

Voter Learning in State Primary Elections¹

Shigeo Hirano

Department of Political Science
Columbia University

Gabriel S. Lenz

Travers Department of Political Science
University of California, Berkeley

Maksim Pinkovskiy

Department of Economics
Massachusetts Institute of Technology

James M. Snyder, Jr.

Department of Government
Harvard University
and NBER

March, 2013

¹We thank participants at the MIT and Columbia Political Economy Breakfast Seminars, Doug Ahler and David Broockman for comments and suggestions. We would also like to thank Abuchi Agu, Jessica Cui, Mason Fitch, Aaron Kaufman, Karen Li, Margaret Mattes, Aidan McCarthy, Laura Mills, Cody Nager, Adria Schulman-Eyink, Jasmine Senior, Tony Valeriano, Nida Vidutis and Miranda Yaver for outstanding research assistance. Hirano and Snyder thank the National Science Foundation for their financial support under grants SES-0617556 and SES-0959200. Pinkovskiy would like to thank the Paul and Daisy Soros Foundation for New Americans for intellectual stimulation, the NSF Graduate Research Fellowship Program and the Institute for Humane Studies for funding.

Abstract

When voters learn about candidates' issue positions during election campaigns, does it affect how they vote? This basic question about voters remains unanswered in part because of a methodological obstacle: learning candidates' issue positions may not only influence voters' vote choice but also their issue positions. To surmount this obstacle, we attempt to answer this question by examining statewide primary elections, which are arguably less vulnerable to this reverse-causation problem because they lack partisan cues and are of much lower salience than presidential elections are. Using both existing polling data and our own panel internet surveys, we find that voters learn about the ideologies of candidates during statewide primary campaigns and that this learning affects their voting decisions in senate and gubernatorial primaries. We fail to find similar results for downballot primaries, raising concerns about voters' ability to make informed judgments for these types of elections.

1. Introduction

Do voters learn about candidates' attributes and issue positions during election campaigns? Does this learning affect how they vote?

The answer to the first question is almost certainly “yes.” Work by numerous scholars shows that voters know more about candidates' relative policy stances and interest group endorsements at the end of campaigns than at the beginning.¹ Summarizing the literature on the 1976 Carter-Ford debates, Sears and Chaffee (1979) found clear evidence of short-term increases in voter knowledge. Alvarez (1997, 2001) finds that voter uncertainty about candidates' issue positions typically falls over the course of presidential and U.S. senate campaigns. Lipsitz (2004) finds that voter knowledge of candidates' issue positions in the 2000 presidential election was higher in battleground states than in non-battleground states. Several studies show that voter knowledge is higher in highly contested U.S. senate and house races than in lopsided races.² Finally, numerous papers find that exposure to news media, advertising, debates, and direct contact with candidates during a campaign is positively related to voter information by election day.³

The answer to the second question is less clear, for two reasons. First, even taken at face value the evidence from previous research is mixed. Many older studies, and some recent ones, find small or insignificant effects. They also find that most voters decide how to vote before campaigns begin in earnest, or that campaigns mainly reinforce preexisting attitudes.⁴ On the other hand, a number of recent studies find relatively large effects.⁵

Second, methodological issues make the question especially challenging. Ruling out omitted-variable biases and reverse causation is difficult. Consider a typical presidential general election. Suppose a voter prefers one of the candidates (or parties) at the beginning

¹See, e.g., research by Patterson and McClure (1976), Zhao and Bleske (1993), Martinez et al. (2001), Rice (2004), Craig et al. (2005), Henderson (2009), and Lenz (2009, 2012).

²Clarke and Evans (1983), Krasno (1994) and Gronke (2000).

³See, e.g., Atkin and Heald (1976), Weaver (1996), Vavreck et al. (2002), Benoit and Hansen (2004), and Druckman (2005).

⁴See, e.g., Lazarsfeld et al. (1948), Berelson et al. (1954), Patterson and McClure (1976), Finkel (1993).

⁵See, e.g., Fan and Tims (1989), Roberts (1992), Bartels (1993), Holbrook (1996), Shaw (1999), Martinez et al. (2001), Iyengar and Simon (2000), and Highton (2006). See also Gelman and King (1993), who argue that the increasing stability of the polls over a campaign is the result of voters learning key facts during the campaign and deciding how to vote on the basis of these facts.

of the campaign for non-policy reasons, such as charisma (or partisan attachment). Suppose also that during the campaign this voter learns about the policy or ideological positions of this candidate (party). Then the voter might adjust her own positions so they are more consistent with those of the candidate (party), or even adopt the candidate’s (party’s) positions outright on some issues. The data would then show a strong relationship between the citizen’s vote choice and the “policy distance” between the voter and the candidate, but this correlation would not reflect any causal effect of policy information on vote choice. This relationship might also grow stronger during the campaign as the voter learns more about the candidate’s issue positions and changes her own positions accordingly. The increasingly strong relationship might look like learning that affects voting decisions, but, again, it would be spurious. These same patterns could also arise if, rather than shifting her policy stances, the voter selectively retains information that supports her prior positive opinion of the candidate and disregards other information. Numerous studies find that individuals do, in fact, tend to adopt policy positions consistent with their partisanship, the positions of their preferred candidate, or other predispositions.⁶

This paper attacks the questions above by focusing on an under-explored type of election: non-presidential primary elections. In particular, we study primary elections for governorships, U.S. senate seats, and state downballot offices. The vast majority of previous research has studied presidential elections, and the small number of non-presidential studies all focus on general elections.⁷

What are the advantages of studying state primaries? First, except for incumbents, most voters probably have very little information about the candidates before campaigning begins. Thus, the potential for learning to occur is large.⁸ Second, the partisan version of reverse-causality discussed above is not likely to be a problem, because the elections are *intra*-party

⁶See, e.g., Berelson et al. (1954), Campbell et al. (1960), Abramowitz (1978), Zaller (1992, 1994), Bartels (2002a, 2002b), Carsey and Layman (2006), and Lenz (2009).

⁷To our knowledge, only one previous paper has examined ideological voting in state primary elections. Tedin and Murray (1981) study the 1978 Texas state races for governor and attorney general. However, they only look at sorting on ideology in a single primary. Their results are consistent with ours – they find sorting in the governor’s race but not in the attorney general race.

⁸In general, this might or might not be an advantage. However, given the strength of the “minimal effects” position in the literature, it is reasonable to look for situations where the potential for learning is large.

contests. Third, the relatively low salience of the offices sought – at least compared to the presidency – means that the candidate-centered version of reverse causality is also unlikely to be a serious problem. How many people in California changed their ideological positions because of their high (or low) opinions of Bill Simon, Pete Wilson, or Leo McCarthy? We guess the number is very small, most likely close to zero, a guess confirmed by panel analyses that we discuss below.

Finally, there are important reasons to study non-presidential primaries for their own sake. One reason is that it is not obvious that we should even have them. Do voters learn enough during the campaigns for these offices to vote in a reasonably informed manner? As noted, state primary elections are low-information affairs, so it would not be surprising if most voters know very little about the candidates even by the end of the campaign. On the other hand, turnout in these elections is low, so perhaps the citizens who do vote are especially interested, attentive, and informed.

Our findings are straightforward. First, using polling data we find evidence consistent with the claim that voters learn candidates’ ideological positions over the course of a campaign and that this learning affects their vote choice. During primary elections with no incumbent, conservative (liberal) respondents increase their support for the conservative (liberal) candidate over the course of the campaigns.

Consistent with this result, we also find evidence that the magnitude of the ideological gap between the candidates, voters’ initial beliefs about the candidates, and the amount of attention given to campaigns appear to also affect the degree to which voters learn about candidates’ ideologies. The amount that voters learn is positively related to the size of the ideological gap between the candidates as well as the degree to which voters had incorrect beliefs about the candidates’ ideologies. We find little effect for “placebo” races – i.e. races with little ideological gap between the candidates – or races with an incumbent – i.e. races where voters are likely to have more information about at least one candidate before the primary campaign. However, these patterns we observe are only significant for gubernatorial and senate races. The pattern is not significant when we focus on downballot races, which is perhaps not surprising given the difference in campaign resources and media attention given

to the different races.

Finally, we confirm our findings using panel surveys in which we recruited participants online and interviewed them before and after 2010 primary races. Although the samples are not representative, the overall patterns they exhibit are similar to those in the polling data; and they allow us to conduct important supplementary analyses. First, since we asked participants about their perceptions of candidates' ideological stances, we were able to measure who learns the positions and who does not. Consistent with the poll results, we find that individuals who learn these positions are indeed those who also shift their support to the candidate who, they have just learned, agrees with them ideologically. Second, the panel data allow us to further rule out the reverse causation alternative. Indeed, we show that ideology influences vote change even when we measure ideology before voters learn. Unlike findings in presidential elections, we find no evidence that participants' views about primary candidates shape their ideologies. Third, the panel findings are also more consistent with a learning explanation than a priming explanation (Iyengar and Kinder 1987; Johnston, Blais, Brady, and Crête 1992), as much of the increase in candidate-voter ideological congruence occurs among voters who learned these ideologies. Of course, priming and learning likely occur simultaneously and both may contribute to the effects we observe.⁹

Taken together, our findings are reassuring, implying that campaigns are fulfilling their role in senate and gubernatorial primaries. In these races, primary voters are learning policy-relevant information and incorporating that information into their voting decisions. On the other hand, our downballot findings are potentially worrisome for democracy, as we fail to find a similar pattern. Voters may therefore be learning too little in these races to make reasonably informed voting decisions.

2. Data and Specifications

2.1. Data Sources

⁹Priming and learning are difficult to distinguish empirically because messages that inform voters about the candidates' ideologies may also prime ideological considerations. Likewise, messages that prime may convey information (see Lenz [2012]). However, both ultimately lead to the same outcome: increased candidate-voter ideological congruence. We return to this point in the conclusion.

Our data comes from three sources. First, we use data from the California/Field Poll for Democratic and Republican primaries in California between 1966 and 2006. The Field Poll is conducted approximately 8 times during every election year in California, with 3-4 of the polls asking questions about primary campaigns to respondents who indicate they were likely to vote in the next election. Besides asking its respondents about their voting intentions in the primaries, the Field Poll also collects a large set of richly-coded covariates for each respondent, including ideology (coded on a 5-, 7- or even 9-point scale), age, race, gender, income, education and area of residence within California. The Field Poll releases all of its unit record data to the public, so we will be able to make maximal use of these covariates.

Our second data source comes from the press releases of the survey research firm Survey USA (hereafter, SUSA). SUSA has conducted election polls since 1992, but we have found publicly available releases of polls for primary elections in a usable format since 2005. The sample frame is likely voters. We look at Republican and Democratic primaries in various states since 2005. The publicly available data consists of cross tabulations of ideology (coded on a 3-point scale) and other covariates (usually age, gender, race, income and region of state are included), so we cannot exploit the data as much as we can for the Field Polls. We study all of the available primaries in either source of data for the offices of governor and senator that do not include incumbents, and where we have at least two polls before the primary date (see Table A.1 for a list of these races). We also study some primaries with incumbents and primaries for downballot offices such as lieutenant governor and attorney general.

Our third data source is panel surveys we conducted during the 2010 primaries. We surveyed likely primary voters in Michigan, Pennsylvania, and New Hampshire before and after their gubernatorial and senate primaries, with 1,201 complete panel respondents. Before Election Day, we recruited participants with online ads and through a consumer internet panel (for details, see Appendix D). We then recontacted participants after the election and reinterviewed 60% to 75% of the original respondents. The surveys asked about perceptions of the candidates' ideologies and about each participant's ideology, vote, party identification, age, race, gender, education, political interests, and political knowledge.

2.2. Specifications

We assume the following simple model. Consider a primary election with only two candidates: one who is more conservative and one who is more liberal. Also suppose that there are only two types of voters in the primary electorate: relatively conservative and relatively liberal. Finally, assume that there have been exactly two polls taken during the primary: an early poll and a late poll. Let $t=0$ denote the early poll and let $t=1$ denote the late poll, and let L_{it} be the indicator for the late poll ($t=1$). Let V_{it} be the preference of respondent i at time t , with $V_{it}=1$ if respondent i prefers the more conservative candidate and $V_{it}=0$ if respondent i prefers the more liberal candidate. Let C_{it} be the respondent's ideology, with $C_{it}=1$ if the respondent is relatively conservative and $C_{it}=0$ if the respondent is relatively liberal. To look at the relative change in support for the conservative candidate among conservatives, we estimate the following regression model:

$$V_{it} = \alpha + \beta_C C_{it} + \beta_L L_{it} + \gamma C_{it} L_{it} + \epsilon_{it} \quad (1)$$

Each coefficient has an intuitive interpretation. The intercept α is the fraction of liberal voters intending to vote for the conservative candidate in the early poll. The coefficient on ideology, β_C , is the amount by which the conservative voters' support for the conservative candidate in the early poll is greater than the liberal voters' support for this candidate. The coefficient on the late indicator, β_L , is the amount by which liberal voters increase their support for the conservative candidate from the early to the late poll. Finally, the interaction coefficient γ is the amount by which conservative voters increase their support for the conservative candidate over the course of the campaign *over and above* the amount by which liberal voters increase their support for the conservative candidate.¹⁰

Our quantity of interest is precisely this coefficient γ . Under the assumption that the only differential trends in support for the conservative candidate among conservatives and liberals arise from voter learning, the coefficient γ also is a measure of how much voters' voting intentions change because of information that they learn during the campaign. If our

¹⁰We use a linear probability model because it does not require normality and homoskedasticity assumptions for consistent estimates and provides coefficients that are relatively easy to interpret.

hypothesis that campaigns inform voters about candidates' positions in ways that change how voters vote is correct, we should observe $\gamma > 0$; that is, conservatives should increase their support for the conservative candidate over the course of the campaign more than liberals do for that same candidate.

2.3. Classifying Ideological Positions and Timing of Polls

In order to implement this specification empirically, we need to match the rich data on primaries that we have, with multiple polls and many ideological shades, to the stylized description of a primary that we present above. In particular, we need to identify candidates as either liberal or conservative. We also need to classify respondents as either liberal or conservative. Finally, we divide the polls between those that are “early” versus those that are “late.”

For all gubernatorial and senate primaries, we classified candidates as either liberal or conservative using information from newspaper articles and interest group endorsements. In most of the races, this distinction is relatively clear. When there were more than two candidates in a race, we focus on the two with the largest ideological difference who were also among the top three vote getters. In Appendix Table A.1 we list the 37 governor and senate primaries analyzed, and the candidates identified as being relatively liberal or conservative.

Even though we were able to easily identify candidates' relative ideological positions from various sources, we confirmed our classification using independent coders of candidates' positions based on newspaper coverage of the primary campaigns. Using an automated process, we collected sentences likely to identify candidates' issue positions from newspapers available on a news aggregating service (Factiva). During the three months prior to the primary election. The sentences collected had to mention at least one of the candidates together with certain key words or phrases related to an issue or ideological position, such as abortion, tax cut, etc. (see Appendix A for the full list).¹¹ For each sentence, the coders attempted to identify the ideological position of candidates mentioned in the sentence on

¹¹These keywords capture issues of broad national interest and may miss state-specific issues. To the extent they do, however, the resulting measurement error will likely lead us to understate learning effects, not overstate them.

a five-point scale based only on the information conveyed within each sentence. Because of limitations in Factiva coverage, this independent coding only covered two-thirds (27 out of 37) of the races we originally coded. Among these two-thirds, however, the independent coders confirmed our liberal / conservative classifications of the candidates in all but one case.¹² We describe the procedure for coding the sentences and computing ideology gaps in Appendix A.

The newspaper sentence coding also provides us a quantification of relative candidate ideological position. We use this measure to divide primaries into a “baseline” sample of races with acute ideological distinctions and a “placebo” sample of races where we would expect little ideological learning because of minimal ideological distinctions between candidates. We considering primaries with an ideology gap of 0.5 or more (on the five-point scale) to be “baseline,” and primaries with gaps less than 0.5 to be “placebo” (the results are similar with other cutoffs). For the races that lack independent ratings, we rely on our classifications.

In our main analyses, we classify voters as either liberal or conservative using their self-reported ideological position in the surveys. For most of our analyses we discard voters in the ideological center of each party in order to focus on voters with relatively distinct self-identified ideological positions. In Democratic primaries, all self-identified conservatives are classed as conservative, and strong liberals (or all liberals on a 3-point scale) are classed as liberal. Weak liberals and moderates are excluded from analyses of the Democratic primaries. In Republican primaries, strong conservatives (or all conservatives on a 3-point scale) are classed as conservatives and all self-identified moderates and liberals are classed as liberal. In these primaries, we are unable to study liberals separately from self-identified moderates, since so few Republican respondents identify themselves as liberal.¹³ We call this the Extremist Only coding and adopt it because it provides predictions about voting decisions without additional assumptions.¹⁴ We examine the robustness of our findings

¹²The only race where we disagreed with the coders is the 2006 Michigan Republican Senator race, between Bouchard and Butler. We code Bouchard as conservative, while the coders code Butler as conservative. The ideology gap in this race is 0.08, which is one of the smallest ideology gaps for all coded races (it is at the 5th percentile of the ideology gaps).

¹³For the polls using a five-point ideology scale, we exclude weak conservatives from the analysis of Republicans. Both weak liberals and moderates are excluded from the analysis of Democrats.

¹⁴Alternatively we could treat the ideological measures as continuous and not exclude moderates. How-

to these decisions, including specifications where we include all respondents. The main substantive findings remain relatively unchanged. These are reported in Appendix B.

Finally, in our main analysis we focus on only the first (early) and the last (late) polls. We would expect the difference between these polls to best reflect the amount of learning that occurs during the campaign. We also examine whether our results are sensitive to the exclusion of the intermediate polls or the variation in the dates of the first and last polls across primary races. More specifically, we conducted additional analyses where we classify all polls 45 days or more before the primary as early, and all polls less than 45 days before the primary as late. The *45-Day-Cutoff* is motivated by a simple analysis of the number of newspaper sentences about primaries per day (discussed above) as a function of the number of days remaining before the primary. The coverage remains flat at a low level averaging 59 sentences per day until about 45 days before the primary, at which point it rises linearly until the primary date. During the week before the primary the average number of sentences per day is 302.¹⁵ These are reported in Appendix B.

Note that we must drop from the analysis all respondents who do not state a preference for any candidates in a poll. Since there are more of these respondents in early polls than in late polls, and since these respondents tend to be less well informed in the early polls, our estimates likely understate the total amount of learning, as well as the overall impact of learning on voters' choices. We can address this issue directly using the data from our panel surveys. We discuss this in more detail in Section 5.

3. Results for Governor and Senator Races Using Polling Data

In this section, we focus on primaries for governor and senator. These tend to be the statewide primary elections most salient to voters. Moreover, given the amount of newspaper coverage and campaign expenditures, voters are more likely to be exposed to information about candidates for these offices as compared to downballot offices over the course of the primary campaign. Thus, we would expect considerable voter learning about the candidates.

ever, since our measures may lack cardinal information about voter and candidate ideal points, we cannot accurately determine when moderates are closer to the conservative or liberal candidates.

¹⁵The point estimate of the slope during the last 45 days is 6.08, with a standard error of 0.48; the regression R-squared is 0.79. We also tried a *30-Day-Cutoff* and the results were substantively unchanged.

Moreover, we use the information in the substantial newspaper coverage of these races to quantify the magnitude of the ideological gap between the candidates. With this measure we examine whether the size of the ideological gaps between candidates corresponds with the amount of information that voters “learn” over the course of the primary campaign.

3.1. Main Results

Do liberal voters shift their support towards liberal candidates and conservative voters shift their support towards conservative candidates over the course of an electoral campaign? To answer these questions, we estimate a saturated version of equation 1 in which we allow all coefficients to vary by primary:

$$V_{ijt} = \alpha_j + \beta_{jC}C_{ijt} + \beta_{jL}L_{ijt} + \gamma_j C_{ijt}L_{ijt} + \epsilon_{ijt} \quad (2)$$

with j indexing primaries.¹⁶ The statistic of interest – how much conservative voters shift towards conservative candidates relative to liberal voters – is the average interaction coefficient:

$$\gamma_{\text{avg}} := \frac{1}{n} \sum_{j=1}^n \gamma_j \quad (3)$$

In Table 1 we present our estimate of γ_{avg} for the 28 “baseline” gubernatorial and senate primary races – i.e. races where campaigns should especially help citizens differentiate between the candidates because they lack incumbent candidates and because the candidates differed ideologically (the independent coders found an ideological difference between the candidates larger than 0.5 on a five-point scale). The table also presents the estimates for races where we would expect electoral campaigns to have less of an effect on the congruence between the ideological position of voters and the ideological positions of the candidates they support – i.e. races where the candidates do not differ in their ideological positions or races with an incumbent.

The first column of Table 1 finds that citizens do indeed appear to learn in these “baseline” races. The average interaction coefficient is positive and statistically significant.¹⁷

¹⁶Note that to run this regression, the sufficient statistics are the cross-tabs of voter ideology across vote choice for each poll since all variables are binary. Hence, we can run this regression equally well for the SUSA data as for the Field Poll data.

¹⁷We report robust standard errors. Since we are estimating a model that is saturated by primary, we

Moreover, the coefficient is large in magnitude: 0.10. Conservatives increase their support for conservative candidates over the course of the campaign by an average of 10 percentage points more than liberals do. The congruence thus increases between the ideological position of the voters and the candidates they support. Figure 1 presents a histogram of the interaction coefficients for each of these races (i.e. γ_j). It reveals that the average 0.10 effect reflects the central tendency of the distribution and is not skewed by outliers.

Table 1: Gubernatorial and Senate Races

	Baseline	Placebo	Incumbents
Δ Voter-Cand. Ideological Congruence (γ_{avg})	0.100** (0.021)	-0.036 (0.034)	0.022 (0.044)
Observations	12,099	4723	2894
Number of Primaries	28	9	5
P-value	3.39e-08	0.774	0.199

Notes: Robust standard errors in parentheses. P-value is for the F-statistic for testing that all interaction coefficients are zero. Dependent variable is support for conservative candidate. Δ Voter-Candidate Ideological Congruence is the average interaction coefficient, γ . This estimates the amount by which conservative voters increase their support for the conservative candidate over the course of the campaign over and above the amount by which liberal voters increase their support for the conservative candidate. ** p<0.01, * p<0.05, + p<0.1

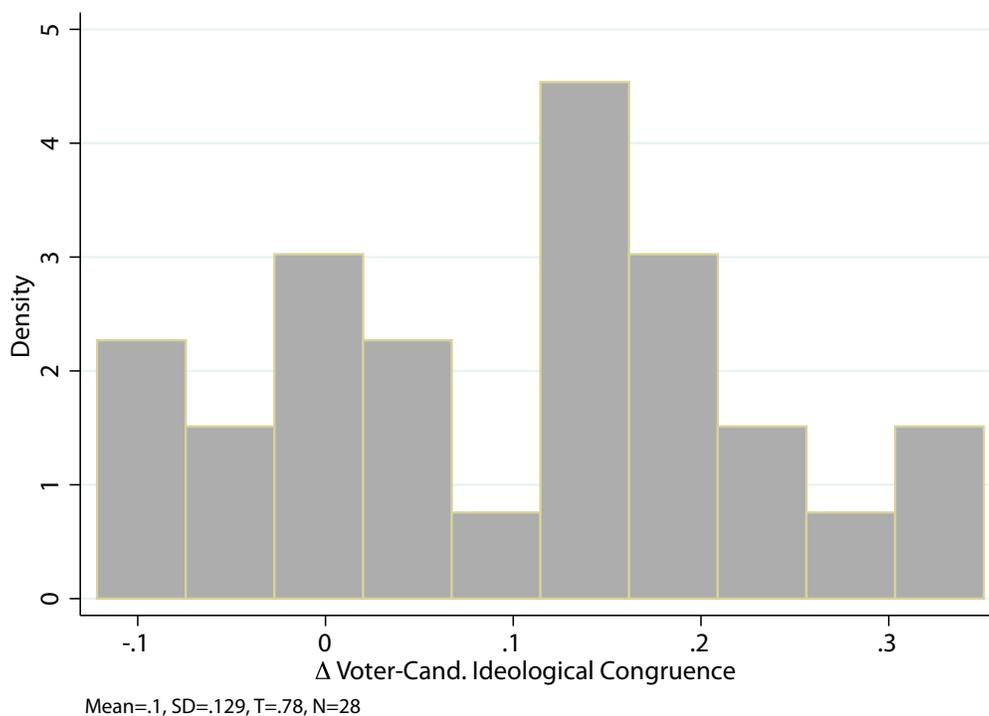
The first estimate uses the Extremist Only classification and the first and last poll, but the estimate is robust to other coding decisions. In results reported in Tables B.1 and B.2, we find a similar effect when we include voters who classify themselves as “moderate” and include all polls but divide them between those taken more or less than 45 days prior to the primary. Not surprisingly, the effect attenuates slightly as we expand the sample to include moderate voters and/or intermediate polls.

According to these findings, voters appear to be learning about candidates’ ideologies in primaries and changing their support accordingly. If these associations do indeed reflect learning, we should observe several other patterns in these data. In particular, we should fail

cannot cluster our standard errors at the primary level. We have estimated models in which we force the interaction effect (the key coefficient of the paper) to be the same across all primaries considered. In this specification, when we cluster our standard errors by primary, the standard errors increase but the coefficient remains significant at the 1% level.

Figure 1

Histogram of Δ Voter-Candidate Ideological Congruence Coefficients



Notes: Δ Voter-Candidate Ideological Congruence is the interaction coefficient γ and estimates the amount by which conservative voters increase their support for the conservative candidate over the course of the campaign *over and above* the amount by which liberal voters increase their support for the conservative candidate. The histogram shows these coefficients for baseline races.

to find the same effect in the “placebo” races – those with minimal ideological differences between candidates. In the second column of Table 1, we present our estimate of γ_{avg} for the nine “placebo” races for governor and senator.¹⁸ For these races, our estimate of the interaction coefficient is close to zero and has the wrong sign, -0.036, indicating little effect of any learning about the ideological positions of the candidates among the voters in these races. The F-test fails to reject the null hypothesis that all the interaction coefficients for all of the races are zero ($p > 0.33$) and the difference between the “baseline” and “placebo” races is significant at $p < 0.01$).¹⁹

¹⁸See Table A.1 for a list of placebo races.

¹⁹Figure 2 shows the interaction coefficients for these races (those to the left of 0.5).

Finally, in the third column of Table 1, we present the results for five races where an incumbent governor or senator was challenged in their primary.²⁰ In these races, we suspect that voters are better informed about the incumbents' ideological positions before the campaigns. As with the "placebo" races, the coefficient is indeed smaller: 0.022. The F-test again fails to reject the null hypothesis that all the interaction coefficients for all of the races are zero (p-value of 0.20) and the coefficient is significantly smaller than the "baseline" estimate ($p < 0.10$).²¹

3.2. More Evidence That It's Learning: Magnitude of the Ideological Gap between Candidates

If these effects indeed reflect learning, then we might expect voters to learn more when the ideological gap is relatively large. In particular, we expect a positive association between the magnitude of the ideological gap between the candidates and the increase in congruence between voters' and candidates' ideological positions. This is precisely what we find.

This association is evident in a simple scatterplot of our coefficient estimates on the interaction between late poll and conservative voter ideology and our measures of ideological gap between the candidates. In Figure 2, we observe that the races with a relatively pronounced ideological gap between the candidates are also those where we find more "learning." One exception is the 1994 California Republican senate primary. In this case, the voters seemed to have identified the ideological position of candidates early in the primary. As is discussed in section 3.3, we find less learning in these types of cases.

To parametrically test whether this relationship exists, we add a triple interaction between the ideological gap measure, the late poll indicator, and the conservative respondent indicator variables. In other words, we run a triple difference specification of the following form, in which we compare conservative and liberal voters, early and late polls, and

²⁰The races included in this analysis are the 1966 California Democratic gubernatorial primary, 1992 California Republican senate primary, 2006 Alabama Republican gubernatorial primary, 2006 Oregon Democratic gubernatorial primary, and the 2007 Kentucky Republican gubernatorial primary.

²¹We should note that races with incumbents may suffer from a similar endogeneity problem as studies of voter learning in general elections. The incumbent may be substantially more influential in his state than any of the challengers and may induce people to modify their ideology to match his own in the reverse-causation model, which would upwardly bias the interaction effect.

behavior among conservatives and liberals can be more plausibly attributed to the existence of ideological gaps between the candidates, and thus, that voters are learning about ideology and using what they learn to shape their vote.

We estimate equation 4 using various methods of constructing g_j . One straightforward way to obtain g_j is to use the ideology gap itself. However, one may be skeptical about the cardinalization assumptions implicit in using the gap directly. Therefore, we also consider setting g_j to be a dummy variable equal to 1 when the ideology gap is above a chosen cutoff. In Table 2 we present the results using a cutoff of 0.5, which is the same cutoff used to identify the “baseline” and placebo samples in Table 1.²²

Table 2: Triple Differences

Gap Computation Period Gap Measure	Whole Period		First Third	
	Continuous	Cutoff 0.5	Continuous	Cutoff 0.5
Δ Voter-Cand. Ideological Congruence (γ_{int})	-0.016 (0.033)	-0.024 (0.032)	-0.043 (0.059)	-0.004 (0.051)
Δ Voter-Cand. Ideological Congruence x Ideological Gap (γ_{gap})	0.095** (0.032)	0.133** (0.039)	0.165** (0.055)	0.163** (0.059)
Observations	12279	12279	6367	6367
Number of Primaries	27	27	14	14
Number of Placebos	0	9	0	4

Notes: Robust standard errors in parentheses. Dependent variable is support for conservative candidate. Δ Voter-Candidate Ideological Congruence is the average interaction coefficient, γ . This estimates the amount by which conservative voters increase their support for the conservative candidate over the course of the campaign over and above the amount by which liberal voters increase their support for the conservative candidate. ** $p < 0.01$, * $p < 0.05$, + $p < 0.1$.

We also varied the period over which we obtain the information from newspapers needed to construct g_j . If candidates manipulate their ideological stances or emphasize different issues over the course of the primary campaign, then sentences taken from newspaper stories and editorials later in the campaign might reflect these changes, and thus g_j might be

²²We use 0.5 because Figure 2 shows that it is a natural break in the data. We also tried a cutoff of 0.7, which allows for more “placebo” primaries. The results do not substantively differ from the results using the 0.5 cutoff.

endogenous. Thus, in addition to constructing g_j using all the sentences describing primary j (as described in the previous section), we also construct a measure of g_j using only the sentences from the first third of our sample for primary j . Focusing on the ideological gap between candidates in the early part of campaign reduces the potential endogeneity problem.²³

Table 3 presents the estimates of γ_{int} and γ_{gap} . Using the candidate ideological gap from either the whole period or the first third of the sample, our hypothesis is confirmed: $\gamma_{\text{gap}} > 0$ and $\gamma_{\text{int}} \approx 0$. In fact, γ_{gap} is very close to our “baseline” estimate in section 3 (it is 10 percentage points), and γ_{int} is always statistically insignificant and rather small. This increases our confidence that the results reflect voter learning. For these estimates, we only use races where we code the ideological difference between the candidates using newspaper sentences for this analysis (see Table A.2 for a list of races for which we were able to quantify the ideological gap between the candidates).

Table 3: Downballot Races

	Baseline	Placebo	Lt Gov & Atty Gen
Δ Voter-Cand. Ideological Congruence (γ_{avg})	-0.006 (0.029)	-0.089 (0.059)	-0.041 (0.040)
Observations	5952	1702	3209
Number of Primaries	26	8	10
P-value	0.474	0.492	.344

Notes: Robust standard errors in parentheses. P-value is for the F-statistic for testing that all interaction coefficients are zero. Dependent variable is support for conservative candidate.

Δ Voter-Candidate Ideological Congruence is the average interaction coefficient γ and estimates the amount by which conservative voters increase their support for the conservative candidate over the course of the campaign *over and above* the amount by which liberal voters increase their support for the conservative candidate. ** p<0.01, * p<0.05, + p<0.1.

Interestingly, using the ideological gap from the newspaper coverage during first third of the sample yields somewhat larger coefficients, but the differences are not substantively

²³In fact, we do not find much evidence that candidates change ideological positions during the primary campaign. We should also note that the substantive findings are not particularly sensitive to the cutoff used to identify the ideological gap. The estimate of the coefficient on the triple interaction variable is about the same if we use the ideological gap from the first half, as opposed to the first third, of the sample.

significant.

3.3. More Evidence That It's Learning: Error Correction

The results in Table 1 where we find no significant “learning” in primary races with an incumbent candidate may reflect the amount of information voters have about the candidates at the outset of the primary campaigns. We might expect that voters are more likely to “learn” about candidates’ ideological positions over the course of the campaign when they start with no perceptions or incorrect perceptions of these positions. Thus, another prediction consistent with learning concerns the relationship between the coefficient on ideology (β_{jC}) and the interaction coefficient (γ_j) in each primary. If conservatives are initially insufficiently in favor of the conservative candidate (low β_{jC}), then we should observe a greater increase in their support for the conservative candidate relative to liberals over the course of the campaign (higher γ_j), and if conservatives are overly in favor of the conservative candidate (high β_{jC}), then we should observe a smaller increase in their support for her relative to liberals over the course of the campaign (lower γ_j).

In Figure 3 we plot our estimates of γ_j against our estimates of β_{jC} for each race. The pattern in this figure is consistent with this prediction. There is a strong negative association between our estimates of these two coefficients. This negative association may be interpreted in the context of an error-correcting model of voter learning, in which voters have a prior perception of candidate ideology and then correct it over the course of the campaign.²⁴

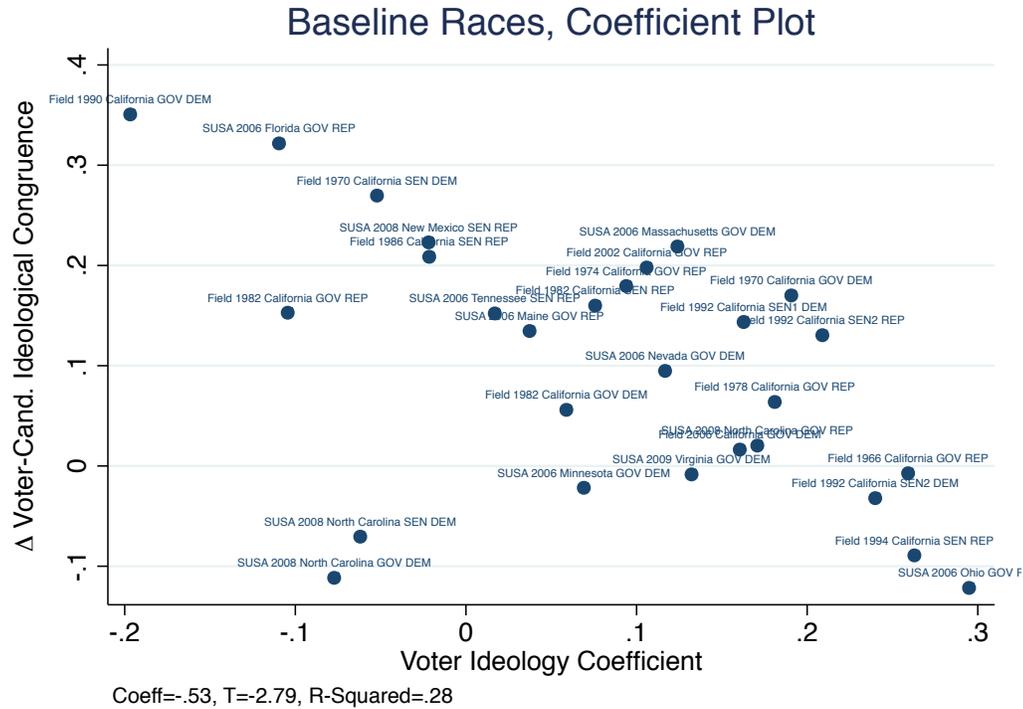
3.4 More Evidence It's Learning: Voter Information Early in the Primary Campaign

For the case of California, the Field polls often ask respondents whether they have “favorable” or “unfavorable” opinions of various primary election candidates. Many respondents are unwilling or unable to state an opinion, especially early in the campaign, and answer “don’t know.” It seems unlikely that informed respondents are actually indifferent, i.e.,

²⁴The two outlier cases are for the 2008 Democratic primaries in North Carolina. The 2008 Democratic primaries may have been unusual since the Obama-Clinton presidential primary drew so much attention that year, especially for states – including North Carolina – that held their state primaries on the same day as the presidential primary. Also, the ideological gaps between the candidates in the North Carolina races are close to the 0.5 cutoff to be considered placebo cases, so it is not particularly surprising that the interaction coefficient is not very large.

Figure 3

Scatterplot of Δ Voter-Candidate Ideological Congruence vs. Ideology Coefficient



have completely neutral opinions about candidates. Rather, respondents who answer “don’t know” probably lack even rudimentary information about the candidates. Such respondents also almost surely do not know the ideological positions of these candidates.

We identified all of the contested primary election races for governor and senator in California over the period 1986 to 2006 for which at least one “early” Field poll exists and in which the respondents were asked the candidate “favorability” items. Averaging across all non-incumbent candidates, 67% of all all respondents answered “don’t know.” For incumbents on the other hand, the percentage of “don’t know” answers was only 10.5%. A few examples are illustrative. In the March 1986 poll, over 70% of those polled had no opinion of Congressman Ed Zschau or political commentator Bruce Herschensohn, the two highest vote-getters in the Republican primary for senator. By contrast, only 5% had no opinion of the incumbent governor, George Deukmejian. In the January 2004 poll, over

80% had no opinion of three of the top four contenders for the Republican nomination (Toni Casey, Rosario Marin, and Howard Kaloogian), and even 34% had no opinion about the clear primary leader, former Secretary of State Bill Jones. By contrast, only 12% of respondents had no opinion about incumbent Senator Barbara Boxer.

3.5 Checking for Changes in the Composition of the Respondents

One potential concern with the above analysis is that the respondents in the Field Polls and SUSA were not the same across the polls within races. Consequently, the composition of the survey respondents in a primary might change substantially over the course of the campaign in a way that could lead to differential trends in vote intentions without there actually being learning. While we address this concern more directly in our internet panel surveys by polling the same respondents over time, we can also examine whether the results are sensitive to including various covariates in the Field Poll and whether there are any noticeable differences in observed covariates across the samples. When we include covariates – such as age, income, religion – our effect becomes slightly larger and more significant than the average interaction effect. We describe the specifications and show the results in Appendix C.

Besides including the covariates in our regressions, we can also check their stability in the primary electorate or (for the Field Poll) among conservatives and liberals separately. Finding stability would reassure us about the main finding, but finding instability does not necessarily invalidate it. Even if changes in the composition of the primary electorate take place, they might offset each other or even work against our hypothesis, and so not explain the differential trends we observe between conservative and liberal voters. We find few changes in the distributions of voter characteristics over the course of the primary campaigns and that those that might have taken place were scattered and unsystematic (see Appendix C for details).

4. Downballot Races Using Polling Data

The above results suggest that voters do learn about candidates' ideological positions

during primary election campaigns that generate relatively large amounts of information, that is, in gubernatorial and senate races. However, the campaigns for many offices do not receive the same amount of attention as these races do. Do voters in relatively low-information races, such as lieutenant governor or attorney general, also learn about their candidates' ideological positions?

In this section, we investigate whether we observe a similar pattern of learning about the ideological positions of candidates for downballot offices. Primaries for downballot offices often involve substantial ideological conflict – e.g., in the 1990 Republican primary for lieutenant governor in California one of the contestants accused the other of being insufficiently conservative during the Republican state convention. However, the amount of information about the candidates made available to voters during the primary campaigns tends to be relatively sparse, as the media often fails to cover these races and the candidates themselves have limited campaign resources. For example, the *Los Angeles Times* published, on average, 295 articles for each gubernatorial primary and 115 articles for each senatorial primary, but only published 44 articles, on average, for an attorney general primary, the downballot primary with the most media exposure. There were even fewer articles for other downballot primaries. Between 2002 to 2007, downballot candidate spending was on average 20% of gubernatorial candidate spending. We might therefore expect that voters would not learn as much about downballot candidates' ideological positions as senate or gubernatorial candidates.

The California field polls included surveys of voter preferences for a number of downballot offices (see Table A.2). Thus, we can estimate whether voters learn about the ideological positions of the downballot candidates using the same saturated model as in the above section. Again, we are mainly interested in the coefficients estimates on the interaction between the late polls and voter ideology, $\hat{\gamma}_j$. Since voters are not exposed to as much information about downballot candidates, we would expect the magnitude of $\hat{\gamma}_{\text{avg}}$ to be smaller for these races as compared to our estimate of this coefficient for the “baseline” gubernatorial and senate races in Table 1.

Because of the scarcity of newspaper coverage of downballot candidates, we could not

have independent raters evaluate candidate ideology by coding newspaper sentences, as we did with top-of-the-ticket primaries. Instead, we used a variety of sources including roll-call voting scores, endorsements, newspaper articles, and various other media sources to classify the downballot candidates in 36 races.²⁵ For several races we were unable to identify the relative ideological positions of the candidates. In some cases, this was due to the candidates having similar ideological positions. More often we were unable to identify the positions due to lack of information about the candidates’ issue positions. Thus, similar to the “baseline” estimate for governors and senators, our “baseline” estimate of γ_{avg} for downballot races excludes the cases where we could not determine the relative ideological positions. We also estimate γ_{avg} for the two types of races that tend to receive slightly more media coverage and campaign resources: lieutenant governor and attorney general.

In the first column of Table 3, we present our “baseline” estimates of γ_{avg} . We find little evidence that voters were learning about the ideological positions of the candidates. The coefficient estimates are close to zero. The estimates are also small in magnitude and not statistically significant when we focus on the “placebo” cases where we had difficulty determining the ideological positions of the candidates.

In the third column of Table 3, we present our estimates of γ_{avg} for the primary campaigns for lieutenant governor and attorney general. Even in these downballot races that receive slightly more media coverage, $\hat{\gamma}_{\text{avg}}$ is still relatively small and not statistically significant. Moreover, the F-test fails to reject the null hypothesis that all the interaction coefficients are zero for each sample of races.²⁶

In contrast with our top-of-the-ticket findings, where voters appear to learn about can-

²⁵When we were unable to use roll-call voting scores or endorsements, we searched various media sources, primarily on-line newspapers, for information to identify the relative ideological positions of the candidates. When the roll-call voting records disagreed with information from other media sources, we based our coding on the roll-call voting records. Thus, while Bowen is often portrayed as having a more liberal position than Ortiz in the media, she had a more moderate roll-call voting position relative to Ortiz.

²⁶In separate analyses, we find little evidence that the lack of a significant interaction effect for downballot offices reflects substantial learning about the candidates’ ideologies prior to the campaign period. Across all down ballot races, our estimates of the conservative main effects, $\hat{\beta}_{jC}$, are distributed around zero. We also fail to find the error correction we observe in senate and gubernatorial races (see Figure 4). Instead, we mostly observe regression to the mean – i.e. races with large main effects on voter conservative ideology tend to be associated with large negative coefficients on the conservative ideology and late poll interactions terms.

didates' ideological positions and act on that knowledge, these results raise concerns about voters' ability to make informed judgments in downballot primaries.²⁷ We return to this point in the conclusion.

5. Panel Survey Analysis

Our analysis above relied on two assumptions that we investigate further by conducting internet panel surveys. First, we assumed that learning the candidates' ideologies lies behind the increased associations we found between ideology and vote choice. Second, we assumed that reverse causation – people adjusting their own ideology to match that of their preferred candidate – is absent. Both assumptions seem plausible in part because of the kinds of races we study: primary races without incumbents. In these contests, voters presumably had much to learn and lacked strong views about the candidates and thus seem unlikely to have adopted (or rejected) candidates' ideologies. Since these races are intra-party, voters were also not necessarily learning their parties' positions (only potential party nominees' positions) and so were unlikely to be learning and then following the ideologies of their parties.

To provide further assurance, however, we conducted our own panel surveys. We interviewed likely voters before and after seven gubernatorial and senate primaries in 2010 (see Table 4). With these surveys, we can test the learning assumption because we can determine who learned the candidates' ideologies and who did not. Since we also measured voters' ideologies earlier in the campaigns, we can be more certain that increased voter-candidate ideological congruence arises, not from voters changing their ideology, but from changing their votes.

Our panel surveys are not probability samples of voters. We recruited participants online through our own ads and through a consumer internet panel.²⁸ Although not perfectly representative, the samples are generally representative of the public on demographics (see Appendix D) and the panel nature of our data nevertheless allows us, to some extent, to vet

²⁷Of course, voters may be learning about nonideological aspects of these candidates.

²⁸We recruited subjects through online ads shown on Facebook and through Decision Analysts (see Appendix D for details). We considered conducting a probability phone survey or a representative Internet panel through firms such as Knowledge Networks or Polimetrix, but the costs were prohibitive, into the hundreds of thousands of dollars, while ours cost only a few thousand dollars.

further the cross-sectional findings. We discuss concerns with self-selection after presenting the findings.

In analyzing these data, we follow the earlier analyses as closely as possible. To classify candidates as liberals or conservatives, we rely on respondents' perceptions and on coding of newspaper sentences, which largely agreed. We present respondents' perceptions in Table 4 and show the newspaper coding in the Appendix E, where we also discuss the minor differences between them and present robustness checks. As shown in Table 4, most of the races we surveyed had a noticeable ideological gap between the major candidates. In the Republican New Hampshire senate primary, for example, our respondents placed the most moderate candidate, Bill Binnie, at 4.0 on a seven-point ideology scale and the conservative candidate, Ovide Lamontagne, at 6.4. Two races, however, had small gaps: the New Hampshire Republican gubernatorial primary and the Pennsylvania Republican senate primary. We exclude these races and the race with an incumbent (Pennsylvania Democratic senate race), leaving five primaries: the Michigan gubernatorial primaries for both parties, the New Hampshire Republican senate primary, and the Pennsylvania gubernatorial primaries for both parties.

Between the early and late interviews, many respondents learned about the candidates' positions. Measuring this learning, however, requires numerous coding decisions, especially in multi-candidate races. We adopt a simple approach and use a minimal definition of knowledge: we code respondents as knowing the positions if they placed at least one conservative candidate to the right of a liberal candidate.²⁹ Using this minimal test, about 45% passed early in the primary and 75% did so after the primary, for a net increase of 30%. This 30% net increase arises from 33% learning and 3% "unlearning."³⁰ The increase in minimal knowledge varies from race to race, as shown in Table 5, and is statistically significant at conventional levels for all the races with an ideological difference between the candidates. In contrast, the bottom half of Table 5 also shows that considerably less learning occurred in the two placebo races and the one race with an incumbent.³¹

²⁹The results are similar with other codings.

³⁰These statistics are for respondents who survive the panel and exclude the two "placebo" races and the race with the incumbent.

³¹Research has produced mixed results as to whether campaigns increase or reduce the knowledge gap

Table 4: Internet Panel Surveys and Perceptions of Ideology

Contest Description	Ideological Placements (7-point Scale)			Vote Coding	N
	Pre-election	Post-election			
Michigan Gov Dem					156
Virg Bernero	2.5	2.2	0		
Andy Dillon	3.9	4.6	1		
Michigan Gov Rep					195
Rick Snyder	5.5	5.2	0		
Mike Cox	5.5	5.7	1		
Mike Bouchard	5.7	5.8	1		
Peter Hoekstra	5.8	6.0	1		
New Hampshire Gov Rep					344
Karen Testerman	5.4	5.5	Placebo		
John Stephen	5.5	5.7			
Jack Kimball	5.7	5.8			
New Hampshire Sen Rep					365
Bill Binnie	4.3	4.0	0		
Jim Bender	5.2	5.2	0		
Kelly Ayotte	5.4	5.5	0		
Ovide Lamontagne	5.9	6.4	1		
Pennsylvania Gov Dem					93
Joe Hoeffel	2.3	2.3	0		
Dan Onorato	3.7	3.7	1		
Jack Wagner	3.8	3.8	1		
Pennsylvania Gov Rep					96
Tom Corbett	5.6	4.9	0		
Sam Rohrer	6.3	6.3	1		
Pennsylvania Sen Dem					163
Joe Sestak	2.1	2.1	Excluded,		
Arlen Specter	4.2	4.2	incumbent		
Pennsylvania Sen Rep					161
Peg Luksik	5.9	5.9	Placebo		
Pat Toomey	6.3	6.3			

Not only did these panel respondents learn, but they also appeared to act on this knowledge. Table 6 shows the findings for the panel respondents, presenting estimates from a series of models. As in the prior analysis, we use the Extremist Only classification of ideology (the results are similar for other classifications).³²

Table 5: Learning About Candidates' Positions

Contest description	% Knowing Positions		Difference (net learning %)	P-value on Difference
	Pre	Post		
Michigan Gov Dem	61.2	82.1	20.9	0.001
Michigan Gov Rep	30.8	72.4	41.6	0.001
New Hampshire Sen Rep	49.8	81.7	31.9	0.001
Pennsylvania Gov Dem	42.4	62.7	20.3	0.001
Pennsylvania Gov Rep	34.6	63.0	28.4	0.001
Placebo/Incumbent Cases:				
New Hampshire Gov Rep	24.3	33.9	9.6	0.002
Pennsylvania Sen Rep	44.9	44.1	-0.8	0.867
Pennsylvania Sen Dem	77.0	86.9	9.8	0.033

Since the sample of panel respondents is less representative than the cross-sectional polls we analyzed earlier, we start by replicating the previous analyses in this panel and assess whether the results are similar. We therefore first treat the panel as if it was two separate cross-sections and estimate a variant of equation 1. Since we have individual-level data, we estimate an individual-level model in a pooled data set, including indicator variables for primary and sample type, and clustering the standard errors at the individual.³³ After

between low information and high information voters (More 1987; Holbrook 2002; Wei and Hindman 2011). In the panel data, the campaign appears to slightly reduce the knowledge gap. Among respondents who answered the four political-knowledge questions correctly, the percent “knowing” the positions increased by 30% on average. For those who correctly answered three or less, the percent increased by 35% ($p = 0.10$ on the difference). Appendix F describes the knowledge scale and shows that the effects of learning on ideological congruence are similar in the lower-knowledge group.

³²For Republican primary voters, we code sixes and sevens as conservative and fours and lower as moderates. For Democratic primary voters, we code ones and twos as liberal and fours and higher as moderates. Other codings produced similar results.

³³Since these analyses primarily assess whether we can replicate the cross-sectional results in the panel data, we do not include fixed effects in columns 1 and 2 of Table 6. When we do include them, however, the results are essentially unchanged.

excluding respondents because they failed to report an ideology or a vote preference, or because of the Extremist Only classification, the sample size in this pooled data set is 716 pre-election and 868 post-election, for a total of 1,584. As shown in Column 1 of Table 6, the cross-sectional and panel findings are reassuringly similar. The interaction between the conservative indicator variable and vote is positive and marginally statistically significant. On average, conservatives increased their support for the conservative candidate over the course of the campaign by 16 percentage points, somewhat larger than the 10 percentage points we found in the cross-sectional surveys. This difference could arise because of differences between the samples, the races we sampled, or just by chance. The important point, however, is that the estimates are not too far apart despite these differences.

To rule out reverse causation, we next present estimates based on measures of ideology only from the earlier interview. Since regression to the mean on the dependent variable can bias downwards coefficients on prior-wave variables (Finkel 1995), we modify the model for panel data:

$$V_{it} = \alpha_j + \gamma C_{it-1} + \rho V_{it-1} + \epsilon_{it} \quad (5)$$

Besides measuring conservativeness at $t - 1$, this model includes lagged vote, V_{t-1} , which addresses this potential bias and transforms the equation into a model of vote change. Since this model examines respondents from both waves, we retain only those who expressed a vote intention and an ideology in both interviews (649 of a possible 868 Extremist Only respondents). Column 2 shows that the cross-sectional estimate in Column 1 remains similar among this subset.

Column 3 presents the estimates of equation 5. The prior conservative coefficient is 0.10 and is statistically significant. It implies that a conservative became 10 percentage points more likely to vote for the conservative candidate over the course of the campaign, a somewhat smaller effect than in column 1. Since reverse causation cannot bias this estimate, we can therefore be more confident that it represents a causal effect of ideology.

We next examine whether learning lies behind this increase. To do so, we compare two groups: those who learned according to the minimal knowledge test (“learners”) and

Table 6: Panel Analyses

	(1)	(2)	(3)	(4)	(5)	(6)
Conservative Indicator	0.14 (0.043)	0.18 (0.050)				
Post-election Indicator	-0.012 (0.027)	-0.013 (0.027)				
Conserv. \times Post	0.16** (0.035)	0.16** (0.035)				
Prior Conserv.			0.10* (0.047)	0.077 (0.069)	-0.17 (0.081)	-0.040 (0.070)
Prior Conserv. \times Learning					0.39** (0.11)	0.43** (0.090)
Learning Indicator					-0.015 (0.094)	-0.053 (0.080)
Prior Vote			0.59 (0.035)	0.53 (0.062)	0.42 (0.11)	0.43 (0.097)
Prior Vote \times Learning					0.010 (0.12)	-0.029 (0.11)
Decision Analyst Sample	-0.076 (0.046)	-0.086 (0.053)	-0.17 (0.054)	-0.17 (0.083)	-0.13 (0.081)	-0.085 (0.075)
Constant	0.27 (0.035)	0.24 (0.040)	0.20 (0.037)	0.21 (0.062)	0.24 (0.074)	0.16 (0.065)
Observations	1,584	1,298	649	323	323	442
R^2	0.178	0.180	0.441	0.274	0.366	0.329
SER	0.45	0.45	0.37	0.42	0.40	0.41

Notes: Notes: Ordinary least squares regression coefficients with robust standard errors in parentheses. Dependent variable is support for conservative candidate. In columns 1 and 2, the standard errors are clustered at the individual. The N = 1584 in column 1 is from 716 pre-election and 868 post-election. Column 6 adds the 119 respondents who failed to report a vote intent pre-election but who did report a vote choice post-election, coding their vote to 0.5. ** p<0.01, * p<0.05, + p< 0.1 (shown only for key coefficients).

those who never passed the test (“never learners”). We focus on these two groups because they lend themselves to clear interpretations: the learners appear to know nothing early, but something late, while the never learners remained ignorant of even the most minimal knowledge. In contrast, drawing clear inferences about the other major group – those who passed the minimal test in both interviews – is hard because it may include individuals who learned more about the candidates’ positions between interviews. We therefore exclude respondents who passed the minimal test in both waves.³⁴ This leaves us with a sample of 323 Extremist Only respondents who learned or never learned. Column 4 shows that the effect of prior ideology is somewhat smaller in this subsample compared to the overall estimate.

To estimate whether the tendency to gravitate towards like-minded candidates arises from learning, we add an indicator for the learners and the interaction between this indicator and prior conservative ideology.³⁵ The never learners are the comparison group (the omitted category). Column 5 shows the results. The interaction between prior conservative and learning is large, 0.39. Compared to conservatives who never learned, it suggests that conservatives who did learn increased their support for the conservative candidate over the course of the campaign by 39 percentage points. As expected, conservatives who never learned failed to shift towards the conservative candidate – if anything, they shifted away. Finally, column 6 adds the 119 respondents who failed to report a vote intent pre-election but who did report a vote choice post-election, coding their vote to 0.5. The estimate of the learning effect is larger: 0.43.

Taken together, these results support the cross-sectional findings, reassuring us that learning lies behind the increased associations between ideology and vote and further ruling out the reverse causation alternative. During primary campaigns, voters appear to learn

³⁴We also exclude respondents who passed the minimal test in their first interview, but failed to do so in their second.

³⁵Specifically, we estimate

$$V_{it} = \alpha_j + \gamma_{nl}C_{it-1} + \omega L_i + \gamma_l C_{it-1}L_i + \rho_{nl}V_{it-1} + \rho_l V_{it-1}L_i + \epsilon_{it}, \quad (6)$$

where L_i is an indicator variable for learning, γ_{nl} captures the increase in voter-candidate congruence among never learners, γ_l captures it among learners, and ρ_{nl} and ρ_l do the same for prior vote.

about the candidates' ideological positions and then act on this newly acquired knowledge, shifting their votes to the candidates who agree with them.

To provide added reassurance, we also directly test for reverse causation, that is, test whether individuals adopt the ideology of their preferred candidate. We estimate models such as those in columns 3-6, but with postelection ideology as the outcome variable. If people were following their preferred candidate, we should see prior vote influencing change in voters' ideologies, especially among learners. In contrast with evidence from national elections (Lenz 2012, ch. 8), the estimates reveal no such findings. The coefficient on prior vote and its interaction with learning is always close to zero and precisely estimated.

One potential concern with these findings is self-selection into the panels. Individuals who care enough to learn the candidates' ideologies and change their votes could be extraordinarily rare in the population – too rare to explain the cross-sectional findings in earlier sections – but extremely likely to self-select into our panels. Although this seems unlikely, the panel samples undoubtedly do overrepresent politically knowledgeable and interested individuals. However, other research on learning candidate positions finds little variation by political knowledge in learning, subsequent vote changes, or subsequent opinion change (Lenz 2009, p. 826; 2012, p. 208). Moreover, our results hold up among high and among not-so-high knowledge individuals in the two states where we included a political-knowledge battery (see Appendix F). Without representative panels, we cannot be absolutely sure. Nevertheless, the panel results combined with cross-sectional findings point to learning as the mechanism, though priming likely occurs as well.

6. Conclusions

Although the conventional wisdom is that U.S. voters are not well informed about politics, a number of studies find that many voters acquire a substantial amount of information about candidates' policy positions and attributes over the course of electoral campaigns, especially for the presidential candidates. What is less well known is whether the information voters gain during campaigns actually affects their voting decisions. Our results suggest that voters learn about candidates during state primary election campaigns and that this information

affects their vote choices. This pattern is most apparent for gubernatorial and senate races with no incumbent candidate.

In downballot races, however, this pattern appears absent – we see little evidence of voter learning about ideology. This could be due to the lack of easily available information for voters. For example, as noted above, many of these races receive little media coverage, and the candidates often have little money for campaign advertising. Voters may also pay little attention to ideology in these races, since it presumably matters less for the job. Does it matter if the state treasurer, auditor, or secretary of state is liberal or conservative? This issue deserves further study, since one possible implication is that the amount of information is too low to allow voters to make meaningful choices; so, perhaps, these offices should be filled by appointment rather than elections.

Our study focuses on the effect of learning about candidate ideology, but voters may act on what they learn about candidates' other attributes as well. Do voters in primaries learn about candidates' positions on salient and divisive issues – e.g., gay marriage inside the Democratic party, or immigration policy inside the Republican party – and vote on the basis of these positions? Voters might also learn about “electability.” If voters are strategic in the sense that they look ahead to the general election, then electability might interact in subtle and interesting ways with candidate ideology. For example, voters who are extremely strategic and well informed might even respond to the events in the primary campaign of the opposing party.

Do our findings reflect learning or priming? The evidence supporting learning is considerable. Given how little voters know about the nonincumbent primary candidates early in the races, learning would have to occur before or in conjunction with any priming (Section 3.4). Our finding that increases in candidate-voter ideological congruence is positively related to the size of the ideological gap between the candidates (Section 3.2) as well as the lack of early congruence (Section 3.3) is also consistent with learning. Finally, while not representative, the panel studies find a tight link between learning the candidates' ideologies and increases in voter-candidate congruence and fail to support the reverse causal pathway (Section 5). Nonetheless, since learning and priming are often deeply entangled – messages

meant to inform voters about candidates' positions on issues may also prime those issues and messages meant to prime issues may also convey information – part of the effect we are observing may be due to priming, as both may lead to ideological congruence. Further disentangling the learning and priming effects is potentially an area for future research.

Finally, our study says nothing about how voters learn. Is it through media coverage? Campaign advertising? Debates? Word-of-mouth? Are there other cues – race, ethnicity, and gender – that help voters identify or make educated guesses about candidates' ideologies? At this point, we cannot say, but future research can address these important questions.³⁶ For example, if campaign advertising by candidates and interest groups proves to be a key channel, this would have important implications for campaign finance reforms.

³⁶Some previous studies find that in nonpartisan elections voters use proxies for party ID to help decide how to vote, when such proxies are available. See, e.g., Smith and Squire (1988) on judicial reconfirmations in California, and McDermott (2005) on candidate occupation designations for downballot California offices.

REFERENCES

- Abramowitz, Alan I. 1978. "Impact of a Presidential Debate on Voter Rationality." *American Journal of Political Science* 22(3): 680-90.
- Alvarez, R. Michael. 1997. "Why Study Only Presidential Campaigns? Statewide Races as a Laboratory for Campaign Analysis." Unpublished manuscript, California Institute of Technology.
- Alvarez, R. Michael. 2001. *Information and Elections*. Ann Arbor, MI: University of Michigan Press.
- Atkin, Charles and Gary Heald. 1976. "Effects of Political Advertising." *Public Opinion Quarterly* 40:2:216-228.
- Bartels, Larry M. 1993. "Messages Received: The Political Impact of Media Exposure." *American Political Science Review* 87: 267-285.
- Bartels, Larry M. 2002a. "Beyond the Running Tally: Partisan Bias in Political Perceptions." *Political Behavior* 24(2): 117-50.
- Bartels, Larry M. 2002b. "The Impact of Candidate Traits in American Presidential Elections." In *Leaders' Personalities and the Outcomes of Democratic Elections*, edited by Anthony King, 44-70. Oxford: Oxford University Press.
- Benoit, William L. and Glenn J. Hansen. 2004. "Presidential Debate Watching, Issue Knowledge, Character Evaluation, and Vote Choice." *Human Communication Research* 30:1:121-144.
- Berelson, Bernard R., Paul F. Lazarsfeld and William N. McPhee. 1954. *Voting: A Study of Opinion Formation in a Presidential Campaign*. Chicago: University of Chicago Press.
- Campbell, Angus, Philip E. Converse, Warren E. Miller, and Donald E. Stokes. 1960. *The American Voter*. Chicago: University Of Chicago Press.
- Carsey, Thomas M., and Geoffrey C. Layman. 2006. "Changing Sides or Changing Minds? Party Identification and Policy Preferences in the American Electorate." *American Journal of Political Science* 50(2): 464-77.
- Clarke, Peter, and Susan H. Evans. 1983. *Covering Campaigns: Journalism in Congressional Elections*. Stanford, Calif.: Stanford University Press.
- Craig, Stephen C., James G. Kane, and Jason Gainous. 2005. "Issue-Related Learning in a Gubernatorial Campaign: A Panel Study." *Political Communication* 22: 483-504.
- Druckman, James N. 2005 "Media Matter: How Newspapers and Television News Cover Campaigns and Influence Voters." *Political Communication* 22: 463-481.

- Fan, David P. and Albert Tims. 1989. "The Impact of the News Media on Public Opinion: American Presidential Elections, 1987-1988." *International Journal of Public Opinion Research* 1:2:151-63.
- Finkel, Steven E. 1993. "Reexamining the 'Minimal Effects' Model in Recent Presidential Campaigns." *The Journal of Politics* 55:1:1-21.
- Finkel, Steven E. 1995. *Causal Analysis with Panel Data*. Thousand Oaks, California: Sage Publications, Inc.
- Frank, Thomas. 2004. *What's the Matter with Kansas?*. New York: Henry Holt and Co.
- Gelman, Andrew, and Gary King. 1993. "Why are American Presidential Election Campaign Polls So Variable When Votes are So Predictable?" *British Journal of Political Science* 23: 409-451.
- Gronke, Paul. 2000. *The Electorate, the Campaign, and the Office: A Unified Approach to Senate and House Elections*. Ann Arbor, MI: University of Michigan Press.
- Henderson, Mike. 2009. "Campaign Learning and Issue Voting: Evidence from the 2000, 2004, and 2008 Elections." Unpublished manuscript.
- Highton, Benjamin. 2006. "Alternative Tests for the Effects of Campaigns and Candidates on Voting Behavior." In *Capturing Campaign Effects*, edited by Henry E. Brady and Richard Johnston. Ann Arbor, MI: University of Michigan Press.
- Hollbrook, Thomas M. 1996. *Do Campaigns Matter?* Thousand Oaks, CA: SAGE Publications.
- Holbrook, Thomas M. 2002. "Presidential Campaigns and the Knowledge Gap." *Political Communication* 19(4): 437-54.
- Iyengar, Shanto, and Donald Kinder. 1987. *News that Matters: Television and American Opinion*. Chicago: The University of Chicago Press.
- Iyengar, Shanto and Adam F. Simon. 2000. "New Perspectives and Evidence on Political Communication and Campaign Effects." *Annual Review of Psychology* 51:1:149-169.
- Johnston, Richard, André Blais, Henry E. Brady, and Jean Crête. 1992. *Letting the People Decide: Dynamics of a Canadian Election*. Stanford, Calif.: Stanford University Press.
- Krasno, Jonathan S. 1994. *Challengers, Competition, and Reelection: Comparing Senate and House Elections*. New Haven, CT: Yale University Press.
- Lazarsfeld, Paul F., Bernard R. Berelson, and Hazel Gaudet. 1948. *The People's Choice: How the Voter Makes up His Mind in a Presidential Campaign*. Second edition. New York: Columbia University Press.
- Lenz, Gabriel S. 2009. "Learning and Opinion Change, Not Priming: Reconsidering the Evidence for the Priming Hypothesis." *American Journal of Political Science* 53: 821-837.

- Lenz, Gabriel S. 2012. *Follow the Leader? How Voters Respond to Politicians' Policies and Performance*. Chicago, IL: University of Chicago Press.
- Lipsitz, Keena. 2004. "The Significance of Rich Information Environments: Voter Knowledge in Battleground States." Paper presented at the annual meeting of the Midwest Political Science Association, Chicago.
- Martinez, Michael D., Stephen C. Craig, and James G. Kane. 2001. "Information and Voting: A Panel Study." Paper presented at the annual meetings of the American Political Science Association, San Francisco.
- McDermott, Monika L. 2005. "Candidate Occupations and Voter Information Shortcuts." *The Journal of Politics* 67:1:201-219.
- Moore, David W. 1987. "Political Campaigns and the Knowledge-Gap Hypothesis." *Public Opinion Quarterly* 51(2): 186-200.
- Patterson, Thomas E. and Robert D. McClure. 1976. *The Unseeing Eye: The Myth of Television Power in National Politics*. New York: G.P. Putnam's Sons.
- Rice, Laurie. 2004. "Learning About Presidential Primary Candidates: Campaign Advertising, Learning, and Recall." Paper presented at the annual meeting of the American Political Science Association, Chicago.
- Sears, David O. and Steven H. Chaffee. 1979. "Uses and Effects of the 1976 Debates: An Overview of Empirical Studies." In *The Great Debates*, edited by Sidney Kraus. Bloomington IN: Indiana University Press.
- Shaw, Daron R. "The Effect of TV Ads and Candidate Appearances on Statewide Presidential Votes, 1988-96." *The American Political Science Review* 93:2:345-61.
- Squire, Peverill and Eric Smith. 1988. "The Effect of Partisan Information on Voters in Nonpartisan Elections." *The Journal of Politics* 50:1:168-79
- Tedin, Kent L. and Richard W. Murray. 1981. "Dynamics of Candidate Choice in a State Election." *The Journal of Politics* 43:2:435-455.
- Vavreck, Lynn, Constantine J. Spiliotes, and Linda L. Fowler. 2002. "The Effects of Retail Politics in the New Hampshire Primary." *American Journal of Political Science* 46(3): 595-610.
- Weaver, David H. 1996. "What Voters Learn from Media." *Annals of the American Academy of Political and Social Science* 546: 34-47.
- Wei, Lu, and Douglas Blanks Hindman. 2011. "Does the Digital Divide Matter More? Comparing the Effects of New Media and Old Media Use on the Education-Based Knowledge Gap." *Mass Communication and Society* 14(2): 216-35.
- Zaller, John. 1992. *The Nature and Origins of Mass Opinion*. New York: Cambridge University Press.

Zaller, John. 1994. "Politics as Usual: The Rise and Fall of Candidate Perot." Manuscript. University of California, Los Angeles.

Zhao, Xinshu and Glen L. Bleske. 1993. "The Effect of 'Horse Race' Reporting in Increasing Voters' Issue Knowledge." Paper presented at the annual meeting of AEJMC, Kansas City.