

ARTICLE

Polarization and state legislative elections

Cassandra Handan-Nader¹  | Andrew C. W. Myers²  | Andrew B. Hall^{3,4}

¹Assistant Professor of Political Science,
Department of Political Science, New York
University, New York, New York, USA

²Ph.D. Candidate, Department of Political
Science, Stanford University, Stanford,
California, USA

³Davies Family Professor of Political
Economy, Graduate School of Business,
Stanford University, Stanford, California,
USA

⁴Senior Fellow, Hoover Institution, Stanford,
California, USA

Correspondence

Andrew C. W. Myers, Department of Political
Science, Stanford University, 616 Jane
Stanford Way, Stanford, CA 94305, USA.
Email: myersa@stanford.edu

Abstract

US state legislatures are critical policymaking institutions that are increasingly polarized, yet data and measurement limitations have prevented researchers from understanding how state legislative elections contribute to this polarization. To address this gap, we construct new measures of candidate ideology based on campaign contributions and roll-call votes, and we use them to offer the first systematic study of the relationship between candidate ideology and electoral outcomes in primary and general elections in state legislatures, 2000–2022. We find that the set of people running for state legislature has polarized substantially in recent decades. More-moderate candidates enjoy a meaningful advantage in contested general elections, but that advantage has declined somewhat in recent years. At the same time, more-extreme candidates are favored in contested primary elections. These new measures and data will allow researchers to build on these basic findings to understand how elections function in lower information, lower salience environments like American state legislatures.

To what extent do primary and general elections for state legislatures advantage more-moderate or more-extreme candidates? While research on American candidate ideology and electoral outcomes focuses mainly on the national level (e.g., Ansolabehere et al., 2001; Canes-Wrone & Kistner, 2022; Canes-Wrone et al., 2002; Hall, 2015), state legislatures provide an ideal laboratory for understanding how the long-studied relationship between candidate ideology and electoral outcomes might extend to lower information, lower salience settings. Such settings are common across the democratic world, and complicate the assumptions behind spatial models of voting that underpin our ideas about the potential advantages of more-moderate candidates (e.g., Downs, 1957). State legislatures are themselves highly consequential and increasingly polarized policymaking bodies (Caughey & Warshaw, 2022; Rogers, 2023; Shor & McCarty, 2011), responsible for disbursing nearly two trillion dollars

in spending and with authority over many salient policy areas including education, healthcare, and election administration.¹ In addition to being critical components of local government in their own right, state legislatures are also the main source of future House and Senate candidates (e.g., Thomsen, 2014), and could therefore be helping to drive national polarization.²

Despite the value of these simple empirical questions about state legislative elections, they have been impossible to answer comprehensively because we lack a measure of candidate ideology for state legislatures that both (a) corresponds closely to legislative polarization as measured by roll-call votes in state

¹ <https://www.urban.org/policy-centers/cross-center-initiatives/state-and-local-finance-initiative/state-and-local-backgrounders/state-and-local-expenditures>.

² For arguments for why a more-extreme state legislative candidate pool could cause an important part of the rising polarization of Congress, see Hall (2019) and Thomsen (2017).

Verification Materials: The materials required to verify the computational reproducibility of the results, procedures, and analyses in this article are available on the *American Journal of Political Science* Dataverse within the Harvard Dataverse Network, at: <https://doi.org/10.7910/DVN/0XQNDNR>.

The Cornell Center for Social Sciences verified that the data and replication code submitted to the AJPS Dataverse replicates the numerical results reported in the main text of this article.

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2025 The Author(s). *American Journal of Political Science* published by Wiley Periodicals LLC on behalf of Midwest Political Science Association.

legislatures and (b) applies to incumbents and non-incumbents alike. Existing measures available for state legislators either capture state legislative voting behavior but apply only to incumbents (Shor & McCarty, 2011), or are meant to represent candidate positions in an ideological space defined primarily by candidates for federal office and subsequent voting behavior in Congress (Bonica, 2014, 2018). As a result, previous work specific to state legislative polarization has largely been limited to studying incumbents. Two important recent books show that more-moderate incumbent state legislators modestly outperform more-extreme incumbents in general elections (Caughey & Warshaw, 2022; Rogers, 2023).³ But the relationship between ideological positions and electoral performance may be different for challengers and candidates running for open seats who do not have the other advantages that incumbents possess. Moreover, to assess whether the advantages of more-moderate or more-extreme candidates have changed over time, we need to make apples-to-apples comparisons of candidates across years, which is not possible if we have to condition our sample only on the incumbents who have survived in each cycle to face another reelection.

The primary contribution of this paper is to offer new data and measures that overcome these obstacles, allowing researchers to study candidate ideology, polarization, and elections in state legislatures in ways not previously possible. We develop two new measures of state legislative candidate roll-call ideology that apply to incumbents and nonincumbents alike. As a baseline model, we use the approach developed in Hall and Snyder (2015) that imputes candidate NP-Scores—a widely used measure of state legislator's roll-call-based ideology from Shor and McCarty (2011)—using the weighted averages of the NP-Scores of incumbents to which that candidate's donors also donated. To improve predictive accuracy, we build off of the ideas in Bonica (2018),⁴ and use a machine learning (ML)-based approach that predicts NP-Scores using campaign donation records. Though they build off prior methodological approaches in the literature, these measures are developed specifically to study state legislative elections in three ways. First, by using a comprehensive source of state contribution records, our measures incorporate more donor information from state candidates' campaigns than measures that must rely on donors that donate to both federal and state campaigns to score state candidates (Bonica, 2014, 2018). Using these data, our models are able to take advantage of cross-state giving to improve predictive accuracy for smaller states. Second, we train

separate ML models for each party to improve the measure's ability to distinguish between candidates within partisan primaries. Third, we develop our models to achieve as much parity in accuracy as possible between winners and losers of elections by using only contributions received before a candidate first takes office, helping to allay concerns that campaign-finance-based scalings could partially be a function of having won office previously. Both of the resulting measures correlate highly with NP-Scores, even within party.

In collaboration with Fourinaies and Hall (2020) and Rogers (2023), we also construct a new data set on state legislative primary elections, collected and digitized from each state's official records, and extensively cleaned and standardized. We merge this with data on general elections from 2000 through 2022 and combine it with CFscores and our candidate ideology scores to form a data set containing the estimated ideological positions and primary- and general election performances of nearly 48,000 candidates for state legislative office. The resulting data set, including our new measures of candidate ideology, will be made publicly available so that researchers can use them freely for the study of state legislative elections.

With these new data, we first show that the polarization of the whole set of candidates seeking state legislative office has risen dramatically over the past two decades. The growing polarization of state legislators tracks the polarization of the set of candidates running for office quite tightly. We argue that who runs for state legislature may therefore be very important for understanding state legislative polarization, despite the focus of existing research on incumbent positioning, and may therefore be important for explaining polarization at the federal level, too. If the entire pipeline of candidates seeking state legislative office is polarizing, this will increase the polarization of congressional candidates, too.⁵

Next, using a panel design that compares over-time changes in the ideological midpoint between candidates within a given district, we show that contested general elections have favored more-moderate candidates, on average. This result is consistent with canonical spatial models despite the low levels of information in state legislative elections. However, this advantage is relatively modest in magnitude and appears to have declined noticeably in recent years. While there may be many explanations for this decline, it is at least consistent with the growing literature on the nationalization of state legislative elections (Abramowitz & Webster, 2016; Hopkins, 2018;

³ For earlier work on the relationship between incumbent ideology and electoral outcomes in state legislatures, see Birkhead (2015), Hogan (2008), and Rogers (2017).

⁴ A similar ML approach is also used in Bonica and Li (2021).

⁵ In fact, Phillips et al. (2024) show that the general trend of state legislators becoming more extreme is a more important part of the explanation for the polarization of the congressional candidate pool than changes caused by redistricting or electoral competitiveness.

Rogers, 2016, 2023).⁶ We bolster these findings with a regression discontinuity design on close primary races, where a more-extreme candidate in the general election is arguably as-if randomly assigned, and find a modest but meaningful vote-share penalty for the more-extreme candidates in the general election.

Next, we take advantage of our large data set to study important sources of variation in these effects. Spatial models of voting, along with prior empirical work, identify at least three important factors that might increase or decrease the advantage of more-moderate candidates: the degree of competition in the district, which we measure using partisan presidential vote share; the level of legislative professionalization, which makes elections more salient to voters and candidates; and the presence of other, more salient offices at the top of the ballot, which we proxy for with on-cycle versus off-cycle elections. We find that the advantage to more-moderate candidates is higher in more-competitive districts, yet surprisingly, we find a relatively precisely estimated non-difference between on-cycle and off-cycle elections, and only weak and very imprecise evidence that more professionalized legislatures feature larger advantages. These results suggest that the more important factor in boosting more-moderate candidates and reducing polarization is electoral competition—which has declined in recent decades as more districts have become especially partisan—rather than the timing of elections or the overall level of importance of the state legislature itself.

A related literature, again focused on the national level, posits a push and pull between primary and general elections, with primary elections favoring more-extreme candidates while general elections favor more-moderate candidates (e.g., Aranson & Ordeshook, 1972; Brady et al., 2007; Hall, 2015). In the final part of the paper, we speak to this literature by offering the first comprehensive estimates of the relationship between candidate ideology and electoral outcomes in contested primary elections. Consistent with this theoretical literature, despite the low levels of information in state legislatures, we find that these elections favor more-extreme candidates, on average, and this advantage has remained large in recent years. These findings build from the analyses presented in Rogers (2023), which shows that incumbents with more-extreme roll-call votes in state legislatures enjoy a modest advantage in primaries.

Taken together, our estimates paint a picture of a changing state legislative system in which more-extreme candidates are increasingly seeking office,

face limited competition, are favored in primary elections, and face relatively small and diminishing penalties in the general election. This pattern shows how state legislatures have polarized over the past two decades, and also helps to explain why the set of people running for Congress has also polarized so much over this same time period. There are many possible explanations for why the state legislative electoral system has evolved in this manner, including changes to the media environment, to the structure of districts and elections, to American political culture, and to the policy agenda facing state legislatures. In the conclusion to the paper, we discuss how our findings help to set up future research on these important questions.

CAMPAIGN-FINANCE-BASED MEASURES OF CANDIDATE ROLL-CALL IDEOLOGY

To assess the electoral roots of roll-call-based polarization, we first need a measure that closely captures how both winning and losing candidates would cast roll-call votes in state legislatures. However, no existing measures of ideology that extend to candidates for state office are optimized explicitly for capturing roll-call voting behavior in state legislatures. Alternative measures such as CF-Scores (Bonica, 2014) and DW-DIME scores (Bonica, 2018) map state candidates to an ideological space defined primarily by federal contributions, and therefore must rely on the subset of donors that donate to both state and federal campaigns to estimate the positions of state candidates. While useful for comparing state and federal candidates on the same scale, these measures only incorporate a small number of state candidates' campaign donors by design, and, when optimized specifically to predict roll-call voting in Congress, extend to a very limited number of state candidates (see Online Appendix A, Table A.1). Moreover, these measures incorporate donations from both before and after a candidate wins election, which makes candidates' estimated ideological positions dependent on their past electoral successes or failures. Hence, our goal in this section is to build new measures that are highly predictive of candidates' subsequent roll-call voting behavior in state legislatures within each party, and are tailored specifically to studying the relationship between ideological positions and winning elections.

Using campaign-finance records to predict NP-scores

We begin with the key target variable that we want to predict, the ideological mappings for state legislators from Shor and McCarty (2011, 2023), called NP-Scores. The most recent version provides scores for 27,629 state legislators between 1993 and 2020.⁷ These

⁶ On the other hand, despite these patterns of nationalization, meaningful amounts of split-ticket voting still occur in state legislative races (Kuriwaki, 2023), especially where information is higher (Moskowitz, 2021). It also remains a puzzle why there has been a more substantial advantage to more-moderate candidates in the early 2000s, when partisanship was still important and voter information was presumably still low in state legislative elections.

mappings are the result of projecting a measure of voting behavior from each state legislature onto a measure of legislator responses to the Project Vote Smart National Political Awareness Test (NPAT). The methodology uses a state-specific ordinary least squares (OLS) regression of a one-dimensional NPAT-based scaling onto a one-dimensional roll-call-based scaling using legislators who have both scores available. Therefore, voting disagreement between legislators within a state on the first roll-call dimension is amplified the more it correlates with the NPAT's first dimension.

Though NP-Scores are the most widely used measure of roll-call-based polarization at the state level, they also have some inherent limitations. First, though the NPAT is designed to measure nationally relevant ideological issue positions, voting behavior in legislatures can be driven by party control in addition to individual ideological leanings (Lee, 2009; Stiglitz & Weingast, 2010). As in other studies of roll-call-based polarization (e.g., Poole & Rosenthal, 2007), we use "roll-call ideology" as a shorthand for the mixture of ideological and partisan factors that might drive roll-call-based polarization. Second, legislators' positions on certain local issues may not align well with their positions on nationalized issues (Anzia, 2021), which could cause the measure to be insensitive to polarization on important local issues in some states. Polarization should therefore be interpreted relative to the national issue space as captured by the NPAT's first dimension, with the understanding that states less polarized on national issues could still be polarized on issues that are not relevant to the two-party axis that characterizes the NPAT's first dimension. Finally, NP-Scores are static by construction throughout a legislator's career, which means they cannot be used to analyze behavioral changes that individual members may make over the course of their careers.

Since these NP-Scores are only available for legislators who won election, we need another set of information to help us predict scores for people who have not served, and may not ever serve, in office. While there are many potential data sources one might use for this purpose—such as the text of candidate speeches or behavior on social media—we build on the supervised learning approach of Bonica (2018) to estimate scores for both incumbents and challengers running for state office using donations to their campaigns. We obtain campaign donations for state legislative candidates from the National Institute on Money in State Politics (NIMSP), which digitizes and standardizes information from campaign-finance reports for all state-office candidates.⁸ The data con-

sist of nearly 30 million transactions for state legislative elections between 1989 and 2022, with comprehensive election coverage beginning in the year 2000. We merge this donation information at the legislator-election-year level to the NP-Scores to obtain a unified set of predictors and outcomes for 20,757 state legislators in at least one election, 11,739 of which have enough donations before they won office to enter the training set. Though our research design only makes within-state comparisons (i.e., between two candidates in the same election), our approach to the data was designed to reduce the variance in predictive accuracy between states as much as possible. To impose a unified threshold on data sparsity, we only produce scores for candidates that received money from at least five donors who also gave to at least five legislators with an NP-score, and to improve predictive accuracy for states with less data, our models take advantage of donors that give to candidates in multiple states to pool information across states.⁹ With this approach, we are able to substantially increase the total number of donors that contribute to a candidate's score relative to the federal approach (median of 16 donors vs. median of 2 donors).¹⁰

Though our basic supervised learning approach follows Bonica (2018), we further customize the prediction problem based on the downstream empirical analyses we wish to conduct. First, because our analysis hinges on producing credible ideology estimates for candidates that do not win office, we develop our models only using donations received *before* candidates win state legislative office for the first time. This avoids biasing the predictive models with information following electoral victories when studying performance in elections (Hall & Snyder, 2015), and improves the generalizability of the models to losers of elections by mimicking the information set that donors have about candidates whom they have not yet observed in state legislative office. As a result, our predicted NP-Scores are static over a legislator's career, matching the construction of NP-Scores, and do not incorporate any donation information from elections in which candidates ran as incumbents. Second, in order to make the assumption that prediction error is equal in expectation between winners and losers more plausible, we assign all legislators in the training data an "out-of-sample" score based on a model that was not trained on their particular donations in any capacity. On top of the cross-validation procedures employed in Bonica (2018), the out-of-sample scores provide an additional

⁷ To obtain some coverage for 2021–2022 elections, we carry forward the NP-Score for incumbents with scores in the 1993–2020 data set.

⁸ See <https://www.followthemoney.org/our-data/about-our-data> (Follow The Money, 2024) for more information on the data. See Online Appendix C for details on our validation of this data set.

⁹ Twenty-two percent of the donors in our modeling data gave to candidates running in multiple states, and we found that pooling information across states reduces the mean squared prediction error by 38%. Online Appendix B reports the average proportion of out-of-state donors and contributions per candidate by state.

¹⁰ We show in Online Appendix C that this is not an artifact of differences in donor identity resolution between the DIME database (Bonica, 2016) and NIMSP.

guarantee that the predicted scores are not overfit to the training legislators. Third, for our ML approach, we train separate models within each party to maximize our ability to distinguish the positions of candidates within partisan primaries.

Since we study primary as well as general elections, we also calculate versions of our scores that scale legislators using only *primary* donations from before they won office. These scores are noisier and apply to fewer legislators, but do not change the substantive findings, so we report the main results using these primary-only scores as a robustness check in Online Appendices G and H.

Baseline model: Contribution-weighted averages

Our baseline model follows the methodology of Hall and Snyder (2015) and McCarty et al. (2006) by predicting candidates' NP-Scores using the contributed-weighted average NP-Score of the candidates to which their donors contribute.¹¹ This straightforward method has been employed in previous work on Congressional candidates (Hall, 2015; Hall & Snyder, 2015; Hall, 2019; Hall & Thompson, 2018), and builds on an extensive literature that uses campaign contributions to scale candidates without ML (e.g., Bonica, 2013, 2014, 2018; McCarty & Poole, 1998; McCarty et al., 2006; McKay, 2008, 2010; Poole & Romer, 1985).

Drawing on the NP-Score and campaign contributions data, we estimate predicted NP-Scores for candidates in two stages. First, we estimate a preference score for all state legislative donors as the average contribution-weighted NP-Score of the incumbents to which a donor contributes. More formally, let \mathbf{X} be an $m \times n$ matrix of campaign contributions, where \mathbf{X}_{ij} is the donation amount from donor j to candidate i , and y_i is incumbent i 's NP-Score. Then donor j 's revealed preference z_j is given by

$$z_j = \frac{\sum_{w \neq i} y_w \mathbf{X}_{wj}}{\sum_{w \neq i} \mathbf{X}_{wj}}, \quad (1)$$

where we leave out candidate i when estimating donor j 's preferences to avoid a feedback loop.¹²

The second step is to impute candidate NP-Scores from the preferences of their donors. Specifically, we calculate each candidate's predicted score as

$$\hat{y}_i = \frac{\sum_j z_j \mathbf{X}_{ij}}{\sum_j \mathbf{X}_{ij}}. \quad (2)$$

For the remainder of the paper, we will refer to these contribution-weighted average scores as Hall–Snyder scores.¹³

Machine learning extension: Random forest regression

We improve upon the predictive accuracy of our baseline model with an ML approach. It is important to note that, while the ML-based scores can provide the best within-party predictions of NP-scores, they do carry some costs. For one thing, ML-based scores may contain difficult-to-understand biases, since the process trades off bias in exchange for reducing variance in the prediction problem. Depending on the nature of these biases, they could affect our downstream estimates of the electoral advantage of different types of candidates. Compounding this issue, ML-based scores are inherently something of a black box, and we have only a limited ability to examine what determines the scores that different candidates receive. For that reason, we use both the baseline and ML scores in all of our election analyses, along with CFscores.

To produce the ML scores, we learn a party-specific mapping $\hat{f}_p(\cdot)$ between the donations a candidate receives before they ever take office, and their subsequent NP-Score that summarizes voting behavior over their entire careers:

$$y_{i,\text{post}} = \hat{f}_p(\mathbf{x}_{i,\text{pre}}) + \epsilon_{i,\text{post}}, \quad (3)$$

where $y_{i,\text{post}}$ is the NP-Score for legislator i in party p that summarizes voting behavior post winning office, and $\mathbf{x}_{i,\text{pre}}$ is a vector of predictors for legislator i before ever winning office for the first time. We learn $\hat{f}_p(\cdot)$ using a random forest regression (Breiman, 2001), an ensemble method that learns a large number of decision trees on bootstrapped samples of the training data by randomly selecting subsets of predictor variables to consider at each split of each tree, and averages predictions across trees to produce a final prediction. We choose the optimal number of predictors to select at each split through 10-fold cross-validation. As in Bonica (2018), we construct donation-based predictors using both standalone donations received from larger donors (represented as dummy variables), and summaries of all donors' preferences using the contribution-weighted average method of the baseline model. These donor summaries are constructed in accordance with the cross-validation scheme to avoid data leakage. Due to limited coverage of candidate

¹¹ A similar method is also used in Caughey and Warshaw (2022).

¹² Note, however, that our results are substantively identical if we include candidate i when estimating donor j 's ideology (i.e., $z_j = \frac{\sum_w y_w \mathbf{X}_{wj}}{\sum_w \mathbf{X}_{wj}}$).

¹³ In Online Appendix Figure E.1, we show that we recover substantively identical Hall–Snyder scores when using an indicator for contributions rather than the actual dollar amount. The results suggest that it is the decision to donate, rather than the donation amount, that primarily drives our ideological scaling.

demographic information, we include state dummies as the only non-donation-based predictor (in addition to, implicitly, the party of the legislator by training separate models by party). Training data legislators are assigned their predicted score from the cross-validation round where they were not used to train the model or build the feature set. Online Appendix B describes in more detail how we constructed the feature set, reports the results of the cross-validation exercise, and provides a summary of which features were most predictive. We learn two separate mappings for Republicans and Democrats to improve prediction accuracy within party and, hence, our ability to measure extremism within partisan primaries.

For the remainder of the paper, we will refer to these ML-based scores as “HMH scores” (for Handan-Nader, Myers, and Hall).

Validating the scalings

To validate these new scalings, we conduct two empirical exercises: We compare them to observed NP-Scores for candidates who eventually take office, and we use them to predict a large data set of roll-call votes cast in all 99 state legislatures from 2010 to 2022. We report the results of the first analyses later and refer the reader to Online Appendix F for results on the second analysis.

Figure 1 compares legislator NP-Scores (on the horizontal axis of each panel) to their predicted score based on four possible scalings: their Hall-Snyder scores (top left panel), their HMH scores (top right panel), their static CFscores (bottom left panel), and their dynamic CFscores (bottom right panel). As the figure shows, while all four scores have relatively high overall correlations (indicating their ability to separate legislators of the two parties), both the Hall-Snyder and HMH scores achieve substantially higher within-party correlations. As expected, the HMH scores that use ML achieve the highest within-party correlations. Because not every candidate can be scored with every model, we report a balance table of candidate characteristics in Online Appendix A. We also report within-state, within-party correlations for the HS and HMH scores in Online Appendix D for completeness, though we note that static differences in predictive accuracy *between* states would not bias our downstream regression results due to the inclusion of state fixed effects. In Online Appendix E, we report that within-state, time-variant trends in prediction error are largely uncorrelated with over-time changes in state-level campaign-finance regulations, and thus these types of changes are unlikely to bias our regression results.

Second, leveraging a panel of 72 million raw roll-call votes, we show in Online Appendix F that our scalings

do the best job of replicating the roll-call classification success of the NP-Scores themselves.

In sum, across these two exercises, we see that both of our new contribution-based scores correlate well with NP-Scores within party and predict roll-call voting effectively. This makes them useful tools for analyzing the relationship between candidate ideology and electoral performance, which we will turn to now.

NEW DATA ON STATE LEGISLATIVE ELECTIONS

In order to provide a comprehensive analysis of candidate ideology and electoral performance in both primary and general elections, we assemble a new data set of state legislative election results. We begin with the State Legislative Election Returns (SLERs) data set from Klarner (2023) that covers all general elections in state legislatures, including full coverage of the years of our study, 2000–2022. Though election data are available before the year 2000, we study elections between 2000 and 2022 due to limited coverage of the donation data in the 1990s.

Next, we construct a comprehensive record of primary election outcomes for 2000–2022 in all relevant states. To do this, we started from partial data on 42 states for the period 2000–2014 from Rogers (2023). We added data on primaries in runoff states collected in Fourinaies and Hall (2020). We then collected the remainder of the data—filling in gaps in the other data sets, adding the remaining states, and extending the data through 2022—from state websites, and cleaned and standardized the resulting combined data set extensively. Overall, almost exactly 50% of the data we use were collected anew for our study, with the other half coming roughly equally from the two sources referenced above. When applicable, our primary data include both first-round and runoff primary-election results. The complete primary data set includes full coverage of all primary elections corresponding to general elections in our sample.

To facilitate meaningful comparisons between candidates, we restrict our analysis data along three margins. First, we focus on Democratic and Republican candidates. Second, we subset our data to include state-chamber-years for which a majority of all available seats are in single-member districts. Finally, we exclude state-chamber-years with nonconventional primary election systems (i.e., top-two and blanket primaries), all special elections, and require each election to send its winner to office for a full term.¹⁴

¹⁴ The latter two restrictions affect few legislators, reducing our sample by approximately .08%.

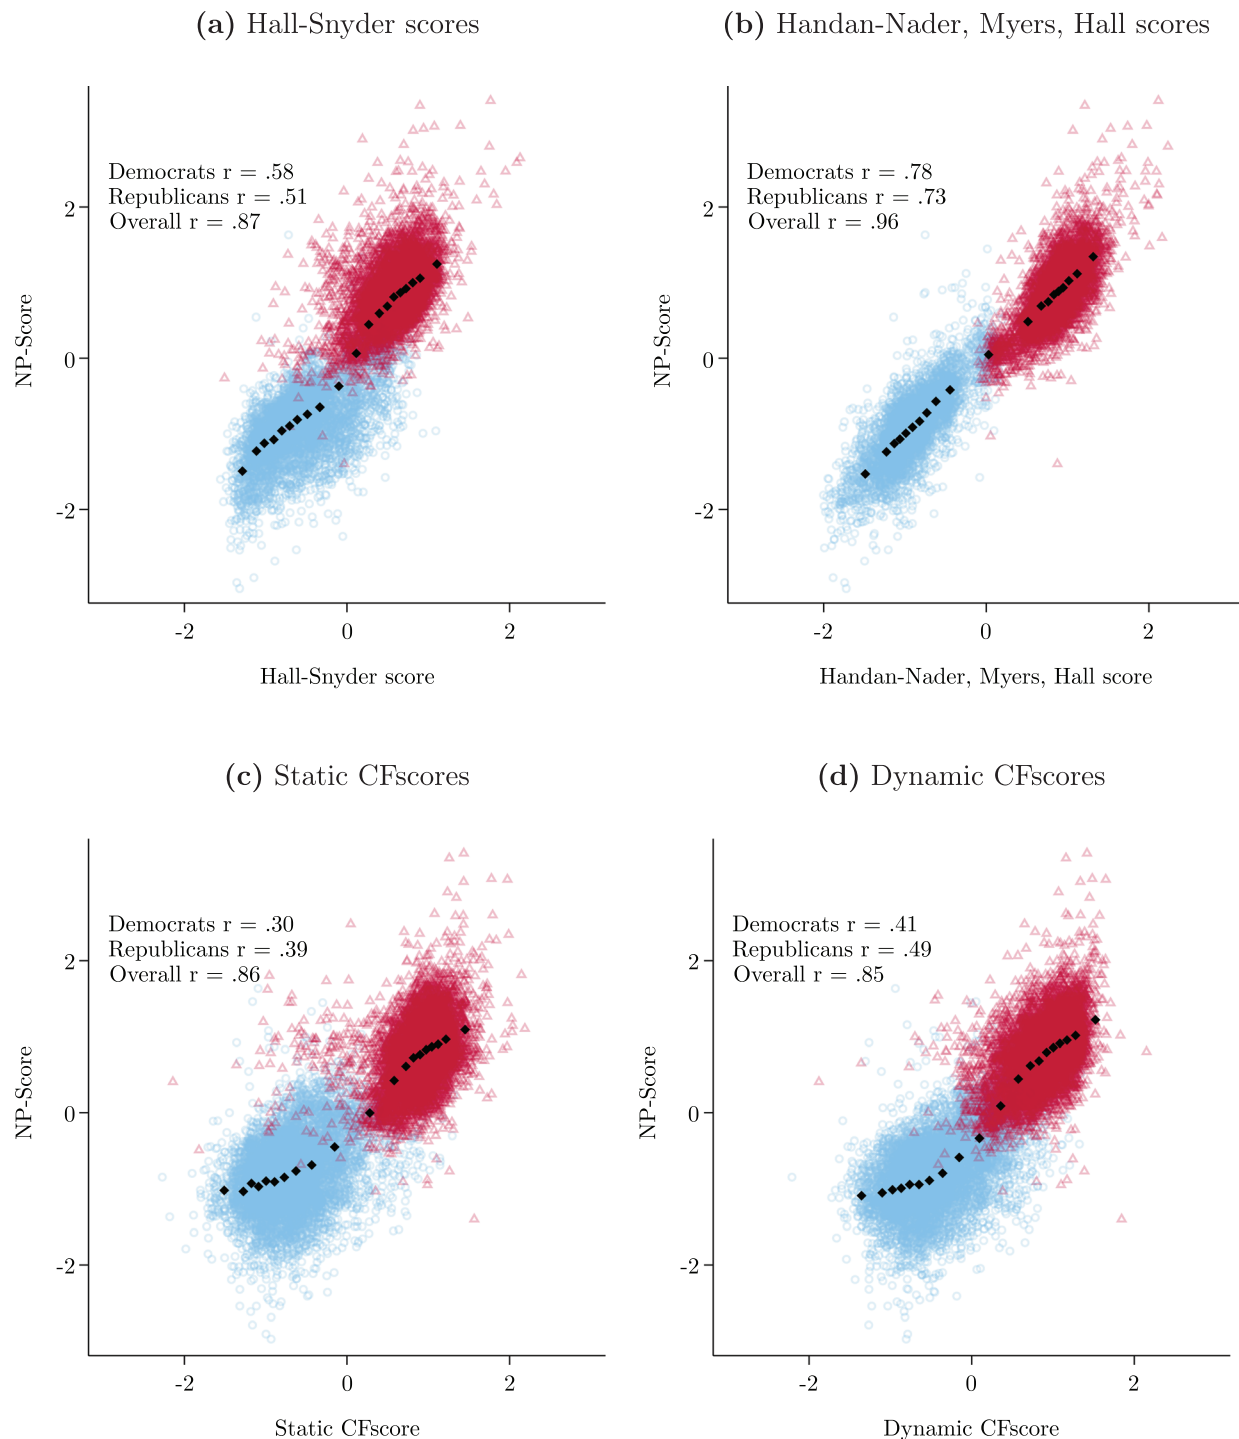


FIGURE 1 Within-party correlation between various scalings and NP-Scores. *Note:* The figures plot roll-call-based NP-Scores against the various campaign-finance-based scalings for Democratic (circle) and Republican (triangle) incumbent legislators. Diamonds represent equal-group-size averages.

We merge the primary and general election data together into a master data set along with the candidate ideology scores validated in section “Campaign-finance-based measures of candidate roll-call ideology.” Due to the restrictions discussed in section “Campaign-finance-based measures of candidate roll-call ideology” that ensure that candidates

have sufficient donation data to receive predicted scores, we are able to assign HMH and Hall-Snyder scores to all candidates within a race in about 10,000 general elections and about 4,000 primary elections, and CFscores to all candidates within a race in about 20,000 general elections and 8,000 primary elections.

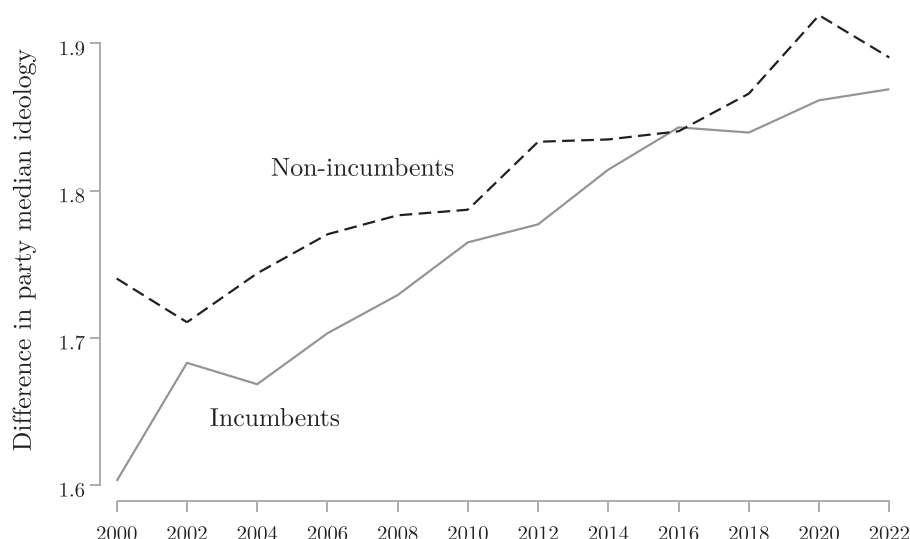


FIGURE 2 Polarization of candidates in state legislatures over time, 2000–2022. *Note:* This figure plots the absolute difference between each party's median incumbent legislator (blue line) and between each party's median nonincumbent candidate (black line), across all states, by year, as measured using Handan-Nader, Myers, Hall scores. Nonincumbent includes both challengers and open-seat candidates.

What kinds of elections enter our sample? Online Appendix A reports characteristics of the elections included in the analysis sample compared to all elections ($n = 63,109$ generals and $n = 79,888$ primaries), all contested elections ($n = 37,335$ generals and $n = 18,362$ primaries), and all competitive elections ($n = 16,450$ generals and $n = 6,883$ primaries). Competitive elections are defined as having a winning margin of 20% or less. As Online Appendix A shows, we are able to capture the majority of competitive races. The races we study are a little less likely to have incumbents than the overall population of competitive races but look very representative of the population in terms of partisanship (measured with Democratic presidential vote share).

POLARIZED CANDIDATE ENTRY IN STATE LEGISLATIVE ELECTIONS

We first use our new data to describe the ideological positions of the people who run for state legislature over time. With relatively low rates of electoral competition, who runs for office becomes especially important in determining the polarization of state legislatures. Figure 2 plots the difference in the median candidate's ideology for each party, using the HMH scores, over time. The plot shows separate lines for the entire set of new candidates in each cycle (i.e., all non-incumbent candidates), and for sitting legislators (i.e., incumbents). To keep the plot easily readable, we omit odd-year elections from it.

As the figure shows, we see a steep increase in the polarization of candidates over time; as legislative

polarization has increased, so, too, has the polarization of the set of people running for office in the first place. The figure also suggests that, though incumbents are less polarized than nonincumbents throughout the study period, the gap between incumbents and nonincumbents appears to have narrowed since 2010. This suggests that, in addition to a steady increase in polarization among who runs for office over the past two decades, there has also been a shift in electoral selection, from a system that weakly favored more-moderate candidates from among the set of candidates to a system that is indifferent between more-extreme and more-moderate candidates. We will formally document this pattern in the analyses below.

GENERAL ELECTIONS AND THE ADVANTAGE OF MORE-MODERATE CANDIDATES

Do contested general elections in state legislatures favor more-moderate candidates, and if so, how much?

This is a classic question in the study of American elections going back to the foundational “median voter theorem” and related ideas about spatial voting explored in Downs (1957), among others. A long debate in political science rages over whether there is, in reality, any advantage to more-moderate candidates as predicted in the spatial model. Behavioral research often argues that voters are not sufficiently informed about candidates and do not have a sufficiently sophisticated view of ideology to favor more-moderate

candidates (e.g., Achen & Bartels, 2016). Research on federal elections, in contrast, consistently documents an electoral advantage for more-moderate candidates (Ansolabehere et al., 2001; Canes-Wrone & Kistner, 2022; Canes-Wrone et al., 2002; Hall, 2015), perhaps because there is a critical mass of informed swing voters, or because campaigns, interest groups, parties, the media, and other elite actors are able to help voters coordinate on more-moderate candidates in the absence of widespread individual-level sophistication. State legislatures present an even stronger challenge to the prediction that more-moderate candidates should be advantaged electorally, however; voters are significantly less informed about state legislators than about national-level politicians (Rogers, 2023), campaigns are significantly less resourced (e.g., Fournaies & Hall, 2014), and there is less media coverage of these races (Moskowitz, 2021). Given these contrasting predictions, it is valuable to see what the data can tell us.

Existing estimates of the advantage to more-moderate candidates in state legislatures do not fully answer this question, because they are forced to focus on incumbents and cannot account for challenger positioning or study open-seat races. To answer this question using our new measures, we first follow the “midpoint” method of Ansolabehere et al. (2001). For each contested election, we compute the distance in ideology between the Democrat and Republican candidates, and we compute the midpoint between their estimated platforms. When this midpoint moves to the right while the distance between the candidates remains constant, it means that the Republican candidate has become more extreme and the Democratic candidate has become more moderate, and vice versa when the midpoint moves to the left while the distance remains constant.

To implement the midpoint method, we estimate regressions of the form

$$Y_{ict} = \beta_1 \text{Midpoint}_{ict} + \beta_2 \text{Distance}_{ict} + X_{ict} + \gamma_i + \delta_t + \epsilon_{ict}, \quad (4)$$

where Y_{ict} represents the Democratic vote share in district i in chamber c at time t . The vector X_{ict} stands in for an optional vector of control variables, and γ_i and δ_t stand in for district-regime and time fixed effects.

The quantity of interest is β_1 , which captures the association between how moderate the Democratic candidate is (when the midpoint between the two candidates shifts right while holding the distance between them equal) and Democratic electoral outcomes. In the original Ansolabehere et al. (2001) approach, the unobserved district median voter’s preferences are held constant by controlling for presidential vote share in the district. Because presidential vote share is not widely available for all state legislative districts, our

preferred specification uses year fixed effects and district fixed effects generated separately for each redistricting period. As a robustness check, for the districts where presidential vote share is available, we report substantively similar results in Online Appendix G.¹⁵

Table 1 presents the results. Each cell in the table reflects a different estimate of β_1 , capturing the relationship between candidate moderation and vote share. The rows show the estimates for different candidate ideology scalings, while the columns are for different regression specifications. The first column is our baseline midpoint specification in Equation (4) using year and district fixed effects (within a districting regime). Because our scalings are not perfect predictors of roll-call voting but rather rely on campaign contributions to estimate candidate positions, columns (2)–(4) subsequently add in different ways to control for possible differences in prediction error between competing candidates due to disparities in the amount of money they raise. In all cells, the midpoint variable is scaled to run from 0 in the race with the leftmost midpoint to 1 in the race with the rightmost midpoint, so that β_1 reflects the predicted change in Democratic vote share for the maximal shift in the midpoint observed in the sample. While we prefer to focus on the first two rows that use our preferred scalings for these analyses—HMH and Hall–Snyder scores—we also present estimates for static and dynamic CFscores (rows 3 and 4) in order to offer a point of comparison to previous work. Because CFscores are not designed to predict NP-Scores, and because they pool contributions before and after successful candidates win office, we do not rely on these estimates for our main results.

Looking down the rows for our baseline specification in column (1), we see that we find a consistently positive coefficient, indicating an advantage for more-moderate candidates. The estimates using static CFscores and HMH scores are roughly half the size of those using Hall–Snyder scores and dynamic CFscores, but are directionally similar. In column (2), we add controls for the total contributions raised in the primary by each candidate, in logs. We focus only on primary contributions in order to avoid a sort of “post-treatment” bias that might occur where a candidate’s ideology both affects their ability to raise money in the general and affects their electoral outcome—such as if a more-moderate candidate is able to raise more money in the general election.¹⁶ These controls are helpful for making sure that our results are not driven by any possible linkage between raising more money and being erroneously scaled as more moderate. As we see, with this control included, all the estimates shrink

¹⁵ We have presidential vote share data for 70% of the state legislative races in our sample.

¹⁶ In Online Appendix G, we re-estimate Table 1 using the HMH scores that only use primary donations as a robustness check for possible “post-treatment” bias in the scores themselves, with substantively similar results.

TABLE 1 Advantage of more-moderate candidates in contested general elections, 2000–2022.

| | Democratic vote share | | | |
|--|-----------------------|------------------|------------------|------------------|
| | (1) | (2) | (3) | (4) |
| HMH score | 0.16** (0.02) | 0.12** (0.02) | 0.13** (0.03) | 0.11** (0.03) |
| Hall-Snyder score | 0.26** (0.02) | 0.19** (0.02) | 0.22** (0.04) | 0.23** (0.03) |
| Static CFscore | 0.14** (0.01) | 0.10** (0.01) | 0.17** (0.02) | 0.14** (0.02) |
| Dynamic CFscore | 0.30** (0.01) | 0.22** (0.01) | 0.32** (0.02) | 0.30** (0.02) |
| District-by-regime FE | Y | Y | Y | Y |
| Year FE | Y | Y | Y | Y |
| Controls for primary contributions | N | Y | N | N |
| Only races with below-median contribution gap | N | N | Y | N |
| Only races with ≥ 10 primary donors per candidate | N | N | N | Y |

Robust standard errors are clustered by district-regime in parentheses. HMH score refers to Handan-Nader, Myers, Hall score.

[†] $p < .10$; * $p < .05$; ** $p < .01$.

Each cell in this table reports an estimate of *Midpoint*, or the association between how moderate a Democratic candidate is (when shifting the midpoint between the Democratic and Republican candidate to the right) and the Democrat's general-election vote share, after holding constant the distance between candidates.

from their prior magnitudes but remain positive. These estimates are corroborated in columns (3) and (4), which use different ways to control for the same source of confounding due to potential disparities in fundraising between competing candidates.

As further robustness checks against confounding due to scaling prediction error, we also report substantively similar results in Online Appendix G when removing states for which the within-state correlations between the HMH and NP-Scores are especially low, and when removing states for which Shor and McCarty (2011) report high vote prediction error in the NP-Scores themselves. In Online Appendix I, we show that the midpoint coefficient shrinks in magnitude but remains positive for both Hall-Snyder and HMH scores as we increase the number of donors each candidate must have to be included in the analysis.

In general, looking across all of the specifications, we see strong evidence for a positive overall advantage to more-moderate candidates. If we focus on the HMH scores estimate in column (2) as our best single estimate, we estimate that shifting from the most-extreme Democratic candidate to the most-moderate predicts a 12 percentage point increase in vote share. Based on the standard deviation of the midpoint variable in this sample, a one standard deviation increase in the midpoint variable would predict a 1.56 percentage point increase in Democratic vote share. This is not nothing, and could certainly matter in a close election, but it does not seem like a very large advantage. For comparison, using the same midpoint approach, Hall (2019) estimates that shifting from the leftmost to the right-

most midpoint corresponds to a 30 percentage point increase in Democratic vote share, an advantage that is more than double this estimate for state legislatures.

Regression discontinuity

The midpoint approach used above has the advantage of using all of our data on contested general elections where we are able to scale both candidates. However, as we discussed, it requires being able to hold fixed the unobserved preferences of the district, which we do either using fixed effects or by controlling for presidential vote. Neither of these is a silver bullet; if districts' political preferences change within redistricting cycles substantially, the fixed effects would fail to capture these trends. A similar issue occurs with presidential vote since it is not observed every year.

Hall (2015) provides an alternative way to hold fixed the preferences of the district by focusing on close primary elections between a more-extreme and a more-moderate candidate, with the idea that this approximates a natural experiment in which the district “randomly” receives one type of candidate or the other. To the extent this natural experiment is valid—an assumption for which we provide evidence below—then the districts that just barely nominate a more-extreme candidate will be otherwise just like those that nominate a more-moderate candidate, on average, including in their overall political preferences. For each contested primary and general election in which we are able to scale at least two candidates, we




FIGURE 3 Effect of extremist nominee on general election vote share in US state legislatures, 2000–2022. *Note:* This figure plots averages of parties' general election vote share across equal-sample-sized bins of the extremist candidate's primary winning margin. Regression lines are from ordinary least squares (OLS) and are estimated on the raw data separately on each side of the discontinuity.

compute the estimated ideological distance between the top two vote-getting candidates.¹⁷ To focus on cases where there is a meaningful ideological distance between the more-moderate and the more-extreme candidate, we then restrict the data to cases where the distance between these top two candidates is at or above the median distance across all cases. We follow standard approaches to estimate the “jump” at the discontinuity that occurs when the more-extreme of the two candidates just barely switches from losing the primary to winning it.

Figure 3 shows the results graphically. As can be seen in both panels, when the more-extreme candidate goes from just barely losing the primary (left side of each plot) to just barely winning (right side of each plot), the party's general election vote share drops noticeably. The size of this drop is meaningful but not huge.

We estimate the size of this drop formally using standard approaches including the optimal bandwidth approach of Calonico et al. (2014). Table 2 presents estimates for four different specifications and all four possible scaling approaches. In the first column, we focus on data in a 10 percentage point window around 50/50 and use only a linear specification of the running variable. In the second and third columns, we include all of the data and use either a third-order or fifth-order polynomial specification of the running variable. Finally, in the fourth column, we use the automated procedure from Calonico et al. (2014).

Looking across the first row, we see that the estimates using the HMH scores range from a 2 percentage point penalty to a 5 percentage point penalty. These estimates grow modestly with Hall-Snyder

scores and CFscores but remain relatively stable across specifications. Looking across the estimates, we find strong evidence for a modest penalty to extremist nominees. While a 2–5 percentage point penalty in vote share is enough to tip close elections, it is small enough to not matter in many cases, too. In comparison, Hall (2019) uses the same RD setup with Hall-Snyder scores and estimates an 8 percentage point effect on vote share.

As is standard with RD analyses, in Online Appendix J, we show that there is no evidence for sorting or for an imbalance that would contribute to these negative estimates. Also in Online Appendix J, we investigate how our estimated effect varies as we change the minimal ideological distance between candidates that is required for a race to enter our

TABLE 2 Effect of extremist nominee on general election vote share, US state legislatures 2000–2022.

| | Party vote share | | | |
|-------------------|-------------------|-------------------|-------------------|-------------------|
| | (1) | (2) | (3) | (4) |
| HMH score | −0.04** (0.02) | −0.04** (0.01) | −0.05** (0.02) | −0.02 (0.02) |
| Hall-Snyder score | −0.05** (0.02) | −0.05** (0.01) | −0.06** (0.01) | −0.06** (0.02) |
| Static CFscore | −0.06** (0.01) | −0.05** (0.01) | −0.05** (0.01) | −0.06** (0.01) |
| Dynamic CFscore | −0.05** (0.01) | −0.04** (0.01) | −0.04** (0.01) | −0.05** (0.01) |
| Polynomial | 1 | 3 | 5 | CCT |
| Bandwidth | .10 | — | — | — |

Robust standard errors are reported in parentheses. HMH score refers to Handan-Nader, Myers, Hall score.

[†] $p < .10$; * $p < .05$; ** $p < .01$.

Each cell in this table reports an estimate of the effect of narrowly nominating a primary extremist on their party's general-election vote share relative to a moderate.

¹⁷ Ideally, we would construct the ideological distance measure using only primary donations, since general donations are post-treatment in this setting. However, due to the sparsity of primary donations for nonincumbents, the primary donation measure is too noisy to cleanly estimate the RD, so we rely on the measure that uses both primary and general donations.

sample (i.e., the cutoff). As the figure shows, we find that the estimated penalty to nominating an extremist increases noticeably as we increase the cutoff and thereby focus on more “intensive” treatments in which the more-extreme candidate is farther away from the more-moderate candidate.

VARIATION IN THE ADVANTAGE OF MORE-MODERATE CANDIDATES

So far, we have established that there is a modest, positive advantage to more-moderate candidates in contested general elections in state legislatures. This estimate pools across roughly 20 years of data and across 47 states. We can learn more about the roots of this modest advantage now by exploring where and when the advantage is larger and smaller.

Declining general election advantage to moderates over time

First, we explore whether the advantage to more-moderate candidates in contested general elections has changed in recent years. As state legislatures have polarized and elections have nationalized, we might suppose that the advantage to more-moderate candidates has gone down in state legislatures. This would be consistent with the argument advanced in Rogers (2016, 2023) that voters in state legislative elections are highly partisan, focus on national races at the top of the ballot, and rarely know much or anything about their state legislative candidates. On the other hand, decades of research at the federal level shows how campaigns, interest groups, parties, the media, and other elite actors can structure elections such that more-moderate candidates are favored even if most voters are unaware of candidate positions. As such, it is not clear whether the advantage to more-moderate candidates has actually declined or not in state legislatures; we need to examine the data directly.

As we discussed in the Introduction, making over-time comparisons requires having access to measures of candidate ideology for not only incumbents, but also for challengers and open-seat entrants, so that we do not confuse over-time changes in conditioning on incumbency with changes to the unconditional advantage to more-moderate candidates. Our new measures and data allow us perform this over-time comparison for the first time.

In Figure 4, we estimate Equation (4) separately for each year, for both the Hall–Snyder scores and the HMM scores. This is only possible when we use presidential vote to control for district preferences, since our main approach uses district fixed effects that require multiple years to be pooled in order to work.




FIGURE 4 General election advantage to moderates over time. *Note:* This figure reports the coefficient on candidates’ *Midpoint* estimated separately by year along with the upper and lower 95% confidence intervals (vertical bars). This figure uses district presidential vote share to hold the median voter constant. HMM Scores refer to Handan-Nader, Myers, Hall Scores.

The figure reveals a relatively steady decline in the size of the coefficient on midpoint over time, indicating that the advantage to more-moderate candidates in contested general elections is shrinking. Although exactly how and when this decline has occurred varies across the two scores, the trends are very similar, and in both cases, the estimate in 2020 is the smallest of the whole time period, and indicates a quite modest advantage.¹⁸

Regression results

Table 3 presents formal regression results for the above heterogeneity tests as well as some additional ones, using the HMM scores. In each column, we re-estimate Equation (4), our basic midpoint approach, and we interact the midpoint variable with a key moderator variable. Consistent with prior work on state legislatures, we focus on moderators that speak to the nationalization of state elections (e.g., Rogers, 2016, 2023), the professionalization of state legislatures (e.g., Birkhead, 2015; Rogers, 2017), and the degree of electoral competition (e.g., Rogers, 2017), as these factors have been hypothesized to contribute to polarization at the state level.

Column (1) of Table 3 simply provides a formal test related to Figure 4. Specifically, we interact the midpoint variable with an indicator for whether the election takes place in 2012 or later. We chose 2012 because it is the first cycle that occurs after the 2010 redistricting cycle, which aligns the way we cut the data with the district-by-regime fixed effects. In the

¹⁸ The year 2022 is excluded from this figure because we lack data on presidential election returns at the legislative district level for elections after the 2020 decennial redistricting process.

TABLE 3 Variation in midpoint coefficient. This table reports estimates of *Midpoint* and its interaction with relevant moderating variables.

| | Handan-Nader, Myers, Hall score | | | | |
|---|---------------------------------|-------------------|------------------|-------------------|------------------|
| | (1) | (2) | (3) | (4) | (5) |
| Midpoint | 0.18** (0.03) | 0.12** (0.02) | 0.12** (0.04) | 0.16** (0.03) | 0.07** (0.03) |
| Midpoint · Year ≥ 2012 | −0.09* (0.04) | | | | |
| Midpoint · Off cycle | | 0.02† (0.01) | | | |
| Midpoint · Odd year | | −0.45** (0.14) | | | |
| Midpoint · Prof. (Squire dynamic average) | | | 0.03 (0.10) | | |
| Midpoint · Prof. (Squire staff) | | | | −0.28** (0.10) | |
| Midpoint · Prof. (Squire salary) | | | | 0.17† (0.10) | |
| Midpoint · Prof. (Squire session length) | | | | −0.06 (0.06) | |
| Midpoint · Competitive | | | | | 0.05** (0.02) |
| N | 6,349 | 6,349 | 6,291 | 6,285 | 6,349 |
| Controls for primary contributions | Y | Y | Y | Y | Y |
| District FEs | Y | Y | Y | Y | Y |
| Year FEs | Y | Y | Y | Y | Y |

Robust standard errors clustered by district in parentheses. *Midpoint* is defined according to Equation 4 and is scaled to run from 0 (most liberal) to 1 (most conservative) for each scaling. *Professionalization* (Prof.) is scaled to run from 0 (least professionalized state) to 1 (most professionalized state). *Competitive* districts are districts where neither party received greater than 70% of the two-party presidential vote share. *Off-Cycle* elections occur in non-presidential election years. *Distance* is included in all models but not reported in the table. Data on professionalization for Louisiana and West Virginia were not available in columns 3 and 4.

† $p < .10$; * $p < .05$; ** $p < .01$.

second row, we see that this interaction is negative and statistically significant, indicating a decline in the advantage post 2010.

In column (2), we explore whether the advantage to more-moderate candidates varies across the types of elections: presidential elections (captured in the main effect on midpoint in the first row of the table), off-cycle elections (the interaction in the third row of the table), and odd-year elections (the interaction in the fourth row). The results indicate that there is no interaction for off-cycle races, meaning that we estimate that more-moderate candidates have similar advantages whether or not they share the ballot with a presidential race. Odd years potentially provide a more-interesting test because in the few states that hold odd-year elections for state legislature, there are often no national races at all on the ballots. Here we find a large negative interaction that is imprecisely estimated, as there are few cases that hold elections in odd years.

In column (3), we interact the midpoint variable with Squire's (2017) measure of state legislative professionalization, as implemented in Birkhead (2015). Professionalization is scaled from 0 (least professionalized) to 1 (most professionalized). In contrast to the finding in Birkhead (2015), we find if anything, a positive though imprecise interaction coefficient indicating that more professionalized state legislatures exhibit somewhat larger advantages for more-moderate candidates. In column (4), we further explore this interaction by decomposing the professionalization measure into its key constituent parts, following Rogers (2017). Consistent with the findings in Rogers (2017), we find a large negative interaction with staff size, indicating that state legislatures with larger staffs exhibit, on average, lower advantage for more-moderate candidates. Coefficients for salary and session length are smaller and noisier, roughly consistent with Rogers (2017) as well. Together, these estimates indicate that there is no simple relation-

ship between legislative professionalization and the advantage to more-moderate candidates, possibly because, as both Birkhead (2015) and Rogers (2017) point out, professionalization bundles together both factors that make elections more salient with factors that may allow legislators to gain voters through other, nonideological dimensions. Sussing out the precise mechanisms between legislative professionalization, candidate ideology, and polarization is beyond the scope of this study, but the new measures and data we provide should prove useful to future research in this direction.

Finally, in column (5), we show that the advantage to more-moderate candidates is stronger in competitive districts. Competitive districts are districts where neither major party averaged greater than 70% of the two-party presidential vote share across the districting regime. As the table shows, the advantage is estimated to be larger in these districts. This is consistent both with the logic of the spatial model, as well as with classic findings for the federal level (Ansolabehere et al., 2001) and previous work looking at state legislative incumbents (Rogers, 2017). This is also important because it suggests that, to the extent there are fewer competitive districts in state legislatures than there used to be, this decline in competition could lead to a decline in the advantage to more-moderate candidates.

PRIMARY ELECTIONS AND THE ADVANTAGE OF MORE-EXTREME CANDIDATES

Having explored the links between candidate ideology and electoral outcomes in general elections, we now turn to estimating the advantage for more-extreme candidates in contested primary elections. A long literature at the federal level documents the conflict between primary and general electorates, with primary electorates thought to prefer more-extreme candidates, while general electorates are thought to prefer more-moderate candidates (e.g., Aranson & Ordeshook, 1972; Brady et al., 2007; Hall, 2015). The mechanisms underlying the advantage of more-extreme primary candidates are not well understood, but may include the activation of ideologically extreme interest groups whose influence is heightened in primaries because they are lower salience (e.g., Bawn et al., 2015), as well as the differential participation of more-extreme voters in primaries (Hill & Tausanovitch, 2016).

Whether this same dynamic is at play in state legislatures is unclear, given the general lack of information and competition in state legislative primaries. If the voters turning out in primaries tend to be

more extreme, state legislative primaries might favor more-extreme candidates like at the federal level; on the other hand, if these voters are turning out to vote on the top-of-ballot primaries and are generally uninformed about their potential state legislative nominees, then we should not expect an advantage for more-extreme candidates. Further, there is no evidence on whether these potential effects have increased, decreased, or remained the same over time.

The one existing analysis of state legislative primary elections and candidate ideology only has access to a measure of incumbent ideology. While Rogers (2023) finds that more-extreme incumbents do somewhat better electorally in primary elections, it is unclear whether the same pattern persists in open-seat primaries, which are essential for sending new incumbents to office.

To measure how extreme primary candidates are in relation to one another, we follow Hall and Snyder (2015) and define

$$\begin{aligned} \text{Relative Centrism}_{ipdt} \\ = |\text{Cand Ideology}_{ipdt} - \text{Most Extreme Ideology}_{pdt}|, \end{aligned} \quad (5)$$

where $\text{Cand Ideology}_{ipdt}$ reflects the ideology score of candidate i running in the contested primary in district d for party p in year t . The variable $\text{Most Extreme Ideology}_{pdt}$ represents the most-extreme candidate running in primary pdt , that is, the candidate with the maximum scaling, in a Republican primary, and the candidate with the minimum scaling, in a Democratic primary. The basic idea here is to give each candidate in a contested primary a score that indicates how much more moderate she is than the most-extreme candidate in the race. This measure is better than using the simple absolute value of the scaling to measure extremism because it deals with cases where Republicans have scalings less than zero or Democrats have scalings greater than zero. This occurs with nontrivial frequency because 0 is an arbitrary value in the candidate scalings.

Armed with this measure, we then estimate regressions of the form

$$Y_{ipdt} = \beta \text{Relative Centrism}_{ipdt} + Z_{pdt} + X_{ipdt} + \epsilon_{ipdt}, \quad (6)$$

where Y_{ipdt} reflects the vote share for candidate i in the primary for party p in district d in year t .¹⁹

¹⁹ Note that the regression discontinuity design from section "Regression discontinuity" does not apply to primary elections because we lack an as-if random assignment mechanism for running in a primary election. We also cannot study subsequent primary election outcomes because doing so would

TABLE 4 Advantage of more-extreme candidates in contested primary elections, 2000–2022.

| | Primary vote share | | | |
|------------------------------------|--------------------|-------------------|-------------------|-------------------|
| | (1) | (2) | (3) | (4) |
| HMH score | −0.17** (0.02) | −0.17** (0.02) | −0.17** (0.02) | −0.17** (0.02) |
| HMH score (Primary donations only) | −0.15** (0.03) | −0.15** (0.02) | −0.16** (0.03) | −0.17** (0.02) |
| Hall-Snyder score | −0.12** (0.02) | −0.12** (0.02) | −0.13** (0.02) | −0.14** (0.02) |
| Static CFscore | −0.13** (0.02) | −0.14** (0.02) | −0.16** (0.02) | −0.18** (0.02) |
| Dynamic CFscore | −0.02 (0.03) | −0.02 (0.03) | −0.07* (0.03) | −0.09** (0.03) |
| District-by-party FE | Y | N | Y | N |
| Party-by-year FE | Y | N | Y | N |
| Number of candidates FE | Y | N | Y | N |
| Race FE | N | Y | N | Y |
| Controls for primary contributions | N | N | Y | Y |

Each cell in this table reports the coefficient on *Relative Centrism* from Equation 6 which is scaled to run from 0 (most extreme) to 1 (most moderate) for each scaling. The sample is restricted to contested primary elections and excludes races in districts where the opposing party received greater than 70% of the two-party presidential vote share. Robust standard errors in parentheses. HMH score refers to Handan-Nader, Myers, Hall score.

† $p < .10$; * $p < .05$; ** $p < .01$.

Each cell in this table reports an estimate of *Relative Centrism*, or the association between more-moderate positions (relative to the most extreme candidate in the primary) and primary-election vote share.

The vectors X and Z stand in for covariates, which we include to address the fact that the relative centrism score in Equation (5) may still be sensitive to differences in the predictive accuracy of the scalings between primary districts and election years. For example, the location of the most extreme candidate within a race could vary more widely in races with comparatively less donation data, producing systematically larger swings in candidate extremism for some districts relative to others. To address this potential source of confounding, we employ two different baseline model specifications in Table 4. Column (1) performs a difference-in-differences in which we compare within-primary-district variation in candidate extremism over time, conditional on the number of candidates in the primary (as vote share decreases mechanically with an increasing number of primary candidates). Column (2) instead includes fixed effects for the specific primary election (i.e., for each state-district-party-year), which makes comparisons only among candidates in a given race. In this latter specification, we do not need to include fixed effects for the number of candidates, since it is fixed within each race. This is arguably the strongest specification, since

it does not require making any cross-district comparisons, but it may be statistically noisier. Columns (3) and (4) supplement these baseline specifications with controls for each candidate's primary contributions as a further robustness check against confounding due to fundraising disparities between competing candidates, as in column (2) of Table 1.

Table 4 presents the results for contested primary elections, excluding districts where the opposing party received greater than 70% of the two-party presidential vote share.^{20,21} Each cell in the table reflects the coefficient on *Relative Centrism* for a different scaling and specification of the regression. Looking at the first two rows that use our preferred scalings, we see a consistent, large, negative coefficient—indicating that more-moderate candidates do worse, on average, in contested primaries. These results are corroborated by the static and dynamic CFscores.

Consider our preferred estimate, which uses race fixed effects, the HMH scores, and no additional controls (since using primary contributions as a control

condition on a post-treatment variable (whether the candidate seeks reelection again). This post-treatment bias could be severe since, in our sample, moderates are nearly 10 percentage points more likely to run for reelection following a close primary election than extremist candidates.

²⁰ In Online Appendix K, we document that all four scalings tend to underestimate extremism for candidates in very uncompetitive presidential districts, particularly for Democratic candidates, likely due to access-seeking behavior on the part of donors. The HMH scores do the best at ameliorating this relationship, but for robustness we exclude very uncompetitive districts from the analysis. Our substantive findings are unchanged when adding the remaining 30% of races.

²¹ In our sample, 39% of primary elections are contested, while 60% of general election races are contested.

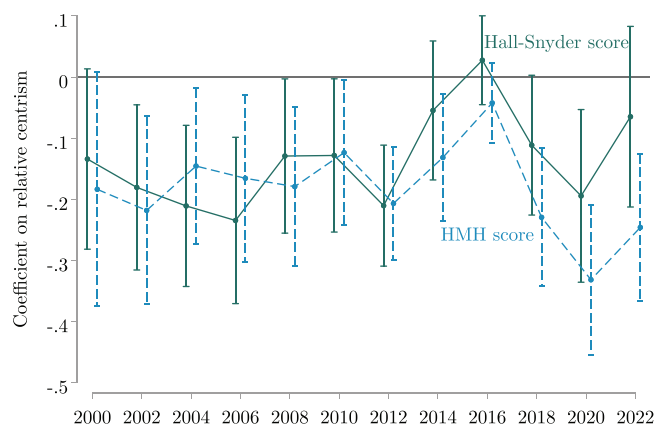


FIGURE 5 Primary-election penalty to extremists over time.
Note: This figure plots estimates from Equation (5) separately for every on-cycle election year. HMH Scores refer to Handan-Nader, Myers, Hall scores.

is arguably post-treatment in this case). Here we estimate that going from the most moderate candidate to the most extreme candidate in a primary predicts a 17 percentage point decrease in primary vote share. This maps to a 1 percentage point decrease in predicted vote share for every one standard deviation shift in relative centrism, which seems like a small but potentially important effect. In Online Appendix H, we report substantively similar results for our preferred specification when removing states with high prediction error, as in earlier robustness checks for the midpoint models. In Online Appendix I, we show that the advantage to extremists shrinks in magnitude but remains statistically different from zero for both HS and HMH scores as we increase the minimum number of unique donors that both candidates in a race must receive a contribution from to be included in the analysis. Finally, in Online Appendix Table H.1, we show that our results are of similar magnitudes after restricting the analysis to races with below-median contribution gaps between candidates and races with at least 20 donors per candidate. In sum, our results corroborate at the state level what has been found at the federal level—more extreme candidates appear to fare better in contested primary elections.

Advantage to extremists in primaries persists over time

Finally, we can also again explore whether these relationships are changing over time. In Figure 5, we plot our estimates for the coefficient on *Relative Centrism* by year, using Equation (6) as our specification. We see a brief decline in the disadvantage for more-moderate

candidates between 2012 and 2016, but in recent years, the disadvantage has returned to being statistically indistinguishable from pre-2012 levels, though the point estimate is lower.²²

CONCLUSION

Understanding how state legislatures have polarized is important both because the state legislatures are themselves highly important policymaking bodies, and because they are the main pathway for candidates to Congress. In this paper, we have offered the first systematic analyses of the links between candidate ideology, electoral competition, and legislative polarization in state legislatures that cover all three stages of the process: candidate entry, primary elections, and general elections. Using new data and new measures of candidate ideology based on campaign contributions, we have established a number of empirical patterns relevant for future work on elections and polarization.

Our study is meant to be only the first key step in what must be a broader effort to understand why state legislative elections work the way that they do. Why are the people running for state legislature themselves so much more polarized than they used to be? Why are more-extreme candidates advantaged in these primaries, and why has their disadvantage in general elections decreased? At the same time, how do state legislative elections sustain *any* advantage for more-moderate candidates, when information on state legislative candidates has always been quite low? These are key questions for future research, and should be aided by the new measures and data that we have assembled to understand state legislative elections.

ACKNOWLEDGMENTS

For data, the authors thank Adam Bonica, Carl Klarner, Steve Rogers, and Jim Snyder. For comments, the authors thank Avi Acharya, Justin Grimmer, Steve Rogers, Chris Warshaw, and participants of the UC Berkeley American Politics Workshop, the Stanford Political Science Workshop, the Stanford GSB Political Economy Theory Lunch, and the Princeton American Politics Seminar. Some of the computing for this project was performed on the Sherlock cluster. We would like to thank Stanford University and the Stanford Research Computing Center for providing computational resources and support that contributed to these research results. This research was supported in

²² In a previous version of the paper, we found that the advantage to more-extreme primary candidates increased in the 2010s relative to earlier years. With our updated data and model specifications, we no longer find that this is the case.

part by funding from the UPS Endowment Fund at Stanford University. For extraordinary research assistance, we thank Elise Kostial, Laniah Lewis, Janet Malzahn, and Anish Sundar.

ORCID

Cassandra Handan-Nader  <https://orcid.org/0000-0003-4868-2509>

Andrew C. W. Myers  <https://orcid.org/0009-0008-5522-8877>

REFERENCES

- Abramowitz, Alan I., and Steven Webster. 2016. "The Rise of Negative Partisanship and the Nationalization of US Elections in the 21st Century." *Electoral Studies* 41: 12–22.
- Achen, Christopher H., and Larry M. Bartels. 2016. *Democracy for Realists: Why Elections Do Not Produce Responsive Government*. Princeton, NJ: Princeton University Press.
- Ansolabehere, Stephen, James M. Snyder, Jr, and Charles Stewart, III. 2001. "Candidate Positioning in US House Elections." *American Journal of Political Science* 45(1): 136–59.
- Anzia, Sarah F. 2021. "Party and Ideology in American Local Government: An Appraisal." *Annual Review of Political Science* 24: 133–50.
- Aranson, Peter H., and Peter C. Ordeshook. 1972. "Spatial Strategy for Sequential Elections." In *Probability Models of Collective Decision Making*, edited by Richard G. Niemi and Herbert F. Weisberg, pp. 298–331. Columbus, OH: Charles E. Merrill.
- Bawn, Kathleen, Knox Brown, Angela Ocampo, Shawn Patterson, John Ray, and John Zaller. 2015. "Social Choice and Coordination Problems in Open House Primaries." Presented at the American Political Science Association Conference, San Francisco, CA, 2016.
- Birkhead, Nathaniel A. 2015. "The Role of Ideology in State Legislative Elections." *Legislative Studies Quarterly* 40(1): 55–82.
- 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): 367–86.
- Bonica, Adam. 2018. "Inferring Roll-Call Scores from Campaign Contributions Using Supervised Machine Learning." *American Journal of Political Science* 62(4): 830–48.
- Bonica, Adam. 2016. DIME Plus. [Computer file]. Stanford, CA: Stanford University Libraries. <https://data.stanford.edu/dime-plus>.
- Bonica, Adam, and Zhao Li. 2021. "Inferring Candidates' Issue-Specific Positions from Itemized Campaign Contributions Using Supervised Machine Learning." <https://www.dropbox.com/s/6uie90whk4n7xdl/Bonica%20and%20Li%20%282021%29.pdf?dl=0>.
- Brady, David W., Hahrie Han, and Jeremy C. Pope. 2007. "Primary Elections and Candidate Ideology: Out of Step with the Primary Electorate?" *Legislative Studies Quarterly* 32(1): 79–105.
- Breiman, Leo. 2001. "Random Forests." *Machine Learning* 45: 5–32.
- Calonico, Sebastian, Mattias D. Cattaneo, and Rocio Titiunik. 2014. "Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs." *Econometrica* 82(6): 2295–326.
- Canes-Wrone, Brandice, and Michael R. Kistner. 2022. "Out of Step and Still in Congress? Electoral Consequences of Incumbent and Challenger Positioning Across Time." *Quarterly Journal of Political Science* 17(3): 389–420.
- 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): 127–40.
- Caughey, Devin, and Christopher Warshaw. 2022. *Dynamic Democracy: Public Opinion, Elections, and Policymaking in the American States*. Chicago, IL: University of Chicago Press.
- Downs, Anthony. 1957. *An Economic Theory of Democracy*. New York, NY: Harper and Row.
- Follow The Money. 2024. "Database of State Legislative Campaign Contributions." <https://www.followthemoney.org/>.
- Fourinaies, Alexander, and Andrew B. Hall. 2014. "The Financial Incumbency Advantage: Causes and Consequences." *Journal of Politics* 76(3): 711–24.
- Fourinaies, Alexander, and Andrew B. Hall. 2020. "How Divisive Primaries Hurt Parties: Evidence from Near-Runoffs in US Legislatures." *Journal of Politics* 82(1): 43–56.
- Hall, Andrew B. 2015. "What Happens When Extremists Win Primaries?" *American Political Science Review* 109(1): 18–42.
- Hall, Andrew B. 2019. *Who Wants to Run? How the Devaluing of Political Office Drives Polarization*. Chicago, IL: University of Chicago Press.
- Hall, Andrew B., and Daniel M. Thompson. 2018. "Who Punishes Extremist Nominees? Candidate Ideology and Turning Out the Base in U.S. Elections." *American Political Science Review* 112(3): 509–24.
- Hall, Andrew B., and James M. Snyder, Jr. 2015. "Candidate Ideology and Electoral Success." https://www.andrewbenjaminhall.com/Hall_Snyder_Ideology.pdf.
- Hill, Seth J., and Chris Tausanovitch. 2018. "Southern Realignment, Party Sorting, and the Polarization of American Primary Electorates, 1958–2012." *Public Choice* 176: 107–132.
- Hogan, Robert E. 2008. "Policy Responsiveness and Incumbent Reelection in State Legislatures." *American Journal of Political Science* 52(4): 858–73.
- Hopkins, Daniel J. 2018. *The Increasingly United States: How and Why American Political Behavior Nationalized*. Chicago, IL: University of Chicago Press.
- Klarner, Carl. 2023. *State Legislative Election Returns, 1967–2022*.
- Kuriwaki, Shiro. 2023. "Ticket Splitting in a Nationalized Era." <https://osf.io/preprints/socarxiv/bvgz3/>.
- Lee, Frances E. 2009. *Beyond Ideology: Politics, Principles, and Partisanship in the US Senate*. Chicago, IL: University of Chicago Press.
- McCarty, Nolan M., and Keith T. Poole. 1998. "An Empirical Spatial Model of Congressional Campaigns." *Political Analysis* 7: 1–30.
- McCarty, Nolan M., Keith T. Poole, and Howard Rosenthal. 2006. *Polarized America: The Dance of Ideology and Unequal Riches*. Cambridge, MA: MIT Press.
- McKay, Amy. 2008. "A Simple Way of Estimating Interest Group Ideology." *Public Choice* 136(1): 69–86.
- McKay, Amy. 2010. "The Effects of Interest Groups' Ideology on Their PAC and Lobbying Expenditures." *Business and Politics* 12(2): 1–21.
- Moskowitz, Daniel J. 2021. "Local News, Information, and the Nationalization of US Elections." *American Political Science Review* 115(1): 114–29.
- Phillips, Connor Halloran, James M. Snyder, Jr, and Andrew B. Hall. 2024. "Who Runs for Congress? A Study of State Legislators and Congressional Polarization." *Quarterly Journal of Political Science* 19(1): 1–25.
- Poole, Keith T., and Thomas Romer. 1985. "Patterns of Political Action Committee Contributions to the 1980 Campaigns for the United States House of Representatives." *Public Choice* 47(1): 63–111.
- Poole, Kieth T, and Howard Rosenthal. 2007. *Ideology and Congress*. New Brunswick, NJ: Transaction.
- Rogers, Steven. 2016. "National Forces in State Legislative Elections." *Annals of the American Academy of Political and Social Science* 667(1): 207–25.

- Rogers, Steven. 2017. "Electoral Accountability for State Legislative Roll Calls and Ideological Representation." *American Political Science Review* 111(3): 555–71.
- Rogers, Steven. 2023. *Accountability in American Legislatures*. Chicago, IL: University of Chicago Press.
- Shor, Boris, and Nolan McCarty. 2011. "The Ideological Mapping of American Legislatures." *American Political Science Review* 105(3): 530–51.
- Shor, Boris, and Nolan McCarty. 2023. "Individual State Legislator Shor-McCarty Ideology Data, April 2023 Update." <https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/NWSYOS>.
- Squire, Peverill. 2017. "A Squire Index Update." *State Politics & Policy Quarterly* 17(4): 361–71.
- Stiglitz, Edward H., and Barry R. Weingast. 2010. "Agenda Control in Congress: Evidence from Cutpoint Estimates and Ideal Point Uncertainty." *Legislative Studies Quarterly* 35(2): 157–85.
- Thomsen, Danielle M. 2014. "Ideological Moderates Won't Run: How Party Fit Matters for Partisan Polarization in Congress." *Journal of Politics* 76(7): 786–97.

Thomsen, Danielle M. 2017. *Opting Out of Congress: Partisan Polarization and the Decline of Moderate Candidates*. Cambridge: Cambridge University Press.

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

How to cite this article: Handan-Nader, Cassandra, Andrew C. W. Myers, and Andrew B. Hall. 2025. "Polarization and State Legislative Elections." *American Journal of Political Science* 1–18. <https://doi.org/10.1111/ajps.12973>