Online Appendix

Intended for online publication only.

Contents

A.1 Additional Statistical Results . . . 38
A.1.1 Characteristics of the RD Sample . . . 38
A.1.2 RD Balance Tests . . . 39
A.1.3 Alternative RD Specifications . . . 40
A.1.4 RD Results Across Minimum Donor Threshold . . . 41
A.1.5 RD Results Controlling For Winner's Total Primary Receipts . . . 42
A.1.6 RD Results Using State Legislative Roll-Call Votes . . . 43
A.1.7 RD Results Omitting Leaners . . . 44
A.1.8 RD Results on CCES Reported Vote Choice . . . 45
A.1.9 RD Results on Independents . . . 46
A.1.10 Test of RD with Placebo Cut-Points . . . 47
A.1.11 McCrary Test . . . 48
A.2 Panel Results . . . 49
A.3 Discussion of Roll-Off . . . 54
A.4 House Elections and Turnout . . . 55
A.5 Details on Datasets . . . 56
A.5.1 Cleaning the Ansolabehere et al. (2007) Election Data . . . 56

A.1 Additional Statistical Results

In this section, we present a variety of additional statistical analyses.

A.1.1 Characteristics of the RD Sample

Table A.1 – Characteristics of Districts in the RD Sample, U.S. House, 2006–2014.

	Full Sample	Sample Near Discontinuity
% Dem	0.35	0.31
Lag Dem Pres Vote Share	0.55	0.53
Safe District	0.41	0.39
% Open Races	0.35	0.41
% Presidential Elections	0.43	0.42
N	234	113

Sample near the discontinuity includes races in which the extremist won or lost the primary by less that 10%.

In order to assess the locality of the RD estimates, we compare cases near the RD threshold—that is, primary races where the top two candidates were within 10% of each other in top two vote share—to the full sample of contested primaries between two ideologically distinct candidates (i.e., cases where the top two candidates' distance, in terms of ideology, was at or above the median distance across all cases). The table above compares the full sample to the close races. As we see, both the full sample and the RD sample are quite evenly split in terms of Democratic and Republican primary elections and in terms of partisanship, as measured using lagged presidential vote. Defining "safe" districts to be those where the lagged presidential vote share was above 60% for one party or the other, we see that the RD sample is slightly skewed towards competitive places, but still contains a large share of safe districts (in contrast to a general-election RD, which would mostly contain competitive and not safe districts). The RD sample is tilted more towards open-seat races, because these are when many of the most competitive primaries occur. And the RD sample appears to be slightly tilted towards midterm races, although not substantially so.

A.1.2 RD Balance Tests

Table A.2 – Differences In Party's Lagged General-Election Turnout, U.S. House, 2006–2014.

	Partisan Share of Turnout, $t-1$			
Extremist Nominee	-0.05	-0.01	-0.04	-0.06
	(0.05)	(0.04)	(0.05)	(0.05)
N	64	127	127	71
Polynomial	1	3	5	CCT
Bandwidth	0.10	–	–	0.11

Robust standard errors clustered by district in parentheses in columns 1-3; standard error in column 4 comes from rdrobust package and is clustered by district. The running variable is the extremist primary candidate's vote share winning margin in the primary.

The key assumption to the RD is that candidates cannot sort across the discontinuity. If such sorting occurs, then the districts that just barely nominate an extremist will be systematically different from those that nominate a moderate. Technically speaking, this will create a discontinuity in district covariates at the RD cutoff. To test for this, we exactly replicate the specifications from the RD table in the paper, except using *lagged* turnout instead of turnout. If a violation of the RD assumption is driving our estimates, we should see that the estimated treatment effect in these lag tests is large and negative. Instead, as the table shows, we find small coefficients and we cannot reject the null of no difference.

A.1.3 Alternative RD Specifications

Figure A.1 – RD Estimates Across Bandwidths and Specifications

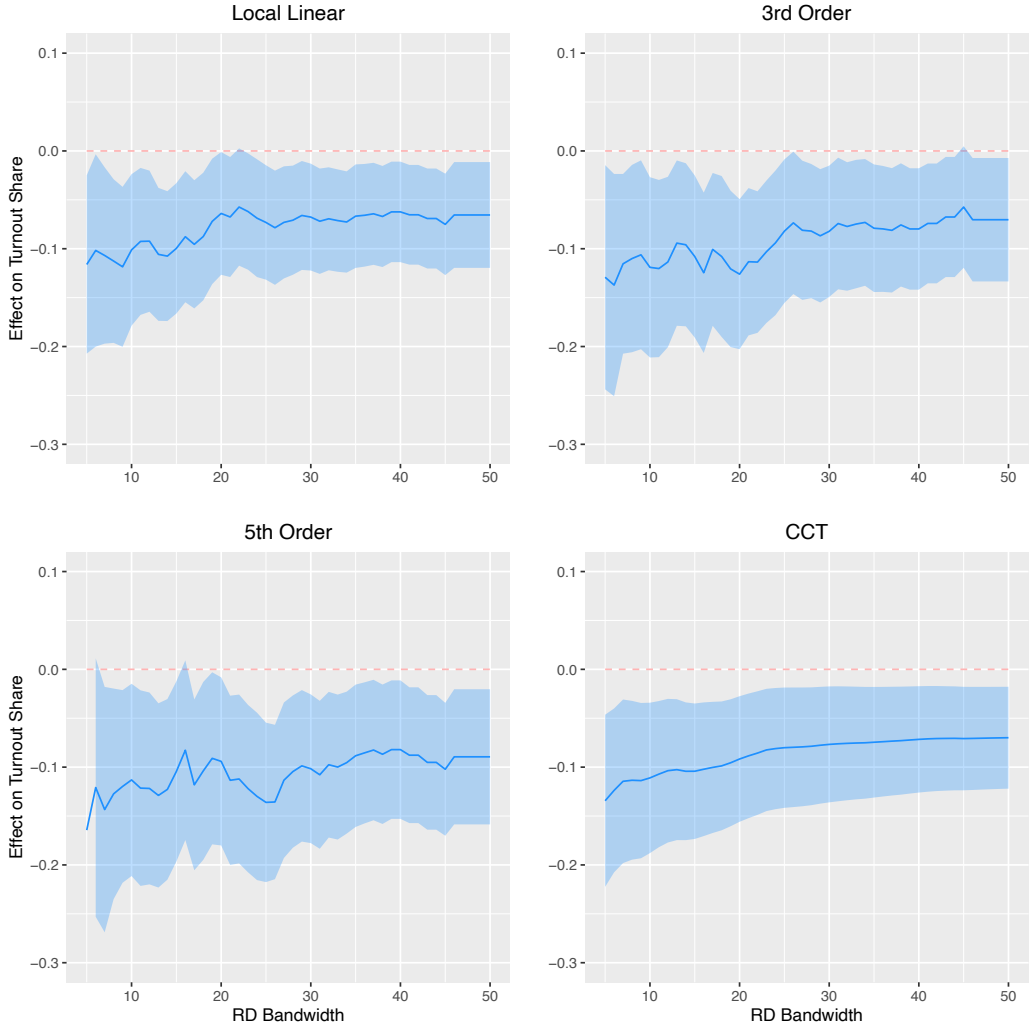

In the paper, we already present the RD at a variety of bandwidths and specifications, including at the "optimal" bandwidth using the CCT procedure. In this section, we present more evidence that our conclusions are robust to the choice of bandwidth and specification. The figure above plots the local linear, 3rd-order and 5th-order polynomial, and CCT estimates across bandwidths ranging from 5 (i.e., races where the vote-share distance between the winning candidate and the 50% threshold was below 5 percentage points) to 50 (i.e., including all races). As the plot shows, we consistently find large, negative results for all plausible choices. Although the local linear estimate attenuates at extremely large bandwidths, we would never use the local linear specification with a bandwidth that large.

A.1.4 RD Results Across Minimum Donor Threshold

Figure A.2 – RD Estimates Across Minimum Donor Cutoff

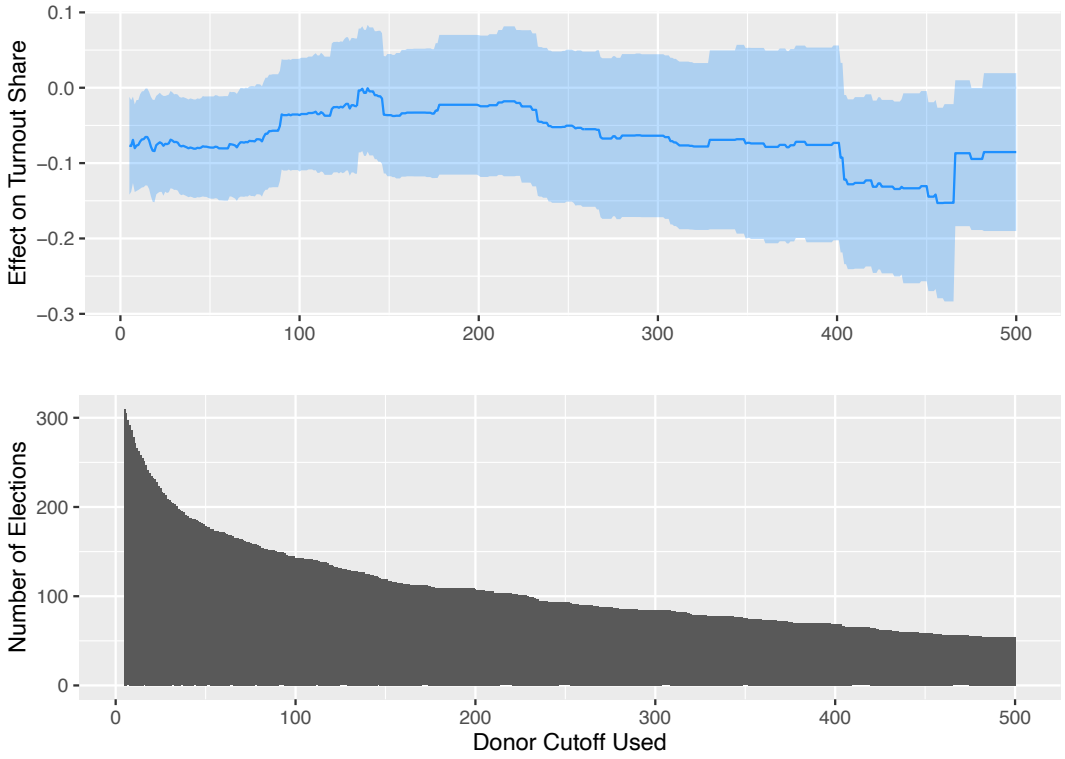

In the paper, we only included in the RD analysis races in which both of the top two candidates in the primary received at least 20 donations over the course of their careers as candidates. to make sure our results do not depend on this arbitrary choice, in the figure above, we re-plot the RD estimate at every donor threshold from 0 to 500. We focus on the 3rd-order polynomial estimates because these are generally the smallest estimates already (so as to be conservative). The plot shows the estimates, 95% confidence intervals, and the sample size. As we see, the estimate does not vary much across cutoffs. Across almost all possible cutoffs, it is large and negative. There are a few regions where the estimate attenuates (around cutoffs of 120 and 220), but these seem to be blips and are at cutoffs where the sample size is already quite small. For any reasonable cutoff (e.g., between 0 and 100) the estimate barely changes at all. Of course, as we increase the cutoff and the sample size decreases, the estimate does become noticeably noisier. But it is clear that the results in the draft are not driven by our choice of 20 as the cutoff.

A.1.5 RD Results Controlling For Winner's Total Primary Receipts

Table A.3 – Effect of Extremist Nominee on Party's General-Election Turnout Controlling for Total Primary Contributions, U.S. House, 2006–2014.

	Partisan Share of Turnout			
Extremist Nominee	-0.09	-0.07	-0.09	-0.10
	(0.04)	(0.03)	(0.04)	(0.04)
Total Contributions	0.01	0.01	0.01	-
	(0.01)	(0.01)	(0.01)	(-)
N	61	128	128	80
Polynomial	1	3	5	CCT

Robust standard errors clustered by district in parentheses in columns 1-3; standard error in column 4 comes from rdrobust package and is clustered by district. The running variable is the extremist primary candidate's vote share winning margin in the primary.

Here, we re-estimate the main RD on party turnout, but controlling for the log total amount of money the bare-winning moderate or extremist raised in the primary. Effects remain substantively large and negative.

A.1.6 RD Results Using State Legislative Roll-Call Votes

Table A.4 – Effect of Extremist Nominee on Party's General-Election Turnout, U.S. House, 2006–2014, Using Shor Legislator Scores.

	Partisan Share of Turnout			
Extremist Nominee	-0.41	-0.04	-0.15	-0.35
	(0.18)	(0.11)	(0.12)	(0.19)
N	26	70	70	27
Polynomial	1	3	5	Kernel
Bandwidth	0.10	–	–	0.10

Robust standard errors clustered by district in parentheses in columns 1-3; standard errors in column 4 come from rdrobust package. The running variable is the extremist primary candidate's vote share winning margin in the primary.

As we discussed in the draft, it is important to make sure the results are not driven by particular scaling we used. In this section, we re-estimate the main RD result on turnout share with a completely different way of scaling candidates. Specifically, we use the Shor and McCarty (2011) data, which scales state legislators in terms of their roll-call votes, to study close primary races between two state legislators. Given these scalings, we carry out precisely the same analysis as before; we focus on primary races where the distance between the top two candidates is at or above the median distance across all the races we have between two state legislators. The table below presents the RD estimates. Unfortunately, sample sizes are too small to draw any strong conclusions. While it is helpful to see that the coefficients are negative, like those reported in the paper, the coefficients are quite fragile across specifications and rarely in the plausible range.

A.1.7 RD Results Omitting Leaners

Table A.5 – Effect of Extremist Nominee on Party's General-Election Turnout, U.S. House, 2006–2014, Leaners Not Included.

	Partisan Share of Turnout			
Extremist Nominee	-0.09	-0.05	-0.08	-0.07
	(0.04)	(0.03)	(0.03)	(0.04)
N	109	228	228	98
Polynomial	1	3	5	CCT
Bandwidth	0.10	–	–	0.09

Robust standard errors clustered by district in parentheses in columns 1-3; standard error in column 4 comes from rdrobust package and is clustered by district. The running variable is the extremist primary candidate's vote share winning margin in the primary.

Our regression discontinuity analysis labeled a respondent a Republican if they reported themselves as "Strong Republican", "Not Very Strong Republican", or "Lean Republican". We used the same convention when dealing with Democratic respondents. We believe this is consistent with best practice, but we also constructed a table of the main results leaving the leaners out of either party. As we expected, the results are quite similar to the results that we included in the main paper; the impact of nominating an extremist on that party's general election turnout share appear to be slightly smaller when leaners are excluded, but it is well within the range of expected results do to noise. The impact of nominating an extremist again appears to have no meaningful effect on that party's general election turnout rate, but meaningfully increases the opposing party's turnout rate, though the effect is quite noisy.

A.1.8 RD Results on CCES Reported Vote Choice

Table A.6 – Effect of Extremist Nominee on Party's General-Election CCES-Reported Vote Share, U.S. House, 2006–2014.

	Vote Share			
Extremist Nominee	-0.10	-0.07	-0.07	-0.12
	(0.04)	(0.03)	(0.04)	(0.04)
N	112	232	232	115
Polynomial	1	3	5	CCT
Bandwidth	0.10	–	–	0.10

Robust standard errors clustered by district in parentheses in columns 1-3; standard error in column 4 comes from rdrobust package and is clustered by district. The running variable is the extremist primary candidate's vote share winning margin in the primary.

As a sanity check on the CCES data, we can also estimate the RD effect of extremist nominees on self-reported vote choice in the CCES, as opposed to using the actual election outcomes themselves. As we see, we find similarly large, negative estimates when we use the CCES self-reported vote choice.

A.1.9 RD Results on Independents

Table A.7 – Effect of Extremist Nominee on Rates of Independent General-Election Turnout, U.S. House, 2006–2014.

	Independents			
Turnout Rate	0.05	0.02	-0.02	0.02
	(0.07)	(0.06)	(0.07)	(0.07)
# Turn Out	0.27	-1.03	-0.65	0.02
	(1.19)	(0.94)	(1.07)	(1.19)
# Registered	-1.34	-2.82	-2.23	-1.09
	(2.03)	(1.51)	(1.79)	(2.09)
N	113	234	234	141
Polynomial	1	3	5	CCT

Robust standard errors clustered by district in parentheses in columns 1-3; standard errors in column 4 come from rdrobust package and are clustered by district. The running variable is the extremist primary candidate's vote-share winning margin in the primary.

Our analysis focuses on voters who at least lean to one party. The choices independents make may also be different when they face an extreme candidate or a moderate. Given the small sample sizes, we are unable to say definitively how pure independents respond to extreme candidates, but positive estimates suggest to us that these voters behave similarly to opposing party voters.

A.1.10 Test of RD with Placebo Cut-Points

Table A.8 – Placebo Estimate Across All Feasible Cut Points.

	Placebo Estimate			
Avg Effect	0.00	-0.00	-0.00	0.00
Rejection Rate	0.01	0.01	0.01	0.15
Polynomial	1	3	5	CCT

Every election in the data with at least 40 elections on both sides of it is used as a possible placebo cut-point. Each side of the true cutoff is evaluated separately. Rejection rates based on 95% confidence interval and the same specifications and standard error estimators as in the main analyses. The running variable is the distance from the placebo cut point.

Following RD best practices, we use our exact RD specifications but using placebo cutoffs where no jump in treatment occurs. We report the average estimated effect and the rejection rate (with 95% confidence) for each of the four specifications. As the average estimates show, we find no evidence of bias, as all four estimates are very close to zero. For the three OLS specifications (columns 1-3), we also find no evidence of over-rejection. For the CCT procedure, we find some over-rejection (0.11 instead of 0.05), perhaps because of the relatively small sample sizes. Since our results in the manuscript do not depend at all on using CCT, this behavior is not overly concerning, though it may be relevant for RD practitioners considering specification choices in the future.

A.1.11 McCrary Test

Figure A.3 – McCrary Test for Running Variable

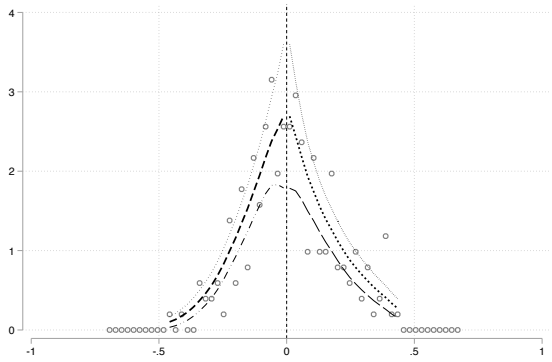

In this section, we perform the McCrary test to investigate whether moderate or extremist candidates systematically sort across the discontinuity. We find no evidence of sorting (log height of 0.27 with a standard error of 0.31.) Figure A.3 presents the output from the Stata program downloaded from `https://eml.berkeley.edu/~jmccrary/DCdensity/`.

A.2 Panel Results

Now we attempt a variety of observational approaches all intended to help hold the underlying features of the district constant. We replicate two main observational strategies from the literature, and add one new adaptation of our own that takes advantage of recent data releases.

In the Canes-Wrone, Brady, and Cogan (2002) specification, we estimate

$$Y_{it} = \beta_0 + \beta_1 Dem\ Extremism_{it} + \beta_2 Rep\ Extremism_{it} + X_{it} + \epsilon_{it}, \tag{2}$$

where Y_{it} is either the two-party vote share or a measure of turnout in district i at time t. We mainly focus, arbitrarily, on the Democratic vote share and turnout share, though in the Appendix we show these effects for Republican outcomes as well. In Canes-Wrone, Brady, and Cogan (2002), the authors use a measure of roll-call extremism that is simply the absolute value of the incumbent's DW-NOMINATE score. To generalize this to all candidates, instead of using DW-NOMINATE, we define $Dem\ Extremism_{it}$ and $Rep\ Extremism_{it}$ as $|Dem\ CFScore_{it}|$ and $|Rep\ CFScore_{it}|$, respectively. Higher values of this variable generally indicate more extreme ideological positioning. To account for the partisanship of the district, Canes-Wrone, Brady, and Cogan (2002) use a variety of controls including presidential vote share. The vector X_{it} stands in for these controls. Although we first present results without controls, when we include them, they are: lagged presidential vote share; the estimated CFScore for the average donor in the district; an indicator for Dem and Rep incumbency, respectively; an indicator for whether the race is an open-seat race; and the number of donations received by the Democratic and Republican candidates, respectively, over all of the campaigns they run in.

The key assumption to the Canes-Wrone, Brady, and Cogan (2002) approach is that DW-NOMINATE scores further from zero correspond to more extreme positions, once the partisanship of the district is accounted for. While this may be a plausible assumption, there is no guarantee that this is the case—if an incumbent is more moderate than the median voter in her district, then by increasing the absolute value of her DW-NOMINATE score or her CFScore, she could actually become less extreme relative to her district. To account for this possibility, we also implement an alternate approach in which we attempt to directly measure how far away candidates are from their districts. In particular, we define $Distance_{pit} = |CFScore_{pit} - AvgDonorCF_i|$, where $CFScore_{pit}$ is the CFScore of party p's general-election candidate in district i in the election at time t and $AvgDonorCF_i$ is the average CFScore of all the donors in district i. We call these variables *Dem Distance* and *Rep Distance*, respectively.

Needless to say, the average donor is not the median voter. Donors are more politically active than voters and appear to hold more extreme political views (Hill and Huber 2015). Despite this issue, the distance between candidates and the average donor is still likely to inform us about how non-median the candidate is. As long as the bias in the average donor score relative to the median voter is relatively fixed across districts, computing candidate distances will cancel the bias out. At a bare minimum, the approach offers a nice alternative to existing approaches—if it comes to similar conclusions, we can be more confident that we are finding the correct answer.

Armed with this distance variable, we then estimate simple equations of the form

$$Y_{it} = \beta_0 + \beta_1 Dem\ Distance_{it} + \beta_2 Rep\ Distance_{it} + X_{it} + \epsilon_{it}, \tag{3}$$

where all variables are defined as before.

Finally, we also use a third approach which avoids attempting to measure the median voter in each district directly but explicitly accounts for the problem that not all shifts away from 0

in CFScores will necessarily indicate extremism. Ansolabehere, Snyder, and Stewart (2001) hold fixed the distance between the candidates and the partisanship of the district and instead look at changes in the *midpoint* between the candidates. No matter where the median voter is located, a shift right in the midpoint between the candidates always benefits the left-wing candidate in the spatial model—so long as the distance between the candidates is held fixed. Again, we use CFScores to operationalize this technique. Using OLS, we estimate equations of the form

$$Y_{it} = \beta_0 + \beta_1 Midpoint_{it} + \beta_2 CandDistance_{it} + X_{it} + \epsilon_{it}, \tag{4}$$

where $Midpoint_{it} = \frac{CFScore_{pit} + CFScore_{-p,it}}{2}$ is the midpoint between the Democratic and Republican candidates' CFScores and $CandDistance_{it} = |CFScore_{pit} - CFScore_{-p,it}|$ is the distance between them. Like in the Canes-Wrone, Brady, and Cogan (2002) approach, we present these results with and without the control variables (the same ones listed above when we discussed the first approach).

Table A.9 presents these estimates on Democratic vote share for all three approaches, with and without optional controls. In all cases, we replicate previous work's finding that more moderate candidates do better, electorally. To aid in interpretability, all explanatory variables are standardized to have mean 0 and standard deviation 1.

In the first two columns, we use the new CFScore approach in which we calculate distance to the average donor in the district. In the first column, we see a substantial association between a candidate being more distant from the average donor and the candidate doing worse, electorally. A one standard deviation increase in the Democratic candidate's distance from the average donor (first row) is associated with an almost 7 percentage-point decrease in the Democratic party's share of the two-party vote. Likewise, a one standard deviation increase in the Republican candidate's distance from the average donor (second row) is associated with an almost 8 percentage-point increase in the Democratic party's share of the two-party vote, or, equivalently, an almost 8 percentage-point decrease for the Republican party. These estimates do become smaller when we control for district and race characteristics, but they remain substantial and very precisely estimated.

Estimates for the other two approaches are highly similar. In the middle two columns, we see that more extreme Democratic candidates, where extreme is now measured using the Canes-Wrone et. al. approach, are associated with lower Democratic vote shares, while more extreme Republican candidates are associated with higher Democratic vote shares. In the final two columns, we see that shifts right in the midpoint between the two candidates—which can be interpreted as making the Republican candidate more extreme and/or the Democratic candidate more moderate—are associated with an increase in Democratic vote share, too. All in all, we find clear evidence that moderate candidates do better than more extreme candidates in U.S. House elections.

We now turn to the estimates on the Democratic party's share of turnout in the general election, estimated from the CCES as discussed in the Data section. Table **??** presents these results in exactly the same format as the previous table.

Here, results all point in the same direction but are not always as precise as the vote share estimates. In the very first column, we see that a one standard deviation increase in the distance between the Democratic candidate and the average donor in the district is associated with almost a 3 percentage-point decrease in the Democratic party's share of turnout in the general election; a corresponding increase in the Republican candidate's distance likewise increases the Democrat turnout share in the general election by more than 5 percentage points. However, these estimates attenuate noticeably when control variables are added. While the resulting estimates are still sizable, we can no longer reject the null of no relationship.

When we turn to the Canes-Wrone estimates, we find more stable estimates, at least for Democratic candidates. In the final two columns, we see that shifts right in the midpoint, corresponding

to more moderate Democratic candidates and/or more extreme Republican candidates, are associated with an increase in the Democratic turnout share. Again, these results are somewhat noisy but, again, in the same direction.

Taken as a whole, the panel estimates *suggest* that more extreme candidates are associated with their parties voters making up a smaller part of general-election turnout. In every case, the estimated coefficients are in this direction, and in most cases, the magnitude of the coefficients is meaningful (around 1 percentage-point per standard deviation, or larger). But the estimates are not particularly precise—of the 10 estimates of interest, 7 are statistically significant at the 0.05 level. In addition, the fact that the estimates move around with the inclusion of control variables makes us hesitant to draw any strong conclusions.

We suspect that the results here are imprecise in large part because of the difficulty of modeling underlying district partisanship. The panel approach leaves us vulnerable to all the usual problems of model dependence, with results dependent on the precise set of control variables included in the regression.

Table A.9 – Candidate Ideology and Electoral Outcomes, U.S. House, 2006–2012.

	Outcome: Dem Vote Percentage (0-100)					
Method:	CFScore		Canes-Wrone et. al.		Ansolabehere et. al.	
Dem Distance	-6.79	-3.66				
	(0.62)	(0.49)				
Rep Distance	7.70	3.23				
	(0.46)	(0.45)				
Dem Extremism			-4.63	-3.00		
			(0.46)	(0.39)		
Rep Extremism			3.47	2.48		
			(0.55)	(0.31)		
Midpoint					5.82	3.95
					(0.53)	(0.42)
Average Donor CFScore		-0.80		-5.06		-4.98
		(0.40)		(0.28)		(0.29)
Lag Presidential Vote		3.29		3.26		3.31
		(0.33)		(0.32)		(0.33)
Open Seat		-2.62		-2.60		-2.65
		(1.15)		(1.13)		(1.17)
Rep Inc		-5.51		-5.42		-5.57
		(1.69)		(1.65)		(1.71)
Dem Inc		-1.21		-1.23		-1.21
		(1.64)		(1.61)		(1.66)
# Donors, Rep		-0.51		-0.52		-0.50
		(0.21)		(0.21)		(0.21)
# Donors, Dem		1.10		1.13		1.11
		(0.38)		(0.39)		(0.38)
Dist Between Cands					-0.41	-0.08
					(0.56)	(0.30)
N	985	978	985	978	985	978

All explanatory variables are standardized to have mean 0 and standard deviation 1. Extremism is defined as the absolute value of the candidate's CFScore. Distance is measured as the absolute distance between each candidate and the district's average donor, in CFScores. Midpoint is average between the Dem and Rep's CFScores in each race. Robust standard errors in parentheses.

Table A.10 – Candidate Ideology and Partisan Turnout, U.S. House, 2006–2012.

	Outcome: Dem Turnout Percentage (0-100)					
Method:	CFScore		Canes-Wrone et. al.		Ansolabehere et. al.	
Dem Distance	-3.21	-1.36				
	(0.44)	(0.43)				
Rep Distance	5.37	0.96				
	(0.53)	(0.64)				
Dem Extremism			-1.16	-1.17		
			(0.42)	(0.35)		
Rep Extremism			0.67	0.81		
			(0.53)	(0.46)		
Midpoint					1.14	1.32
					(0.50)	(0.48)
Average Donor CFScore		-2.51		-3.95		-3.89
		(0.53)		(0.43)		(0.43)
Lag Presidential Vote		4.78		4.76		4.79
		(0.47)		(0.46)		(0.46)
Open Seat		1.32		1.34		1.31
		(1.28)		(1.30)		(1.28)
Rep Inc		2.62		2.69		2.59
		(1.85)		(1.88)		(1.84)
Dem Inc		1.87		1.86		1.87
		(1.79)		(1.82)		(1.78)
# Donors, Rep		-0.15		-0.16		-0.15
		(0.22)		(0.22)		(0.22)
# Donors, Dem		-0.42		-0.40		-0.42
		(0.24)		(0.24)		(0.24)
Dist Between Cands					-0.29	-0.14
					(0.46)	(0.32)
N	968	961	968	961	968	961

All explanatory variables are standardized to have mean 0 and standard deviation 1. Extremism is defined as the absolute value of the candidate's CFScore. Distance is measured as the absolute distance between each candidate and the district's average donor, in CFScores. Midpoint is average between the Dem and Rep's CFScores in each race. Robust standard errors in parentheses.

A.3 Discussion of Roll-Off

The partisan turnout rate effects could also mask effects on voter roll-off. That is, while random roll-off would not affect how we interpret our results, it is possible that the nomination of an extremist could increase or decrease the rate at which some voters cast votes for higher offices but not for the House, which would change the interpretation of our estimates. For roll-off to produce the negative effect of extremists on the party's turnout share and vote share, if there were no effects on turning out and actually voting for the House, extremist nominees would have to significantly decrease roll-off in their own party and/or increase it in the opposing party. It is hard to see why this would be the case, though it would be an interesting behavioral mechanism for the results if it did occur. Since roll-off is typically found to be around 5% in US House races (e.g., Wattenberg, McAllister, and Salvanto 2000), and we do not expect large differential roll-off rates, we suspect that roll-off is not a main mechanism for our results.

A.4 House Elections and Turnout

A potential concern with the results in the paper is that we document substantial turnout effects based on House races, despite the fact that House races are rarely at the top of the ballot. Conventional wisdom often suggests that most of turnout should be based on top-ballot races, particularly the president and to a lesser extent Senator or Governor. In the paper, we discussed some reasons to believe relatively large turnout effects within House races. Here, we present some results to support the overall claim. In particular, here we document that House turnout varies systematically with the nature of the House race—in particular, in relation to its competitiveness. We do not claim that these analyses are causal; we only claim that the clear variation we show in House turnout suggests that there is more to turnout than only presidential elections.

In particular, we estimate equations of the form

$$\textit{Log Total Votes}_{it} = \beta \textit{House Vote Margin}_{it} + \gamma_i + \delta_t + \epsilon_{it}, \tag{5}$$

where $\textit{Log Total Votes}_{it}$ is the log of the total number of votes cast for the House race at time t in district i. The variable $\textit{House Vote Margin}_{it}$ measures the competitiveness of the House race in terms of the winner's vote margin (so, smaller means more competitive). For interpretability we scale this in terms of 10s of percentage-points, so that a one-unit shift is a 10 percentage-point change in the winning vote margin. Finally, γ_i and δ_t stand in for district and year fixed effects, respectively.

Table A.11 presents the results. Whether we look over a long time period (from 1980 to present) or a short one to match the sample in our main analyses (2004 to present), we see a systematic relationship between competitiveness and turnout. This relationship is equally large in presidential vs. midterm election cycles. On average, a 10 percentage-point increase in the winner's vote margin in a House race is associated with a 6-8% decrease in turnout. We take this as evidence that House-specific electoral features have room to change electoral turnout, even despite the salience of races higher up on the ballot.

Table A.11 – Turnout and Competitiveness in U.S. House Races.

	1980-2016		2004-2016	
	Midterms	Presidentials	Midterms	Presidentials
	Log Total Votes	Log Total Votes	Log Total Votes	Log Total Votes
House Vote Margin	-0.08	-0.07	-0.06	-0.07
	(0.00)	(0.00)	(0.01)	(0.01)
N	3,367	3,800	851	1,276
District Fixed Effects	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes

Robust standard errors clustered by district in parentheses. Vote margin scaled in 10 percentage-point increments. Log Total Votes is log of total votes cast in House race.

A.5 Details on Datasets

A.5.1 Cleaning the Ansolabehere et al. (2007) Election Data

We use the Ansolabehere et al. (2007) election data as the source of primary, general, and special election outcomes in US House races. The data serves two main purposes by providing information on the two-party vote share which we use an outcome explained by candidate positioning, while also providing data on close primary elections and their outcomes for the forcing variable in the regression discontinuity design. We limited the outcome data to cases in which two candidates of opposing parties ran in the election that determined the office holder – either a special or general election. We limited the primary election data on the RD forcing variable to the top two candidates in terms of vote share within their party's primary.