

Evaluating the Minority Candidate Penalty with a Regression Discontinuity Approach

September 23, 2022

Abstract

Do parties face an electoral penalty when they nominate candidates of color? We use a regression discontinuity design with state legislative election data from 2018 and 2020 to isolate the effect of nominating a candidate of color on the party's general election performance. Using this approach with real-world data heightens external validity relative to existing racial penalty studies, which are largely supported by surveys and experiments. We find no evidence that candidates of color are disadvantaged in state legislative general elections, relative to narrowly-nominated white candidates from the same party. These findings challenge leading explanations for the underrepresentation of racial/ethnic minority groups, with implications for candidate selection across the United States.

1 Racial Bias and Election Outcomes

Many political scholars and practitioners believe that parties pay a significant penalty in the voting booth when they nominate racial and ethnic minority candidates, particularly in places with heavily-white electorates (Kinder and Dale-Riddle 2012; Lewis-Beck, Tien and Nadeau 2010). There are good reasons to believe this. Decades of research into white racism suggest that many white voters are prejudiced against candidates of color in experimental and survey settings, and that this prejudice results in a dearth of minority representation in majority-white electoral districts. This research is incredibly consequential, influencing everything from redistricting, to candidate ambition, to party leaders tasked with recruiting individuals to run for office (Canon and Posner 1999; Doherty, Dowling and Miller 2019).

However, these electoral penalties may be overstated or obscured, especially in an era of partisan and racial polarization (Kam 2007). And while candidates of color may lose votes from some out-partisan voters, they may simultaneously get a bigger boost from some co-partisans, either in turnout or vote choice (Chudy 2021; Chudy, Piston and Shipper 2019; Jackman and Vavreck 2010; Agadjanian et al. 2020; Stout 2020; Tesler and Sears 2010). In this research note, we offer new evidence to help adjudicate between these contrasting narratives about racial penalties.

Early research into white racial bias and elections in the United States focused on the paucity of minority legislators outside of majority-minority districts (Canon and Posner 1999; Lublin 1999). This work is foundational in redistricting disputes and has been codified into representational legal frameworks for decades (*Page v. Bartels* 2001; *Georgia v. Ashcroft* 2003; *Bartlett v. Strickland* 2009). However, selection bias limits our ability to draw useful inferences from research that only looks at election winners (Juenke 2014), and work that includes both candidates and officeholders finds little evidence of bias penalizing Black or Latino general election candidates that is easily separable from party (Voss and Lublin 2001; Juenke and Shah 2016; Fraga, Shah and Juenke 2020). Furthermore, as candidates of color do not uniformly emerge in all types of districts, it is empirically difficult to distinguish

candidate supply versus voter demand effects using observational designs (Branton 2009) or to separate the effects of party from district racethnic factors (Juenke and Shah 2016).

Given the limitations of observational work, some researchers have turned to experimental research that tests for racial bias in a controlled setting. Experimental work shows white voters' candidate preferences are heavily influenced by racethnic considerations. For example, information search processes differ for many white voters, harming the electoral prospects of candidates of color (Ditonto 2020; Crowder-Meyer et al. 2020; Valentino, Hutchings and White 2002). Partisanship can also work to either minimize the impact of racism (Kam 2007; Sigelman et al. 1995) or reinforce prejudice (Stout 2020; Tesler 2013). Similar disagreements abound regarding the interaction of candidate racethnicity and information like ideology, candidate quality, or incumbency (Andersen and Junn 2010; Kam 2007). The difficulty of replicating the complex and dynamic information environment of political campaigns limits the external validity of the evidence from these experimental studies.

In this paper, we use a regression discontinuity (RD) design that leverages closely-fought primary elections yielding either white or non-white party nominees. We follow previous work using close elections to examine otherwise-similar districts where different candidates are chosen (Fraga and Hassell 2021; Broockman 2014). RD designs emphasize real world external validity while improving our ability to test causal claims. They are particularly valuable tools to help assess the disparate methodological evidence in this literature.

Similar to Bucchianeri (2018)'s work on candidate gender, we leverage close primary winners to estimate the causal effect of party primary voters nominating either a white candidate or a candidate of color. Because these nominees emerge from very similar districts and primary elections near the cutpoint, we can more precisely and confidently estimate any differences in the subsequent general election two-party vote share. While we take care to contextualize the limitations of this design (De la Cuesta and Imai 2016; Marshall 2021) the results offer a unique and significant contribution to the racial penalty literature.

2 States as Laboratories of (Racial) Representation

State legislative elections are ideal for testing theories of bias because there are thousands of partisan elections in the states every two years compared to hundreds of congressional races, thus creating statistical power for the kinds of racial/ethnic comparisons that are difficult to make using congressional election data. For this paper, we focus on primary elections where a white candidate narrowly defeated, or was narrowly defeated by, a candidate of color. The dataset consists of 391 elections (from 36 states) in which the top two candidates in a primary were of different racethnicities, such that the outcome of the primary election determined whether the party’s nominee was white or a candidate of color.¹

The regression discontinuity (RD) design provides a test of what happens to a party’s general election voteshare when a candidate of color [white candidate] wins a competitive primary and the party advances them to a competitive general election, relative to a white candidate [candidate of color]. While other work has explored what happens when candidates of color (Juenke and Shah 2016) and women of color (Fraga, Shah and Juenke 2020) emerge more generally, our analysis goes one step further in trying to parse out the general election *effect* of nominating candidates of color after close primary elections. In sum, we examine the local average treatment effect of a *barely*-winning (in the primary) white candidate or candidate of color on parties’ general election outcomes using real-world data.

We focus on general elections where the top two vote-getters in the preceding primary election were of different racial/ethnic backgrounds. That is, the candidate who received the party’s nomination via a primary election was of a different race/ethnicity from the candidate who came in second place in the primary. The forcing variable in this design is the “minority

¹These races are drawn from a broader dataset developed by the Candidate Characteristics Collaborative (C3), a group of scholars who cooperatively coded the race, ethnicity, and gender of state legislative candidates in 2018 and 2020 (Fraga, Juenke and Shah 2021). The dataset provides racial and ethnic information on thousands of candidates for office.

candidate primary win margin”: what share did the minority candidate win (or lose) by? As an example, if there were two candidates and the candidate of color won 55% of the vote and the white candidate 45%, this variable would take the value 0.1. In practice, most instances where there is a difference in the race/ethnicity of the top two vote-getters occur when one of the candidates is [non-Hispanic] White. This aligns with previous literature and our substantive interest in examining the impact of nominating a candidate of color relative to a white candidate.²

The dependent variable is the two-party general election vote share for the party that had the contested interracial primary. Thus, if the Democrats were the party with the contested interracial primary, our dependent variable is the two-party Democratic vote share in the subsequent general election. In the case of a contested interracial Republican primary, the dependent variable is instead the inverse (i.e., the Republican share of the two-party vote). For uncontested general elections, the general election two-party vote share is not meaningful and thus these observations are excluded from the analysis (though see the Online Appendix for analyses that impute vote counts for these races).


3 Results

We begin by graphically displaying the comparison being made in the RD design. Figure 1 presents a binned scatterplot of the data described above.³ Observations to the left of the vertical line marking the cutoff are primary elections where the minority primary win margin is negative: that is, the minority candidate lost the primary election and did not proceed to the general. To the right of the cutoff line are cases in which the minority candidate won

²See the Online Appendix for additional details about our election data, including discussion of the partisan breakdown of matchups, how we dealt with multimember and nonpartisan primaries, and the exclusion of uncontested elections.

³Figure and all RD analyses produced with **R** package `rdrobust` (Calonico et al. 2021).

Figure 1: Illustration of the RDD setup, showing full range of the dataset



the primary and became their party's nominee in the election. The y-axis displays the mean values of our outcome measure, the party's general election vote share, for each bin. The red line shows a polynomial function fit to the data on each side of the cutpoint. The light gray dots plot each individual observation (each election) in the dataset.

The apparent discontinuity at the cutpoint indicates that when a party narrowly nominates a non-white candidate for the general election, instead of a white candidate, the party wins a slightly *larger* voteshare in the general election that follows (not a smaller share as racial-penalty accounts would indicate). The intuition of the RD design is that observations proximate to the cutpoint should be very similar on all other observable and unobservable dimensions, and indeed, district-level characteristics are smooth across the cutpoint with no apparent discontinuities. The main difference is that the party's voters either do or do not nominate the white candidate, and this graphical analysis suggests that parties choosing a minority nominee in the primary may fare *better*, not worse, in the general election.

Table 1 presents estimates from several regression discontinuity specifications.⁴ The first column shows a simple local linear approach that fits lines to the data on either side of the cutpoint within an automatically-selected bandwidth. The second adds district-level covariates,⁵ which are not needed for identification but may increase precision of the estimates. The third column uses a third-order local polynomial regression to model the data on either side of the cutpoint, again using an automatically-selected bandwidth. These estimates vary slightly in size, but yield similar overall conclusions: a slight positive effect, and thus if anything an electoral *advantage* to parties that nominate the minority candidate in a cross-racial primary. While the positive effect is not significantly different from 0, we can rule out substantial electoral penalties: the 95% confidence interval for the estimates in column 1, for example, exclude minority candidate electoral penalties of one percentage point or larger.

Table 1: RDD Estimates of the Effect of Nominating a Minority Candidate on General-Election Voteshare

Minority nominee	0.07 (-0.01,0.16)	0.01 (-0.02,0.06)	0.03 (-0.01,0.08)
Polynomial	p=1	p=1	p=3
Covariates		X	X
Effective Sample Size	207	129	264
Bandwidth	0.27	0.16	0.35

These estimates cut against the conventional wisdom that parties will pay an electoral penalty should they nominate a minority candidate, at least in districts where there are closely-fought cross-racial primaries. We examine the robustness of these findings in the Online Appendix. First, in Section A3, we explore many plausible RD specifications, varying the observations included, bandwidth-selection approaches, kernels, polynomial order, and bias-correction decisions. The vast majority of these specifications yield positive estimates

⁴All estimates report bias-corrected estimates with robust standard errors (Calonico, Cattaneo and Farrell 2020; Calonico et al. 2021).

⁵Party, district population, an indicator for open-seat races, and the primary party’s voteshare in the district in the 2016 presidential election.

like those shown in Figure 1, and never document a significant negative effect. All of our evidence points against an electoral penalty when parties nominate minority candidates.

Second, we address questions about whether the sample of elections used here is dominated by majority-minority or heavily Democratic districts where an electoral penalty is especially unlikely. Section A1 of the Online Appendix demonstrates that the districts in our sample, though more diverse and slightly more Democratic than average legislative districts, are not extreme cases. Fewer than one quarter of the districts in the sample are “majority-minority” districts, and most are competitive between the parties. An analysis limited to Republican primaries does not provide any stronger evidence of an electoral penalty.

Last, we note that our estimates do not measure a causal effect of candidate race itself, but of the selection of a candidate of a given race (Hall 2015; Marshall 2021). The RD design does not change any individual candidate’s identity; it compares otherwise-similar elections that yielded nominees of different races. This means that candidate characteristics *other than race* could also vary across the cutpoint, if these characteristics are *correlated* with race. We view this possibility not as a threat to inference, but as an important part of the interpretation of these estimates: the value of using real-world data is that we can see how actual candidates fare in actual elections, so if some characteristics tend to vary with race we view that as an important way to draw out the real-world implications of our study. Accordingly, we briefly examine discontinuities in other candidate characteristics in Appendix Figure A5. We show that when parties narrowly nominate a candidate of color, they are choosing a whole bundle of characteristics: those narrowly-winning minority candidates are also more likely to be newcomers running for a seat, and are slightly less likely to be women. In the Online Appendix we also discuss how these correlations reveal potential mechanisms by which candidate race could affect election outcomes. Further, we consider threats to the interpretation of these estimates as a causal effect of nominating candidates of color, such as the possibility of strategic discrimination by primary voters leading to “compensating differentials” in candidate quality (Marshall 2021).

4 Implications

Our unique data and design allow for a novel test of the “racial penalty” story, and our results find no support for this perspective in recent state legislative elections. In this set of real-world elections where parties narrowly nominate candidates of color, parties do not face electoral penalties in the general election. These findings are more in line with recent work on the partisan effects and implications of racial polarization, suggesting that an out-partisan racial penalty may be overwhelmed by a larger boost from co-partisans (Agadjanian et al. 2020; Stout 2020; Tesler and Sears 2010).

Questions remain about the specific processes of candidate emergence that yield the elections studied here. We focus on a set of primaries where multiple candidates of various racethnic backgrounds emerged; not all districts currently see this sort of intra-party competition. Our results describe what happens in actual multiracial elections right now, but it would be worth revisiting these results if there were substantial shifts in party elites’ candidate recruitment practices or other election dynamics.

These results challenge the idea that parties consistently face an electoral penalty when they nominate racial and ethnic minority candidates (Doherty, Dowling and Miller 2019; Fraga and Hassell 2021; Stephens-Dougan 2021). Such a perspective may have kept party elites from promoting and supporting candidates of color, thus reinforcing the bigotry that undermines the fortunes of ambitious and talented candidates of color. While racial bias is a fact of life in the American electorate, polarization has shifted the landscape of descriptive representation in the United States. We show that racial and ethnic minority candidates are not just viable but could potentially boost the electoral prospects of a party that chooses to nominate them.

References

- Agadjanian, Alexander, John M. Carey, Yusaku Horiuchi and Timothy J. Ryan. 2020. “Disfavor or Favor? Assessing the Valence of White Americans’ Racial Attitudes.” *Assessing the Valence of White Americans’ Racial Attitudes* (September 28, 2020) .
- Andersen, David J. and Jane Junn. 2010. “Deracializing Obama: White voters and the 2004 Illinois US senate race.” *American Politics Research* 38(3):443–470.
- Branton, Regina P. 2009. “The importance of race and ethnicity in congressional primary elections.” *Political Research Quarterly* 62(3):459–473.
- Broockman, David E. 2014. “Do female politicians empower women to vote or run for office? A regression discontinuity approach.” *Electoral Studies* 34:190–204.
- Bucchianeri, Peter. 2018. “Is running enough? Reconsidering the conventional wisdom about women candidates.” *Political Behavior* 40(2):435–466.
- Calonico, Sebastian, Matias D Cattaneo and Max H Farrell. 2020. “Optimal bandwidth choice for robust bias-corrected inference in regression discontinuity designs.” *The Econometrics Journal* 23(2):192–210.
- Calonico, Sebastian, Matias D. Cattaneo, Max H. Farrell and Rocio Titiunik. 2021. *rdrobust: Robust Data-Driven Statistical Inference in Regression-Discontinuity Designs*. R package version 1.0.4.
- Canon, David T. and Richard A. Posner. 1999. *Race, redistricting, and representation: The unintended consequences of black majority districts*. University of Chicago Press.
- Chudy, Jennifer. 2021. “Racial sympathy and its political consequences.” *The Journal of Politics* 83(1):122–136.
- Chudy, Jennifer, Spencer Piston and Joshua Shipper. 2019. “Guilt by association: White collective guilt in American politics.” *The Journal of Politics* 81(3):968–981.
- Crowder-Meyer, Melody, Shana Kushner Gadarian, Jessica Trounstine and Kau Vue. 2020. “A different kind of disadvantage: Candidate race, cognitive complexity, and voter choice.” *Political Behavior* 42(2):509–530.

- 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.
- Ditonto, Tessa. 2020. “The mediating role of information search in the relationship between prejudice and voting behavior.” *Political Psychology* 41(1):71–88.
- Doherty, David, Conor M. Dowling and Michael G. Miller. 2019. “Do Local Party Chairs Think Women and Minority Candidates Can Win? Evidence from a Conjoint Experiment.” *Journal of Politics* 84(4):1282–1297.
- Fraga, Bernard, Eric Gonzalez Juenke and Paru Shah. 2021. “Candidate Characteristics Cooperative (C3) 2018 Data.”
URL: <https://doi.org/10.7910/DVN/VHAPHV>
- Fraga, Bernard and Hans Hassell. 2021. “Are minority and women candidates penalized by party politics? Race, gender, and access to party support.” *Political Research Quarterly* 74(3):540–555.
- Fraga, Bernard, Paru Shah and Eric Gonzalez Juenke. 2020. “Did women and candidates of color lead or ride the democratic wave in 2018?” *PS: Political Science & Politics* 53(3):435–439.
- Hall, Andrew B. 2015. “What happens when extremists win primaries?” *American Political Science Review* 109(1):18–42.
- Jackman, Simon and Lynn Vavreck. 2010. Obama’s Advantage: Race, Partisanship, and Racial Attitudes in Context. In *Annual Meeting of the Midwest Political Science Ass’n*.
- Juenke, Eric Gonzalez. 2014. “Ignorance is bias: The effect of Latino Losers on models of Latino representation.” *American Journal of Political Science* 58(3):593–603.
- Juenke, Eric Gonzalez and Paru Shah. 2016. “Demand and supply: Racial and ethnic minority candidates in White districts.” *Journal of Race, Ethnicity, and Politics* 1(1):60–90.
- Kam, Cindy D. 2007. “Implicit attitudes, explicit choices: When subliminal priming predicts candidate preference.” *Political Behavior* 29(3):343–367.

- Kinder, Donald R and Allison Dale-Riddle. 2012. *The End of Race?: Obama, 2008, and Racial Politics in America*. Yale University Press.
- Lewis-Beck, Michael S., Charles Tien and Richard Nadeau. 2010. "Obama's Missed Landslide: A Racial Cost?" *PS: Political Science & Politics* 43(1):69–76.
- Lublin, David. 1999. *The paradox of representation*. Princeton University Press.
- Marshall, John. 2021. "Can Close Election Regression Discontinuity Designs Identify Effects of Winning Politician Characteristics? Working Paper."
- Sigelman, Carol K., Lee Sigelman, Barbara J. Walkosz and Michael Nitz. 1995. "Black candidates, white voters: Understanding racial bias in political perceptions." *American Journal of Political Science* pp. 243–265.
- Stephens-Dougan, LaFleur. 2021. "The Persistence of Racial Cues and Appeals in American Elections." *Annual Review of Political Science* 24:301–320.
- Stout, Christopher T. 2020. *The Case for Identity Politics: Polarization, Demographic Change, and Racial Appeals*. University of Virginia Press.
- Tesler, Michael. 2013. "The return of old-fashioned racism to White Americans' partisan preferences in the early Obama era." *The Journal of Politics* 75(1):110–123.
- Tesler, Michael and David O. Sears. 2010. "Obama's Race: The 2008 Election and the Dream of a Post Racial America."
- Valentino, Nicholas, Vincent Hutchings and Ismail White. 2002. "Cues that matter: How political ads prime racial attitudes during campaigns." *American Political Science Review* 96(1):75–90.
- Voss, D. Stephen and David Lublin. 2001. "Black incumbents, white districts: An appraisal of the 1996 congressional elections." *American Politics Research* 29(2):141–182.

Online Appendix for: Evaluating the Minority Candidate Penalty with a Regression Discontinuity Approach

A1 Description of Elections Used in the Analysis

We use data on state legislative candidates and outcomes in our regression discontinuity analysis. As noted in the main text, we begin with information on primary election outcomes where the race of the top vote winners differed, where the winning margin serves as our RD's running variable. In systems that hold runoff elections to choose between the top two candidates in a multi-candidate primary, we use the runoff results to calculate the winning margin, rather than the first round. Some of the 36 states we examine have multi-member state legislative districts or use primary systems where multiple candidates advance to the general election from a single primary ballot. We exclude systems with non-partisan "top two" primary systems that advance the two highest vote-getters regardless of party, as they do not produce party nominees. In the case of multi-member districts, we focus on the "last in" and "first out" candidates when calculating the primary win margin, comparing the candidate who most narrowly won nomination to the person who most narrowly lost it.

Table A1 reports the total number of multiracial primary elections we use in the analysis, excluding uncontested general elections as well as races omitted due to missingness from other sources¹, by year and party; there are more Democratic primaries than Republican ones, but both parties are well-represented across both years of the dataset. Furthermore, in each year the breakdown is about three-quarters lower-chamber elections (State House) and one-quarter upper (State Senate).

¹These are mainly cases where the running variable, "minority candidate primary win margin (over white candidate)," is missing because none of the top candidates in a contested multiracial primary are white.

Table A1: Election counts

	Democratic	Republican
2018	141	52
2020	152	46

We note that the districts included in the RD sample are not exclusively “majority-minority districts” (fewer than one-quarter have a white population under 50%), nor are they overwhelmingly safe Democratic districts. Figure A1 explores the partisan composition of the districts in our sample by plotting the density of Democratic two-party presidential voteshare in 2016 for the state legislative districts in the RDD sample, and for the rest of the US. Presidential voteshare data used here come from MIT’s Election and Data Science Lab (they rely on aggregating precinct-level voteshare data to the level of the state legislative district) and have some missingness, such that about 9 in 10 of the districts in our sample are successfully matched to district presidential-vote data. Figure A1 indicates that districts in our sample are slightly more Democratic than average, but are generally not “safe” Democratic districts in the sense of having very high Democratic vote shares. Rather, the most common values for both our sample of districts and for the rest of the US fall within the competitive zone of 40-60% Democratic voteshare.

In addition to the distribution of districts included, some readers may wonder whether the effects differ across parties: could it be that Democrats (the majority of the sample) do not face an electoral penalty when nominating a minority candidate, but that Republicans do? We note the limitations of heterogeneity analysis; given the size of the sample used for the main study, any subgroup analysis is going to be somewhat underpowered. But we present a version of the main table here that restricts the sample to the 120 Republican primaries included in the main dataset. The estimates included in Table A2 are noisy and variable, but they do not suggest that Republicans face a substantial electoral penalty.

Figure A1: Comparing RD sample districts to the rest of the US

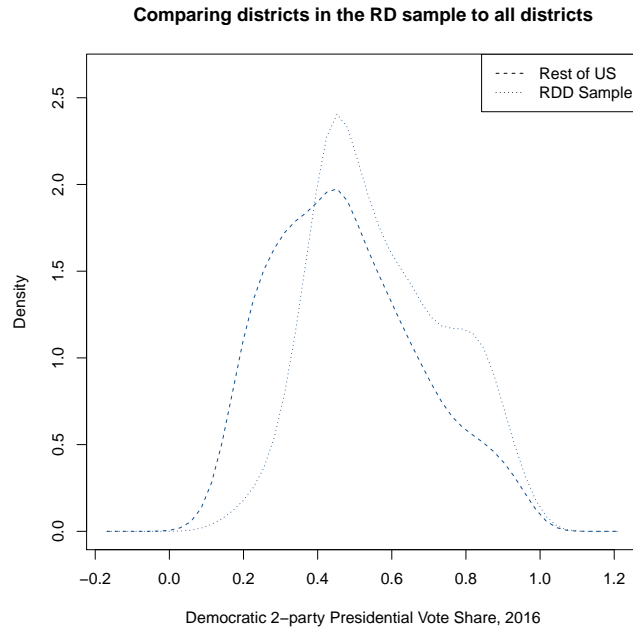


Table A2: RDD Estimates of the Effect of Nominating a Minority Candidate on General-Election Voteshare, Republican primaries only

Minority nominee	0.23 (0.1)	0 (0.04)	0 (0.05)
Polynomial	p=1	p=1	p=3
Covariates		X	X

A2 Alternative Treatment of Uncontested Elections

Our main specifications omit uncontested elections, considering an election as uncontested if at least one of the two major parties does not have a candidate on the general election ballot. These uncontested general elections are not informative about the question asked in the paper. In most cases, the opposing party had no nominee for the general election even *before* the primaries took place, so there is no logical way that one party’s nomination of a white candidate, for example, could change the presence of the other party’s candidate on the ballot.

Nevertheless, in this section we explore an alternative approach to uncontested elections, setting the two-party vote share to 100% for winning parties (rather than missing) and to

0% for non-contesting parties in cases of uncontested general elections. This approach adds additional observations to the sample, but also introduces noise (as we are simply imputing vote counts in elections that did not occur). Table A3 reproduces Table 1 from the main paper using this alternative outcome measure. The specifications vary in size and direction, but still do not clearly point to a substantial electoral penalty.

Table A3: RDD Estimates of the Effect of Nominating a Minority Candidate on General-Election Voteshare, with alternative voteshare measure

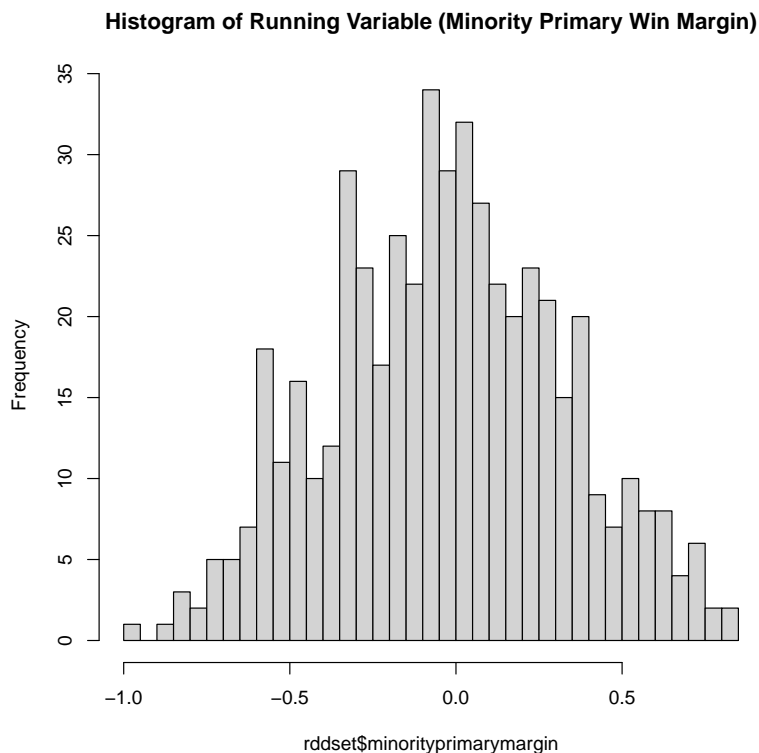
Minority nominee	0.05 (0.08)	-0.02 (0.04)	-0.03 (0.05)
Polynomial	p=1	p=1	p=3
Covariates		X	X

A3 RDD Validity

In this section, we test the validity of the regression-discontinuity assumptions and explore robustness to alternate specifications.

We begin by looking for evidence of sorting around the cutpoint. If units were able to select which side of the cutpoint they landed on, the RD setup would not be valid. We present a simple histogram of the running variable (Figure A2), which doesn't show any notable lopsidedness around the cutpoint. We also ran a test for continuity of the density functions for control and treatment units around the cutoff using the `rddensity()` package in R (Cattaneo, Jansson and Ma 2021), which failed to reject the null hypothesis of no manipulation at the cutoff.

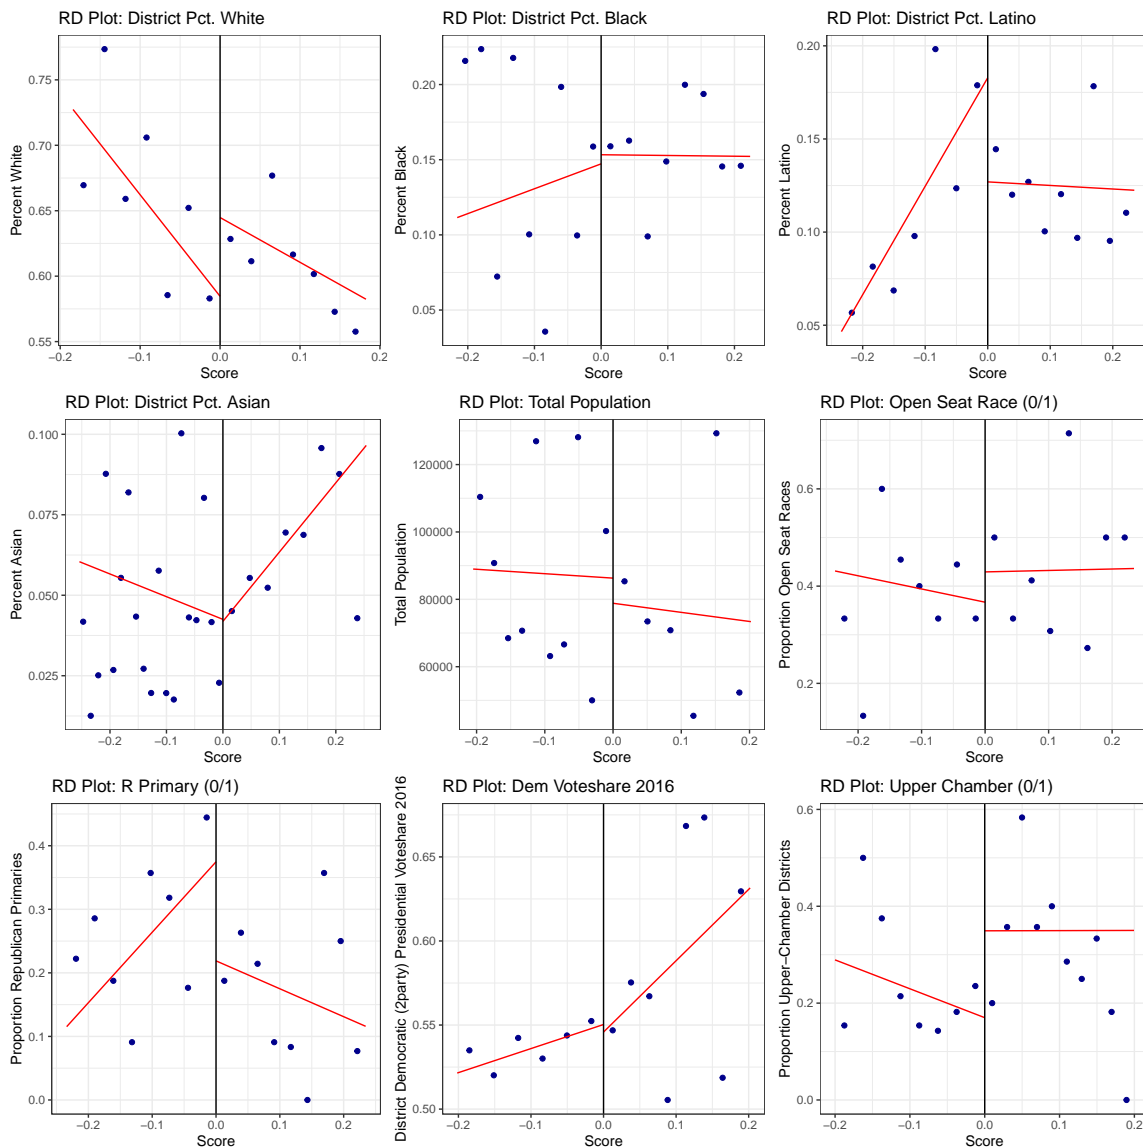
Figure A2: Looking for sorting around the cutpoint



We next turn to background covariates, which ought to look smooth across the cutpoint: districts where minority candidates narrowly win the primary should not look systematically

different from districts where they lose. To illustrate this smoothness, we graphically display the equivalent of our main RDD specification (local linear with “mserd” automated bandwidth selection), except that in this case we use a variety of pre-election covariates as the outcome variable. These can be thought of as a sort of placebo test: a significant “treatment effect” on these pre-election covariates would indicate a problem with the RD setup. We observe no such pattern in Figure A3: discontinuities at the cutpoint are relatively small, and none are significant (p-values range from .21 to .96).

Figure A3: Examining smoothness of background covariates across the cutpoint

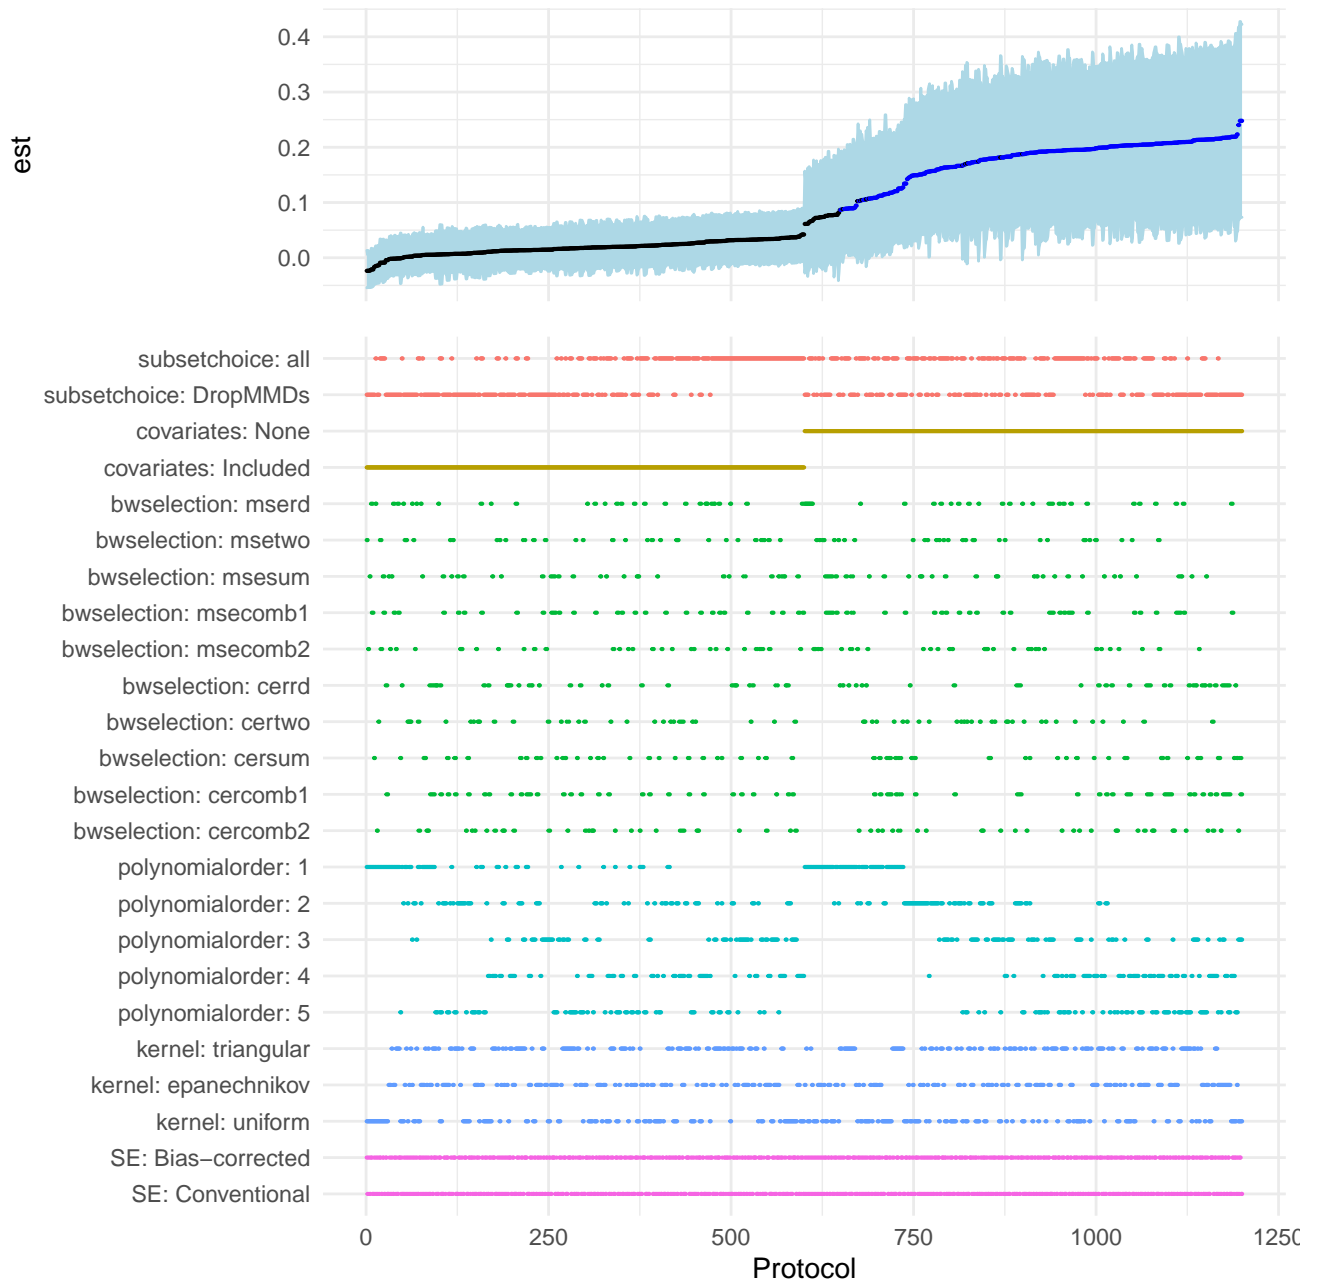


Finally, we present a plot that explores many plausible specifications for the main RD estimates presented in the paper. In estimating the electoral penalty/benefit to a party that nominates a minority candidate, our main approach in the paper uses all “factory default” settings from the `rdrubust()` package: a local linear specification with automated bandwidth selection using the “`mserd`” option and bias-corrected estimates with robust standard errors. In this plot, we begin to vary some of the data-management and analytic decisions made in those main estimates in order to explore how much those choices mattered for the estimates presented. The plot presents RD estimates from 1,200 specifications that vary these decisions. The top part of the plot displays the distribution of point estimates from those models along with 95% confidence intervals. The bottom section of the plot indicates which specification choices are associated with which estimates.

The choices vary as follows. First, the plot varies the data included (“`subsetchoice`”) to rely either on the full analysis dataset from the main paper, or to omit the small number of multi-member districts that were included in our dataset and focus on single-member districts. We include specifications that do not include covariates, as well as specifications that incorporate available covariates (total district population, primary party, whether an election is an open-seat race, party’s voteshare in the district in the 2016 presidential election). We also include specifications that rely on all available automated bandwidth-selection (“`bwselection`”) approaches in the `rdrubust()` package, as well as varying the order of the polynomial used to construct the point estimates (from 1, local linear, to 5). We include specifications using all available kernel choices from the `rdrubust()` package. We also vary the types of estimates presented: our main table presents estimates with robust bias-corrected standard errors (Calonico, Cattaneo and Farrell 2020) but in this plot we also include estimates with conventional standard errors. Each set of choices in the plot (“`outcomechoice`,” “`subsetchoice`,” etc.) interacts with all possible values for all other choices, for a total of 1,200 specifications.

The estimates across these specifications are fairly consistent with the paper’s main esti-

Figure A4: Exploring many plausible specifications for the main RDD



mate. Of the 1,200 specifications, only a few dozen yield negative point estimates (all non-significant), with the largest negative estimate showing less than three percentage points of disadvantage for minority candidates. The estimates range from $-.024$ to $.25$ in size (-2 to 25 percentage points' general-election advantage for minority nominees), and they vary in

precision. Most of these estimates rule out even small negative effects: over half of them have 95% confidence intervals that exclude estimates of one percentage point.

It does not appear that bandwidth-selection methods, kernel choice, or bias-correction decisions make much difference for the estimates (these choices are spread across the range of estimates). In general, point estimates look smaller and more precisely-estimated when including covariates, which makes sense given the limited sample size available for this RD analysis.

Inspired by Stommes, Aronow and Sävje (2021), we also consider the statistical power of this design. Our main analysis presents null findings, rejecting claims of an “electoral penalty” for parties that nominate minority candidates. These findings are only compelling to the extent this analysis is powered to detect evidence of such a penalty should it exist. We use the `rdpower()` package to perform ex-ante power calculations for the design presented in the main paper. A local-linear specification including background covariates (as presented in column 2 of the main paper table) has 80% power to detect effects of size .0555 or greater. Such effects (a 5.6-percentage point electoral penalty for nominating a minority candidate) correspond to a cohen’s d of .36, a small-to-medium effect size. Being able to rule out effects of this size is substantively meaningful, especially as some experimental work on electoral penalties for nonwhite candidates finds much larger effect sizes. In one recent study, a nationally-representative sample showed about a fifteen-percentage point electoral penalty for nonwhite candidates (in the “face-saving” treatment condition, which the paper suggests is most similar to real-world election conditions) (Krupnikov, Piston and Bauer 2016). In another, respondents in a low-information condition showed huge racial penalties across groups: a 22 percentage point disadvantage for Black candidates and 10 points for Latino or Asian candidates (all compared to white candidates), though these penalties shrank somewhat with the provision of party information (Crowder-Meyer, Gadarian and Trounstein 2020)(Figure 2).

A4 Relevance of Other Candidate Characteristics


Below we examine how other observable candidate characteristics may also change with candidate race/ethnicity. If we saw, for example, that characteristics like incumbency tended to vary along with winners' race, that could provide an important understanding of the mechanisms by which nominating minority candidates was providing parties with an advantage in the elections we study.

Accordingly, we briefly examine the other candidate characteristics available to us in the dataset. Figure A5 presents two plots similar to our main graphical presentation in Figure 1. In this case, rather than focusing on our main outcome of interest, we instead ask whether nominating a white candidate yields different nominee characteristics on gender or incumbent status. The plot on the left side examines incumbency, and finds a positive discontinuity at the cutpoint: when parties narrowly nominate a white candidate, they are more likely to get a nominee who is also the incumbent for the legislative seat in question. The plot on the right side of the figure examines candidate gender: parties nominating a white candidate for the general election appear slightly more likely to have nominated a woman.² So it appears that when parties narrowly nominate a candidate of color, they are choosing a whole bundle of characteristics: those narrowly-winning minority candidates are also more likely to be newcomers running for a seat, and they may also be less likely to be women.

The correlation between gender or incumbency and the close nomination of a White candidate reveals potential mechanisms by which candidate race matters for election outcomes. Understanding a person's race to be a "bundled treatment" (Sen and Wasow 2016) from the perspective of estimating causal effects, the fact that barely-nominated white candidates are more likely to be incumbents *who performed worse than a barely-nominated candidate of color* indicates that minority candidates may be more likely to mount a formidable challenge

²Running the RDD specification with "party's nominee is a woman" as the outcome measure yields negative estimates, but they are statistically indistinguishable from zero.

Figure A5: Examining other candidate characteristics across the cutpoint



to white incumbents who are strong on paper, but underperform in the general election campaign. Or, a potential mechanism could be a divisive primary campaign that depresses general election turnout among minority voters.

Another possibility raised by Marshall (2021) is that of “compensating differentials” or of candidates winning the nomination in different ways depending on race. For example, it could be the case that primary voters concerned about candidate electability would strategically discriminate against candidates of color because they feared that these candidates would lose in the general election if nominated (Bateson 2020; Green, Schaffner and Luks N.d.) If this kind of strategic discrimination took place, it could mean that barely-nominated minority candidates would be of higher quality than barely-nominated white candidates (since they had passed through this discriminatory filtering process), and thus could be expected to do better in general elections simply because of the quality differential introduced by our RD’s focus on narrowly-contested elections. Such a pattern would yield statistically-valid estimates, but would only teach us about candidate outcomes in districts with this specific pattern of strategic primary discrimination. If candidate recruitment or selection patterns

changed, so too could the general-election vote patterns observed here.

We cannot rule out all possible patterns of compensating differentials because there is limited pre-primary data available on state legislative candidates. However, we do not think such a pattern of “Jackie Robinson effects” is likely in these sorts of low-information elections. For one thing, we note the pattern seen in Figure A7: narrowly-nominated minority candidates are substantially *less* likely to have already been in state legislative office, which is the opposite of what we would expect if there were a quality differential being introduced by the research design.

Second, additional analyses are inconsistent with stories of strategic discrimination by primary voters. One prominent story about compensating differentials imagines that primary voters are strategic and accurately predict that candidates of color will face bias in the general election. Thus, they hold candidates of color to a higher standard in the primary election, ensuring that any candidates of color that make it to the general election will be more qualified than white nominees, and thus counterbalancing the racial bias of general-election voters. This form of forward-looking strategic discrimination could yield patterns equivalent to the ones shown in our main analyses even if nominees of color did face an electoral penalty in the general. However, we note that this kind of strategic discrimination should only be likely to happen in certain types of elections: specifically, primary voters should be especially likely to engage in strategic discrimination in contexts where the general election is likely to be closely contested. In districts that are “safe” for either party and where a given party’s nominee is *ex ante* either very likely or very unlikely to win the general election, there is little reason to engage in strategic discrimination as it is unlikely to make a difference. Thus, we re-run our main RD analysis limiting to primaries in safe partisan districts (defined as those where one party’s presidential nominee won more than 60% of the vote in the 2016 election, though a cutoff of 65% or 70% yields equivalent conclusions). Table A4 reproduces the analysis from Table 1 of the main paper within this subset, and finds similar general-election vote patterns (though noisier given the smaller sample). If strategic discrimination

were driving the main estimates, we should expect to see negative point estimates in Table A4: without strategic primary discrimination “masking” the general election bias, we should be able to see any electoral penalty faced by nominees of color. We do not see such a pattern even in the subset of elections where strategic primary discrimination should be least likely.

Table A4: RDD Estimates of the Effect of Nominating a Minority Candidate on General-Election Voteshare, Safe-seat races only

Minority nominee	0.16 (-0.01,0.36)	0.19 (0.03,0.41)	0.37 (0.02,0.71)
Polynomial	p=1	p=1	p=3
Covariates		X	X
Effective Sample Size	88	84	99
Bandwidth	0.23	0.21	0.26

References

- Bateson, Regina. 2020. “Strategic discrimination.” *Perspectives on Politics* 18(4):1068–1087.
- Calonico, Sebastian, Matias D Cattaneo and Max H Farrell. 2020. “Optimal bandwidth choice for robust bias-corrected inference in regression discontinuity designs.” *The Econometrics Journal* 23(2):192–210.
- Cattaneo, Matias D., Michael Jansson and Xinwei Ma. 2021. *rddensity: Manipulation Testing Based on Density Discontinuity*. R package version 2.2.
URL: <https://CRAN.R-project.org/package=rddensity>
- Crowder-Meyer, Melody, Shana Kushner Gadarian and Jessica Trounstine. 2020. “Voting can be hard, information helps.” *Urban Affairs Review* 56(1):124–153.
- Green, Jon, Brian Schaffner and Sam Luks. N.d. “Strategic Discrimination in the 2020 Democratic Primary.” . Forthcoming.
- Krupnikov, Yanna, Spencer Piston and Nichole Bauer. 2016. “Saving face: Identifying voter responses to black candidates and female candidates.” *Political Psychology* 37(2):253–273.
- Marshall, John. 2021. “Can Close Election Regression Discontinuity Designs Identify Effects of Winning Politician Characteristics? Working Paper.”.
- Sen, Maya and Omar Wasow. 2016. “Race as a ‘Bundle of Sticks’: Designs that Estimate Effects of Seemingly Immutable Characteristics.” *Annual Review of Political Science* 19.
- Stommes, Drew, PM Aronow and Fredrik Sävje. 2021. “On the reliability of published findings using the regression discontinuity design in political science.” *arXiv preprint arXiv:2109.14526* .

