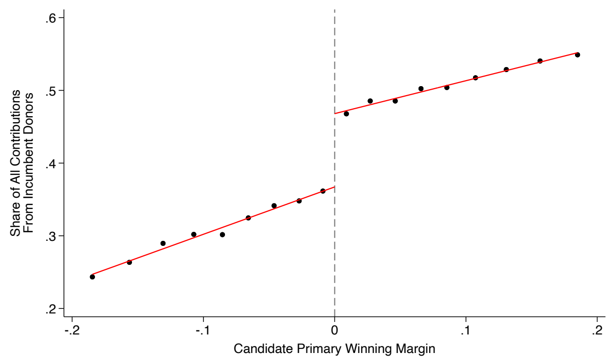

(b) Share of **All** Contributions
From Corporate PACs

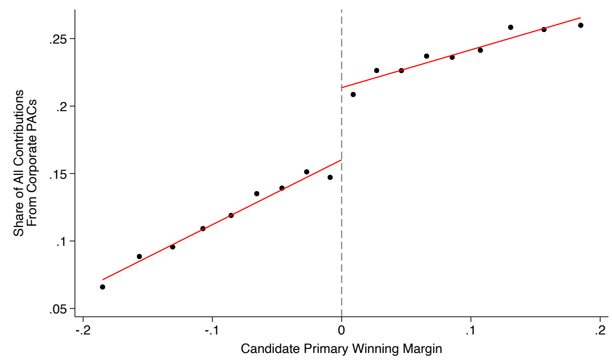

(c) Share of **Primary-Election** Contributions
From Incumbent Donors

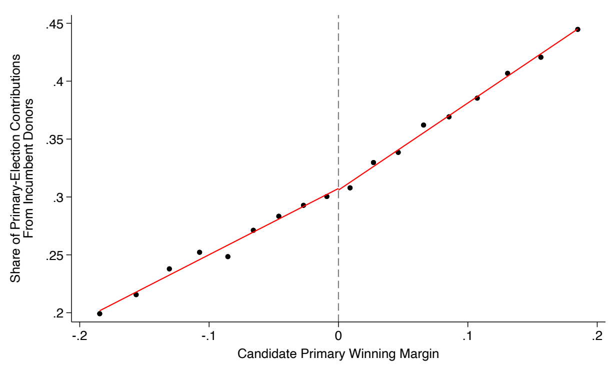

(d) Share of **Primary-Election** Contributions
From Corporate PACs

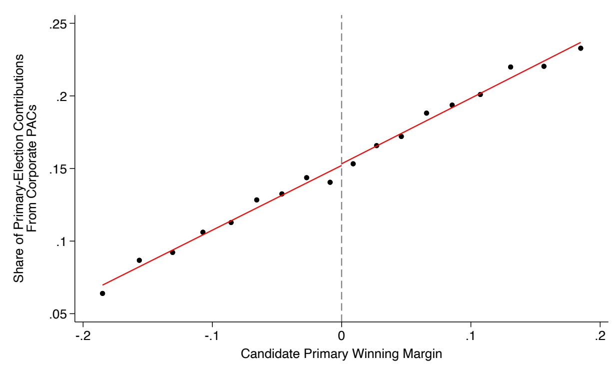

include only primary-election contributions. After making this restriction, the discontinuity disappears entirely, indicating that the data restrictions I introduce address concerns about post-treatment bias.

Scalings Conflate Moderation with Fundraising Success

A second concern is that campaign finance-based scalings may conflate ideological moderation with fundraising success if donors contribute on the basis of candidates' non-ideological characteristics. This concern is particularly acute when jointly scaling candidates based on primary- and general-election fundraising or including contributions from access-seeking PACs. However, since I omit general-election contributions and contributions from access-seeking PACs from my preferred ideological scaling, I focus in this section on concerns related to primary-election fundraising from individual donors.

While it is difficult to test this concern directly, I begin by running a series of simulations

Figure A.3 – Simulated Effect of Altering Candidates' Primary-Election Fundraising Success. This figure plots the results from a series of simulations where candidates' primary-election fundraising success is increased or decreased, holding fixed the underlying ideological space. Altering primary-election fundraising success does not meaningfully affect candidates' estimated *Primary-Only Scaling*.

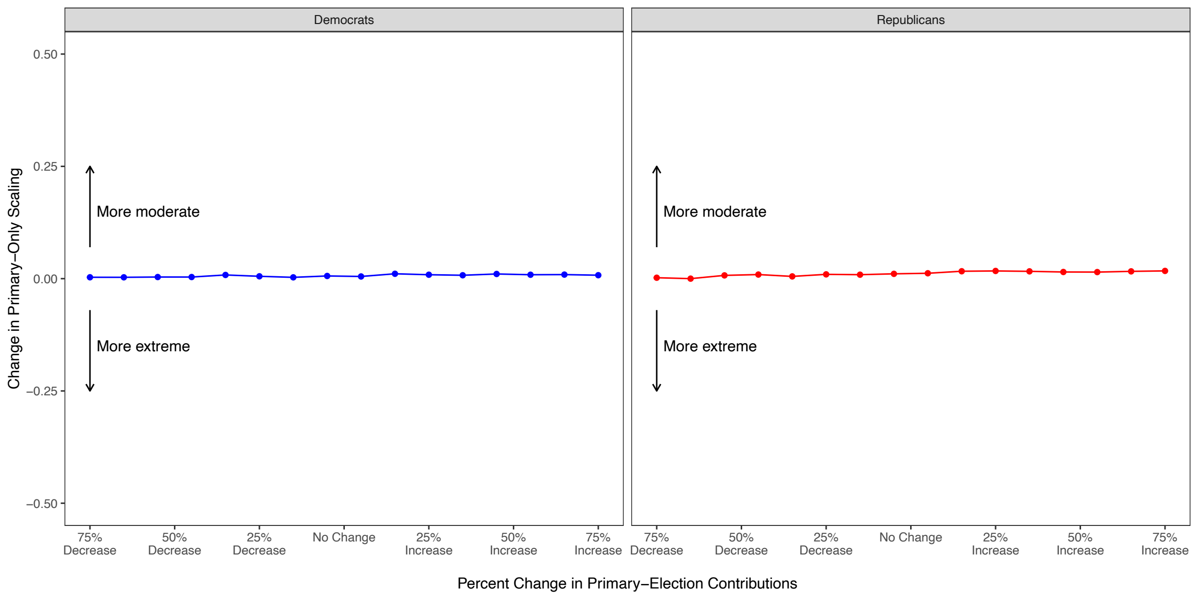

that evaluate whether altering a candidate's primary-election fundraising success systematically affects their estimated ideology. Specifically, I bootstrap candidate i's primary-election contribution matrix—increasing or decreasing their number of contributions by factors between 5% and 75%—while holding fixed all other candidates' primary-election contributions.[3] Using this modified contribution matrix, I calculate the full set of candidate ideology scalings following the methodology described in Section and extract candidate i's scaling (henceforth, *Bootstrapped Primary-Only Scaling*$_i$). I then calculate the change in candidate i's scaling caused by altering their fundraising success as

$$\textit{Scaling Change}_i = \textit{Primary-Only Scaling}_i - \textit{Bootstrapped Primary-Only Scaling}_i,$$

where *Primary-Only Scaling*$_i$ is candidate i's true *Primary-Only Scaling*.[4]

Figure A.3 plots the average value of *Scaling Change* separately for Democratic and Republican candidates. I find that permuting candidates' primary-election contributions does not meaningfully affect their estimated ideological positions. The estimated change in candidates' ideological positions is less than .01 across sample size factors.[5]

[3]I iteratively bootstrap an individual candidate's contributions, rather than the universe of scalable candidates at the same time, in order to hold the underlying ideological space fixed. Because this process is computationally costly, I conduct this analysis for a random sample of 1000 candidates.

[4]For clarity, I set the polarity of the change in scalings such that larger values indicate more-moderate positions.

[5]This difference is equal to less than 2% of the standard deviation of the true *Primary-Only Scaling*.

Figure A.4 – Empirical Relationship Between Primary-Election Fundraising and Candidates' *Primary-Only Scaling.*

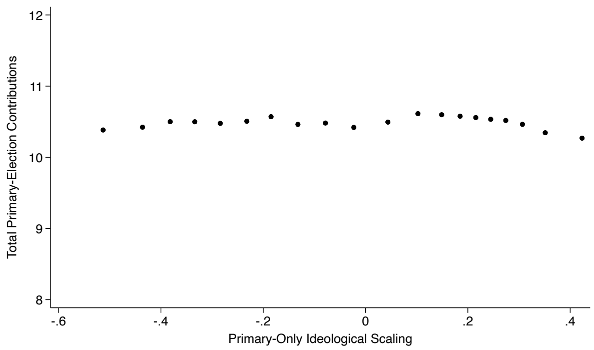

(a) Total Primary-Election Contributions

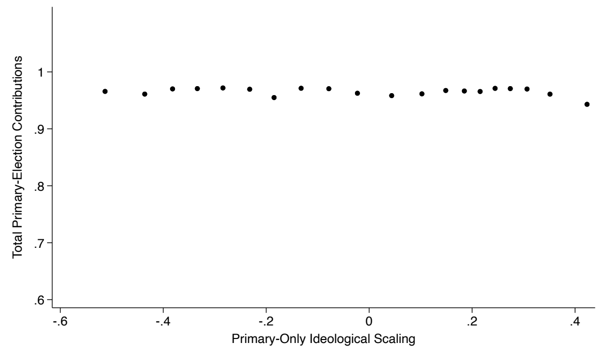

(b) Share of Total Primary-Election Contributions

To further probe whether my measurement of candidate ideology is endogenous to candidates' primary-election fundraising success, Figure A.4 plots binned averages of candidates' total primary-election fundraising and share of primary-election fundraising (vertical axis) across their estimated ideology using my *Primary-Only Scaling* (horizontal axis). As the figure illustrates, there is no evidence that candidates' positions estimated using the *Primary-Only Scaling* are endogenous to their primary-election fundraising success.

Finally, if the restrictions I impose on the scaling process in section three address the concerns documented in this section, there should be no causal effect of winning the primary election on a candidate's estimated ideology. Figure A.5 tests this prediction using a candidate-level RD in primary elections. The running variable in Figure A.5 is a candidate's primary-election winning margin and the outcome is their *Primary-Only Scaling*. I plot estimates separately for Democrats (in blue) and Republicans (in red). As Figure A.5 depicts, I find that there is no effect of winning a primary election on a candidate's estimated ideology scaling.

Evidence that Strategic Donating and Post-Treatment Contributions Would Bias Estimates

Having documented forms of strategic donating and post-treatment bias in the contribution matrix and how my preferred scaling method addresses them, I now provide direct evidence of how failing to account for these biases would affect my results. To to so, I create a second version of the scalings introduced in the main text that use primary- and general-election contributions from all donors to scale candidates (henceforth, the *Unrestricted Scaling*). These *Unrestricted Scalings* correlate with NP-Scores at very similar rates to the *Primary-Specific Scalings* ($r = .92$ overall, .71 for Democrats, and .72 for Republicans).

First, I compare primary candidates' designation as relative moderates or extremists using the *Primary-Specific Scaling* and *Unrestricted Scaling*. Table A.1 reports the results. The rows in Table A.1 report candidates' classifications using the *Primary-Specific Scaling*, while columns report candidates' classifications using *Unrestricted Scaling*. As is apparent, using

Figure A.5 – Effect of Winning the Primary Election on Candidates' *Primary-Only Scaling* in Congress, 1980-22, and State Legislatures, 1996-22. Winning the primary election does not alter candidates' *Primary-Only Scaling*.

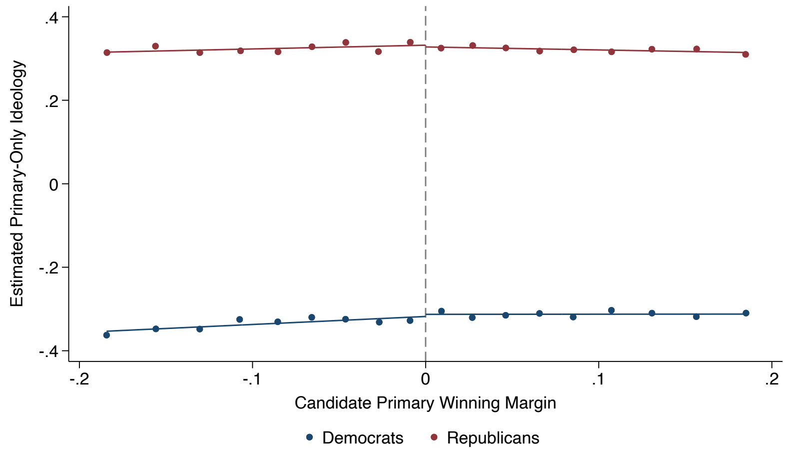

general-election contributions to scale candidates significantly affects primary candidates' relative positioning. Using the *Unrestricted Scaling* would cause the researcher to "flip" 17% of primary candidates' moderate and extremist designations, relative to the *Primary-Specific Scaling* ($1859/10861 \approx .17$).

Second, to evaluate whether these "flips" are consequential, Table A.2 replicates Table 1 using *Unrestricted Scalings*. The estimates across Table A.2 are negative and significant, indicating that my substantive conclusions would be unchanged using *Unrestricted Scalings*. However, the estimates using *Unrestricted Scalings* are significantly larger than the estimates when using *Primary-Specific Scalings*. For example, column one indicates that using the *Unrestricted Scalings* would inflate my coefficient estimate by roughly 35% in comparison to *Primary-Specific Scalings* (-8 vs -6 percentage points).

Overall, Tables A.1 and A.2 suggest that the scaling correction I employ meaningfully addresses concerns about strategic donating and post-treatment bias on my estimates. In Appendix E, I show that my estimates using the *Primary-Specific Scaling* are very similar to estimates obtained using NP-Scores, an ideological scaling that is entirely distinct from campaign contributions.

Table A.1 – Top Two Primary Candidates' Moderate/Extremist Classifications Using *Primary-Specific Scalings* and *Unrestricted Scalings* in Congress, 1980-22, and State Legislatures, 1996-22. Table reports candidates' classifications as relative moderates or extremists using *Primary-Specific Scalings* (rows) and *Unrestricted Scalings* (columns).

	Unrestricted Scaling Scaling Classification		
Primary-Specific Scaling Classification	Moderate	Extremist	N
Moderate	9002	1859	10861
Extremist	1859	9002	10861
N	10861	10861	21722

Note: Sample is restricted to contested primary elections where top two candidates have both a *Primary-Specific Scaling* and *Unrestricted Scaling*. The unit of analysis is the individual candidate.

Table A.2 – Effect of Nominating the Extremist Primary-Election Candidate on their Party's General-Election Contribution Share Using *Unrestricted Scaling* in Congress, 1980-22, and State Legislatures, 1996-22. RD estimates of the effect of nominating the extremist candidate on their party's share of general-election contributions are approximately 35% larger when using *Unrestricted Scalings*.

	Share of Total General Election Contributions			
	(1)	(2)	(3)	(4)
Extremist Primary Win	-0.08	-0.09	-0.09	-0.09
	(0.02)	(0.03)	(0.02)	(0.02)
N	3,145	6,519	3,200	6,519
Specification	Linear	Cubic	CCT	IW
Spline	Yes	Yes	-	-
Bandwidth	.10	-	0.10	-

Note: Robust standard errors clustered by district are reported in parentheses. The running variable is the extremist candidate's win margin in the primary election. Spline indicates that the regression function was fit separately on either side of zero. Cubic refers to a third-order polynomial regression. CCT refers to the method from Calonico, Cattaneo, and Titiunik (2014). IW refers to the method from Imbens and Wager (2019).

B RD Balance Tests

The key identifying assumption that underlies my regression discontinuity design is that districts that narrowly nominate a relative moderate candidate are, in the limit, identical to districts that narrowly nominate the extremist candidate (Imbens and Lemieux, 2008; Lee and Lemieux, 2010). In other words, there must be no district-level sorting at the discontinuity. In Table B.1, I test for any chance imbalances in my sample by estimating Equation 1 where the outcome is the party's fundraising totals in the previous election cycle. If the "no sorting" assumption holds, these estimates should be null, indicating that, in districts where the more-moderate candidate barely wins, the party fundraised no better in the prior election than in districts where the more-extreme candidate was nominated. The coefficients in Table B.1 are all exceedingly small, indicating that there is no evidence of bias. Further, using the standard McCrary (2008) manipulation test, Figure B.1 shows that I fail to reject the null hypothesis of no jump at the discontinuity (p-value = .595).

Finally, Table B.2 tests for chance imbalances in two additional relevant variables: the party's lagged presidential and legislative vote shares. As the table shows, I find no evidence of an imbalance in these variables that would contribute to the estimates reported in the main article.

Table B.1 – Effect of Nominating the Extremist Primary-Election Candidate on Lagged General-Election Contribution Share in Congress, 1980–22, and State Legislatures, 1996–22.

	Lagged Share of Total General Election Contributions			
	(1)	(2)	(3)	(4)
Extremist Primary Win	-0.00	-0.00	0.02	0.00
	(0.03)	(0.03)	(0.04)	(0.04)
N	1,701	3,388	1,409	3,388
Specification	Linear	Cubic	CCT	IW
Spline	Yes	Yes	-	-
Bandwidth	.10	-	0.08	-

Note: Robust standard errors clustered by district are reported in parentheses. The running variable is the extremist candidate's win margin in the primary election. Spline indicates that the regression function was fit separately on either side of zero. Cubic refers to a third-order polynomial regression. CCT refers to the method from Calonico, Cattaneo, and Titiunik (2014). IW refers to the method from Imbens and Wager (2019).

Figure B.1 – Density of the Running Variable Using McCrary (2008) Test in Congress, 1980–22, and State Legislatures, 1996–22.

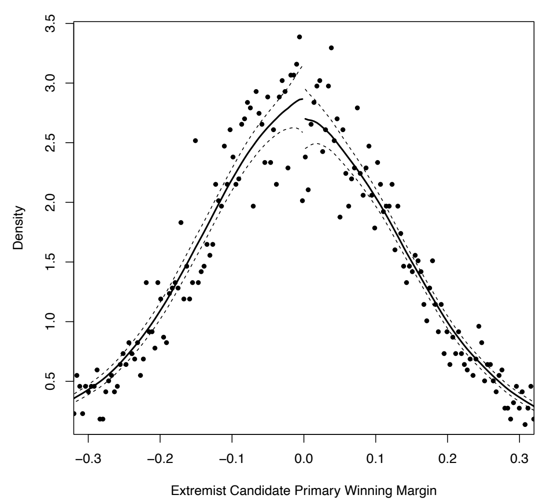

Table B.2 – Effect of Nominating the Extremist Primary-Election Candidate on Lagged Party Presidential and Legislative Vote Share in Congress, 1980–22, and State Legislatures, 1996–22. Districts that narrowly nominate the extremist primary-election candidate do not differ in terms of prior support for their party's presidential candidate (columns 1-4) or legislative candidate (columns 5-8).

	Party's Lagged Presidential Vote Share				Party's Lagged Legislative Vote Share			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Extremist Primary Win	-0.00	0.00	-0.00	-0.00	0.00	0.02	0.01	0.01
	(0.01)	(0.01)	(0.01)	(0.01)	(0.03)	(0.03)	(0.03)	(0.03)
N	1,643	3,315	1,742	3,315	1,848	3,675	2,200	3,675
Specification	Linear	Cubic	CCT	IW	Linear	Cubic	CCT	IW
Spline	Yes	Yes	-	-	Yes	Yes	-	-
Bandwidth	.10	-	0.11	-	.10	-	0.12	-

Note: Robust standard errors clustered by district are reported in parentheses. The running variable is the extremist candidate's win margin in the primary election. Spline indicates that the regression function was fit separately on either side of zero. Cubic refers to a third-order polynomial regression. CCT refers to the method from Calonico, Cattaneo, and Titiunik (2014). IW refers to the method from Imbens and Wager (2019).

C Characteristics of Moderate and Extremist Bare-Winners

As Marshall (2022) notes, my politician characteristic RD design identifies the aggregate effect of candidate ideology and all other candidate-level characteristics that differ between the two types of barely-winning candidates (i.e., compensating differentials). Studying this bundled treatment is appropriate for evaluating the consequences of primary voters' electoral selection, where all differences between candidate types matter (Hall, 2015). To understand the underlying mechanisms, however, it is important to examine whether moderate and extremist candidates differ on observable non-ideological characteristics. In Figure C.1, I test whether barely-winning moderate and extremist candidates systematically differ in terms of incumbency status, prior office-holding experience, gender, and race. I find no significant differences on these characteristics.

Figure C.1 – Characteristics of Moderate and Extremist Bare-Winners. This figure plots the difference in the probability that moderate and extremist bare-winners are an incumbent, have previous office-holder experience, are female, and are non-white. Data on candidate characteristics is from Porter and Treul (2024) and is limited to candidates for Congress.

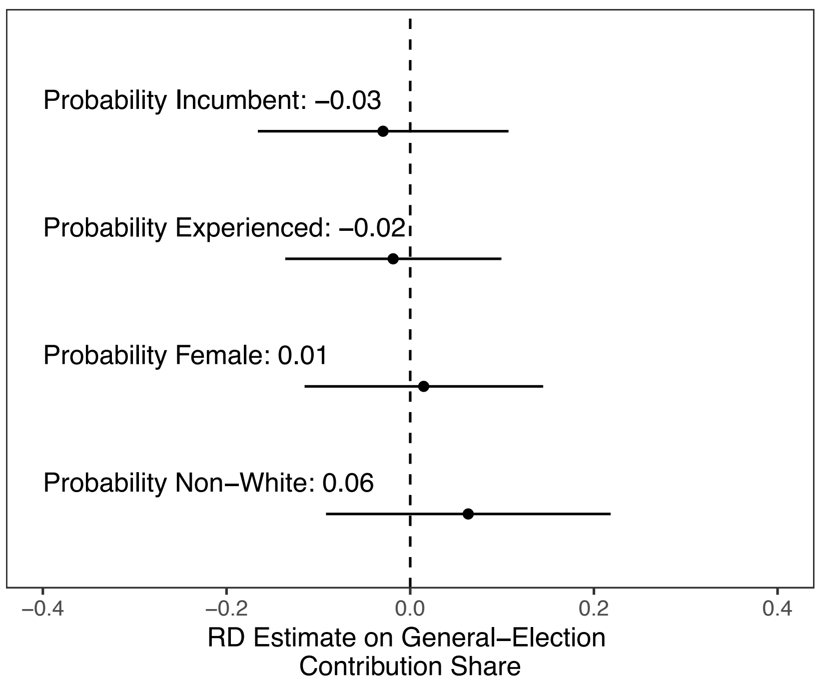

D RD Estimates by Corporate Industry

Figure D.1 – Effect of Nominating the Extremist Primary-Election Candidate on their Party's General-Election Contribution Share by Corporate Industry in Congress, 2000-22, and State Legislatures, 1996-22. The penalty to extremist primary nominees is similarly sized across all 10 corporate industries defined by NIMSP and the FEC. This figure reports estimates using a cubic specification of the running variable.

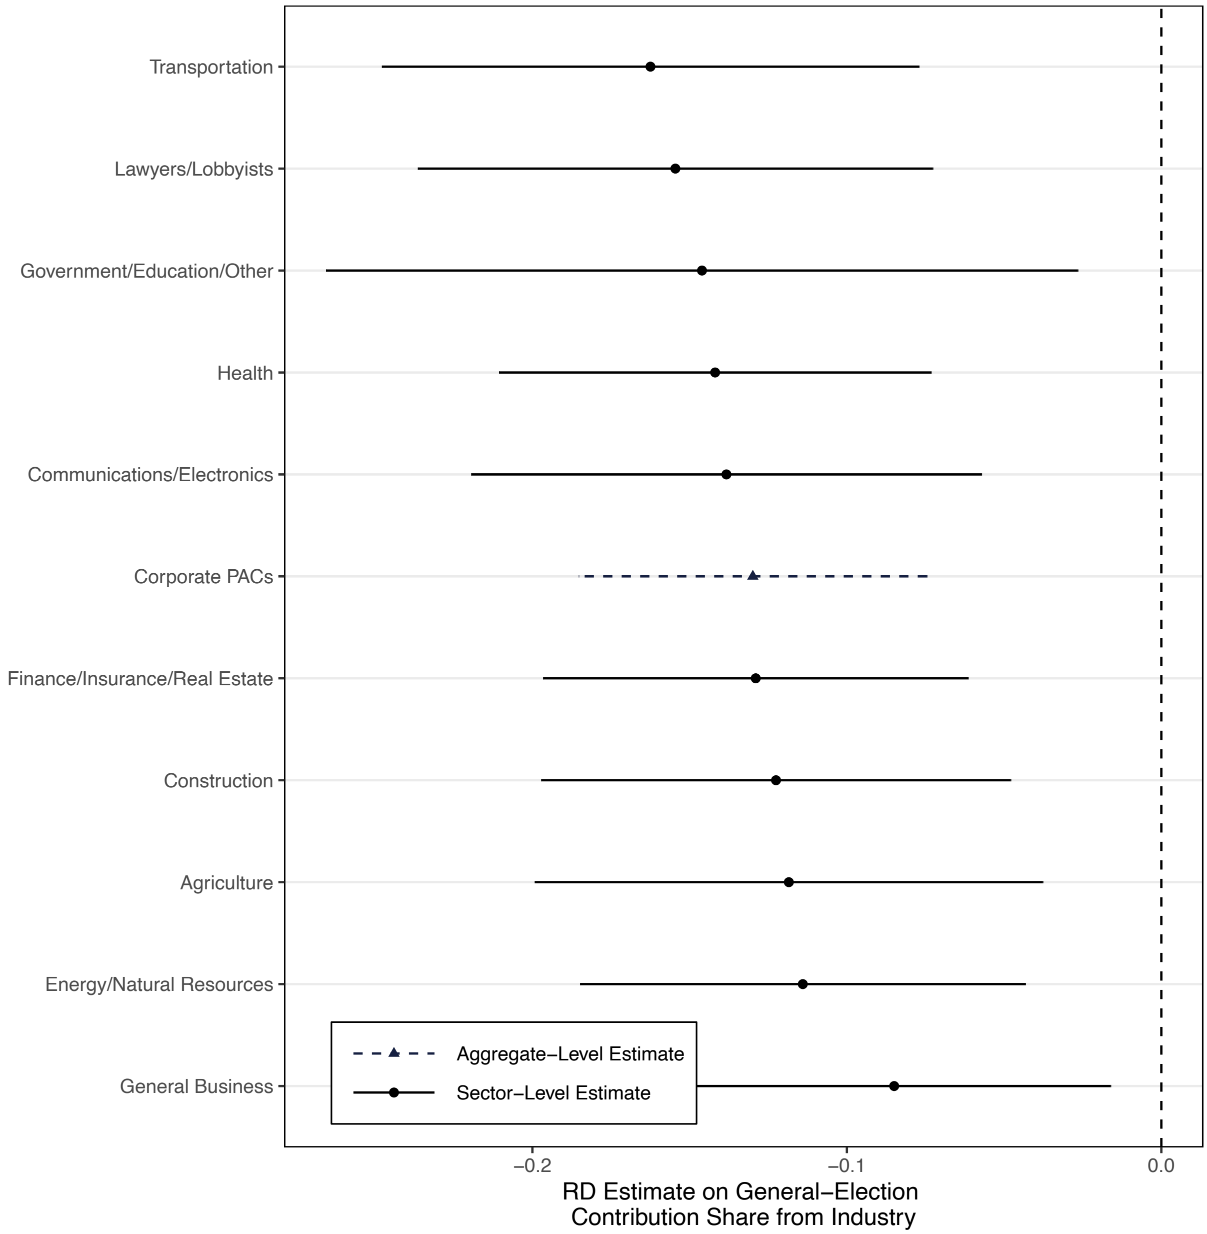

E Replicating Results Using State Legislative Roll-Call Voting Records

To ensure that my results are not an artifact of the contribution-based scaling, I replicate the panel-based analysis from the main text using a measure of candidate ideology that is independent of campaign contributions. This measure draws on the state legislative roll-call voting records of prior, current, or future state legislators who face another candidate with a state legislative roll-call voting record, either in a congressional or state legislative election. The results are plotted in Figure E.1. As the figure illustrates, my estimates are highly similar using this alternative scaling, although the coefficients are estimated imprecisely due to the small sample size.

Figure E.1 – Comparison of Midpoint Estimates Using Campaign Finance-Based and Roll Call-Based Scalings in Congress, 1980-22, and State Legislatures, 1996-22. This figure compares midpoint estimates using *Primary-Specific Scalings* and the NP-Scores of prior, current, or future state legislators who face another candidate with a state legislative roll-call voting record. Estimates are transformed to the RD scale. This figure uses Democratic presidential vote share to hold the district median constant.

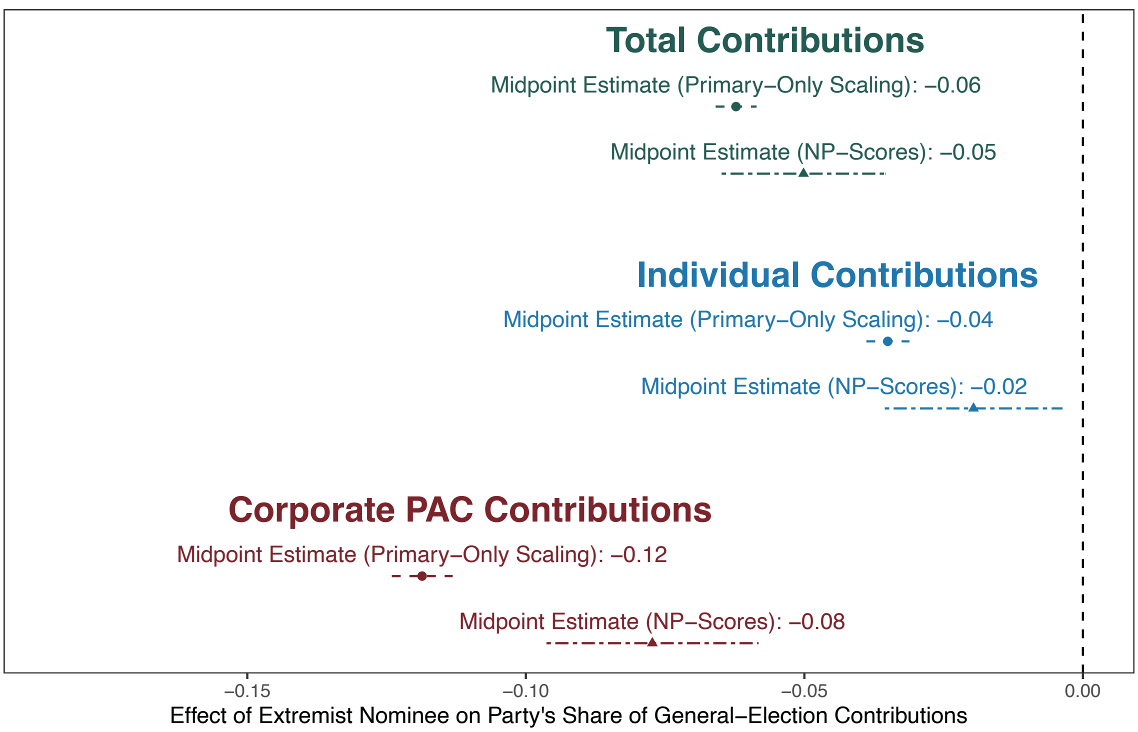

F Additional RD Estimates for Individuals and Corporate PACs

Table F.1 – Effect of Nominating the Extremist Primary-Election Candidate on their Party's General-Election Contribution Share from Individual Donors and Corporate PACs in Congress, 1980-22, and State Legislatures, 1996-22. The close primary nomination of the extremist candidate causes a 11-13 percentage point decrease in that party's share of general-election contributions from corporate PACs and 4-5 percentage point decline among individual donors.

	Share of General-Election Contributions From Corporate PACs				Share of General-Election Contributions From Individuals			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Extremist Primary Win	-0.12	-0.13	-0.11	-0.11	-0.04	-0.05	-0.05	-0.05
	(0.03)	(0.03)	(0.03)	(0.03)	(0.02)	(0.02)	(0.02)	(0.02)
N	2,598	5,106	2,845	5,106	2,588	5,091	3,155	5,091
Specification	Linear	Cubic	CCT	IW	Linear	Cubic	CCT	IW
Spline	Yes	Yes	-	-	Yes	Yes	-	-
Bandwidth	.10	-	0.11	-	.10	-	0.13	-

Note: Robust standard errors clustered by district are reported in parentheses. The running variable is the extremist candidate's win margin in the primary election. Spline indicates that the regression function was fit separately on either side of zero. Cubic refers to a third-order polynomial regression. CCT refers to the method from Calonico, Cattaneo, and Titiunik (2014). IW refers to the method from Imbens and Wager (2019).

G RD Estimates Over Time

Table G.1 – Financial Penalty Imposed on Extremist Primary Nominees Over Time Using the RD. The close primary nomination of an extremist causes a 7 percentage point decline in their party's share of general-election contributions, but this penalty has declined significantly in recent years.

	Share of Total General Election Contributions	Share of Total General Election Contributions
	(1)	(2)
Extremist Primary Win	-0.07	-0.10
	(0.02)	(0.03)
Extremist Primary Win $\times$ Year $\geq$ 2002		0.04
		(0.02)
N	5,223	5,223
Specification	Cubic	Cubic
Spline	Yes	Yes

Note: Robust standard errors clustered by district are reported in parentheses. The running variable is the extremist candidate's win margin in the primary election. Spline indicates that the regression function was fit separately on either side of zero. Cubic refers to a third-order polynomial regression.

H RD Estimates Using Only In-State Donors to Scale Candidates

Table H.1 – Effect of Nominating the Extremist Primary-Election Candidate on their Party's General-Election Contribution Share in Congress, 1980-22, and State Legislatures, 1996-22, Using Only In-State Contributions to Scale Candidates. This table replicates Table 1 after scaling candidates only on the basis of contributions from in-state donors.

	Share of Total General Election Contributions			
	(1)	(2)	(3)	(4)
Extremist Primary Win	-0.06	-0.06	-0.06	-0.06
	(0.02)	(0.02)	(0.02)	(0.02)
N	2,608	4,952	3,022	4,952
Specification	Linear	Cubic	CCT	IW
Spline	Yes	Yes	-	-
Bandwidth	.10	-	0.12	-

Note: Robust standard errors clustered by district are reported in parentheses. The running variable is the extremist candidate's win margin in the primary election. Spline indicates that the regression function was fit separately on either side of zero. Cubic refers to a third-order polynomial regression. CCT refers to the method from Calonico, Cattaneo, and Titiunik (2014). IW refers to the method from Imbens and Wager (2019).

I RD Estimates Using Scalings that Adjust for Protest Voting

Table I.1 – Effect of Nominating the Extremist Primary-Election Candidate on their Party's General-Election Contribution Share in Congress, 1980-22, and State Legislatures, 1996-22, Using Fowler-Lewis Scores to Anchor U.S. House Incumbents' Positions. This table replicates Table 1, except that incumbents' roll-call voting records in the U.S. House are measured using scalings that account for protest voting from Fowler and Lewis (2024).

	Share of Total General Election Contributions			
	(1)	(2)	(3)	(4)
Extremist Primary Win	-0.05	-0.06	-0.06	-0.06
	(0.02)	(0.02)	(0.02)	(0.02)
N	2,662	5,170	2,996	5,170
Specification	Linear	Cubic	CCT	IW
Spline	Yes	Yes	-	-
Bandwidth	.10	-	0.11	-

Note: Robust standard errors clustered by district are reported in parentheses. The running variable is the extremist candidate's win margin in the primary election. Spline indicates that the regression function was fit separately on either side of zero. Cubic refers to a third-order polynomial regression. CCT refers to the method from Calonico, Cattaneo, and Titiunik (2014). IW refers to the method from Imbens and Wager (2019).

J Donor Partisanship Over Time

This section evaluates a potential competing explanation for the decline in the financial penalty to extremist nominees among corporate PACs. Instead of reallocating funds to increasingly-competitive extremist nominees, corporate PACs could have become increasingly partisan in their donating.

To probe this possibility, Figure J.1 plots the share of individual and corporate donors who contributed to at least one candidate from both parties, conditional on making at least five contributions in a given election cycle.[6] The red line in the figure shows that the contribution-weighted share of corporate PACs that contribute to both parties has remained remarkably constant at 94% during the period of study. These results indicate that corporate PACs have not become more partisan in recent years. Individual contributors, in contrast, have become substantially more partisan. In 2000, roughly 35% of individual contributors donated to candidates of both parties (again, weighted by contribution amount), but that number has declined to 5% in 2022.

Figure J.1 – Share of Donors Who Contributed to At Least One Democratic and One Republican Candidate in Congress, 1980-22, and State Legislatures, 2000-22. This figure plots the share of individual and corporate donors who contributed to at least one candidate from both parties, conditional on making at least five contributions in a given election cycle. Results are weighted by contribution amount.

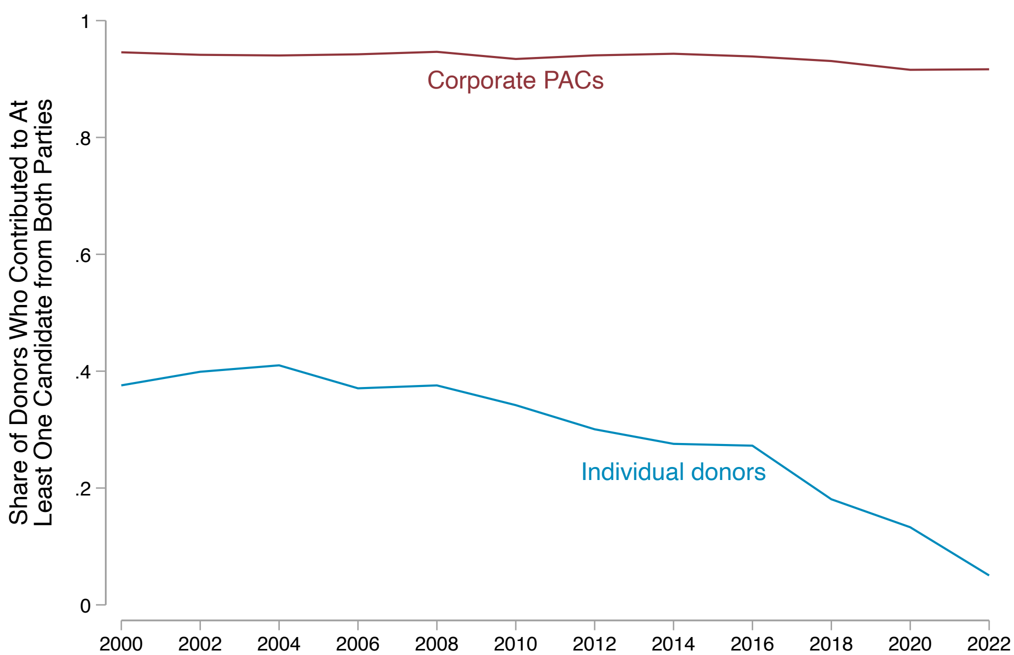

[6]The results are highly similar across a variety of cutoffs.

K Attenuation Bias Simulations

Another alternate explanation for the decline in the financial penalty to extremist nominees is attenuation bias. Specifically, a combination of an increase in measurement error and a decrease in within-party heterogeneity could cause the financial penalty to extremists to appear to decline when its true value remains constant.[7] This section evaluates this possibility using a simulation.

Let the true midpoint regression be expressed as

$$Y_{dt} = \beta_0 + \beta_1 Midpoint_{dt} + \varepsilon_{dt}, \tag{1}$$

where Y_{dt} is the Democratic candidate's share of general-election contributions and $Midpoint_{dt}$ is the true midpoint between Democratic and Republican candidates.[8] Let the variance of $Midpoint_{dt}$ be σ^2_M, and assume that the observed midpoint is given by

$$\widehat{Midpoint}_{dt} = Midpoint_{dt} + \nu_{dt}, \tag{2}$$

where $\nu_{dt} \sim (0, \sigma^2_\nu)$. The observed financial penalty to extremists is estimated as

$$Y_{dt} = \tilde{\beta}_0 + \tilde{\beta}_1 \widehat{Midpoint}_{dt} + u_{dt}. \tag{3}$$

Then the classical error-in-variables model indicates that $\tilde{\beta}_1 = \lambda\beta_1$, where $\lambda = \frac{\sigma^2_M}{\sigma^2_M + \sigma^2_\nu}$.

Clearly, either an increase in σ^2_ν (measurement error) or a decrease in σ^2_M (decreasing party heterogeneity) would increase attenuation bias. Fortunately, it is possible to estimate each of these quantities over time for candidates who are ultimately elected. To measure σ^2_ν, I calculate the variance of the difference between candidates' *Primary-Only Scaling* and their true roll-call voting score (i.e., DW-NOMINATE or NP-Score). To measure σ^2_M, I calculate the variance of legislators' DW-NOMINATE or NP-Score. Using these quantities, I calculate λ_t, the attenuation factor in year t.

Finally, I simulate how much of the decline in the financial penalty to extremists documented in Figure 8 that over-time change in λ_t would explain. Specifically, I fix the baseline value of financial penalty to extremist nominees at its observed value in 2000, and then estimate the counterfactual penalty as the product of the baseline value and λ_t. This quantity is plotted in red in Figure K.1.[9] Because I estimate that the penalty has declined slightly more among corporate PACs than individual donors, I conservatively focus on corporate PACs in Figure K.1. The black line in the figure plots the observed penalty to extremist primary nominees.

As is apparent in Figure K.1, attenuation bias does not explain a meaningful proportion of the decline in the financial penalty to extremists. In this counterfactual scenario, the estimated penalty to extremists remains roughly constant and is far from the observed decline

[7]Note, however, that a decrease in within-party ideological heterogeneity absent measurement error would not bias my estimates.

[8]For simplicity, I omit γ_i, the distinct's Democratic presidential vote share. Results are highly similar when including this variable, because measurement error is not correlated with the district's underlying ideological composition.

[9]As in the main text, I transform the midpoint coefficients to match the RD scale.

plotted with the black line. In fact, it'd require a set of λ_t that are nearly ten times the observed λ_t to fully explain away the observed decline in the financial penalty to extremist nominees.

Figure K.1 – Counterfactual Financial Penalty to Extremist Primary Nominees. This figure plots the counterfactual financial penalty to extremist primary nominees where the baseline penalty is fixed at its original value in 2000 and changes are caused by shifts in the signal-to-noise ratio, λ_t.

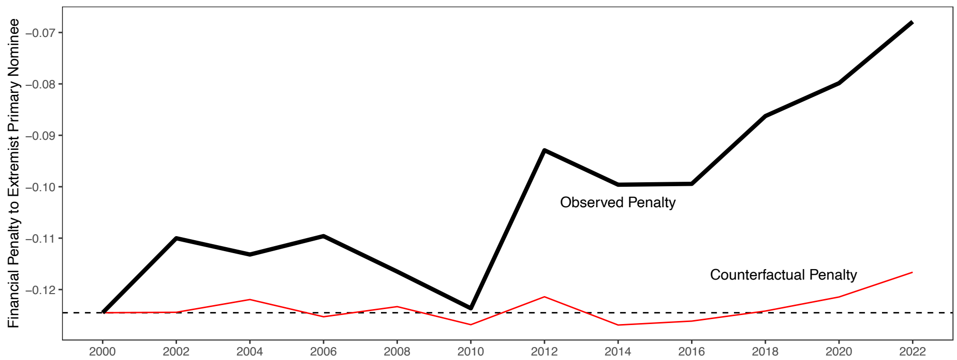

L Roll-Call Classification Exercise

To further validate my preferred ideological scaling, this section uses candidates' *Primary-Specific Scaling* to predict the outcome of observed roll-call votes in Congress and state legislatures.

Data on roll-call votes in Congress was downloaded from Voteview (Lewis et al., 2024). This dataset includes the universe of roll-call votes cast for the years 1980-2023 and data on roll-call voting in 2024 through September 1st. In total, this includes 12 million roll-call votes.

State legislative roll-call data was assembled from two sources. First, data for the near-universe of roll-call votes cast in all 99 state legislative chambers between January 1st, 2010 and September 1st, 2024 was collected from www.Legiscan.com. This dataset consists of 60.8 million individual votes. I supplement this dataset with 11.2 million roll-call votes for the years 2000-2009 from Fouirnaies and Hall (2022) for a varying panel of 21 states.[10] All together, this roll-call dataset encompasses 72 million distinct votes. Following Bonica (2014, 2018) and Poole (2007), I remove lopsided roll calls with margins greater than 97.5% and omit abstentions and missed votes. Table L.1 reports the total number roll-call votes in this dataset by level and year.

To evaluate the predictive ability of my ideological scaling and other measures of candidate ideology, I calculate the optimal cutting point between "yea" and "nay" votes following Poole (2007). Specifically, for every roll-call in our dataset, I find the maximally-classifying point in one-dimensional space that predicts "Yea" votes on one side and "Nay" votes on the other. Leveraging these cutpoints, I impute predicted roll-call votes and compare the result to the true votes cast.

Table L.2 reports the classification rates and aggregate proportional reduction in error (APRE) for the primary-specific scaling and, for comparison, Static CFscores, an indicator for party, and scalings derived directly from incumbents' roll-call voting in office (DW-NOMINATE for members of Congress and NP-Scores for state legislators).[11] DW-NOMINATE and NP-Scores are estimated using roll-call votes themselves and represent a theoretical upper-bound on classification rate, while static CFscores are estimated using the full contribution matrix (i.e., primary- and general-election contributions). I find that the *Primary-Specific Scaling* predicts 89.5% of state legislative roll-call votes correctly ($APRE = .716$), outperforming CFscores and an indicator for party, and closely behind DW-NOMINATE and NP-Scores themselves (91.1%; $APRE = .759$). In sum, despite restricting the size of the training contribution matrix, I am still able to consistently recover candidates' ideological positioning.

[10] I include the unbalanced panel of states from 2000-2009 in my main analyses to evaluate the predictive capacity of my scalings over an extended time frame. The results in Table L.2 are very similar if I instead focus on the years for which I have a balanced panel.

[11] $\text{APRE}_i = \frac{\sum_{j=1}^{J}\{\text{minority vote}_j - \text{classification errors}_{ij}\}}{\sum_{j=1}^{J}\text{minority votes}_j}$ for scaling i and roll call j. This quantity measures the extent to which a given scaling improves upon the naive prediction that every legislator always votes with the majority.

Table L.1 – Number of Congressional and State Legislative Roll-Call Votes Included in Roll-Call Prediction Sample.

Year	Overall	Congress	State Legislatures	Year	Overall	Congress	State Legislatures
1980	315,742	315,742	–	2003	1,808,744	339,465	1,469,279
1981	202,296	202,296	–	2004	1,162,502	257,096	905,406
1982	245,565	245,565	–	2005	1,749,748	326,389	1,423,359
1983	252,648	252,648	–	2006	1,155,209	261,662	893,547
1984	205,754	205,754	–	2007	1,851,129	554,794	1,296,335
1985	228,355	228,355	–	2008	1,227,608	319,183	908,425
1986	230,797	230,797	–	2009	2,302,626	467,924	1,834,702
1987	253,249	253,249	–	2010	2,527,895	315,142	2,212,753
1988	232,348	232,348	–	2011	5,142,218	431,903	4,710,315
1989	190,199	190,199	–	2012	4,207,630	306,161	3,901,469
1990	253,321	253,321	–	2013	5,209,044	308,007	4,901,037
1991	213,039	213,039	–	2014	4,005,615	279,056	3,726,559
1992	231,964	231,964	–	2015	5,786,226	337,515	5,448,711
1993	298,676	298,676	–	2016	4,342,870	284,653	4,058,217
1994	249,419	249,419	–	2017	6,252,840	338,575	5,914,265
1995	437,149	437,149	–	2018	4,863,361	241,009	4,622,352
1996	227,096	227,096	–	2019	6,510,474	346,421	6,164,053
1997	304,035	304,035	–	2020	3,756,663	137,408	3,619,255
1998	262,188	262,188	–	2021	6,470,780	246,070	6,224,710
1999	301,777	301,777	–	2022	5,025,915	277,911	4,748,004
2000	815,548	290,518	525,030	2023	6,843,060	289,276	6,553,784
2001	1,593,291	257,550	1,335,741	2024	4,513,115	128,433	4,490,115
2002	882,478	235,085	647,393				

Table L.2 – Percent of Congressional and State Legislative Roll-Call Votes Classified Correctly. The *Primary-specific scaling* predicts roll-call votes better than CFscores or a naive indicator for party, and nearly as well as scalings derived directly from incumbents' roll-call voting records (DW-NOMINATE/NP-Scores).

Scaling	Overall	Congress	State Legislatures
DW-NOMINATE/NP-Scores	0.911	0.904	0.910
	(0.759)	(0.764)	(0.751)
Primary-Specific Scaling	0.895	0.895	0.897
	(0.716)	(0.720)	(0.713)
Static CFscore	0.886	0.891	0.882
	(0.696)	(0.734)	(0.658)
Party	0.857	0.845	0.850
	(0.587)	(0.500)	(0.584)

Note: Aggregate proportional reduction in error (APRE) reported in parentheses. Table is ordered by overall classification rate.