

Who Punishes Extremist Nominees? Candidate Ideology and Turning Out the Base in U.S. Elections*

Andrew B. Hall[†]
Department of Political Science
Stanford University

Daniel M. Thompson[‡]
Department of Political Science
Stanford University

November 27, 2017

Abstract

Political observers, campaign experts, and academics alike argue bitterly over whether it is more important for a party to capture ideologically moderate swing voters or to encourage turnout among hardcore partisans. The behavioral literature in American politics suggests that voters are not informed enough, and are too partisan, to be swing voters, while the institutional literature suggests that moderate candidates tend to perform better. We speak to this debate by examining the link between the ideology of congressional candidates and the turnout of their parties' bases in U.S. House races, 2006–2014. Combining a regression discontinuity design in close primary races with survey and administrative data on individual voter turnout, we find that extremist nominees—as measured by the mix of campaign contributions they receive—suffer electorally, largely because they *decrease* their party's share of turnout in the general election, skewing the electorate towards their opponent's party. The results help show how the behavioral and institutional literatures can be connected. For our sample of elections, turnout appears to be the dominant force in determining election outcomes, but it advantages ideologically moderate candidates because extremists appear to activate the opposing party's base more than their own.

*For helpful discussion, the authors thank Avi Acharya, Bob Erikson, Jim Fearon, Anthony Fowler, Stephen Pettigrew, Kevin Quinn, Ken Shotts, Brad Spahn, Danielle Thomsen, and participants of the MIT American Politics Conference and the Emory Institutions and Lawmaking Conference. For data, the authors thank Shigeo Hirano and Jim Snyder. For guidance using voter file data, the authors especially thank Brad Spahn. All remaining errors are the authors' sole responsibility.

[†]Andrew B. Hall is an Assistant Professor in the Department of Political Science at Stanford University (andrewbenjaminhall@gmail.com, <http://www.andrewbenjaminhall.com>).

[‡]Daniel M. Thompson is a Ph.D. Student in the Department of Political Science at Stanford University.

“The key data is this, and it’s important to re-emphasize if only to shut up the useless, overpaid political consultants who idiotically babble about ‘moving to the center’ or ‘compromising with the other side’...What matters is turning out our voters. That’s it. The Democrats win when we fire up and turn out our base.”

—Blog post on Daily Kos¹

“Democrats cannot win elections without capturing the votes of independent-minded swing voters.”

—Commentary in the Wall Street Journal²

1 Introduction

The current state of American politics, characterized by high degrees of legislative polarization, brinksmanship, and gridlock (e.g., McCarty, Poole, and Rosenthal 2006), has raised new questions about the interplay of ideology and electoral success in U.S. elections. The study of candidate ideology and electoral performance in U.S. elections can be crudely divided into two literatures which seem fundamentally at odds with each other. On one side is what we might call the institutional literature, which uses election data to suggest that there is an electoral advantage for moderate candidates (e.g., Ansolabehere, Snyder, and Stewart 2001; Canes-Wrone, Brady, and Cogan 2002; Erikson 1971; Erikson and Wright 2000; Hall 2015). This literature is often associated with the idea that candidates must appeal to so-called “swing voters” in order to win office. On the other side is what we might call the behavioral literature, which uses survey evidence to suggest that many voters are uninformed and ideologically inconsistent, casting doubt on whether swing voters are relevant, or whether they even exist at all (e.g., Campbell et al. 1960; Converse 1964; Lenz 2012; Miller and Stokes 1963). This latter literature is often associated with the idea that turnout among the parties’ bases determines election outcomes.³ In the strongest version of this claim, voters are rigid partisans, “campaigns consist in large part of reminding voters of their partisan identities—‘mobilizing’ them to support their group at the polls,” and more moderate candidates do no better than more extreme candidates because “election outcomes are essentially random choices among the available parties” (Achen and Bartels 2016: 311-312).

¹<http://www.dailykos.com/story/2014/11/5/1342347/-CRUSH-the-GOP-don-t-compromise-with-em-how-to-win-in-2016-and-what-not-to-do>.

²<http://www.wsj.com/articles/SB10001424052702304252704575155942054483252>

³Hill (2016) is one interesting paper at the intersection of these literatures. The paper combines administrative data on individual voter turnout with precinct-level vote returns to estimate swing voting. Under the assumptions of the paper’s model, partisan turnout is estimated to determine roughly twice as much of the vote share as swing voters are.

This disagreement rages on in the popular press, too, where pundits and campaign practitioners debate the relative merits of hypothetical moderate candidates who capture swing voters or ideologically committed candidates who “fire up the base.” Consider, as a recent example, the 2016 Democratic presidential primary race between Hillary Clinton, widely thought to be the more pragmatic and moderate candidate, and Bernie Sanders, the farther left progressive candidate.⁴ Thom Hartmann, the progressive radio host, argued that Sanders would make a better general-election candidate than Clinton, saying: “the truth is, Democrats win when voter turnout is high. And voter turnout is high when voters have real progressive candidates and a truly progressive platform to support.”⁵ This is a sentiment shared by many progressives, as the colorful Daily Kos quote in the epigraph above suggests. Former member of Congress Barney Frank disagreed with this idea, penning an op-ed telling readers that “Wishful thinking won’t win the White House.” Frank continued, “[Sanders’s] very unwillingness to be confined by existing voter attitudes...is both a very valuable contribution to the democratic dialogue and an obvious bar to winning support from the majority of these very voters in the near term.”⁶ Similar arguments have been made for Republican candidates and for elections up and down the ballot. While the Clinton-Sanders example concerns the presidency, an office whose elections differ fundamentally from the legislative elections we study, it helps highlight the more general argument about turning out the base vs. appealing to swing voters. Although the argument has gone on for decades, empirical evidence to resolve it has been scant. Do more extreme legislative candidates actually galvanize their own bases, and, if so, do they actually perform better than moderate candidates in elections?

In this paper, we attempt to answer these questions, and in so doing, we help connect the behavioral and institutional literatures. In particular, we combine election data with survey data from the Cooperative Congressional Election Study (CCES), as well as comprehensive administrative data on voters from Catalist, to show that more extreme candidates do worse, electorally, in part because they decrease their party’s share of turnout in the general election. As a result, turnout among the parties’ bases is a mechanism for the advantage of moderate candidates, rather than a phenomenon at odds with selecting candidates of certain ideologies. *Extremist candidates*

⁴These are only relative terms, and convey no normative opinion on our part. Many do not consider Clinton’s platform to be particularly moderate, even if less extreme than Sanders; those who agree with Sanders’ positions may not think his views are “extreme” in an absolute sense, also.

⁵<http://www.thomhartmann.com/blog/2016/02/if-you-want-win-go-progressive>

⁶<http://www.politico.com/magazine/story/2015/07/why-progressives-shouldnt-support-bernie-120484>

do worse because, contrary to rhetoric, they fail to galvanize their own base and instead encourage the opposing party's base to turn out more, on average. Put another way, the possibility that voters are tribal partisans, even if true, does not imply that more moderate candidates cannot outperform extremist candidates. In a hypothetical world of rigid partisanship, turnout still allows elections to have ideological consequences.

Our results therefore suggest that the institutional and behavioral literatures on elections are both right, and both wrong. Moderate candidates *do* perform better, electorally, but their success depends largely on their ability to encourage their partisans to turn out—or, similarly, to discourage the other party's partisans from turning out. As we will show, the estimates from our regression discontinuity design are consistent with the possibility that the bulk of the vote-share penalty of extremist nominees comes from changes in turnout (though our tests cannot prove whether or not there are also swing voters at play). A key point we want to emphasize is that these effects on turnout need not be the result of individual voters carefully processing information about candidates' positions—for example, the effects could instead be the result of purely elite-driven behavior filtered through campaigns (e.g., endorsements, fundraising). We will not take a stance one way or another on whether voters are informed or rational. The key takeaway from our results is that, however individual voters process information, their decision to turn out to vote is systematically related to whether a more extreme or more moderate candidate is running—an observation that helps to revise both the institutional and behavioral literatures on American elections.

We also link our study to a long political economy literature on candidate ideology and turnout (Adams and Merrill 2003; Adams, Dow, and Merrill 2006; Downs 1957; Glaeser, Ponzetto, and Shapiro 2005; Hinich and Ordeshook 1969; Peress 2011). We make the simple point that, in a typical spatial model where voters have concave utility (and thus convex disutility), so that they receive increasingly large amounts of disutility from candidates further and further from their ideal point, it is the *opposing* party's voters who should care the most about a party nominating an extremist instead of a moderate. While the party's base voters may prefer a more extreme nominee, the opposing party's base voters dislike this extreme nominee more than the party's base likes him or her. This is even more true if we consider more behavioral models of voter preferences, like loss aversion, in which case voters might especially dislike opposing-party extremists. Our empirical

results are broadly consistent with ideas like these. While the party’s base may or may not turn out more in response to a more extreme nominee, extremist nominees do seem to galvanize turnout in the opposing party.

Documenting these effects empirically requires overcoming a number of issues of omitted variables and selection bias, as well as problems of measurement. Places where more extreme candidates win nomination and stand for office in the general election are sure to vary, systematically, from places where more moderate candidates are nominated. To address this issue, we implement a regression discontinuity design based on close primary elections. This approach ensures, under plausible assumptions, that districts in which a more extreme candidate barely wins nomination and runs in the general election are otherwise identical, in expectation, to those where a more moderate candidate barely wins the primary and runs in the general. Although the identifying assumption of the RD has been challenged in U.S. House general elections (Grimmer et al. 2012; Caughey and Sekhon 2011), our tests, in line with those of Hall (2015), strongly suggest the design is valid in U.S. House primaries.

To implement the RD, we require a way to scale candidates. Previous studies have mainly relied on roll-call votes, which indicate the ideology of incumbents but cannot tell us anything about candidates who have never served in office. We follow recent work in using the mix of campaign contributions that candidates receive as an indicator of ideology (e.g., Bonica 2013, 2014; Hall 2015; Thomsen 2014, 2016). These scalings are invaluable because they inform us not just about incumbents but about candidates who have never served in office. No other measure is available for studying questions about incumbents and challengers at scale. That said, recent work, focusing primarily on the CFScore scalings from Bonica (2014), raises questions about what the scalings actually measure (Hill and Huber 2015; Tausanovitch and Warshaw N.d.). Tausanovitch and Warshaw (N.d.), in particular, shows that CFScores do not correlate well with DW-NOMINATE scores—a conventional measure of roll-call based ideology for sitting legislators from Poole and Rosenthal (1985)—within party. We take these concerns seriously, and we address them in several ways. First, we present all of our main results using the contribution-based scalings developed in Hall and Snyder (2014) and used in Hall (2015). These scalings have a simple interpretation that emphasizes their substantive meaning irrespective of roll-call based ideology, and they also correlate fairly well with DW-NOMINATE within party. Finally, we also show the robustness of our results to the use

of a variety of contribution-based scalings, which run the gamut from not using information about roll-call votes at all, to being setup explicitly to predict roll-call based scalings.

The remainder of the paper is organized as follows. In the following section, we discuss the similarities and differences between the institutional and behavioral literatures on elections and ideology, and we briefly discuss formal theories of candidate ideology and turnout. Following that, we discuss the data we use to study these topics empirically. In the next section, we present regression-discontinuity results that consistently show that extremist candidates do worse, electorally, in part because they skew general-election turnout away from their own party. Finally, we conclude by discussing how our findings help revise and connect behavioral and institutional studies of elections and enhance our understanding of ideology, electoral performance, and turnout.

2 Theoretical Background: Swing Voters and Turning Out the Base

A large literature on American political institutions finds that legislative candidates measured to have more moderate positions do better electorally (e.g., Ansolabehere, Snyder, and Stewart 2001; Canes-Wrone, Brady, and Cogan 2002; Erikson and Wright 2000; Hall 2015). An equally large literature on American voter behavior, however, finds that many voters know little about candidates' ideologies and do not report consistent ideological views (e.g., Broockman 2016; Campbell et al. 1960; Converse 1964; Lenz 2012; Miller and Stokes 1963). Furthermore, in surveys, voters do not seem to express support for candidates based on their ideological proximity (e.g., Tausanovitch and Warshaw 2015). For this reason, the political science literature has long viewed the idea that moderates do better electorally and the idea that individual voters are uninformed and highly partisan as fundamentally at odds. Consider this passage from a recent op-ed by Professors Christopher Achen and Larry Bartels in *The New York Times*:

Decades of social-scientific evidence show that voting behavior is primarily a product of inherited partisan loyalties, social identities and symbolic attachments. Over time, engaged citizens may construct policy preferences and ideologies that rationalize their choices, but those issues are seldom fundamental.

That is one key reason contemporary American politics is so polarized: The electoral penalty for candidates taking extreme positions is quite modest because voters in the political center do not reliably support the candidates closest to them on the issues.⁷

Building off of these arguments, some scholars and campaign practitioners have concluded that elections are much more about “turning out the base”—encouraging partisans who will vote for the party to turn out and vote—than they are about altering ideological positions to appeal to swing voters.⁸ Reviewing precisely these kinds of arguments in 2014, but with a logic that could apply in any year, Professor Lynn Vavreck wrote for *The New York Times* that “The 2014 fight is not over swing voters. It’s for partisans.”⁹

Each of these literatures has a clear strength and a clear weakness. The strength of the institutional approach is that it speaks confidently to overall election results. Because these papers study actual elections, they can make clear points about which types of candidates actually win real elections. The weakness, of course, is that finding this aggregate evidence does not give us direct information about the mechanisms underlying it. Although institutional scholars have been tempted to take these empirical patterns as evidence that moderate voters sway elections, there is, to our knowledge at least, no direct evidence for this idea. The strength of the behavioral literature is that it can examine these mechanisms directly. The main takeaway is that, generally speaking, many voters do not seem informed enough to make systematically ideological choices in whom they support in the voting booth. This forces us to revise electoral theories that depend on swing voters in important ways. The weakness to this approach is that it is difficult to extrapolate from surveys to actual elections. While many respondents may struggle to articulate coherent positions on a survey, we do not know how the electorate as a whole does or does not process ideological information. Even a relatively small number of well-informed voters, for example, could determine most election outcomes if the less informed voters are relatively evenly split between the two parties. One aim of our study is to address these weaknesses by combining data on election outcomes with survey

⁷<http://www.nytimes.com/2016/05/23/opinion/campaign-stops/do-sanders-supporters-favor-his-policies.html>, Accessed August 15, 2016

⁸We are focused here on understanding how candidate ideology affects turnout and swing voters—but it is worth noting that this debate also relates to the electoral strategies of parties. A large literature studies the strategic dynamics of parties targeting swing voters or core supporters (Cox and McCubbins 1986; Cox 2009).

⁹<http://www.nytimes.com/2014/04/23/upshot/the-myth-of-swing-voters-in-midterm-elections.html>, Accessed August 15, 2016.

data on individual voters, so that we can attempt to assess both aggregate electoral outcomes and individual-level mechanisms in a single analysis.

The second point we want to make is that the main findings in the two literatures need not be at odds with one another. Moderate candidates may perform better, electorally, even in the extreme hypothetical where swing voters do not exist, and even if most or all voters are not well informed, if turnout is a mechanism for the support of moderate candidates. In particular, extremist candidates could turn off their own party’s voters from turning out to vote and/or galvanize voters in the opposing party to turn out and vote. Although voter attention is one plausible mechanism for effects of this kind, this idea does not depend on voters being highly informed and learning the ideology of each party’s candidate—perhaps, for example, a party coordinates more get-out-the-vote campaigning when it has nominated a more moderate candidate vs. when it has nominated an extremist. Perhaps newspapers cover moderates more favorably than extremists, inducing supporters to turn out more—or, conversely, perhaps newspapers slam extremists, exciting voters in the opposition party. There are many possible reasons that candidates of varying ideology could affect turnout and, as a result, could induce electoral effects even if there are few or no swing voters. As a first contribution, our focus is on understanding the consequences of extreme vs. moderate candidates in the general election, rather than on identifying which of these mechanisms is most responsible for those consequences.


To explain the idea of how candidates of varying ideology can affect turnout in more detail, we turn now to considering formal theoretic ideas about candidate ideology and voter turnout.

2.1 Formal Theories of Turnout

Although the formal theory literature on elections largely builds on the median-voter theorem’s basic foundation, a significant literature studies the question of candidate ideology and turnout directly (e.g., Adams and Merrill 2003; Adams, Dow, and Merrill 2006; Downs 1957; Glaeser, Ponzetto, and Shapiro 2005; Hinich and Ordeshook 1969; Hortala-Vallve and Esteve-Volart 2011; Peress 2011). We will not test the predictions of these models directly, but we will make one simple point about how moderate vs. extremist nominees might affect turnout in these theoretical models.

Suppose, as is standard, that voters have concave utility over policy positions, i.e., convex disutility, so that they dislike positions farther from their ideal point more and more. This assumption

Figure 1 – Candidate Ideology and Turnout. With a typical concave utility function, the nomination of an extremist over a moderate affects voters in the opposing party more than voters in the nominee’s party. Specifically, voter j , in the opposing party, loses more utility than voter i gains when voter i ’s party nominates an extremist instead of a moderate.



is not the only way to construct a model in which opposing-party voters turn out more in response to extremist nominees, but it is a helpful and standard example. Now imagine a typical voter in each party considering turning out when the Republicans nominate a moderate vs. when they nominate an extremist. Figure 1 presents the situation graphically. The figure considers a base voter in each party—voter j is a Democrat; voter i is a Republican. Each voter is located at their bliss point on the ideological spectrum. The upside-down parabolas represent each voter’s utility, which is zero at its maximum when a candidate is at the voter’s ideal point, and is decreasing away from the ideal point—and, crucially, following normal assumptions, the rate at which it decreases goes up as the distance grows.

Suppose that the Republican party nominates either a moderate Republican candidate (labeled on the plot as the moderate) or an extremist (also labeled on the plot). Clearly, voter i is closer to the extremist than the moderate. Her utility is higher at the extremist’s point than at the moderate’s, so she may be more willing to turn out in the general when the party nominates the extremist. Conversely, though, voter j loses more utility than voter i gains when we consider the switch from the moderate to the extremist. Voter j ’s utility function makes this clear; because it gets steeper as we get further from voter j ’s bliss point, and because candidates in the other party are naturally farther away from j than they are from i , voters in the opposing party are likely to

be more horrified by an extremist nominee than the party’s own base is pleased by one. If this fear drives voter turnout, then extremists will galvanize voters in the opposing party more than those in their own party’s base.

Another reason extremist nominees may galvanize turnout in the opposing party has to do with “rational inattention,” an area of growing interest in political economy. Matějka and Tabellini (2017), for example, offers a model in which rational, forward-looking voters consider whether to invest effort in becoming informed about candidates. A key factor in their decision is what is at stake in the election—that is, it is more valuable for voters to be informed if the potential policy outcomes induced by the election are more consequential to them. Although it is not a specific prediction put forward in the paper, opposing-party extremist voters may have a particularly strong incentive to become informed when the opposing party nominates an extremist. We suspect that rationally inattentive voters are a promising recent literature to which our empirical tests will be relevant.

The spatial model massively abstracts from the actual electoral process. As we have presented it, it omits all of the non-ideological factors that candidates possess, as well as the effects of media, campaigns, and more. We have also assumed that utility functions are quadratic rather than, for example, Gaussian, or asymmetric as in prospect theory (Carroll et al. 2013; Quattrone and Tversky 1988; Brady and Ansolabehere 1989).¹⁰ Further, some research questions the ability of voters to process information about candidate ideology directly (e.g., Ahler, Citrin, and Lenz 2016; Tausanovitch and Warshaw 2013) leaving the task of translating ideology into voter behavior to elites or factions (Snyder and Ting 2002), campaigns (Franklin 1991), or some other mechanism. We do not claim our simple sketch captures all of these possibilities, but it helps illustrate the simple point that, while extremists might galvanize their own base, they could just as easily, if not more easily, galvanize voters in the opposing party. Having laid this idea out, we now turn to our empirical analyses of candidate ideology, electoral outcomes, and turnout.

¹⁰The illustrative model produces roughly the same insight with quadratic loss as with prospect theory. The same effects can be recovered if voters have Gaussian loss functions, but the location of the voter and candidates matters more in this case.

3 Data on Elections, Ideology, and Turnout in the U.S. House

To investigate these topics, we combine three main datasets. The first is a dataset containing all U.S. House primary and general elections, 1980–2014, compiled from primary sources for a series of papers including Ansolabehere et al. (2007) and extended to later years by the same authors.

The second dataset comes from the Cooperative Congressional Election Study (CCES) (Schaffner and Ansolabehere 2015) and measures partisan turnout in U.S. House general elections for the years 2006–2014. For each respondent in the CCES, we rely on the validated turnout indicator rather than the self-reported turnout. For this variable, each respondent’s turnout status is validated using administrative data compiled by Catalist. We define Democrats to be any respondent who reports identifying with the Democratic party in the 7-point party ID item, including strong Democrats, not very strong, and leaners, and we define Republicans analogously. This choice is consistent with the finding that leaners and weak partisans behave in nearly identical ways (Abramowitz 2012; Keith et al. 1992; Pew Research Center 2014).¹¹ We can then compute two types of turnout rates from the CCES data: the share of validated voters who identify as Democrats, and the share of respondents who turn out broken out by party. This first turnout share variable tells us whether the electorate in the district skews towards one party or another in a given election; the second tells us whether members of a given party turned out more or less than usual. For example, imagine a district in which 250 Democrats and 250 Republicans are sampled, and 100 Democrats and 200 Republicans have validated turnout. In this district, the Democratic share of validated votes is $100/(100+200) = 33\%$, the Democratic turnout rate is $100/250 = 40\%$, and the Republican turnout rate is $200/250 = 80\%$.

Response bias from the fact that taking surveys like the CCES is voluntary is certainly a concern. Indeed, validated turnout rates in the CCES are substantially above the population turnout rate (Ansolabehere and Hersh 2012). To attempt to account for these kinds of issues, we do two things. First, we only make comparisons *across* CCES respondents (e.g., turnout for Republicans vs. for Democrats in the same race), to difference out the fixed part of this bias. Second, and more importantly, we replicate our main results using Catalist voter-file data, directly. This alternative data contains full voter files, removing all response bias (although it has other, unrelated issues

¹¹In the Appendix, we report all relevant results using only strong and weak partisans and find that the estimates are consistent with our main findings.

that we discuss below). We find extremely similar point estimates using this data, suggesting that survey response bias does not affect our conclusions.

The third dataset contains the ideological scalings of congressional candidates and comes from Hall and Snyder (2014). The Hall-Snyder scalings are created in two stages. First, donors are given an imputed DW-NOMINATE score (hereafter, Nominatate score) based on the roll-call scalings of the incumbents to whom they donate. Donors with more positive scalings are thus donors that donate more to farther right incumbents; donors with more negative scalings donate more to far left incumbents. In the second stage, candidates' scores are imputed based on the donors from whom they received contributions. A farther left candidate, in the Hall-Snyder scaling, is one who receives more donations from donors who tend to support farther left incumbents, for example.

Details on how we cleaned and merged these datasets can be found in the Appendix.

3.1 Considering Issues with Contribution-Based Candidate Scalings

Recent literature has raised questions about what exactly contribution-based scalings reveal about candidate positions.¹² In particular, Tausanovitch and Warshaw (N.d.) shows that the CFScores developed in Bonica (2014) do not correlate well with estimates of candidate ideology built from roll-call votes (like the Nominatate scalings from Poole and Rosenthal (1985)). The most important point for us to make is that the Hall-Snyder scalings we use actually correlate fairly well with Nominatate, within party; as Hall and Snyder (2014) shows, the Hall-Snyder scores correlate at 0.61 with Nominatate scores for Democrats and at 0.53 for Republican Nominatate scores. Since the Hall-Snyder scalings are built off of Nominatate (though not intended to predict them, necessarily), it is not surprising that they correlate with it more highly than CFScores, which use only the candidate-donor matrix to extract a dimension of ideology.

This raises a broader point about scaling candidates. No scaling technique can isolate ideology, since it is such a complicated object with different meanings in different contexts. Roll-call votes indicate one important aspect of the concept, the part related to the incumbent's formal behavior in the legislature.¹³ Sometimes, like when we are specifically studying roll-call votes—one of the

¹²Hill and Huber (2015) show a low correlation between individual donor's issue positions and their CFScore scalings. Although this is an important issue, since we only use candidate-level scalings rather than donor scalings, we do not focus on this critique here.

¹³In some cases, we may be interested in ideology as represented by roll-call vote behavior on high stakes votes or for some other set of votes (Clinton 2006).

most important representative acts a legislator performs—it makes sense to focus on measures that explicitly attempt to predict roll-call votes for candidates who have not cast them. In other contexts, the goal might not be specifically to predict roll-call votes. When voters, elites, and the media consider the ideology of potential representatives, they may consider a variety of factors. How will a potential candidate shape legislation? What issues will she prioritize? Whose voices will she listen to? These activities include roll-call voting, but also many other items that are distinct from roll-call voting. The mix of campaign contributions a candidate receives may tell us something about who is extreme and who is moderate—in the precise sense of what kinds of donors her portfolio is weighted towards—that is valuable and somewhat distinct from roll-call voting. This is why it is not clear that the goal of contribution-based scalings should be to perfectly recover the Nominat scale. This is clearest to us when considering the Hall-Snyder scalings. As described above, these scalings convey substantive information whether or not they correlate with Nominat scores. It is coherent to talk about an extremist—a candidate who raises most of her money from donors who tend to donate to extreme members of the legislature—in the context of these scalings.

That being said, it is important to understand whether our results depend on our choice of scaling method and the degree to which it accords with Nominat scores. After presenting our results using Hall-Snyder scores, which have a medium correlation with roll-call scores within party, we then present our results with three other possible scalings: CFScores and Dynamic CFScores, which have very low correlations with Nominat within party; and DW-DIME scores, which are constructed to be as highly correlated with Nominat within party as possible. Although these alternative scalings have some problems for our design, because they use post-treatment contributions, we think they provide a useful range of estimates across the range of within-party correlation with roll-call based scores. As we show, we find results that are similar in magnitude across all the scalings. In the end, our results do not seem to depend on whether or not we use Nominat scores as the ground truth for predicting candidate ideology.

4 Effects of Extremist Nominees on Turnout: RD Evidence

Do more extreme candidates do worse in general elections? To isolate the effect of extremists on general-election turnout, we follow Hall (2015) in focusing on close primary elections between a

more extreme candidate and a more moderate one. The logic is that, while in general districts that nominate more or less extreme candidates will systematically differ from one another, among the set of districts that have close primary races between an extremist and moderate, the ones that “as-if” randomly see the extremist win will be otherwise identical, in expectation, to the ones where the moderate wins.

Consider again the rhetoric concerning more moderate vs. more extreme primary candidates. We want to evaluate the consequences of nominating an extremist instead of a more moderate candidate in the general election. Is it true that extremists increase turnout among the base? Or is it the case that more moderate candidates successfully secure swing voters? More moderate and more extreme candidates surely differ from each other in other ways beyond ideology, but by systematically studying these types of choices, we can learn about the electoral and representational consequences of voters’ decision to nominate more extreme or more moderate candidates, whatever the myriad mechanisms underlying these consequences. While the RD estimates the effect of a “bundled” treatment, the bundle of all candidate characteristics correlated with being an extremist, these are precisely the bundles voters are constrained to choose between, and thus the quantity estimated is directly relevant.

As this discussion makes clear, the goal is *not* to isolate the “causal effect” of candidate *positions*, themselves. In fact, it is not even clear that there is such a thing as a causal effect of candidate positions. Like gender or race, ideology is probably a deep characteristic built throughout one’s life; it almost certainly violates the “no causation without manipulation” maxim. Although one can imagine altering the candidate positions that voters see in a survey or field experiment, this does not isolate actual changes in candidate positions. If candidates with different ideologies differ on other dimensions, then voters may make inferences about these other attributes when they are told that a candidate holds certain positions (e.g., Gerber, Gooch, and Huber 2017).

In order to implement the RD design, we need to identify extremist and moderate primary candidates. First, we limit the data to contested primary races where the top two candidates in the race both have raised enough money to have a Hall-Snyder score. Some of these races occur between two candidates who are very distinct, ideologically, where it is clear that one candidate is the extremist and one the moderate. But in other races, the candidates are closer together, ideologically. Following Hall (2015), we define a primary race as being one between an extremist

Table 1 – Effect of Extremist Nominee on Party’s General-Election Vote Share and Victory, U.S. House, 2006–2014.

	Vote Share				Victory			
Extremist Nominee	-0.12 (0.04)	-0.07 (0.03)	-0.10 (0.04)	-0.14 (0.04)	-0.38 (0.17)	-0.34 (0.12)	-0.42 (0.14)	-0.51 (0.17)
N	111	230	230	125	111	230	230	114
Polynomial	1	3	5	CCT	1	3	5	CCT
Bandwidth	0.10	–	–	0.11	0.10	–	–	0.10

Robust standard errors clustered by district in parentheses in columns 1-3 and 4-7; standard errors in columns 4 and 8 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.

and a moderate if the ideological distance between the candidates, in terms of their contribution-based scalings, is at or above the median distance between the top two primary candidates across all races in the sample.


Specifically, we estimate equations of the form

$$Y_{ipt} = \beta_0 + \beta_1 \text{Extremist Win}_{ipt} + f(V_{ipt}) + \epsilon_{it}, \quad (1)$$

where Y_{ipt} is one of our outcome variables measured for party p in district i at time t , usually either party p ’s vote share or turnout share. $\text{Extremist Win}_{ipt}$ is an indicator variable for whether the extremist candidate wins the primary for party p in district i at time t . The quantity of interest is thus β_1 , which reflects the RD estimate of the effect of an extremist nomination. Finally, the function $f(V_{ipt})$ is the function used to control for the extremist’s vote-share winning margin in the primary, V_{ipt} .

In this paper, we use this design to study turnout effects rather than vote-share effects like in the original paper. That being said, we first replicate the Hall (2015) estimates on vote share. We do this replication because our subsequent analyses will be using a smaller dataset that only includes election years since the beginning of the CCES in 2006, while the original paper used data from 1980-on. We will also refer back to these vote-share estimates later in interpreting the turnout estimates. Table 1 presents the estimates. As is standard in RD studies, we report estimates at a variety of specifications and bandwidths for $f(V_{ipt})$. In the first column, we use a local linear OLS estimated separately on each side of the discontinuity, using only observations where the primary

Figure 2 – The Effect of Extremist Nominees on Validated Partisan Turnout in the General Election. U.S. House, 2006–2014. When a party nominates an extremist in its primary, general-election turnout skews towards the opposing party.



Note: The dependent variable (on the vertical axis) is the share of all CCES respondents who turned out in a given general election who reported being strong, not very strong, or leaning supporters of the party holding the primary election (each observation is a district-year-primary). Points are averages in equal-sample-sized bins of the horizontal axis variable. Regression lines are from OLS estimated on raw data (not binned), separately on each side of the discontinuity. Shaded area represents 95% confidence interval.

winner won by 10 percentage points or less. In the second column, we use all the data with a 3rd-order polynomial of the running variable. In the third column, we increase this polynomial to 5th-order. In the fourth column, we use the automated bandwidth selection and kernel estimation from Calonico, Cattaneo, and Titiunik (2014). These first four columns are all on general-election vote share; the final four columns replicate these specifications for electoral victory (an indicator) instead of vote share.

As the table shows, we find large effects on vote share and electoral victory, consistent with Hall (2015). Nominating an extremist drastically reduces the party’s electoral fortunes in the general election. Although the estimate does move around from specification to specification, likely because of the reduced sample size when only focusing on data from 2006–2014, it is consistently negative and statistically significant.

Having documented the effect of extremist nominees on vote share, we now examine their effect on partisan turnout. Following typical practices, we begin by studying the RD graphically. Figure 2 plots binned averages of the extremist candidate’s winning margin in each primary race, on the

Table 2 – Effect of Extremist Nominee on Party’s General-Election Turnout, U.S. House, 2006–2014.

	Partisan Share of Turnout			
Extremist Nominee	-0.10 (0.04)	-0.06 (0.03)	-0.08 (0.04)	-0.09 (0.04)
N	107	224	224	110
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 `rdr` package and is clustered by district. The running variable is the extremist primary candidate’s vote share winning margin in the primary.

horizontal axis, against the general-election vote share of the primary winner, on the vertical axis. When the winning margin is above 0, to the right of the vertical line in the plot, the extremist candidate from among the top two primary candidates wins the race, and the party fields an extremist in the general election. When the winning margin is below 0, to the left of the vertical line in the plot, the moderate wins and stands in the general election instead.

The plot also includes OLS lines fit to each side of the discontinuity, separately. These lines, as well as their 95% confidence intervals, are estimated from the underlying data, not the binned averages. As the lines show, there appears to be a noticeable drop at the discontinuity; when the party goes from just barely nominating a moderate to just barely nominating an extremist, its share of turnout in the general election drops by roughly 6 percentage points in the graph.

Table 2 presents the formal estimates, using the same specifications as in the vote share RD. No matter the specification, we find strong, negative effects of extremist nominees on their party’s share of turnout in the general election.¹⁴ Consider the first column. Here, we estimate that nominating an extremist instead of a moderate causes a party’s voters to constitute 10 percentage-points less of turnout in the general election. Looking across the columns, the estimates range from -6 to -10 percentage points. All together, we find strong evidence that extremist nominees depress their party’s share of turnout in the general election, on average.

¹⁴In the Appendix, we re-estimate the RD adding a control for the log of total contributions the bare-winning moderate or extremist raises in the primary. The results are similar to those presented here.

It is worth restating this finding in the context of the RD, to make clear the things that are *not* driving this effect. We can think of the effect in the context of a hypothetical identical district. In one world, this district receives a more extreme nominee—let’s say it’s a Democratic extremist. In the other world, this district receives a more moderate Democratic nominee. Everything else, including the identity of the Republican opponent and the partisanship of the district, are exactly the same across these two worlds.¹⁵ In the case where the extremist now runs in the general election, our findings say that the pool of people who turn out to vote in the general will have a lower proportion of Democratic voters—that is, Democrats will make up a smaller share of the electorate—than in the case where the moderate runs in the general election instead.

These effects are substantial, and we might wonder how simply changing the identity of a U.S. House nominee could alter turnout so much. Below, we offer some analyses on possible mechanisms, and we suggest that campaigning effects—the ways moderate nominees run their general-election campaigns, and the support they get from party elites—probably drive some of the effect. Consistent with this possibility, our effect estimates are quite similar to the estimates of the aggregated effects of presidential campaigns on local turnout. While individual get-out-the-vote (GOTV) interventions only seem to affect turnout by roughly a percentage point, on average (Green, McGrath, and Aronow 2013), Enos and Fowler (N.d.) estimates that the stacked effects of many such interventions over the course of a presidential campaign causes approximately an 8 percentage-point increase in turnout, squarely in the range of our estimates.

Additionally, our own analyses suggest that House-specific features of the electoral landscape can alter turnout more than might be expected, if one thinks of turnout as mainly driven by presidential elections. To show this, in the Appendix we use data on district-level House vote from 1980 to 2016. We use the winning margin in the race to predict total votes cast. We include district and year fixed effects to account for the possibility that some districts and some election years have both higher turnout and more competitive races. Within districts, when a House race is more competitive, more people turn out, even relative to average turnout across the country in that year. This holds even when we only estimate the regression with presidential election years,

¹⁵This might be violated if the Republican primary came after the Democratic primary in this example. Since partisan primaries for Congress were held on the same day in all cases in our sample, runoffs are the only cases that might introduce a cross-party effect on primaries. There are two such cases. While this could just be understood as another consequence of the treatment, we have re-run our analyses without these cases and the point estimates and standard errors are unchanged to the second decimal point.

and when we zoom up only on recent elections. The relationship is large, too. For example, using only the presidential elections 2004, 2008, 2012, and 2016, we estimate that a 10 percentage-point increase in the House election win margin is associated with roughly a 7 percentage-point reduction in total turnout, relative to turnout in a hypothetical tied race, on average (this association is larger if we include earlier presidential elections, too.) These are substantial shifts in turnout that are not driven by presidential voting.

In interpreting our effects, we should also keep in mind that the underlying “intervention” we are studying, changing the identity of a party’s nominee from a moderate to an extremist, is quite powerful, even if we are studying House rather than presidential elections. GOTV effects are estimated conditional on the identity of the candidates in the race, while the effects we estimate here change who the candidate is. Since candidates appeal to voters in different ways, and run different campaigns, it is plausible that our design could estimate turnout effects that are much larger than the effects of a single GOTV intervention.

4.1 Validating the RD’s Identifying Assumption

For the RD results above to be valid, and not confounded by unobserved differences in the potential outcomes of the districts that nominate extremists vs. those that nominate moderates, it must be the case that candidates cannot “sort” across the discontinuity. Past research has raised questions about the validity of this assumption in U.S. House general elections, offering evidence that candidates who just barely win differ in observable ways (most notably, in terms of incumbency status) from those that barely lose (Grimmer et al. 2012; Caughey and Sekhon 2011; Snyder 2005). However, Eggers et al. (2015) presents evidence that these concerns do not typically apply to most electoral RD designs, finding no evidence of such sorting in any other electoral case.¹⁶ Hall (2015), using the same close primaries RD design as the present study, finds no evidence of sorting, either. Following these studies, in the Appendix we present balance tests that suggest there is no sorting in close U.S. House primary elections. Specifically, we re-estimate the RD specifications above but using *lagged* turnout share as the dependent variable (since this is our main outcome variable of

¹⁶For further discussion of these issues, also see de la Cuesta and Imai (2016), Erikson and Rader (2013), and Skovron and Titunik (2015).

interest). We find no evidence that treated and control districts differ from each other, on average, in terms of lagged turnout share.¹⁷

4.2 Replicating the RD With Administrative Data on Turnout

The CCES data is useful because it provides a wealth of additional information about individual voters, such as their self-reported ideology and partisanship, which we will exploit below. Before doing so, though, we need to make sure that our results are not driven by some of the undesirable aspects of the CCES sample. In particular, we might worry that the set of people who are willing to answer an exhaustive survey about politics are especially engaged individuals. Effects documented on this unusual sample of people may not generalize to the whole electorate, especially since the number of voters per district in the CCES sample is not particularly large.

Table 3 – Effect of Extremist Nominee on Party’s General-Election Turnout, U.S. House, 2006–2014.

	Partisan Share of Turnout (Catalist)			
Extremist Nominee	-0.09 (0.05)	-0.05 (0.04)	-0.08 (0.05)	-0.12 (0.06)
N	110	227	227	101
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.

To address this issue, we now re-estimate the RD, but using data on turnout by party from Catalist, which assembles voter files for all fifty states. We estimate turnout by party for 2014 using a one-percent random sample of the Catalist database drawn in August 2015; for 2012 using a one-percent random sample of the Catalist database drawn in September 2013; and for 2006 through 2010 using a one-percent random sample drawn in November 2011. A few important peculiarities of the samples bear mentioning. First, both samples are random samples of registered voters in a given

¹⁷Generally, researchers present balance tests on additional variables in addition to the lagged outcome variable (although this is obviously the most important to show). Because Hall (2015) already presents an exhaustive set of balance tests using the same data and same design, finding no evidence of imbalance, we do not do so here.

state. This means that voters who move across state borders are overrepresented in our data. The 2013 sample includes all observations for any sampled individual, so we can re-weight the turnout statistics to remove this bias for the 2012 estimates, but we are unable to use this technique in earlier years. Second, 20 states do not include party information as a part of their voter registration process. For these states, we use a modeled Democratic party propensity variable that Catalyst constructed to classify registrants in these states as Republicans and Democrats. Lastly, Catalyst began gathering voter files between 2006 and 2008, so some voters who moved between 2006 and 2008 may be missing in the 2006 analysis.

Table 3 presents the results. Although the estimates are noisier, probably due to some of the data issues discussed above, we obtain very similar estimates when we use the full voter file instead of the CCES sample.

4.3 RD Results Similar Across Possible Scaling Techniques

We now examine the robustness of the results to using alternative contribution-based scalings. Table 4 replicates the RD estimates for every available contribution-based scaling. The first row reproduces the results from before, using the Hall-Snyder scalings and using the same four RD specifications we have used throughout. The second row uses the vanilla CFScore scaling from Bonica (2014). The third row uses Dynamic CFScores, which are the same as CFScores except estimated separately for each year, so that candidate scalings can vary over time. The final row uses the DW-DIME scaling from Bonica (2017). This scaling uses machine-learning techniques to predict Nominat scores as well as possible, and is extremely highly correlated with Nominat within party. All three of these alternative scalings may suffer from some post-treatment bias, because they use contributions that candidates receive after winning office, unlike the Hall-Snyder scores. In addition, the DW-DIME scalings are likely to be less precise because they implicitly draw information from a smaller set of donors (those that are algorithmically found to be most predictive of roll-call votes). But they all provide useful differences from the Hall-Snyder scores and are useful for assessing the robustness of our results.

As the table shows, we tend to find similar results with all of the measures. In all 16 specifications across all 4 scalings, estimates are negative and substantively meaningful (the smallest estimate is -0.04, which we would argue is still a fairly large effect in the electoral context.) Estimates appear

Table 4 – Effect of Extremist Nominee on Party’s General-Election Turnout Across Scalings, U.S. House, 2006–2014.

	Partisan Share of Turnout			
Hall-Snyder	-0.10 (0.04)	-0.06 (0.03)	-0.08 (0.04)	-0.09 (0.04)
CF Score	-0.14 (0.04)	-0.06 (0.03)	-0.10 (0.04)	-0.12 (0.04)
Dynamic CF	-0.08 (0.04)	-0.06 (0.03)	-0.10 (0.04)	-0.12 (0.04)
DW-DIME	-0.04 (0.05)	-0.04 (0.04)	-0.06 (0.05)	-0.04 (0.06)
Polynomial Bandwidth	1 0.10	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. Each pair of rows represents the effect estimates from a different measure of ideology.

to be larger with CFScores than with Hall-Snyder scores, and smaller with DW-DIME scores (the DW-DIME scores are also about 20% noisier, in terms of standard errors), but the overall pattern is fairly clear.

In the Appendix, we also attempt to validate our approach by re-estimating the RD design with an entirely separate measure of candidate ideology that does not depend on donations. We use NP-Scores, based on state legislative roll-call votes Shor and McCarty (2011), and we focus on close House primary races between two state legislators. Unfortunately, there are simply not enough cases to generate a stable RD estimate, and the coefficient estimates vary massively across bandwidths. Although the lack of stability brings with it some implausibly large effect sizes, all of the RD estimates in this alternative approach are also negative. As such, we view this as a weak but useful validation exercise.

Given these robustness checks, we conclude that our findings are not driven by our choice of scaling from among the available options. The fact that we find similar estimates across the spectrum from no within-party correlation with Nominato to high within-party correlation with

Nominate suggests to us that we are capturing something meaningful about extremist and moderate candidates.

4.4 Considering the Local Nature of the RD

In order to identify the effects of extremist nominees, the RD necessarily focuses on close primary races, which make the resulting estimates “local” to this context. In the present study, we have good reasons to think that this local estimate is more generalizable than is often possible. As Hall (2015) shows, close primary elections occur in a variety of districts. Unlike general-election RDs, which by definition only take place in competitive districts, primary election RDs are able to utilize both competitive districts and safe districts, because the dominant party in safe districts often has competitive primaries. In the Appendix, we offer some descriptive evidence to suggest that the RD sample is similar to the broader set of contested primary races, so we suspect that our estimates are not particularly local. Finally, since we also present similar results using the panel analysis (see Appendix), which relies on a much broader set of elections, we are relatively confident in the generalizability of the RD estimates.

5 Who Punishes Extremists? Evaluating Mechanisms.

5.1 Is the Penalty to Extremists Driven Mostly By Turnout and Not Swing Voters?

How much of the overall electoral penalty that extremists receive comes purely from the overall turnout effect, and how much comes from voters switching sides? There is no way to answer this conclusively without observing individual vote choices, but we can conduct a suggestive test. If the RD-estimated effects on partisan turnout share are very different from the effects on vote share, then we might have more reason to think swing voting is occurring.

To make these ideas clear, suppose for starters that all voters are either Democrats or Republicans (i.e., no independents, an issue we discuss below.) Suppose we observe a large effect on vote share and no effect on turnout share. This would directly rule out the hypothesis that all voters are rigid partisans, because if all voters were rigid partisans changes in vote share would have to accord with changes in turnout share. If, on the other hand, the effects are quite similar, we might

suspect that turnout and not swing voting is the main explanation. This would not be conclusive because compositional changes in the numbers of swing voters turning out in each party could coincidentally lead to similar changes in turnout and vote share, but it would be suggestive because in a hypothetical world with no swing voting, where every voter is a rigid partisan, the *only* pattern consistent with a vote share estimate is for there to be an equivalent change in turnout share.

Figure 3 plots the RD estimates on vote share and partisan turnout side-by-side for each specification. As a reminder, the outcome variable for the turnout analysis here is the partisan turnout share, the fraction of Democratic voters among all Democratic and Republican voters who cast votes. Since this is the same denominator as for vote share, effects for turnout share and for vote share are directly comparable.

As we see, the estimates are similar. Though we stress that this is not a dispositive test, it is consistent with the hypothesis that most of the effect of extremist nominees on electoral outcomes is driven by changes in turnout, with many voters fixed in their partisan loyalty. The first estimate in the plot, which uses a 3rd-order polynomial that tends to be the most conservative estimate across specifications in the paper, shows highly similar point estimates on turnout and on vote share, with extremist nominees estimated to cause roughly a 6 percentage-point decrease in the turnout share and a 7 percentage-point decrease in vote share. Recall our hypothetical example from the Data section, where we supposed that 300 voters turned out, 100 of which were Democrats, so that the Democratic turnout share was 33%, and, supposing that all partisans vote for their own party, the Democratic vote share would also be 33%. Now suppose that this hypothetical election occurred with a moderate nominee. The point estimate on turnout share here says that, on average, we would expect the Democratic turnout share to fall to 27%, a decrease of 6 percentage points, and we would expect the Democratic vote share to fall to 26%, a decrease of 7 percentage points.

Subsequent estimates are similar, though the gap between the two effects does grow, particularly when we use CCT, which uses the least amount of data and is therefore the most variable in small samples. In all cases, the turnout effect is large in magnitude, and except for the outlier estimate of -0.15 on vote share in the CCT specification, the gap between the estimates is never more than 2 percentage points. As such, we conclude that the turnout of partisan voters appears to be a very important component of the overall vote-share penalty to extremist nominees. More than zero

swing-voting is no doubt occurring, but pundits and campaigners may be right that the turnout of the bases is particularly important.

Purely independent voters, voters who do not lean towards one party or the other, deserve separate comment. As a reminder, we compute our partisan turnout share variable including leaners as members of each party, but we omit pure independents entirely; as such, they do not contribute to the turnout variable. On the other hand, we cannot remove them from the vote share estimate, so if they act as swing voters they may induce variation in the vote share variable that is separate from changes in the partisan share of turnout. Given our results, this could mean that the similarity in the turnout and vote-share estimates is a coincidence, with the vote-share penalty actually driven by swing voting among independents. This is not really a separate problem from what we already discussed above—namely, that in general we cannot rule out any level of swing voting. Because we find turnout share effects similar to vote-share effects, and because pure independents are a smaller portion of the sample, we conclude that a significant portion of the vote-share estimate is likely driven by the change in partisan turnout.


In the Appendix, we also estimate turnout effects for pure independents, who represent roughly 10% of CCES respondents. Given the small sample sizes, we are unable to say definitively how these voters respond to extreme candidates, but positive estimates of the effect of extremist nominees on the rate of independent voter turnout suggest to us that these voters behave similarly to opposing party voters, and are galvanized to turn out when extremists are nominated.¹⁸

5.2 Effects in Midterms vs. Presidential Years

We have documented large overall turnout effects from nominating extremists in U.S. House races. These effects could indicate individual voters making informed decisions about the candidates, or they could instead indicate coordinated campaigning effort among elites and parties that depends on the identity of the nominee. Isolating either of these mechanisms is extremely difficult and will require follow-up work in the future. That said, there is one simple test we can run that may be somewhat informative. If the penalty is larger in midterm years and smaller in presidential years, it might suggest a voter-centric mechanism. In presidential years, voters are more likely to base

¹⁸We also discuss the possibility that roll-off could change the interpretation of these results to some degree in the online appendix.

Figure 3 – RD Effect of Extremist Nominees on Vote Share and Partisan Share of Turnout, U.S. House, 2006–2014. The estimated vote-share penalty to extremists is similar to the estimated effect of extremists on their party’s share of turnout in the general, suggesting that much of the penalty to extremists may be driven by differential turnout.



Note: Estimates are labeled based on the RD specifications from Table 2.

their decision to turn out on the presidential race, while in midterm races voters might be more likely to turn out based on the identity of the House candidates themselves. If we still find a large penalty in presidential years, then, it would suggest the overall results are not as much due to voter choices, directly, since we would expect those choices not to hinge on the identities of the House candidates. On the other hand, if the penalty is much larger in midterm years, and smaller or null in presidential years, it would suggest that the results could be mainly due to voter choices directly.

Table 5 replicates the RD estimates for midterm and presidential years, respectively. Cutting the data in half results in relatively unreliable estimates, so we do not place too much weight on these analyses, but the patterns are interesting. Effects look smaller in presidential years, but not by much (the sample sizes are not sufficient to provide an interesting test of the null that the effects are the same across the two contexts.) For three out of four specifications, we find quite large turnout effects in presidential years. These speculative results suggest to us that much of the turnout penalty might result from coordinated campaign efforts among parties and elites who might be more willing to allocate scarce resources to moderates rather than extremists.

Table 5 – Effect of Extremist Nominee on Party’s General-Election Turnout in Midterm and Pres Years, U.S. House, 2006–2014.

	Partisan Share of Turnout							
	Midterm				Presidential			
Extremist Nominee	-0.10 (0.06)	-0.10 (0.05)	-0.10 (0.05)	-0.10 (0.05)	-0.07 (0.06)	-0.02 (0.05)	-0.07 (0.05)	-0.06 (0.07)
N	62	129	129	83	44	95	95	41
Polynomial	1	3	5	CCT	1	3	5	CCT

Robust standard errors clustered by district in parentheses in columns 1-3 and 4-7; standard errors in columns 4 and 8 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.

5.3 Extremists Appear to Galvanize the Opposing Party’s Base

Who turns out in response to extremist vs. moderate nominees? We have seen that extremist nominees seem to skew the electorate away from the party. This could be due to a number of possible compositional effects. For example, extremists could increase turnout in their own base but increase turnout in the opposing party’s base more, or extremists could depress turnout in their own base but depress it less (or increase it) in the opposing party’s base. Moreover, extremist nominees could affect party identification or registration—because significant numbers of voters register in between the primary and the general—and/or it could affect who chooses to turn out among those already registered or who already identify with a particular party.

Although we are inevitably limited by the amount of data we have, we can carry out some simple tests to see which party’s base seems to be affected more by extremist nominees. Specifically, we re-estimate the RD with three new outcome variables, estimated for each of the two parties. The first is the rate at which members of each party turn out—so, in an observation corresponding to a Republican primary, for example, this variable would be the fraction of all CCES respondents identifying with the Republican party in that district and year who are validated to have turned out, and the other would be the analogous fraction for Democratic voters in the district. We stress that this rate variable is different from the turnout share variable we have been using as the outcome variable thus far in the paper. Recall once more our example from the Data section. There we supposed that 100 Democrats and 200 Republicans turned out, so that the Democratic turnout share was 33% ($100/(100+200)$). But, we supposed that there were 250 members of each party

Table 6 – Effect of Extremist Nominee on Rates of Partisan General-Election Turnout, U.S. House, 2006–2014.

	Own Party				Opposing Party			
Turnout Rate	0.03 (0.07)	0.07 (0.05)	0.03 (0.06)	0.01 (0.07)	0.09 (0.07)	0.10 (0.05)	0.10 (0.06)	0.09 (0.06)
# Turn Out	-5.75 (5.14)	-6.22 (4.10)	-5.70 (4.67)	-5.14 (5.73)	1.83 (5.81)	0.89 (4.28)	0.33 (5.13)	2.18 (5.80)
# Identify w/ Party	-9.81 (5.91)	-12.28 (4.60)	-10.69 (5.39)	-7.97 (7.39)	-1.96 (6.95)	-4.48 (5.29)	-5.59 (6.34)	-1.77 (6.94)
N	110	229	229	138	110	229	229	146
Polynomial	1	3	5	CCT	1	3	5	CCT


Robust standard errors clustered by district in parentheses in columns 1-3 and 5-7; standard errors in columns 4 and 8 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. In the first row, the dependent variable is the total number of CCES respondents in the party who turned out to vote divided by the total number of respondents identifying with the party. The second and third row dependent variables are the numerator and denominator of this ratio, respectively.

sampled, so the Democratic turnout rate would be 40% (100/250). As we discussed in the Data section, we define Democrats as those on the 7-point scale who report being strong, not very strong, or leaning Democrats (and likewise for Republicans). The two remaining variables are simply the numerator and denominator of this turnout rate—the number of respondents who turn out for that party in the district, and the number of respondents in the district who identify (at least weakly) with that party, respectively.

Table 6 presents the results. Though we are limited by the sample size and cannot draw strong conclusions, opposing party voters appear to turn out at higher rates than own party voters. That is, when a party nominates an extremist, voters of the other party turn out at a higher rate than the party’s own voters do. We cannot reject the null that the effects are the same, but the differences are consistent across all four specifications and are substantively large. For example, in the first and fifth columns, we see that the nomination of an extremist is estimated to increase own-party turnout by 3 percentage points, but is estimated to increase opposing party turnout by 9 percentage points.

To investigate specifically whether the effect on opposing-party turnout is bigger than on own-party turnout, we can also re-estimate the RD with the difference in the rate of turnout between opposing-party voters and own-party voters as the outcome variable. To convey the findings from

Figure 4 – Difference in Effect of Extremist Nominee on Turnout Rate, Opposing-Party Voters Minus Own-Party Voters. Across bandwidths, we see consistent though imprecise evidence that extremist nominees affect opposing-party turnout more than own-party turnout.



Note: Line reflects RD estimates using 3rd-order polynomial of running variable. Shaded area represents 95% confidence interval from robust standard errors clustered by district.

this exercise succinctly, Figure 4 plots the resulting RD estimate across all possible RD bandwidths. We use the 3rd-order polynomial specification for this because it is generally the most conservative specification in our analyses, and in this instance, as can be seen in Table 6, it finds the smallest difference in effects and thus biases us against finding a large difference. As the resulting figure shows, we find consistent, though still somewhat imprecise, evidence that extremist nominees affect opposing-party turnout more than own-party turnout.

The second row of Table 6 shows estimated effects on the number of voters from each party who turn out to vote, the numerator for the rate variable estimated in the first row. Here we find negative estimates for the party holding the primary, and small but positive estimates for the opposing party—all insufficiently large or precise to reject the null of no effect. In the third row, we estimate effects on the number of respondents who identify with the party in the district, the denominator for the rate variable estimated in the first row. Fewer voters in both parties appear to identify with a party after an extremist nominee, though effects look quite a bit larger for the

party holding the primary (first four columns.) Again, we cannot reject the null of no difference because of the limited sample sizes. Nevertheless, the negative estimates across both parties, and the fact that they are more negative for the party holding the primary, does seem consistent with the possibility that extremist nominees wage less effective, less active campaigns, with less support for party elites, than moderate nominees do.

Putting these results together, it appears that the overall effect of extremist nominees on turnout is driven by a concentrated group of opposing party voters who turn out at higher rates in response to extremists.

6 Conclusion

This paper engages with a longstanding debate over the relative strengths of extreme legislative candidates, thought to boost turnout among their party's base, and moderate candidates thought to attract hypothetical moderate swing voters. Academics, political practitioners, and pundits have long disagreed about how candidates of varying ideologies perform electorally, as well as about the mechanisms that determine these electoral outcomes. While the institutional literature suggests that more moderate candidates tend to do better, survey-based behavioral work has gone so far as to suggest ideology is entirely irrelevant, with moderate candidates receiving no electoral advantage whatsoever. Resolving these disagreements is key for understanding how the electoral process contributes to the gridlock and polarization among our elected officials.

We have added to this debate by examining how general-election voters respond to the nomination of more or less extreme candidates. Using several different empirical strategies, we have found consistent evidence that extremist nominees do poorly in general elections in large part because they skew turnout in the general election away from their own party and in favor of the opposing party. On the one hand, this result seems to be in keeping with the institutional literature's findings that moderate candidates outperform, electorally speaking, and contrary to claims in the behavioral literature that candidate ideology is irrelevant. On the other hand, the results suggest that much of moderate candidates' success may actually be due to the turnout of partisan voters, rather than to swing voters who switch sides. In fact, our regression discontinuity estimates are consistent with the possibility that the bulk of the vote-share penalty to extremist nominees is the

result of changes in partisan turnout. Seen in this light, the results are more consistent with the behavioral literature’s focus on turnout than they are with the institutional literature’s theoretical focus on swing voters. As such, we see this paper as helping to link the behavioral and institutional literatures together, suggesting that moderate candidates do possess an electoral advantage, but that this advantage may depend heavily on turnout-based mechanisms.

The paper can also help revise formal theories that consider this same question of turning out the base. We have made the simple point that, under the standard kinds of utility functions these theories suppose voters possess, voters in the opposite party should actually have a stronger incentive to turn out when an extremist is nominated, because the extremist is so far away from them, ideologically. Our follow-up analyses on turnout among base voters and among self-identified extremist voters are consistent with our claim. Extremists appear to do little to galvanize turnout in their own party—although the precision of these estimates does not rule out the possibility that they increase own-party turnout—but they do seem to galvanize the opposing party.

In considering our results, we should be clear that they do not resolve the debate over how informed or “rational” voters are. It is entirely possible that extremist nominees increase the opposing party’s turnout because voters in the opposing party are aware of the extremist, are turned off by the extremist’s positions and other attributes, and turn out to vote as a result. But it is equally possible that other mechanisms are at play. Elites, party organizations, campaigns, and the media all play a role. Voter turnout is not only the result of individuals’ rational calculations but, we know, responds to the strategic activities of canvassers, advertisers, and others (e.g., Green and Gerber 2008). Separating out these mechanisms is, we suspect, a very promising avenue for future research. But whatever the mechanism, more extreme candidates in our sample do worse than more moderate candidates in U.S. House elections in large part because they fail to galvanize their own base and instead encourage voters in the opposing party to turn out and vote against them.

References

- Abramowitz, Alan. 2012. "Forecasting in a Polarized Era: The Time for Change Model and the 2012 Presidential Election." *PS: Political Science & Politics* 45(04): 618–619.
- Achen, Christopher H., and Larry M. Bartels. 2016. *Democracy for Realists*. Princeton University Press.
- Adams, James, and Samuel Merrill, III. 2003. "Voter Turnout and Candidate Strategies in American Elections." *Journal of Politics* 65(1): 161–189.
- Adams, James, Jay Dow, and Samuel Merrill, III. 2006. "The Political Consequences of Alienation-Based and Indifference-Based Voter Abstention: Applications to Presidential Elections." *Political Behavior* 28(1): 65–86.
- Ahler, Douglas J, Jack Citrin, and Gabriel S Lenz. 2016. "Do Open Primaries Improve Representation? An Experimental Test of California's 2012 Top-Two Primary." *Legislative Studies Quarterly* 41(2): 237–268.
- Ansolabehere, Stephen, and Eitan Hersh. 2012. "Validation: What Big Data Reveal about Survey Misreporting and the Real Electorate." *Political Analysis* 20(4): 437–459.
- Ansolabehere, Stephen, James M. Snyder, Jr., and Charles Stewart, III. 2001. "Candidate Positioning in US House Elections." *American Journal of Political Science* 45(1): 136–159.
- Ansolabehere, Stephen, John Mark Hansen, Shigeo Hirano, and James M. Snyder, Jr. 2007. "The Incumbency Advantage in U.S. Primary Elections." *Electoral Studies* 26(3): 660–668.
- Bonica, Adam. 2013. "Ideology and Interests in the Political Marketplace." *American Journal of Political Science* 57(2): 294–311.
- Bonica, Adam. 2014. "Mapping the Ideological Marketplace." *American Journal of Political Science* 58(2): 367–386.
- Bonica, Adam. 2017. "Inferring Roll Call Scores from Campaign Contributions Using Supervised Machine Learning." Working Paper.
- Brady, Henry E, and Stephen Ansolabehere. 1989. "The Nature of Utility Functions in Mass Publics." *American Political Science Review* 83(1): 143–163.
- Broockman, David E. 2016. "Approaches to Studying Representation." *Legislative Studies Quarterly* 41(1): 181–215.
- Calonico, Sebastian, Mattias D. Cattaneo, and Rocío Titiunik. 2014. "Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs." *Econometrica* 82(6): 2295–2326.
- Campbell, Angus, Philip E. Converse, Warren E. Miller, and Donald E. Stokes. 1960. *The American Voter*. University of Chicago Press.
- Canes-Wrone, Brandice, David W. Brady, and John F. Cogan. 2002. "Out of Step, Out of Office: Electoral Accountability and House Members' Voting." *American Political Science Review* 96(1): 127–140.

- Carroll, Royce, Jeffrey B Lewis, James Lo, Keith T Poole, and Howard Rosenthal. 2013. "The Structure of Utility in Spatial Models of Voting." *American Journal of Political Science* 57(4): 1008–1028.
- Caughey, Devin M., and Jasjeet S. Sekhon. 2011. "Elections and the Regression Discontinuity Design: Lessons from Close US House Races, 1942–2008." *Political Analysis* 19(4): 385–408.
- Clinton, Joshua D. 2006. "Representation in Congress: Constituents and Roll Calls in the 106th House." *The Journal of Politics* 68(2): 397–409.
- Converse, Philip. 1964. "The Nature of Belief Systems in Mass Publics." In *Ideology and Discontent*, ed. David E. Apter. New York: Free Press.
- Cox, Gary W. 2009. "Swing Voters, Core Voters, and Distributive Politics." In *Political Representation*, ed. Ian Shapiro, Susan C. Stokes, Elisabet Jean Wood, and Alexander S. Kirshner. Cambridge University Press pp. 342–357.
- Cox, Gary W, and Mathew D McCubbins. 1986. "Electoral Politics as a Redistributive Game." *The Journal of Politics* 48(2): 370–389.
- de la Cuesta, Brandon, and Kosuke Imai. 2016. "Misunderstandings About the Regression Discontinuity Design in the Study of Close Elections." *Annual Review of Political Science* 19: 375–396.
- Downs, Anthony. 1957. *An Economic Theory of Democracy*. New York: Harper and Row.
- Eggers, Andrew, Anthony Fowler, Jens Hainmueller, Andrew B. Hall, and James M. Snyder, Jr. 2015. "On the Validity of the Regression Discontinuity Design for Estimating Electoral Effects: Evidence From Over 40,000 Close Races." *American Journal of Political Science* 59(1): 259–274.
- Enos, Ryan D., and Anthony Fowler. N.d. "Aggregate Effects of Large-Scale Campaigns on Voter Turnout." *Political Science Research and Methods*. Forthcoming.
- Erikson, Robert, and Kelly Rader. 2013. "Much Ado About Nothing: RDD and the Incumbency Advantage." Working Paper.
- Erikson, Robert S. 1971. "The Electoral Impact of Congressional Roll Call Voting." *American Political Science Review* 65(4): 1018–1032.
- Erikson, Robert S., and Gerald C. Wright. 2000. "Representation of Constituency Ideology in Congress." In *Continuity and Change in House Elections*, ed. David W. Brady, John F. Cogan, and Morris P. Fiorina. Stanford University Press.
- Franklin, Charles H. 1991. "Eschewing Obfuscation? Campaigns and the Perception of US Senate Incumbents." *American Political Science Review* 85(04): 1193–1214.
- Gerber, Alan, Andrew Gooch, and Greg Huber. 2017. "Evaluations of Candidates' Non-policy Characteristics from Issue Positions: Evidence about Valence Spillover." Working Paper.
- Glaeser, Edward L, Giacomo AM Ponzetto, and Jesse M Shapiro. 2005. "Strategic Extremism: Why Republicans and Democrats Divide on Religious Values." *The Quarterly Journal of Economics* 120(4): 1283–1330.
- Green, Donald P., and Alan S. Gerber. 2008. *Get Out the Vote: How to Increase Voter Turnout*. Brookings Institution Press.

- Green, Donald P., Mary C. McGrath, and Peter M. Aronow. 2013. "Field Experiments and the Study of Voter Turnout." *Journal of Elections, Public Opinion, and Parties* 23(1): 27–48.
- Grimmer, Justin, Eitan Hersh, Brian Feinstein, and Daniel Carpenter. 2012. "Are Close Elections Random?" Presented at the 2012 Annual Meeting of the Midwest Political Science Association. <http://stanford.edu/~jgrimmer/CEF.pdf>.
- Hall, Andrew B. 2015. "What Happens When Extremists Win Primaries?" *American Political Science Review* 109(1): 18–42.
- Hall, Andrew B., and James M. Snyder, Jr. 2014. "Candidate Ideology and Electoral Success." Working Paper.
- Hill, Seth J. 2016. "Changing Votes or Changing Voters? How Candidates and Election Context Swing Voters and Mobilize the Base." Working Paper.
- Hill, Seth J., and Gregory A. Huber. 2015. "Representativeness and Motivations of the Contemporary Donor: Results from Merged Survey and Administrative Records." *Political Behavior* pp. 1–27.
- Hinich, Melvin J, and Peter C Ordeshook. 1969. "Abstentions and Equilibrium in the Electoral Process." *Public Choice* 7(1): 81–106.
- Hortala-Vallve, Rafael, and Berta Esteve-Volart. 2011. "Voter Turnout and Electoral Competition in a Multidimensional Policy Space." *European Journal of Political Economy* 27: 376–384.
- Keith, Bruce E, David B Magleby, Candice J Nelson, Elizabeth A Orr, and Mark C Westlye. 1992. *The Myth of the Independent Voter*. Univ of California Press.
- Lenz, Gabriel S. 2012. *Follow the Leader? How Voters Respond to Politicians' Performance and Policies*. University of Chicago Press.
- Matějka, Filip, and Guido Tabellini. 2017. "Electoral Competition with Rationally Inattentive Voters." Working Paper.
- McCarty, Nolan M., Keith T. Poole, and Howard Rosenthal. 2006. *Polarized America: The Dance of Ideology and Unequal Riches*. MIT Press Cambridge, MA.
- Miller, Warren E., and Donald E. Stokes. 1963. "Constituency Influence in Congress." *American Political Science Review* 57(1): 45–56.
- Peress, Michael. 2011. "Securing the Base: Electoral Competition Under Variable Turnout." *Public Choice* 148(1): 87–104.
- Pew Research Center. 2014. "Political Polarization in the American Public."
- Poole, Keith T., and Howard Rosenthal. 1985. "A Spatial Model for Legislative Roll Call Analysis." *American Journal of Political Science* 29(2): 357–384.
- Quattrone, George A, and Amos Tversky. 1988. "Contrasting Rational and Psychological Analyses of Political Choice." *American Political Science Review* 82(3): 719–736.
- Schaffner, Brian, and Stephen Ansolabehere. 2015. "2010-2014 Cooperative Congressional Election Study Panel Survey." doi:10.7910/DVN/TOE8I1, Harvard Dataverse, V6.

- Shor, Boris, and Nolan McCarty. 2011. "The Ideological Mapping of American Legislatures." *American Political Science Review* 105(3): 530–51.
- Skovron, Christopher, and Rocío Titiunik. 2015. "A Practical Guide to Regression Discontinuity Designs in Political Science." Working Paper.
- Snyder, Jason. 2005. "Detecting Manipulation in US House Elections." Unpublished Manuscript, University of California, Los Angeles. <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.335.6505&rep=rep1&type=pdf>.
- Snyder, Jr., James M., and Michael M. Ting. 2002. "An Informational Rationale for Political Parties." *American Journal of Political Science* 46(1): 90–110.
- Tausanovitch, Chris, and Christopher Warshaw. 2013. "Measuring Constituent Policy Preferences in Congress, State Legislatures, and Cities." *Journal of Politics* 75(2): 330–342.
- Tausanovitch, Chris, and Christopher Warshaw. 2015. "Does the Spatial Proximity Between Legislators and Voters Affect Voting Decisions in U.S. House Elections?" Working Paper.
- Tausanovitch, Chris, and Christopher Warshaw. N.d. "Estimating Candidates' Political Orientation in a Polarized Congress." *Political Analysis*. Forthcoming.
- Thomsen, Danielle M. 2014. "Ideological Moderates Won't Run: How Party Fit Matters for Partisan Polarization in Congress." *The Journal of Politics* 76(7): 786–797.
- Thomsen, Danielle M. 2016. "Gender and Ideology in U.S. House Elections: Reassessing the Conventional Wisdom." Working Paper.
- Wattenberg, Martin P, Ian McAllister, and Anthony Salvanto. 2000. "How Voting is Like Taking an Sat Test: An Analysis of American Voter Rolloff." *American Politics Quarterly* 28(2): 234–250.

Online Appendix

Intended for online publication only.

Contents

A.1	Additional Statistical Results	36
A.1.1	Characteristics of the RD Sample	36
A.1.2	RD Balance Tests	37
A.1.3	Alternative RD Specifications	38
A.1.4	RD Results Across Minimum Donor Threshold	39
A.1.5	RD Results Controlling For Winner’s Total Primary Receipts	40
A.1.6	RD Results Using State Legislative Roll-Call Votes	41
A.1.7	RD Results Omitting Leaners	42
A.1.8	RD Results on CCES Reported Vote Choice	43
A.1.9	RD Results on Independents	44
A.1.10	Test of RD with Placebo Cut-Points	45
A.1.11	McCrary Test	46
A.2	Panel Results	47
A.3	Discussion of Roll-Off	52
A.4	House Elections and Turnout	53
A.5	Details on Datasets	54
A.5.1	Cleaning the Ansolabehere et al. (2007) Election Data	54

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	0.34	0.31
Lag Dem Pres Vote Share	0.54	0.53
Safe District	0.41	0.38
% Open Races	0.33	0.38
% Presidential Elections	0.41	0.40
N	230	111

Sample near the discontinuity includes races in which the extremist won or lost the primary by less than 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	-0.02 (0.06)	-0.00 (0.05)	-0.02 (0.05)	-0.04 (0.05)
N	65	128	128	66
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 `rdrubust` 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

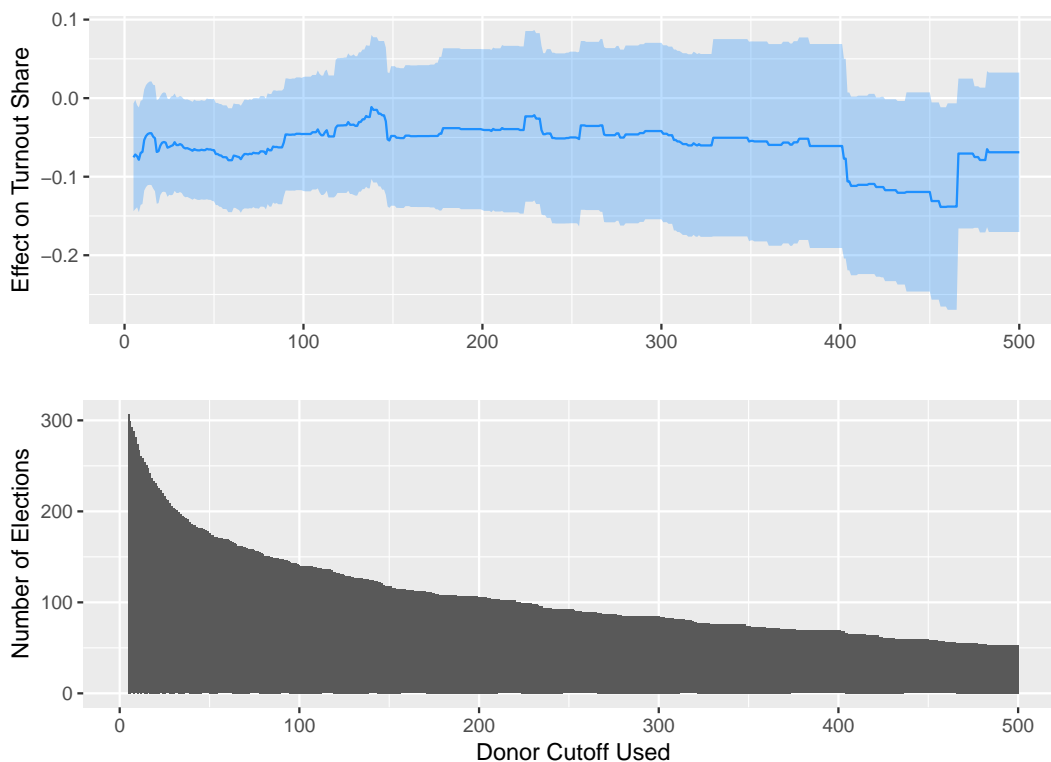
Figure A.1 – RD Estimates Across Bandwidths and Specifications



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	-0.07 (0.04)	-0.06 (0.03)	-0.07 (0.04)	-0.07 (0.04)
Total Contributions	0.01 (0.01)	0.01 (0.01)	0.01 (0.01)	- (-)
N				
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 (though they do lose statistical significance in three specifications, barely.)

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	-0.47 (0.17)	-0.08 (0.11)	-0.24 (0.12)	-0.30 (0.20)
N	25	69	69	21
Polynomial	1	3	5	Kernel
Bandwidth	0.10	–	–	0.08

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	-0.07 (0.04)	-0.04 (0.03)	-0.06 (0.03)	-0.05 (0.04)
N	107	224	224	97
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 (0.04)	-0.06 (0.03)	-0.06 (0.04)	-0.11 (0.05)
N	110	228	228	121
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.

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	0.08 (0.08)	0.04 (0.06)	0.01 (0.07)	0.05 (0.08)
# Turn Out	0.68 (1.14)	-0.64 (0.92)	-0.30 (1.04)	0.47 (1.13)
# Registered	-1.49 (2.08)	-2.79 (1.55)	-2.58 (1.83)	-1.39 (2.08)
N	111	230	230	131
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	0.01	0.00	0.00	-0.00
Rejection Rate	0.01	0.01	0.00	0.11
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

$$Y_{it} = \beta_0 + \beta_1 \text{Dem Extremism}_{it} + \beta_2 \text{Rep Extremism}_{it} + X_{it} + \epsilon_{it}, \quad (2)$$

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 $\text{Dem Extremism}_{it}$ and $\text{Rep Extremism}_{it}$ as $|\text{Dem CFScore}_{it}|$ and $|\text{Rep CFScore}_{it}|$, 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 $\text{Distance}_{pit} = |\text{CFScore}_{pit} - \text{AvgDonorCF}_i|$, 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

$$Y_{it} = \beta_0 + \beta_1 \text{Dem Distance}_{it} + \beta_2 \text{Rep Distance}_{it} + X_{it} + \epsilon_{it}, \quad (3)$$

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

$$Y_{it} = \beta_0 + \beta_1 \text{Midpoint}_{it} + \beta_2 \text{CandDistance}_{it} + X_{it} + \epsilon_{it}, \quad (4)$$

where $\text{Midpoint}_{it} = \frac{\text{CFScore}_{pit} + \text{CFScore}_{-p,it}}{2}$ is the midpoint between the Democratic and Republican candidates' CFScores and $\text{CandDistance}_{it} = |\text{CFScore}_{pit} - \text{CFScore}_{-p,it}|$ 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 A.10 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, 5 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.

Method:	Outcome: Dem Vote Percentage (0-100)					
	CFScore		Canes-Wrone et. al.		Ansolabehere et. al.	
Dem Distance	-6.78 (0.62)	-3.67 (0.49)				
Rep Distance	7.72 (0.47)	3.23 (0.45)				
Dem Extremism			-4.62 (0.46)	-3.00 (0.39)		
Rep Extremism			3.45 (0.55)	2.48 (0.31)		
Midpoint					5.79 (0.53)	3.95 (0.42)
Average Donor CFScore	-5.07 (0.28)		-0.80 (0.40)			-4.98 (0.29)
Lag Presidential Vote		3.26 (0.32)	3.29 (0.33)			3.31 (0.33)
Open Seat	-2.59 (1.13)		-2.61 (1.15)			-2.63 (1.16)
Rep Inc	-5.42 (1.65)		-5.51 (1.69)			-5.57 (1.71)
Dem Inc	-1.23 (1.61)		-1.21 (1.64)			-1.21 (1.66)
# Donors, Rep	-0.52 (0.21)		-0.51 (0.21)			-0.50 (0.21)
# Donors, Dem	1.13 (0.39)		1.10 (0.39)			1.11 (0.38)
Dist Between Cands					-0.42 (0.56)	-0.08 (0.30)
N	982	978	982	978	982	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.

Method:	Outcome: Dem Turnout Percentage (0-100)					
	CFScore		Canes-Wrone et. al.		Ansolabehere et. al.	
Dem Distance	-2.81 (0.45)	-0.99 (0.42)				
Rep Distance	5.30 (0.53)	0.34 (0.61)				
Dem Extremism			-0.85 (0.43)	-0.86 (0.34)		
Rep Extremism			0.55 (0.54)	0.30 (0.46)		
Midpoint					0.86 (0.51)	0.74 (0.46)
Average Donor CFScore	-2.93 (0.42)		-2.13 (0.52)			-2.90 (0.42)
Lag Presidential Vote	6.27 (0.45)		6.28 (0.45)			6.28 (0.45)
Open Seat	1.87 (1.42)		1.86 (1.41)			1.85 (1.40)
Rep Inc	3.50 (2.07)		3.45 (2.05)			3.44 (2.05)
Dem Inc	2.33 (2.01)		2.33 (1.99)			2.33 (1.99)
# Donors, Rep	0.04 (0.25)		0.04 (0.25)			0.05 (0.25)
# Donors, Dem	-0.45 (0.24)		-0.46 (0.24)			-0.46 (0.24)
Dist Between Cands					-0.19 (0.45)	-0.26 (0.31)
N	952	952	952	952	952	952

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

$$\text{Log Total Votes}_{it} = \beta \text{House Vote Margin}_{it} + \gamma_i + \delta_t + \epsilon_{it}, \quad (5)$$

where $\text{Log Total Votes}_{it}$ is the log of the total number of votes cast for the House race at time t in district i . The variable $\text{House Vote Margin}_{it}$ 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 as 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.