

Issues or Affinity? How Voters Decide in Primary Elections*

John A. Henderson[†]
Dept. of Political Science
Yale University

Logan Dancey[‡]
Dept. of Government
Wesleyan University

Stephen N. Goggin[§]
Dept. of Political Science
San Diego St. University

Geoffrey Sheagley[¶]
Dept. of Political Science
University of Georgia

Alexander G. Theodoridis^{||}
Dept. of Political Science
University of California, Merced

May 20, 2019

Word Count: 9,586

*For valuable comments we thank Brian Box, Jacob Hacker, Greg Huber, Travis Johnston, David Jones, Kelly Rader, and Rob Van Houweling, as well as seminar participants at Yale. This work was funded by generous research support from the University of California, Merced and Yale University, and by National Science Foundation Award #1559125. Replication data happily provided upon request.

[†]<john.henderson@yale.edu>, <http://jahenderson.com/>, Institution for Social and Policy Studies, Yale University, 77 Prospect Street, New Haven, CT 06520-8209

[‡]<ldancey@wesleyan.edu>, <http://ldancey.faculty.wesleyan.edu/>, Government Department, Wesleyan University, 238 Church Street, Middletown, CT 06459

[§]<stephen.goggin@sdsu.edu>, <http://sgoggin.org/>, Department of Political Science, San Diego State University, 5500 Campanile Drive, San Diego, CA 92182

[¶]<geoff.sheagley@uga.edu>, <https://sites.google.com/view/geoffreysheagley>, Department of Political Science, University of Georgia, 180 Baldwin Hall, Athens, GA 30602

^{||}<atheodoridis@ucmerced.edu>, <http://www.alexandertheodoridis.com>, Department of Political Science, University of California, 5200 North Lake Rd., Merced, CA 95343

Abstract

While central to theories of polarization, surprisingly little evidence shows that partisans choose primary candidates based on divergent policies, above other considerations. In this study, we investigate how voters weigh various factors when deciding in primary elections, using novel conjoint experiments that fully manipulate the demographic, biographical and issue information of legislative candidates. We recruit 4,093 participants in the 2016 Cooperative Congressional Election Study, randomly assigning half to either rate same-party candidates on an ideological scale or indicate their preferences in a primary-like contest. We find voters perceive *and* choose candidates along coherent ideological lines, with greater support given to candidates taking *party-consistent* issue positions. We find evidence of ‘affinity’ voting on shared racial, gender or religious characteristics, though these effects are smaller, and unlike policy preferences, attenuate when primed to consider electability. Our results underscore the importance of policy divergence in driving vote choice in primaries.

1 Introduction

Primary elections are often cited as a major source of ideological divergence in American politics (Aldrich 1983; Aldrich and Alvarez 1994; Brady et al. 2007). Accordingly, policy-motivated voters screen out relative centrists in favor of more ideological candidates who are expected to faithfully advance the policy goals of the party, perhaps even risking electability (Rickershauser and Aldrich 2007). In spite of their centrality to theories of representation, relatively little evidence has been marshaled showing that voters choose primary candidates based on their issue positions (Bawn et al. 2017; Hirano et al. 2010). A central challenge in studying vote choice in primaries is that the field of candidates is typically distorted by strategic considerations that minimize policy disagreement among them. This is further compounded by the absence of powerful partisan signals voters depend on to make policy evaluations (Arceneaux and Vander Wielen 2017; Bolsen et al. 2014; Klar 2014; Marietta and Barker 2007). Lacking clear issue differences, primary voters may rely on other factors like demographics to base their decisions, using these to infer legislative priorities or otherwise differentiate candidates (Dolan 2008; McDermott 1998), which can mask the importance of ideological divergence.

A major alternative sees primary elections not as contests over policy, but between competing coalitions organized around group interest or ‘affinity’ (Bawn et al. 2012; Cohen et al. 2008; Campbell et al. 2011; Dolan 1998; Grossmann and Hopkins 2016; Jackman and Vavreck 2010; McConnaughy et al. 2010). Instead of demographics serving primarily as policy signals, partisans are expected to choose candidates who identify with or share membership in the different groups that make up the parties. Considerable research has uncovered evidence of such *affinity voting* in both primary and general elections, focusing mainly on the importance of shared race or ethnicity (Branton 2009; Boudreau et al. 2019; McConnaughy et al. 2010), gender (Badas and Stauffer 2019; Dolan 1998, 2008; Lawless and Pearson 2008; McDermott 1998), religious denomination or religiosity (Campbell et al. 2011, 2012), and economic or so-

cial class (Carnes and Lupu 2016). It is difficult to disentangle whether such voting is meant to promote the policy priorities of particular groups or to achieve broader symbolic or empowerment goals (e.g., expanding group representation or influence in the party), and these are not necessarily exclusive. Regardless of aim, partisans would decide using the sociodemographic cues made available about candidates, and not their platform of issue positions.

Though there is renewed scholarly interest in the politics of candidate nomination (Bawn et al. 2017; Boatright 2013; Cohen et al. 2008), few studies examine the way competition over policy or non-policy differences influences how votes decide in primaries. To better understand the competing considerations voters weigh in primary elections, we utilize a conjoint experimental approach that manipulates a large number of demographic, biographical, personality, and policy attributes of hypothetical candidates. The design allows us to assess the relative impact of multiple considerations on primary election choices simultaneously, and as a function of the policy attitudes or attributes of voters. Drawing on a sample of 4,093 respondents in the 2016 Cooperative Congressional Election Study, we randomly assign half of our participants to ‘rate’ pairs of same-party candidates by indicating which is more conservative (Republican candidates) or more liberal (Democratic candidates). The other half of our participants indicate their preferences between pairs of candidates from the same party in a primary election setting. Competing candidates are randomly generated from the set of demographic and policy attributes, providing us a measure of how those factors influence both the *ideological perceptions* and *nominating preferences* voters have about party politicians.

The results of our ideological scoring conjoint provide strong evidence that voters perceive candidates in coherent ideological terms, at least in aggregate. Importantly, these perceptions mirror expectations about the sorts of attributes and policy positions exhibited by real liberal and conservative politicians competing for office (Ahler and Sood 2017; Goggin et al. 2019; Rothschild et al. 2019), demonstrating a voter competence perhaps at odds with the more skeptical work on voter sophistication (e.g., Achen and Bartels 2016). Ideological signals gen-

erated by non-policy information (gender, race) are significantly muted relative to the policy items (endorsements, issues) in the experiment. This suggests that issue voting uncovered in our primary election frame is most likely due to policy signals from the issue items, and not ideological spillover from non-policy characteristics. Put differently, policy voting is made easier when voters can clearly observe ideological differences between alternatives, instead of relying on demographics alone as signals. From these results, we generate an ideology scores for each candidate evaluated in the primary election vignette, allowing a test for the impact of ideological perceptions on candidate choices generally, and when electability concerns are (randomly) primed.

Results from the primary election comparisons indicate that party identifiers base a substantial portion of their decisions of who to nominate on policy grounds. By far, the largest predictor of support is whether candidates take positions on a set of issues that are ideologically consistent with the attitudes of the party's voter base. Similar, but weaker, vote predictors include endorsements by party-affiliated interest groups. We then explore how policy support varies among subsets of ideological and cross-pressured partisans. We find that candidates do increasingly worse as they take more out-party positions, with losses among ideologues nearly always outpacing any gains from more centrist (or cross-pressured) co-partisans. Confirming past research, we also find evidence of affinity voting along gender, race, and religious cleavages, especially for Democrats. Yet, the impact of these and other non-policy dimensions is typically half that found for the policy information communicated in the experiment. We cannot dismiss this smaller effect is due to the way demographic signals are presented in a survey as opposed to how candidates' gender, race or ethnicity are emphasized in elections (e.g., Dolan 2010; Hopkins 2009). Nevertheless, the findings show that primary voters are capable of issue voting in general, and perhaps more likely to evaluate candidates' policy differences when these are included alongside group affinity cues.

Finally, we assess how priming electability alters voters' preferences over competing party

candidates, and observe how this interacts with policy or group voting. To do this, we randomly inform participants that the winner of the primary election will go on to compete in either a Safe Democratic or Safe Republican race, or that the general election will be a Toss Up. Consistent with other work on the group bases of the parties (Grossmann and Hopkins 2016), we find that Democrats appear to moderate their preferences when choosing candidates in competitive districts, while Republicans do not. In exploring this strategic moderation further, we find it does *not* vary based on policy attitudes among Democratic voters, but rather on whether or not voters share the *same non-policy* attributes as candidates. The biggest shift we uncover is a relative reduction in co-party support for black and Hispanic politicians among *non-white* Democrats, who choose white and co-racial candidates at similar rates when general elections are toss ups. Both Democrats and Republicans appear otherwise unwilling to endorse policy-moderates even at the risk of losing a general election. Overall, this study sheds new light on the competing factors that voters balance when nominating legislative candidates and clarifies the relative importance of policy and group affinity in these decisions.

2 Prior Work on Primary Election Voting

Many scholars perceive voting in primaries to involve better-informed, policy-demanders punishing candidates who tack too far to center. However, the existing evidence for this broad theoretical perspective is light given the prominent role primaries are thought to play in driving off-center policy-making in the U.S. (Aldrich 1983; Bawn et al. 2017; Hirano et al. 2010). The notable absence of party cues in primary elections provides an opening for party voters to evaluate candidates on other grounds, including policy and ideological agreement. But in practice, few primary contests produce major competition along policy lines, often focusing on differences in biography, personality or electability. Of course, the latter is observationally equivalent with strategic moderates refusing to run due to high costs and dim prospects, but

also accords with the view that primaries are essentially issue-free affairs.

Other researchers are skeptical that voters rely at all on policy when evaluating candidates, including in primary elections. In the absence of party cues, ideological reasoning is complicated (e.g., Arceneaux and Vander Wielen 2017; Bolsen et al. 2014; Klar 2014; Marietta and Barker 2007). And while primary voters are often thought to be better equipped to judge candidates in ideological terms than the mass public, recent work has found differences between primary and general election participants to be rather minimal (Sides et al. 2019). Like other voters, examples abound of partisans ignoring candidate policy positioning in presidential nomination contests (Lau 2013; Stone et al. 1992). Further, election scholars in recent years have grown particularly pessimistic about voter sophistication generally, questioning *whether* even primary-voting partisans possess meaningful policy attitudes or the ability to use these when selecting candidates (Achen and Bartels 2016). Rather, voters are thought to understand politics in terms of group interest or difference. Indeed, group status is often more readily discernible than policy disagreement, allowing those with shallow issue attitudes and limited political sophistication a lens to evaluate politicians.

Instead of ideological or policy agreement then, voters may alternatively rely on shared group affinities – like race, gender, or religion – in primary election contests (Branton 2009; Campbell et al. 2012; Jackman and Vavreck 2010). A growing body of research highlights the general importance of group affinity in influencing voter behavior across different electoral contexts (Badas and Stauffer 2019; Boudreau et al. 2019; Branton 2009; Campbell et al. 2011, 2012; Carnes and Lupu 2016; Dolan 1998, 2008; Jackman and Vavreck 2010; Lawless and Pearson 2008; McConnaughy et al. 2010; McDermott 1998). Shared identity may play a particularly important role in candidate nomination, where internal battles over the policy and non-policy priorities of the parties are often most intensely fought (Bawn et al. 2012,

2017; Grossmann and Hopkins 2016).¹ This view accords with theories of voters as ‘blind partisans’ motivated by social identity and group attachment, not policy preference (Achen and Bartels 2016). Though one need not abandon policy aspirations when choosing on shared identity. For instance, group-based cues may be more informative about the future policy actions of candidates, especially on issues that advance the ‘linked fate’ of racial, ethnic or other minority partisans (e.g., Branton 2009; McConnaughy et al. 2010).

The extent to which voters rely on ideology or shared group membership may also vary across voters, parties and election contexts. Though polarization accounts expect primary electorates to select ideologically divergent politicians, most primary voters are cross-pressured on some policies (Sides et al. 2019). An open question is whether these voters favor (dyadic) policy congruence over partisan consistency when given a choice. Cross-pressured voters may naturally select primary candidates who agree with them more on policy. Yet, some partisans may vote ‘strategically’ or give co-partisans leeway, discounting their own policy demands to advance the broader legislative or electoral goals of the party. A recent collection of findings also suggests that Republican and Democratic voters and activists view conflict in American politics differently, with the former peering through the lens of ideology (small vs. large government), but the latter reasoning along group (wealthy vs. middle class) interests (Grossmann and Hopkins 2016). This asymmetry has important consequences for primary voting – Republicans would be expected to choose candidates who are better aligned with conservative ideological thinking, while Democrats would prefer those advancing the goals of

¹In the 2008 Democratic presidential primary, for example, though Barack Obama and Hillary Clinton took almost identical policy positions, the contest was immensely heated, splitting Democrats along race and gender lines in candidate support (Jackman and Vavreck 2010). Campbell et al. (2012) find that Republicans were also split in the 2008 primary, with views about Mormonism shaping impressions of Mitt Romney against the rest of the field.

particular interests, and not a more general commitment to policy liberalism.

Finally, researchers have explored whether context, especially *candidate viability* in a general election, impacts who partisans support in primaries. Much classic work finds that primary voters weigh electability when deciding (Abramowitz 1989; Stone et al. 1992), though these studies rely on voter evaluations of viability, which likely proxies their preferences. Recent experimental work has sought to avoid this by manipulating information about candidate electability, mainly through randomizing early polling support (Rickershauser and Aldrich 2007; Simas 2016). Consistent with prior findings, this research suggests that voters may be “willing to sacrifice ideological representation for an increase in the likelihood of supporting the eventual winner” (Simas 2016, p. 286). Yet, a separate concern is that using ‘horse race’ polling to signal viability may be overly strong or interpreted deterministically, leading to overestimation (e.g., poll deficits as signs of lost causes). These may also center attention away from contextual factors about a district that impact electability (e.g., differentials in party registration), and towards individual qualities that predict success, but that voters also highly prefer (e.g., competence). Given this, manipulating information about the closeness of a district would better identify the impact of electability on voter choices, without (in)directly influencing perceptions of candidate quality or desirability.

Altogether, one expectation is that voters will favor policy-moderates in primaries when they believe the general election will be competitive. Another is that voters will avoid nominating candidates with particular attributes (e.g., gender, religion) perceived to hurt electability regardless of their positions (Doherty et al. 2019; Krupnikov and Piston 2015; McDermott 1998). We design our experiment with the goal of assessing how voters weigh trade-offs between electability and policy or non-policy preferences. An important feature, the design can also assess whether this moderation emerges on features shared by respondents, or on other dimensions, illuminating which candidates different primary voters are willing to abandon.

3 Data and Conjoint Design

We utilize novel conjoint experiments to assess which candidate characteristics shape vote choices in primaries, and whether the general election environment impacts these decision. Conjoint experiments are increasingly common in political science (Hainmueller et al. 2013), and have been utilized to study attitudes about immigration (Hainmueller and Hopkins 2015), voter associations with the party label (Goggin et al. 2019), and elite primary election strategy (Doherty et al. 2019). The design generates hypothetical primary candidates from a large array of relevant attributes. This ability to simultaneously vary several factors makes a conjoint design particularly well suited to our analysis of primary election voting.

Our study was embedded in two modules ($N = 4,093$) of the 2016 Cooperative Congressional Election Study (CCES).² Half of our respondents were asked to ‘rate’ four separate pairs of candidates by indicating which in each pair is more conservative (for Republican candidates) or more liberal (for Democrats). The other half was asked to indicate their choices for one of two candidates from the same party, presented as four separate pairs of competitors in a primary election context. In this latter *primary election* frame, party identifiers had a 3/4 probability of evaluating pairs of in-party candidates, and a 1/4 probability of evaluating pairs of out-party candidates.³ Participants in the *ideology rating* frame were assigned to rate Democratic or Republican candidate-pairs with equal (1/2) probability. A portion of respondents (2,482) completed both rating and preference tasks, one on the pre-election wave and the other on the post-election wave, thus performing 8 total paired comparisons. The remaining

²YouGov fields the CCES online in the weeks prior to and after Election Day. We do not use sampling weights in analysis, consistent with the recommendations of Miratrix et al. (2018).

³Once assigned within the *same* pre- or post-election survey, participants only evaluated candidates of the same party. Throughout, we include leaners as partisans, with pure Independents assigned to evaluate either Democratic or Republican candidates at equal (1/2) probability.

respondents (1,611) only completed one of the two tasks during the pre-election wave, thus performing 4 total paired comparisons. Using results from the ideological rating frame, we score candidates on a left-right dimension as the average of coefficients associated with the attributes in their profile. This measure can capture the ‘total’ effect of ideological signals on choices in the primary election arm. Stratifying on self-reports of general or primary vote participation recovers identical results, replicating minimal differences found elsewhere (Sides et al. 2019).⁴ Thus, we include responses for all participants in the analysis below.

In both frames, the characteristics of each candidate in the pair were randomly generated from a set of demographic and policy factors. Candidate gender, race, religion, occupation, personality type, endorsement information, and top-two issue positions were randomly assigned. Table I in the Online Appendix contains a list of full factors and the levels for each factor. Respondents under both frames were told that the information they were receiving in each profile was from a survey filled out by each candidate. Examples of the introductory screens and paired choice interfaces can be found in the Online Appendix as well.⁵ In addition to eliciting ideological perceptions or preferences, we also asked respondents to rate their favorability towards each candidate, and to indicate which is most similar to themselves.⁶

In designing the conjoint experiment, we took care to ensure that candidates seemed realistic while still providing enough variation for voters to make distinctions between competitors. As Table I in the Online Appendix shows, some characteristics (e.g., most group endorsements, certain issue positions) were party specific, so as to preclude conflicting or confusing information. Special attention was devoted to designing policy positions. We included twenty

⁴See Online Appendix Figures IX and X for results stratified by primary or general electorates.

⁵These are presented in Figures I – IV in the Online Appendix.

⁶All outcomes and candidate profiles appear on the same page. Given wording (‘most similar’, ‘favorable’) we doubt any anchoring in responses is biased to uncover issue over affinity effects.

priorities, each constituting either a liberal or conservative position drawn from one of ten issue areas. Three issue areas were party specific – reproductive rights/pro-life, gun control/gun rights, and LGBT rights/traditional marriage – so that Democrats only take the liberal position, and Republicans only the conservative one. Along the other dimensions, Democrats and Republicans could take either the liberal or conservative position on an issue, but not both. A key feature in the experimental design is that partisans are able to evaluate a broad array of candidates, including many types that would likely select themselves out of real primary competition due to expectations of defeat. Thus, we can observe the incentives ‘off-equilibrium’ candidates may face due to the ideological or non-policy motivations of primary voters.⁷

4 Evidence of (Non-)Policy Voting in Primary Elections

Given the limited evidence that primary voters weigh policy when deciding, our central task is to evaluate whether the issue and policy items in our experiment move voter perceptions and preferences of candidates in ways that meaningfully accord with their attitudes. Of course, there are many ways voters can choose on the basis of policy, including the use of non-policy attributes that signal ideological differentiation. Yet, it is exceedingly difficult in general to distinguish whether voters use those items as (unbiased) proxies for policy, or because they

⁷To allay concerns about the realism of the attributes, we show six randomly generated candidate profiles in Online Appendix Figure V(b), arrayed from left to right given their overall ideological score. Above this, in Figure V(a), are densities of the summary ideological locations for all Democratic (blue) and Republican (red) candidates. Both parties are scaled on the same dimension up to μ , the average difference in ratings between Democratic and Republican candidates with all-baseline attributes. We set μ to zero for presentation, though this does not influence any analysis as we restrict all comparisons within the same party.

indicate shared affinity.⁸ Throughout, we consider evidence of policy voting to be limited to voters being swayed to support one candidate over another due to greater agreement or consistency *on the issues*, either overall or on particular items. In that regard, we doubt issue items proxy non-policy attributes voters care about, especially the particular demographic cues presented alongside the issue positions. In reverse, even if some non-policy items convey additional ideological information above that contained in the endorsements or issue priorities, while simultaneously impacting preferences, we refrain from interpreting this as evidence of policy voting. However, the extent to which voters perceive additional ideological information from the demographic cues *can* provide an upper limit to how much these non-policy factors would impact policy voting, when included alongside the other items.

4.1 Ideological Signaling from Candidate Attributes

We first investigate which of the policy and non-policy attributes influence the relative spatial perceptions of same party candidates in the ideology rating frame. If none of the items, and especially the issue priorities, systematically move ideological evaluations of politicians, we should be skeptical that voters are capable of reasoning effectively about issues or able to choose based on ideological or policy differences. Such incoherence would not necessarily preclude that individual voters select candidates who advance particular issues over others, though it does indicate an inability to grasp how issues relate to each other. This would seriously undermine the claim that primary voters can constrain politicians to off-center positions. Thus, a pre-condition of divergent policy voting is whether or not primary voters possess coherent perceptions of the ideological information signaled by a candidate's positions, with

⁸If factors like race or gender impact ideological perceptions when policies are included, these attributes likely capture information about issues excluded from our experiment, a positive (or negative) valence about those characteristics, or partisan bias (Ahler and Sood 2017).

minimal signaling emerging from non-policy characteristics. Absent such coherency, we can discount that voters consider ideological differences when choosing candidates.

As described above, the multi-factor design allows us to estimate the relative impact of each of the randomized factors on liberal and conservative ratings of same-party candidates. These results are presented in Figure 1, which pools results across Democratic and Republican candidates and party identifiers.⁹ Larger (absolute) coefficients indicate attributes that have greater marginal influence on perceptions a candidate is ideologically extreme, relative to a series of baseline categories centered at zero. Since participants provide multiple comparisons, standard errors throughout our analysis are clustered on the respondent. Additionally, the two issue priorities are combined into a single issue indicator for estimation.

As shown in the figure, the factors with the largest impact on ideological perceptions are the twenty policy positions and the endorsing interest groups. Endorsements of Democratic candidates by left-leaning groups, including reproductive rights, civil rights, environmental, gun control and labor organizations lead both Democratic- and Republican-identifiers to perceive endorsed candidates as more liberal, relative to those supported by business or veterans' groups. We uncover a similar finding for Republican endorsements from right-leaning groups with the most conservative endorsements coming from the Tea party, Christian and gun rights organizations. Similarly, candidates taking liberal or conservative issue positions are rated accordingly in the pairwise evaluations. Positions like defending LGBT rights, expanding government assistance, and providing illegal immigrants a path to citizenship rate as the most liberal, in contrast with defending traditional marriage, protecting the unborn, and defending gun rights, which are the most conservative.

⁹We recover virtually identical estimates stratifying by party of candidate and respondent, indicating considerable agreement about the ideological signals of our experimental attributes.

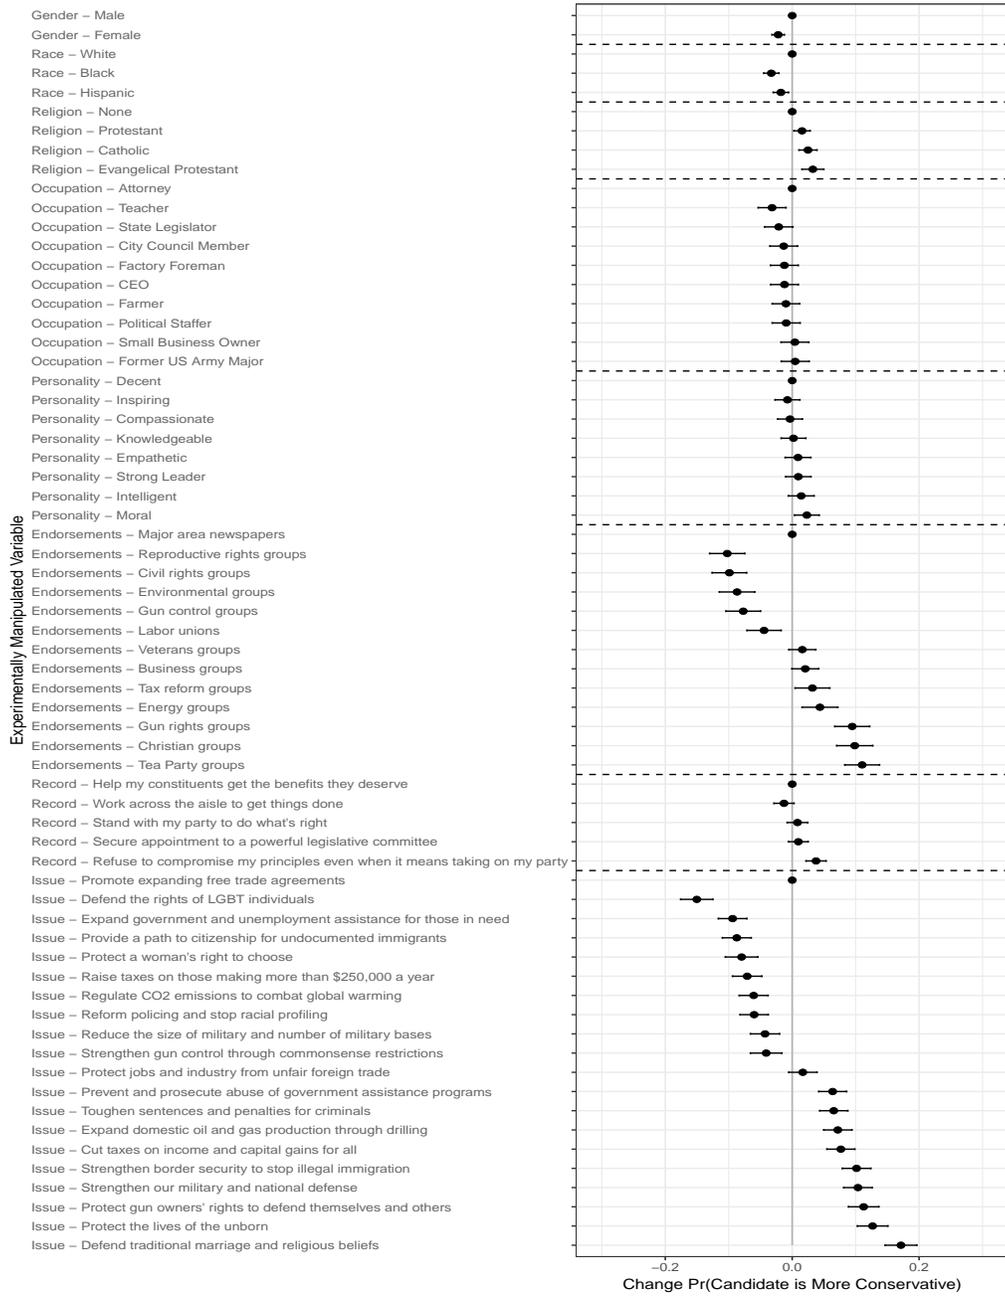


Figure 1: **Overall Rating of Candidate Ideology:** Estimates are OLS, regressing relative ideological ratings on each factor, pooling over all respondents evaluating Democratic and Republican candidates. Standard errors are clustered on the respondent, with error bars displaying 95% confidence intervals. Estimates with no error bars are the excluded levels of each experimental factor. Issues are coded as present if they were in either the first or second policy priority for the candidate. All variables are coded 0-1, with the dependent variable coded as 1=More Conservative (Less Liberal), 0=Less Conservative (More Liberal).

In contrast, the personality and occupation non-policy attributes have virtually no ideological discrimination. However, gender, race, ethnicity and religion clearly signal some ideological information to voters. Women, minority politicians, non-Evangelicals and non-Catholics tend to be viewed as more liberal (less conservative) than white men of Christian faith. This finding replicates other work suggesting that policy-minded voters may use demographic attributes, at least in part, to make inferences about within-party ideological differences (and not just partisanship) when assessing competing candidates (Goggin et al. 2019; Lawless and Pearson 2008; McDermott 1998). Though again, we cannot discern whether these demographic effects emerge from inferences about how these affinities correlate with other policies not included in the design, a projection bias from those with shared or different attributes, or some other source (e.g., Ahler and Sood 2017; Rothschild et al. 2019).

Importantly, the ideological signals from demographic cues are substantially muted in comparison to policy attributes. This latter finding is crucial as it indicates voters organize politicians, attributes and positions in meaningful ideological terms, including an apparent (or implicit) awareness that non-policy attributes are weak predictors of policy orientation once issue positions are clarified. More generally, this finding of policy coherence is especially notable, not only in delimiting the sort of ideological inconsistency we should expect of voters (Achen and Bartels 2016), but also that it arises when the party label is fixed. Respondents display a meaningful grasp for how policy battles split, rather than just separate parties, even if their attitudes do not necessarily map onto the same cleavages. In sum, our results verify the minimal conditions for policy voting hold.

4.2 Candidate Attributes Shaping Voter Support

Next we assess how policy and non-policy factors included in our study shape vote choice in a primary context. Figure 2 displays results from the candidate choice conjoint experiment.

Most respondents evaluated same-party candidates, however 25% of our sample saw candidates from the opposite party. We focus our initial attention on the former primary-like electoral setting. The first column of Figure 2 shows the marginal effect of each factor on vote choice for Democratic and Republican respondents who evaluated same-party candidates. (Across the figures we order how attributes appear based on their ideological discrimination uncovered above in Figure 1.) Here we find strong evidence that policy positions shape preferences between co-party politicians. Democrats are more likely to support Democratic candidates who adopt traditionally liberal policies, like protecting LGBT rights and raising taxes on the wealthy, while Republicans favor co-party candidates who hold typical conservative positions, like defending gun rights and border security.

Generally, the above issue voting emerges from politicians taking party *consistent*, rather than extreme positions. There is little apparent benefit (or cost) for taking the most ideologically off-center positions (e.g., LGBT rights, government assistance, pro-life, traditional marriage), relative to other partisan issues. Yet, two interesting exceptions to party-consistency stand out. Republicans give greater support to candidates who promise to raise taxes on the wealthy compared to those promoting free trade, though less than the other conservative policies. Somewhat similarly, Democrats are less supportive of candidates advancing a long-standing goal of liberals to reduce the size of the military. Group endorsements also shape vote choices in expected ways. Democrats (civil rights, pro-choice, environment) and Republicans (Tea Party, Christian, gun rights) earn greater support from being linked with stereotypical interest organizations that work alongside each party. Finally, we uncover asymmetrical evidence about the impact of candidates' major legislative goals. Democrats negatively evaluate those who 'stand with their party' or seek committee appointments, relative to helping constituents. Republicans appear equally indifferent to all such aims.

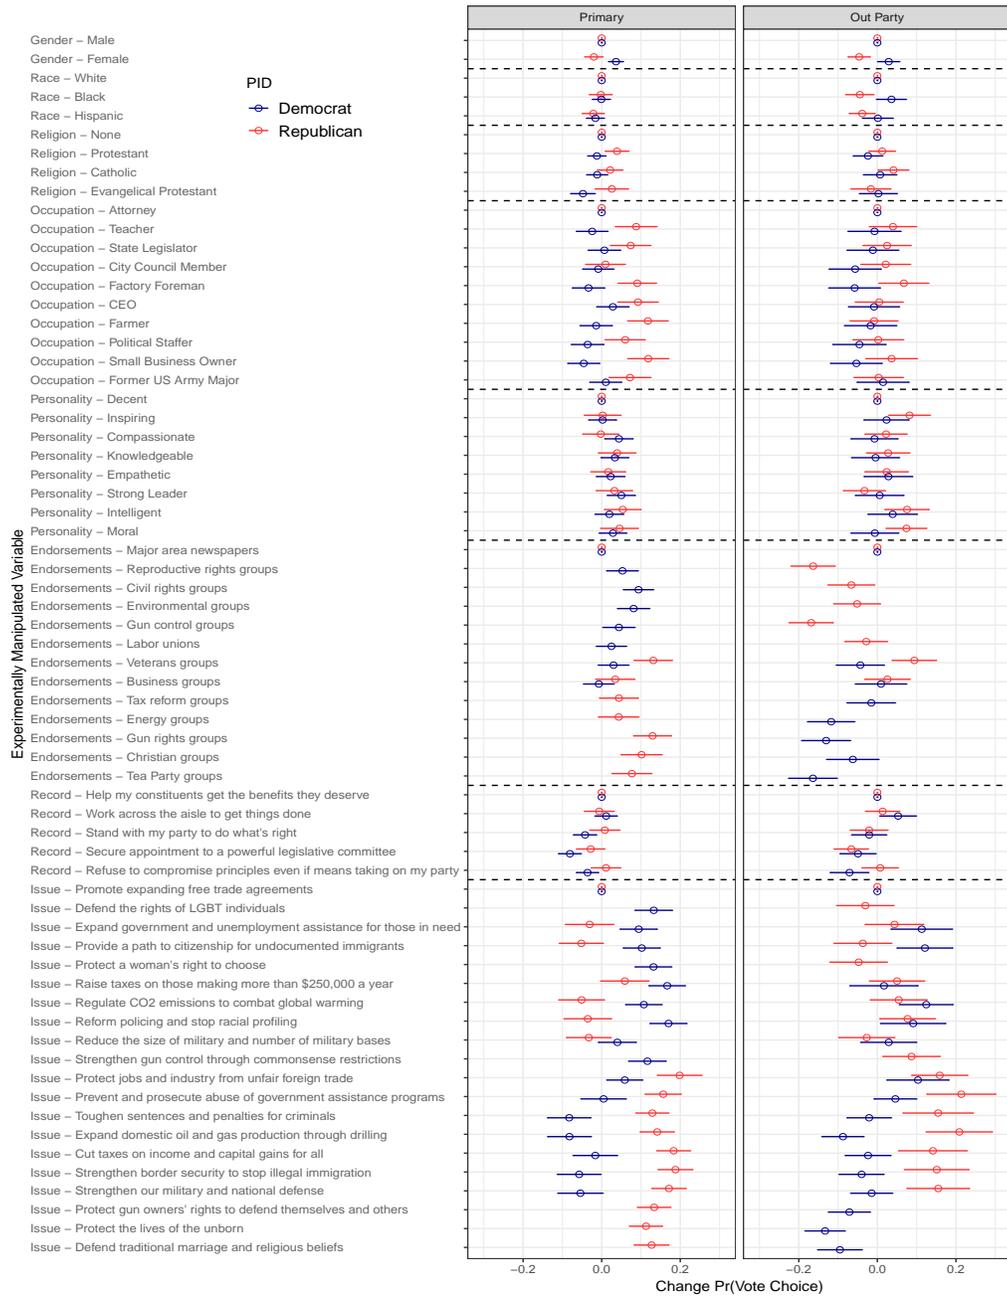


Figure 2: **Impact of Candidate Attributes on Support in Primary Election Settings:** Estimates are OLS, regressing vote support on each factor, stratifying by Democratic (blue) and Republican (red) PID for in- and out-party contests. Standard errors are clustered on the respondent, with error bars displaying 95% confidence intervals. Estimates with no error bars are the excluded levels of each experimental factor. Issues are coded as present if they were in either the first or second policy priority for the candidate. All variables are coded 0-1, with the dependent variable coded as 1=Supported Candidate, 0=Unsupported Candidate.

Turning to non-policy attributes, we see little evidence that personality traits drive support among co-partisans. We observe some occupation effects. Republicans are less supportive of attorneys and city council members, while Democrats dislike small business owners and political staffers. Of the core demographic cues, there is little evidence that race or ethnicity drives *overall* support for candidates in either party. Gender noticeably matters, with Democrats preferring female co-partisans, and Republicans penalizing them. This is consistent with how female candidates are rated in our experiment, as significantly more liberal (less conservative) than their male co-partisans (Lawless and Pearson 2008; Thomsen 2019). Though this may reveal polarized preferences for women in office (Dolan 2010; Sanbonmatsu 2002). Voters also have polarized views of Evangelicals, with Democrats clearly disfavoring them (though evidence of a Republican preference is weaker). Both parties see Evangelicals as more conservative, which may partly drive the differences in candidate support (Campbell et al. 2011).

Mirroring the primary election condition, out-party estimates in the second column of Figure 2 indicate that voters prefer opposition partisans to be moderate and atypical. Democratic voters are again more likely to support female Republican candidates, while Republicans prefer male Democrats. Here we see voters clearly diverging in their support for minority candidates. Democratic voters favor black Republicans. Conversely, Republican identifiers choose white Democrats at much higher rates than both Hispanic and black politicians. Issues operate in analogous fashion. Out-partisans reward politicians who defect from party orthodoxy by taking counter-typical policy positions, and penalize both ideological consistency and stereotypical group endorsements. Similar to other observational and experimental work, our findings suggest that cross-over primary voting is largely sincere, and not a strategic effort to endorse the least electable among partisan opponents (e.g., Blackwell and Calcagno 2019).¹⁰

¹⁰A general concern in candidate choice experiments is that these miss some of the strategic voting that comes with real electoral competition. One implication is that we could over-

As described above, an important factor potentially driving support for candidates is their shared affinity or policy agreement with voters. The above models are estimated unconditionally in each sample of Democratic or Republican party identifiers, all of whom see some differences between themselves and the candidates on offer in the vignettes. We turn here to an analysis of the impact each conjoint item has among those who share the attribute or policy attitude as measured in the CCES. Demographic measures for gender, race and religion are from self-reports on those survey items. For the most part, issue attitude questions in the CCES do not precisely map onto the policies or endorsing groups included in the conjoint experiments. Instead, we take all survey questions that gauge opinion on each of the ten issue areas, and scale these to measure separate liberal or conservative attitudes for each policy area.¹¹ For instance, those with liberal attitudes on the environment are assumed to support the liberal environmental statement (regulate CO2) and oppose the conservative one (expand

estimate the importance of sincere preferences, which may (or may not) distort the relative weight given to (non-)policy cues. We address this by priming district competitiveness to see how electability alters in- and out-party choices. We find this has modest effects on the former and no impact on the latter, which affirms other research on the question (Blackwell and Calcagno 2019; Rickershauser and Aldrich 2007).

¹¹See Tables II and III in the Online Appendix for details on items used to measure agreement and affinity. Personality and occupation items are excluded as the CCES did not include reliable measures of either. To construct policy agreement indicators, we assign each relevant CCES question to an issue area corresponding to the experiment (in Table II). On each of these questions, we set the conservative position to be 1 and the liberal position as -1. We average over item responses in each issue area to produce an issue-specific liberal-conservative scale. We set policy agreement to 1, when the candidate takes the liberal (or conservative) position and the participant has an overall liberal/negative (or conservative/positive) score on the same issue, and 0 otherwise. Given possible measure error (e.g., issues differ in number

domestic oil) in our experiment. Though imperfect, any measurement error from the use these issue-specific summaries is likely to attenuate the strength of policy voting we recover relative to shared non-policy voting, as the latter are measured with much greater accuracy.

These results are shown in Figure 3. In the models we *can* interpret effects for the baseline category variables, as these coefficients all estimate interactions – the impact of having the same characteristic or policy attitude as the candidate, compared to either not having the feature or that feature not appearing. Reflecting work on co-ethnicity and co-religiosity the largest non-policy effects of shared affinity emerge for black Democrats and Evangelical or Catholic members of both parties. We find inconsistent evidence that shared policy attitudes are driving the impact of endorsements on candidates, though this may reflect non-congruence between how we map attitude questions to endorsing groups. We again find strong evidence for policy agreement driving support for candidates taking party-consistent positions. Democratic identifiers who take liberal positions on the different issues areas reward Democratic candidates who do so as well, and punish those who moderate. This is again mirrored by Republicans who favor co-party candidates who offer party-consistent positions, and disfavor those who deviate. Yet, we find little evidence of sustained policy support for Democratic (Republican) candidates taking conservative (liberal) positions among cross-pressured partisans with similar out-party attitudes on each issue. We explore this result below, but generally uncover weak evidence that centrist (or cross-pressured) partisans reward policy moderation in primaries, illuminating an

and polarizing nature of questions), we caution against interpreting the agreement findings across the issues, focusing rather on overall differences between consistent and inconsistent issues, and between issues and affinities. Notably, if measurement is attenuating the small-to-null effects of shared attitudes on party-inconsistent issues for one party, it must also be attenuating the much larger consistency effects for the other party. Hence improving measurement would not change these relative effects.

important way two-stage elections can constrain politicians and drive polarization.

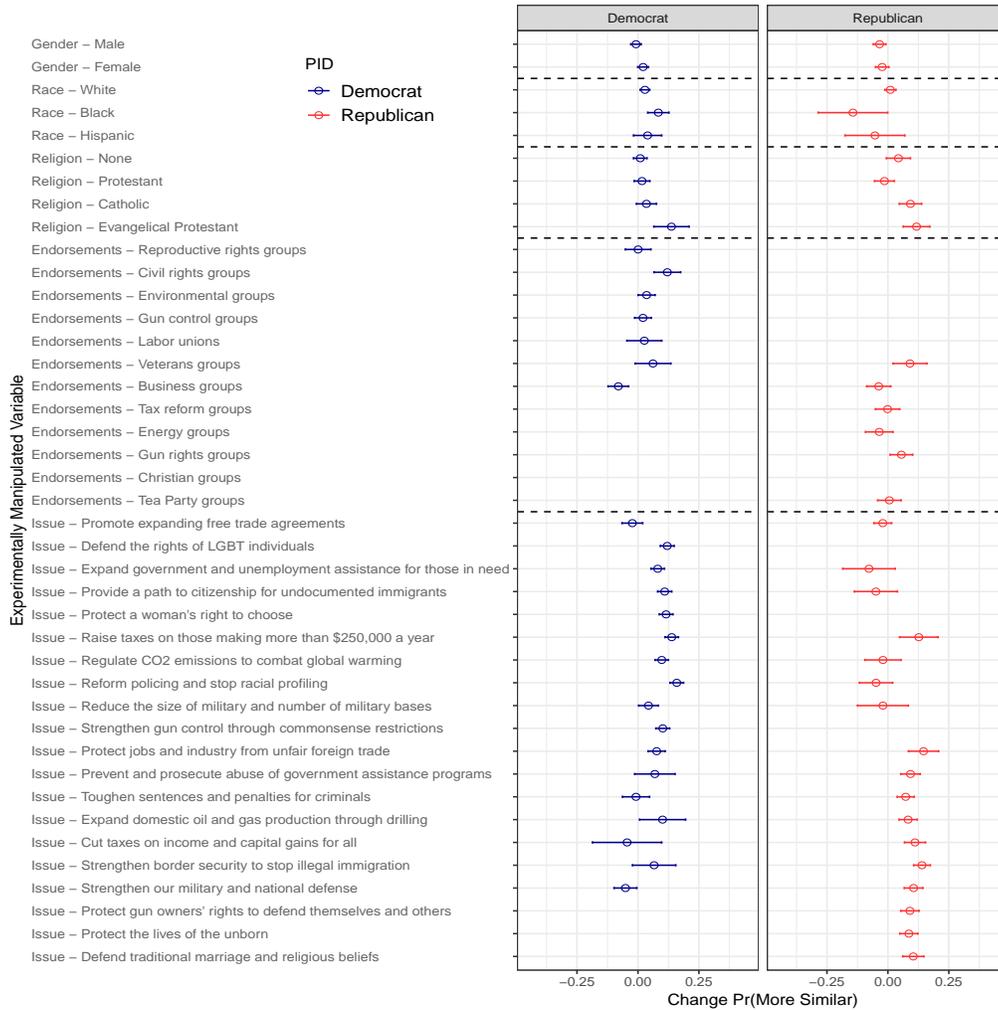


Figure 3: **Impact of Shared Attributes and Policy Agreement on Primary Vote Support:** Estimates are OLS, regressing vote support on an interaction between factors and indicators measuring shared attributes or issue agreement between participants and candidates. Endorsements are interacted with issue-specific measures of liberal/conservative attitudes, and not group membership indicators. Results are stratified by Democratic (blue) and Republican (red) PID for in-party contests. Standard errors are clustered on the respondent, with error bars displaying 95% confidence intervals. Factors without natural affinity or agreement items are excluded. No levels of the included factors are excluded since these effects are identified through the interactions. All interactions are coded 0-1, with the dependent variable coded as 1=Supported Candidate, 0=Unsupported Candidate.

4.3 Any Moderating Effect of Electability?

The above findings indicate that partisans substantially decide on the basis of issue agreement in primary election settings, though almost exclusively along party consistent policies. Shared affinities also matter in these decisions, but tend to be smaller in magnitude and less systematic. We now turn to whether concerns about viability alter voter preferences in primaries. To assess how voters weigh electability, we first utilize a summary measure of candidate ideology, scored using coefficients from the ideological ratings in Figure 1. We include this score in an interaction model, to estimate how the impact of candidate ideology influences voter choices overall, and for each level of a district competitiveness prime. We inform each participant that the series of primary contests they evaluate will eventually result in a nominee running in a Safe or Toss Up general election. (We also include a control condition excluding competitiveness information.) From past work on candidate viability, we expect voters to endorse policy moderates to run in toss up races. Though there is less work on how minority politicians fare in electorally competitive settings, we have some reason to suspect partisans favor non-minority politicians due to concerns about racial or gender backlash (Krupnikov and Piston 2015; Lawless and Pearson 2008). Both follow the expectation that a candidate's attributes or positions have greater chance of deciding close general election outcomes.

Initial results from this analysis are presented in Figure 4. Positive (negative) coefficients indicate more conservative (liberal) candidates are chosen in each experimental condition (toss up, safe, control), for both in- and out-party comparisons. The findings partially confirm our predictions above. Democrats, but not Republicans, are more likely to endorse (apparent) ideological moderates to run in toss up, relative to safe Democratic or Republican seats (and compared to the no-information condition). Yet, Republicans appear to be unswayed by this sort of electability concern. We again see little evidence of strategic behavior by out-partisans supporting policy extreme opponents to help their preferred candidates win in competitive

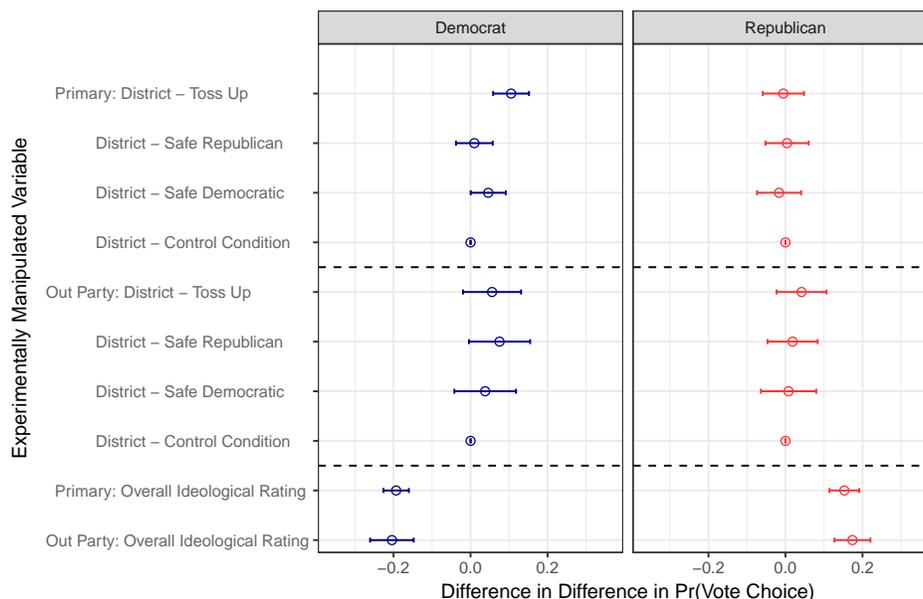


Figure 4: **Moderating Effect of Competitiveness on Support for Ideological Centrists:** Estimates are OLS, regressing vote support on the interaction between estimated candidate ideology and a district competitiveness prime (0=Control, 1=Safe Democratic, 2=Safe Republican, 3=Toss Up), for in- and out-party contests. Impact of overall candidate ideology and vote support are also presented. Results are stratified by Democratic (blue) and Republican (red) PID. Standard errors are clustered on the respondent, with error bars displaying 95% confidence intervals. Estimates with no error bars are the excluded experimental factor. The dependent variable is coded as 1=Supported Candidate, 0=Unsupported Candidate.

racers. In Figure 5, we break apart this competition effect item-by-item. In the figure, we present coefficients for the impact of going from a toss-up to uncompetitive race, interacted with each conjoint attribute. A positive coefficient means candidates with that attribute are supported more frequently in the competitive than the non-competitive conditions, with negative estimates capturing the reverse. As can be seen, the ‘moderation’ observed above is driven mainly by Democratic voters reducing their relative support for black and Hispanic candidates, rather than picking centrists with inconsistent issue positions. Figure 5 also stratifies results by white and non-white (black or Hispanic) respondents. Perhaps surprising, we see this competition effect is driven by *minority* Democratic voters adjusting their choices given an apparent expectation that co-racial politicians will face steeper challenges in a close

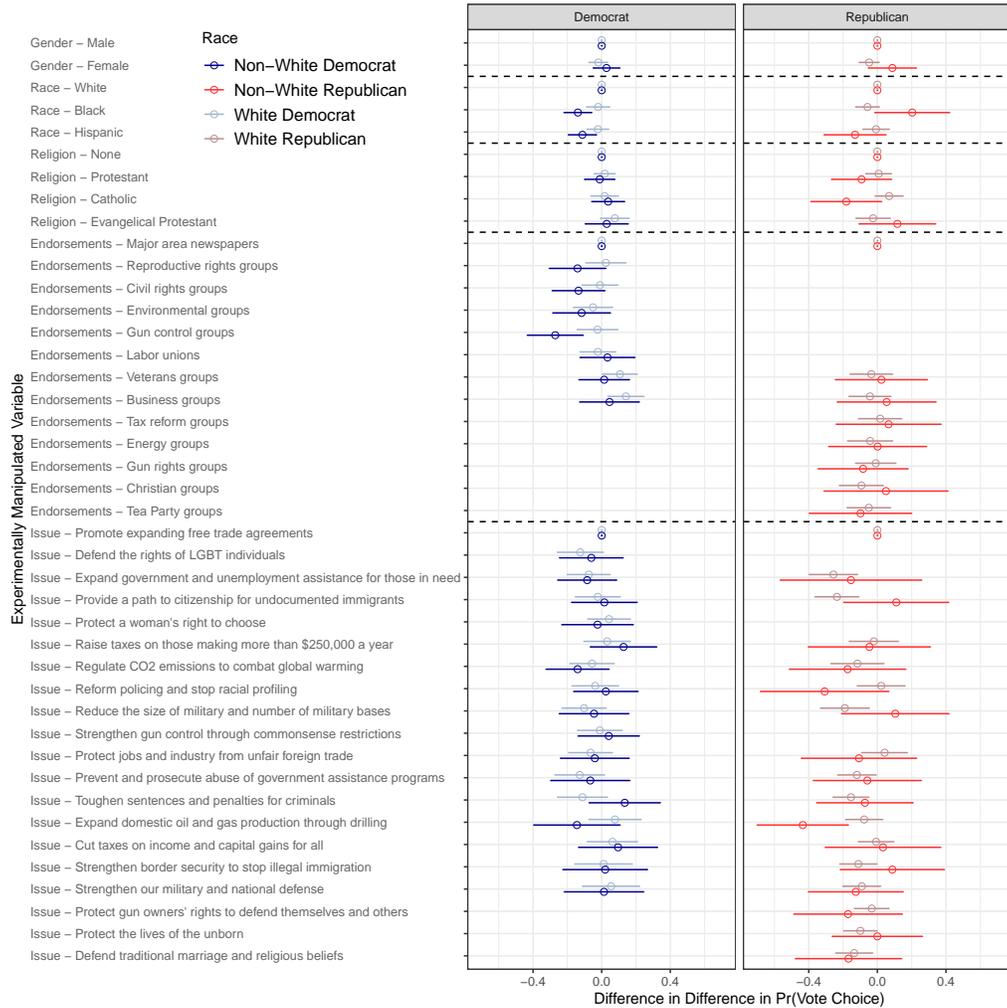


Figure 5: **Moderating Effect of Competitiveness on Primary Vote Support, Stratified by Race of Participant:** Estimates are OLS, regressing vote support on the interaction between each factor and a district competitiveness prime (0=Safe, 1=Toss Up), for in-party contests. Results are stratified by White and Non-White voters by Democratic (blue) and Republican (red) PID. Standard errors are clustered on the respondent, with error bars displaying 95% confidence intervals. Estimates with no error bars are the excluded levels of each experimental factor. Issues are coded as present if they were in either the first or second policy priority for the candidate. All variables are coded 0-1, with the dependent variable coded as 1=Supported Candidate, 0=Unsupported Candidate.

general election.¹² Notably, there is no such dampening effect among white Democrats or

¹²This effect is rather modest. In the toss-up condition, black and Hispanic voters support white (49.2%) and non-white (49.6%) candidates at virtually identical rates, while in safe

Republican voters for racial minority candidates, nor for female candidates of either party.

A striking aspect of the findings is that electability does not moderate the issue voting of partisans in primary-election contexts. In spite of clear expectations drawn from past research, including survey and lab experiments, voters seem unwilling to endorse policy moderates when primed to consider how this might impact general electoral success.¹³ Combined with the above findings, this result suggests that primary voting centers substantially around issue-demanders rewarding party-consistent positions with little willingness to sacrifice policy for electability. We do observe asymmetry among the parties in a willingness to alter candidate support given concerns about candidate viability, with Democrats adjusting who they choose. But this involves a perhaps unexpected form of ‘moderation,’ minority voters penalizing co-racial and co-ethnic politicians putatively due to non-policy considerations.

5 Who Rewards or Penalizes Policy Moderation in Primaries?

We show above that partisans are more likely to support co-party politicians who take ideologically consistent positions on important issues. The effects appear stronger than group-based demographic voting, and persist when voters are primed to think about competitiveness in a follow-up general election. In combination, the evidence suggests an important way

districts this shifts to 46.6% and 51.7%, or a relatively change of 4.7 percentage points.

¹³Identifying strategic behavior in surveys is challenging. Nevertheless, we see non-policy adaption to a competition prime, and other scholars uncover policy moderation using weak-to-no incentives in survey and lab experiments (Blackwell and Calcagno 2019; Simas 2016). An alternative is to see if voters from safe or competitive districts support different candidates in our experiment, having been exposed to different strategic environments. Using Cook Report ratings for 2012 to 2016 House races, we find that the competitiveness of a district where a voter resides does not correlate with their vote choices in our experiment.

primaries can elicit polarization – when voters are given clear choices about the (non-)policy differences between candidates, they punish moderates *even when* electability is at stake.

We explore this finding in more detail here. First, we replicate the above analysis, stratifying on whether party identifiers are ideologically consistent or cross-pressured using a summary attitude score. This summary is produced averaging over the same CCES attitude questions along the ten policy areas used in the above issue agreement analysis in Figure 3. ‘Cross-pressured’ voters are those whose attitude score is towards the center of their party’s distribution, while ‘consistent’ partisans have scores to the extremes, using the median score by PID as cutoff.¹⁴ Figure 6 presents results from this analysis. The major finding is that policy inconsistent partisans exhibit far less discrimination in candidate preferences based on any policy or non-policy attributes. Core candidate demographics like gender, race or religion, appear to elicit little variation in candidate preferences by both Democratic and Republican identifiers. Looking at coefficients for the policy items, these voters do reward more ideologically consistent candidates. However, the magnitude of policy voting is rather muted. This is in clear contrast with the policy *and* non-policy voting displayed by consistent partisans. For these voters, we see virtually all the major effects emerging as uncovered in the above analysis. Consistent Democrats show some preference for women and minority candidates, while punishing Evangelicals. Consistent Republicans choose less on the basis of demographics, though somewhat favor Catholic and mainline Protestants over Evangelical politicians. We again see strong evidence of policy voting among consistent partisans, who clearly reward politicians for taking party-consistent issue positions where they mostly agree.

Next, we look at how overall levels of candidate support change when politicians take two,

¹⁴We recover similar results using the 25th percentile as cutoff to determine cross-pressured voters, and in using coefficients in Figure 1 to produce a similar voter ideology score based on available CCES attributes and issue attitudes items.

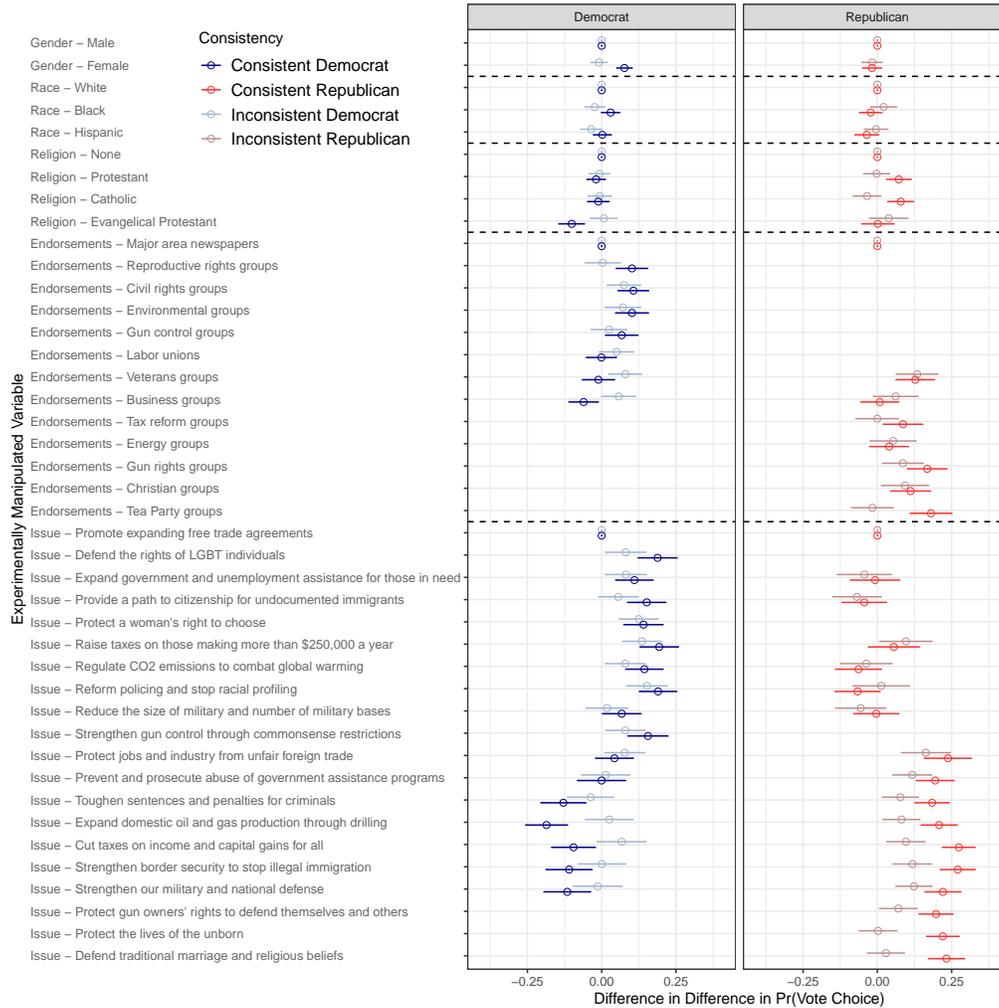


Figure 6: **Moderating Impact of Voter Consistency on Primary Vote Support:** Estimates are OLS, regressing vote support on the interaction between each factor and a measure of voter policy consistency, based on a summary score of liberal-conservative attitudes on ten issues from questions in the CCES. *Inconsistent* partisans have summary scores to the center of the party median, and *Consistent* partisans have scores in the extremes. Results are stratified by Democratic (blue) and Republican (red) PID for in-party contests. Standard errors are clustered on the respondent, with error bars displaying 95% confidence intervals. Estimates with no error bars are the excluded levels of each experimental factor. Issues are coded as present if they were in either the first or second policy priority for the candidate. All variables are coded 0-1, with the dependent variable coded as 1=Supported Candidate, 0=Unsupported Candidate.

one or zero party-consistent positions for partisan voters who agree with all, some or none of those presented policies. For this, we use the above policy affinity measures featured in Figure 3 (and described in Online Appendix Table II), which capture whether participants agree with

Table 1: **Rate of Candidate Support by Levels of Agreement and Number of Party Consistent Positions**

	Agree Both	Agree One	Agree None	Avg. %
Primary				
Both Consistent	0.598	0.521	0.425	0.563
One Consistent	0.464	0.404	0.380	0.420
None Consistent	0.327	0.313	0.198	0.285
Out Party				
Both Consistent	0.524	0.449	0.347	0.415
One Consistent	0.538	0.559	0.482	0.544
None Consistent	0.705	0.608	0.450	0.654

Note: Cell values are percent support for the candidate taking 2, 1, or 0 consistent issues among primary and out-party voters who agree with both, one, or none of the positions. Avg. % is the weighted percent of support over the three levels of agreement for each level of issue consistency.

each of two issue positions randomly included in the candidate profiles shown to them. Thus, this agreement measure is based solely on the two particular positions presented for each candidate profile, ignoring all other information about voter attitudes (e.g., cross-pressured or not on any other issues). Candidate consistency is determined by whether Democratic candidates take the liberal position, and Republicans the conservative one, on the issue items presented. Results are presented in Table 1, pooling Democrats and Republicans (with Independents excluded). Cells in the table display rates of support for candidates at decreasing levels of party consistency, for each level of issue agreement among voters. The last column (Avg. %) shows the overall rate of support for candidates at each level of policy consistency, that is, averaging over the (weighted) proportion of voters across the levels of agreement.

The main aim is to see how vote support changes when candidates take out-party positions, including among voters who are similarly cross-pressured, and thus agree with those counter-typical policies. Not surprising, voters are more likely to support candidates they increasingly

agree with on presented issues. This holds for both in- and out-party comparisons, and at each level of candidate consistency. Yet, candidates always perform worse among co-partisans when taking counter-typical positions, and the loss from this inconsistency nearly always out-paces gains made in appealing to cross-pressured co-partisans. Voters who agree with co-partisan candidates taking two counter-typical (none consistent) positions, support them at a rate of 33%. However, candidates can always do better than this, even among those who disagree with them, when they take an additional party-consistent position. For instance, co-partisans support 38% of politicians taking one typical and one counter-typical issue, and 43% of those taking two party-consistent positions, even when they disagree on both policies. All things equal, taking two policy consistent positions is clearly the optimal strategy to maximize vote support, at least among this sample of partisans on these particular issues.

Rates of support are also shown for party identifiers assigned to both the in- and out-party experimental condition to illuminate how much electoral benefit may arise when defecting on policy to appeal to out-partisans in open-primary settings. In the out-party condition we uncover a similar finding as above – support grows for politicians who take additional counter-typical positions, especially for out-partisans who agree with those policies. An interesting question is, under what conditions would a candidate do better in a primary by advancing one or two out-party policies? Clearly, if candidates are competing entirely for co-partisans in a closed primary (again ignoring Independents), we show they cannot do better by moving to center. In an open-primary setting, what proportion of voters would have to be out-partisan for politicians to do better overall in defecting on one or two issues, given the above findings?

A way to answer this is to compare the overall support (Avg. %) for candidates taking two, one and zero consistent positions, combining the in- and out-party results at varying weights. Define the proportion of in- and out-party voters as $1 - p$ and p . Overall support for politicians taking two consistent positions is given by $0.563(1 - p) + 0.415p$, for one consistent policy $0.420(1 - p) + 0.544p$, and so on. We solve for p , using the inequality $0.420(1 - p) + 0.544p >$

$0.563(1-p) + 0.415p$. As seen, more than 53% of primary voters must be out-partisan for there to be electoral advantages to taking one counter-typical policy given the experimental findings (with identical results for zero consistent positions). Many factors unique to the experiment drive this particular estimate, warranting caution in how far we extrapolate the result to real world settings. Nevertheless, the finding points to a substantial obstacle facing candidates who would like to tack to center, due either to sincerely held moderation or a strategic calculation to appeal to cross-pressured co-partisans and out-party voters in primary or general elections.

6 Discussion

The findings from our candidate conjoint experiments indicate that policy considerations may play a central role in primary election decisions. Counter to the growing skepticism about democratic accountability, voters in our study are able to draw sensible, coherent conclusions about the ideological orientations of competing candidates from their issue positions and non-policy attributes. In particular, policy rather than demographics provide the clearest signals of ideological difference among co-party candidates. Further, we find that voters use issue positions when deciding which candidates to support in our hypothetical primary match-ups. And while policy agreement between candidates and voters increases the likelihood of support, candidates always do better when taking party consistent positions rather than appealing to cross-pressured co-partisans or out-party voters. We uncover some evidence that voters choose candidates with whom they share group membership along gender, race and religious lines, though policy agreement substantially outweighs group affinity in driving voter preferences in our experiments. Indeed, preferences for policy divergence persist even when voters are primed to consider candidates' electability, though the impact of shared affinity returns to baseline in competitive races.

Our study has important implications for how primary elections can drive polarization and

policy divergence through candidate selection. There is ongoing disagreement about whether, and under what conditions, primary electorates select ideologically extreme or party consistent politicians, and thus screen out relative centrists (e.g., Brady et al. 2007; Hirano et al. 2010). A core concern in much past research is that the strategic entry and exit decisions of politicians makes it difficult to isolate the effect that primary elections have on candidate competition, including how incumbents (re)position to defeat or stave off challengers, or how primary voters would evaluate a fuller range of potential contenders. As consequence, primary contests tend to feature ideologically similar candidates. Observational analyses are generally unable to disentangle whether this lack of ideological diversity is due to voters being unmotivated by policy differences or their policy-motivations deter moderates from contesting primaries at all. Our design makes an important contribution to the study of vote choice in primaries. By presenting candidates with meaningful differences on issues that could emerge, but are often muted in real primaries, we can observe how this policy differentiation impacts preferences, alongside important non-policy considerations. Thus, an important feature of our study is the ability to see how primary voters would evaluate (non-)policy differences among the field of candidates that might run when strategic entry costs are low. Moreover, this can help explain how such costs distort the types of challengers who ultimately contest primaries.

Additionally, the findings help us understand how partisans may behave as the strategic forces behind primaries change, and especially as these heighten intra-party conflicts. Ongoing nationalization and polarization of the American parties is reducing general election competition across districts, while simultaneously increasing competitive pressure within primaries. This may lower the threshold for policy divergent or demographically diverse candidates to enter, something perhaps on display in recent presidential nomination battles. One prediction from our findings is that these contests will largely pivot around which candidates are seen advancing party consistent policies on the issues. Given our issue items, we cannot speak to the impact of advancing ideologically *extreme* policies on partisan support. Though our findings

suggest that a promising way to differentiate is by *prioritizing* (old or new) issues that garner the greatest support among partisans. Hence an implication is that primary competition may contribute to polarization by constraining party elites to the policy consistent (rather than extreme) positions that unify and energize co-partisans.

Our experimental design naturally leaves out many important questions that should be pursued in future research. One such question is how voters learn the information required to make policy judgments between candidates in relatively low-information primaries. Even when positions differ, politicians do not always emphasize this in primary campaigns, which often focus on general electability or candidate competence (Boatright 2013). Primary contests are frequently uncompetitive as well, leaving voters with little incentive to learn about the ideological orientations of challengers. Our findings clearly suggest that when made aware of the policy priorities of primary candidates, voters are capable of making ideological evaluations, even when the party label is fixed and thus uninformative. Put simply, voters have a coherent and realistic vision of what makes a candidate more liberal (conservative), and a capacity and willingness to use that information when deciding. Our findings may generalize less readily to cases where voters are uninformed of the issues that divide co-partisans.

In a similar vein, we may under- or over-estimate the impact of non-policy demographics. Surveys can mask negative affect towards candidates with particular attributes due to an effort by respondents to avoid revealing socially undesirable attitudes. Potentially, partisans respond more favorably to certain candidates in our experiment than they would in a real primary, though the prevalence of this is on the decline and lessened in online survey mode (e.g., Hopkins 2009). Yet, we suspect such an effect would attenuate policy voting, as cognitive effort devoted to providing socially appropriate responses could distract from the issue positions. Another important difference is we do not show images of candidates or ask respondents to consider actual politicians – race, gender or religion may matter differently when these are visualized or embodied by real individuals. Finally, we structure primary elections as binary decisions

between relatively comparable alternatives instead of requesting a ranking or first-choice in multidimensional, multi-candidate contests. Some research suggests the cognitive complexity of the latter task would reduce support for minority candidates (Crowder-Meyer et al. 2019). To the extent binary evaluations differ from how partisans decide, we again could overstate levels of support for minority candidates, or otherwise miss how attributes drive preferences.

This study joins a body of emerging research on candidate selection and vote choice in primary elections. Our aim is to assess the importance of issue and affinity voting in a particular information environment, where policy and non-policy factors are presented alongside each other in a questionnaire format. In such a setting, we show that clear policy differences elicit a pattern of primary voting consistent with moderates being deterred or punished in real elections, thus offering additional explanation for how primaries can yield off-center politics. Notably, this design is meant to replicate some, but not all important elements of primary elections. Future work could extend our approach to investigate how different aspects of primaries, including the way information is presented or the attributes featured, impact co-partisan vote choice. Indeed, doing so can help uncover whether stronger or richer signals drive consistent and durable preferences for candidates with shared affinities, and whether issue voting persists even as competing factors become increasingly prominent and realistic.

References

- Abramowitz, Alan I. 1989. “Viability, Electability, and Candidate Choice in a Presidential Primary Election: A test of Competing Models.” *Journal of Politics* 51(4):977–992.
- Achen, Christopher H. and Larry M. Bartels. 2016. *Democracy for Realists: Why Elections Do Not Produce Responsive Government*. Princeton: Princeton University Press.
- Ahler, Douglas J. and Gaurav Sood. 2017. “The Parties in Our Heads: Misperceptions About Party Composition and Their Consequences.” *Journal of Politics* 79(3).
- Aldrich, John. 1983. “A Downsian Spatial Model with Party Activism.” *American Political Science Review* 77(4):974–90.
- Aldrich, John H. and R. Michael Alvarez. 1994. “Issues and the Presidential Primary Voter.” *Political Behavior* 16(3):289–317.
- Arceneaux, Kevin and Ryan Vander Wielen. 2017. *Taming Intuition: How Reflection Minimizes Partisan Reasoning and Promotes Democratic Accountability*. Cambridge University Press.
- Badas, Alex and Katelyn Stauffer. 2019. “Voting for Women in Nonpartisan and Partisan Elections.” *Electoral Studies* 57(1):245–255.
- Bawn, Kathleen, Martin Cohen, David Karol, Seth Masket, Hans Noel and John Zaller. 2012. “A Theory of Political Parties: Groups, Policy Demands and Nominations in American Politics.” *Perspectives on Politics* 10(3):571–597.
- Bawn, Kathleen, Stephanie DeMora, Andrew Dowdle, Mark Myers, Spencer Hall, Angela Ocampo, Shawn Patterson and John Zaller. 2017. “Ideology and Vote Choice in House Primaries.” <https://uakron.edu/bliss/state-of-the-parties/papers/bawn+etal.pdf>.

- Blackwell, Calvin and Peter T. Calcagno. 2019. "Party-Crashers or Wallflowers? The Lack of Strategic Voting in Experimental Primaries." *Applied Economics Letters*, Forthcoming.
- Boatright, Robert G. 2013. *Getting Primaried: The Changing Politics of Congressional Primary Challenges*. University of Michigan Press.
- Bolsen, Toby, James N. Druckman and Fay Lomax Cook. 2014. "The Influence of Partisan Motivated Reasoning on Public Opinion." *Political Behavior* 36(2):235–262.
- Boudreau, Cheryl, Christopher S. Elmendorf and Scott A. MacKenzie. 2019. "Racial or Spatial Voting? The Effects of Candidate Ethnicity and Ethnic Group Endorsements in Local Elections." *American Journal of Political Science* 63(1):5–20.
- 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).
- Branton, Regina. 2009. "The Importance of Race and Ethnicity in Congressional Primary Elections." *Political Research Quarterly* 62(3):459–473.
- Campbell, David E, John C Green and Geoffrey C. Layman. 2011. "The Party Faithful: Partisan Images, Candidate Religion, and the Electoral Impact of Party." *American Journal of Political Science* 55(1):42–58.
- Campbell, David E., John C. Green and J. Quin Monson. 2012. "The Stained Glass Ceiling: Social Contact and Mitt Romneys 'Religion Problem'." *Political Behavior* 34(2):277–299.
- Carnes, Nicholas and Noam Lupu. 2016. "Do Voters Dislike Working-Class Candidates? Voter Biases and Descriptive Underrepresentation of the Working Class." *American Political Science Review* 110(4):832–844.
- Cohen, Marty, David Karol, Hans Noel and John Zaller. 2008. *The Party Decides: Presidential Nominations Before and After Reform*. Chicago: University of Chicago Press.

- Crowder-Meyer, Melody, Shana Kushner Gadarian, Jessica Trounstine and Kau Vue. 2019. "A Different Kind of Disadvantage: Candidate Race, Cognitive Complexity, and Voter Choice." *Political Behavior*, Forthcoming.
- Doherty, David, Connor Dowling and Michael Miller. 2019. "Do Party Chairs Think Women and Minority Candidates Can Win? Evidence from a Conjoint Experiment." *Journal of Politics*, Forthcoming.
- Dolan, Kathleen. 1998. "Voting for Women in the 'Year of the Woman'." *American Journal of Political Science* 42(1):272–293.
- Dolan, Kathleen. 2008. "Is There a 'Gender Affinity Effect' in American Politics?" *Political Research Quarterly* 61(1):79–89.
- Dolan, Kathleen. 2010. "The Impact of Gender Stereotyped Evaluations on Support for Women Candidates." *Political Behavior* 32(1):69–88.
- Goggin, Stephen N., John A. Henderson and Alexander G. Theodoridis. 2019. "What Goes with Red and Blue? Mapping Partisan and Ideological Associations in the Minds of Voters." *Political Behavior*, Forthcoming.
- Grossmann, Matt and David A. Hopkins. 2016. *Asymmetric Politics: Ideological Republicans and Group Interest Democrats*. Oxford University Press.
- Hainmueller, Jens and Daniel J Hopkins. 2015. "The Hidden American Immigration Consensus: A Conjoint Analysis of Attitudes Toward Immigrants." *American Journal of Political Science* 59(3):529–548.
- Hainmueller, Jens, Daniel J Hopkins and Teppei Yamamoto. 2013. "Causal Inference in Conjoint Analysis: Understanding Multidimensional Choices via Stated Preference Experiments." *Political Analysis* 22(1):1–30.

- Hirano, Shigeo, James M. Snyder Jr., Stephen Ansolabehere and John Mark Hansen. 2010. "Primary Elections and Partisan Polarization in the U.S. Congress." *Quarterly Journal of Political Science* 5(2):169–191.
- Hopkins, Daniel J. 2009. "No More Wilder Effect, Never a Whitman Effect: When and Why Polls Mislead about Black and Female Candidates." *Journal of Politics* 71(3):769–781.
- Jackman, Simon and Lynn Vavreck. 2010. "Primary Politics: Race, Gender, and Age in the 2008 Democratic Primary." *Journal of Elections, Public Opinion and Parties* 20(2):153–186.
- Klar, Samara. 2014. "Partisanship in a Social Setting." *American Journal of Political Science* 58(3):687–704.
- Krupnikov, Yanna and Spencer Piston. 2015. "Racial Prejudice, Partisanship, and White Turnout in Elections with Black Candidates." *Political Behavior* 37(2):397–418.
- Lau, Richard R. 2013. "Correct Voting in the 2008 U.S. Presidential Nominating Elections." *Political Behavior* 35(2):331–355.
- Lawless, Jennifer L. and Kathryn Pearson. 2008. "The Primary Reason for Womens Underrepresentation? Reevaluating the Conventional Wisdom." *Journal of Politics* 70(1):1–24.
- Marietta, Morgan and David C Barker. 2007. "Values as Heuristics: Core Beliefs and Voter Sophistication in the 2000 Republican Nomination Contest." *Journal of Elections, Public Opinion and Parties* 17(1):49–78.
- McConaughy, Corrine M., Ismail K. White, David L. Leal and Jason P. Casellas. 2010. "A Latino on the Ballot: Explaining Co-ethnic Voting Among Latinos and the Response of White Americans." *Journal of Politics* 72(4):1199–1211.
- McDermott, Monika L. 1998. "Race and Gender Cues in Low-Information Elections." *Political Research Quarterly* 51(4):895–918.

- Miratrix, Luke W., Jasjeet S. Sekhon, Alexander G. Theodoridis and Luis F. Campos. 2018. “Worth Weighting? How to Think About and Use Weights in Survey Experiments.” *Political Analysis* 26(3):275–291.
- Rickershauser, Jill and John H Aldrich. 2007. “‘It’s the Electability, Stupid’ – or Maybe Not? Electability, Substance, and Strategic Voting in Presidential Primaries.” *Electoral Studies* 26(2):371–380.
- Rothschild, Jacob, Adam Howat, Richard Shafranek and Ethan Busby. 2019. “Pigeonholing Partisans: Stereotypes of Party Supporters and Partisan Polarization.” *Political Behavior* 41(2).
- Sanbonmatsu, Kira. 2002. “Gender Stereotypes and Vote Choice.” *American Journal of Political Science* 46(1):20–34.
- Sides, John, Chris Tausanovitch, Lynn Vavreck and Chris Warshaw. 2019. “On the Representativeness of Primary Electorates.” *British Journal of Political Science*, Forthcoming.
- Simas, Elizabeth N. 2016. “The Effects of Electability on U.S. Primary Voters.” *Journal of Elections, Public Opinion and Parties* 27(3):1–17.
- Stone, Walter J, Ronald B Rapoport and Alan I Abramowitz. 1992. “Candidate Support in Presidential Nomination Campaigns: The Case of Iowa in 1984.” *Journal of Politics* 54(4).
- Thomsen, Danielle. 2019. “Ideology and Gender in U.S. House Elections.” *Political Behavior*, Forthcoming.