Do Primary Elections Exacerbate Congressional Polarization?

Short title: Primary Elections and Congressional Polarization

Anthony Fowler, Harris School of Public Policy, University of Chicago, 1307 East 60th St, #2031, Chicago, IL 60637; anthony.fowler@uchicago.edu.

Shu Fu, School of International and Public Affairs, Shanghai Jiao Tong University, No. 1954 Huashan Rd, Shanhai, China; fushu@sjtu.edu.cn.

Abstract

Do primary election incentives cause elected officials to take more extreme or partisan positions? We study this question for members of the U.S. Congress by leveraging variation in primary election dates across states. Implementing differences-in-differences designs that account for idiosyncratic differences between each member in each Congress and each bill by party, we test whether members vote differently before or after their state's primary election date. Members of Congress cast more ideologically extreme votes before they have secured their party's nomination, but the substantive magnitude of this effect is small, explaining approximately one percent of congressional polarization. We further find that the polarizing effect of primary elections is greater in the Senate, smaller on party-priority legislation, greater for more moderate members; and smaller in states utilizing nonpartisan primaries.

Keywords: primary elections, polarization, Congress, legislative politics

Supplementary material for this article is available in the online appendix. Replication files are available in the *JOP* Dataverse (https://dataverse.harvard.edu/dataverse/jop). The empirical analysis has been successfully replicated by the *JOP* replication analyst.

Members of the U.S. Congress do not closely represent the preferences of their constituents, with Republicans positioned consistently to the right and Democrats to the left of the median voter in their respective constituencies (e.g., Bafumi and Herron 2010). The ideological differences between congressional Democrats and Republicans are substantively large, normatively troubling, and increasing over time (see McCarty 2019). Paradoxically, voters prefer more ideologically moderate candidates, so both parties could win more seats if they were willing to field more moderate candidates (Hall 2015, 2019). Why, therefore, are our elected officials so extreme relative to their constituents, and why is there so much polarization in Congress? One common explanation is the fear of being "primaried." Perhaps incumbent members of Congress worry that if they are too moderate, a more extreme candidate will challenge and defeat them in their party's primary election.

In this paper, we test whether and to what extent primaries exacerbate congressional polarization as claimed. More specifically, we examine whether members of Congress would vote differently if they were less worried about primary competition. We implement a staggered differences-in-differences design that exploits the fact that different states hold their primary elections at different times. Analyzing approximately 8 million roll-call votes taken in Congress from 1995 to 2022, we test whether members vote in more partisan or ideologically extreme ways before the primary election date. Our design allows for idiosyncratic differences between each member in each two-year Congress and between each bill by party.

We find that members become less partisan and more ideologically moderate after their primary date has elapsed. However, the magnitudes of these effects are substantively small. Our estimates imply that if members always voted as if they had just won their primary election, the average difference in voting behavior between Democrats and Republicans in Congress would be approximately 1 percent smaller. Additionally, we find that the polarizing effects of primary elections are greater in the Senate than the House, greater for members who are less likely to seek reelection

again, smaller for party-priority legislation, and weaker in states that utilize a non-partisan primary system.

Related Literature

A large literature shows that general elections in the U.S. have a moderating effect. All else equal, more moderate candidates are more likely to win general elections (Ansolabehere, Snyder, and Stewart 2001; Canes-Wrone, Brady, and Cogan 2002; Fowler 2020; Fowler et al. 2023; Hall 2015; Hall and Thompson 2018), although the electoral advantages of moderation may have declined in recent years (Canes-Wrone and Kistner 2022; Utych 2020). Other studies have found little evidence that members of Congress moderate their positions in response to electoral incentives (Fowler and Hall 2016; Poole 2005; Stone 1980), suggesting that most of the moderating effects of general elections are through selection rather than incentives. If general elections have a moderating effect, and if voters prefer moderate politicians, why are elected officials so polarized?

Pundits, commentators, and scholars often blame primary elections. Specifically, they argue that incumbents must cast partisan or extreme votes to deter or defeat primary challengers. For example, Anderson, Butler, and Harbridge-Yong (2020) argue that elected officials often reject compromises they prefer because they fear that doing so would cause them to lose primary elections. The authors acknowledge that primary electorates do not necessarily select extreme or uncompromising candidates, but nevertheless, the fear of being primaried causes all incumbents to take more extreme positions or to reject compromise. Reviewing the evidence on congressional primary elections, Drutman (2021, 33) concludes that "incumbent members of Congress fear a primary challenge, and adjust to avoid one".

This account has some theoretical grounding. For example, Palfrey (1984) shows that office-motivated candidates might diverge significantly from the median voter's preference if they are worried that they will be flanked by a more extreme candidate. In Palfrey's model, two competing candidates strategically select their platforms, and a third candidate enters and selects the platform that maximizes her vote share. Although the model does not explicitly incorporate primary elections, the purportedly polarizing impact of primaries aligns with the polarizing force introduced by the third candidate in the model. However, this model assumes that voters sincerely support the candidate closest to them. If voters are strategic, primaries will not necessarily have a polarizing effect because forward-looking primary voters may have an incentive to nominate the more moderate candidate who has a greater chance of winning the general election.

As Anderson, Butler, and Harbridge-Yong (2020) acknowledge, there is little evidence that primary elections systematically select more extreme or polarized politicians. As expected, primary-election voters are more ideologically extreme than general-election voters (Hill 2015), although the difference is modest (Sides et al. 2020). When states switched from party conventions to direct primaries, the ideological extremism of their members of Congress, if anything, decreased (Hirano et al. 2010). Primary elections appear to select higher-quality candidates (Hirano and Snyder 2014, 2019) but not ideologically extreme candidates (Cowburn 2022; Hirano et al. 2010). Relatedly, more open primaries do not seem to produce more or less extreme politicians (McGhee et al. 2014). Furthermore, voters and especially donors appear to behave strategically in primary elections (Hall and Snyder 2015), suggesting that any polarizing effect of primaries may be offset by forward-looking individuals aiming to secure victory for their preferred party in the general election. However, these findings do not preclude the possibility of incentive effects whereby politicians' fears of losing primary elections—well-founded or not—cause them to take more extreme positions.

Several studies have attempted to directly test whether primary competition changes the behavior of incumbents in Congress. Boatright (2013) finds that competitive primaries do not affect roll-call voting, although he relies on DW-NOMINATE scores, which are not well-suited to this question because they do not allow members' ideological scores to vary idiosyncratically over time (see Theriault 2016). Jewitt and Treul (2019) find that members of the majority party—but not the minority party—who defeated an extreme challenger in the previous Congress are less likely to vote with their party. Barton (2023) finds that members with an extreme primary challenger are more likely to co-sponsor partisan bills. Meyer (2022) finds that U.S. Senators with a primary challenger are more likely to vote with the party whip in the months leading up to the primary election. Analyzing one election cycle in the House, Burden (2001) finds that the estimated ideology of members did not meaningfully change after the primary election season. However, most of these studies do not study within-member variation in primary competition, few leverage the fact that primary election dates vary across states, and none account for the differences between bills that are voted on before vs. after primary election dates.

Our study builds upon this literature by assessing whether members of Congress change their roll-call voting in response to the primary election calendar. Unlike Hirano et al. (2010), our focus is not on comparing primaries with an alternative nominating system. Rather, we study the effect of the increased or decreased salience of primary competition that comes with the passing of the primary election date. When pundits and scholars argue that primaries exacerbate polarization, they are typically not comparing primaries to alternative nominating institutions. Rather, they contend that incumbents are more extreme than they would otherwise be if they weren't worried about their primary challengers. Therefore, our study evaluates this empirical claim as directly as possible by testing whether incumbents change their behavior when they are freed from concerns about a primary challenger and can shift their focus to the general election.

Theoretical Expectations

In congressional primary and general elections, candidates are faced with two distinct but overlapping electorates (Jacobson 2013). If primary elections exacerbate polarization as many scholars and observers claim, we would expect members of Congress to be less likely to cast ideologically extreme or partisan votes after their primary election date has passed. [1] In other words, we would expect members of Congress to vote in a more extreme or partisan way until their primary campaigns end. Of course, primary elections need not exacerbate polarization since primary voters and self-identified partisans are not typically more ideologically extreme than incumbent members of Congress (e.g., Bafumi and Herron 2010), and since primary voters may be forward-looking (e.g., Hall and Snyder 2015).

We utilize two different measures of roll-call voting that are described in the next section—ideologically extreme voting and partisan voting. If primary competition changes the way members cast their roll-call votes, would we expect a bigger effect for extreme or partisan voting? The answer surely depends on what kind of voting behavior members believe is likely to motivate primary voters. In the most famous examples of veteran members of Congress losing primary elections such as Joe Crowley (D-NY14) and Eric Cantor (R-VA7), their opponents were not motivated by the fact that the incumbents weren't partisan enough. If anything, perhaps they were viewed as too partisan. Instead, their opponents appear to have been motivated by the ideological moderation of the

[1] This prediction is unclear for members who lose their primary election. For them, we would predict that they should shift toward their preferred personal ideology, which could either be more or less extreme. However, very few members (on average, approximately 4 per election year in the House) lose their primary election, so these cases are not likely to meaningfully influence our estimates. Excluding members who lost their primary election has little bearing on our estimates.

incumbents. Therefore, if these anecdotes reflect a larger phenomenon, we would expect to see greater effects on ideologically extreme voting than on partisan voting.

Data and Design

We utilize data from Voteview.com on congressional roll-call votes taken in the U.S. Congress between 1995 and 2022 (Lewis et al. 2022) and from the Federal Election Commission on the congressional primary election dates by state between 1996 and 2022.[2] We also utilize information on contested primaries and retirements as inferred from official election returns and verified using Wikipedia and Ballotpedia. We assemble a data set in which every row is a member by roll-call vote. We exclude those who are not explicitly members of the Democratic or Republican parties. To avoid any potentially conflating effect of members losing reelection, we exclude votes taken after the general election date in every Congress.

Two states warrant special consideration. Louisiana uses a runoff system with no partisan primaries, so we classify the primary date as the date of the general election. In Connecticut, party nominees are often selected at the party conventions, which take place before the typically uncontested primary elections. Therefore, we exclude Connecticut from our main analyses, although in Appendix Table A4, we separately test for the effects of the party convention dates in Connecticut. This leaves us with more than 7 million member-votes for our U.S. House analyses and more than 800,000 member-votes for our U.S. Senate analyses.

[2] In most cases, the election dates are the same for the House and Senate seats within the same state, but when they differ, we account for this and match an incumbent to the relevant date for their state and chamber.

Congressional primary election dates vary considerably across states. The earliest primary elections are typically in early March of the election year, and the latest primary election dates (ignoring Louisiana) are typically in September, meaning there is a six-month period every two years during which some members of Congress have already had their primary election and others are still competing in their primary campaign. Figure A1 in the Appendix visually illustrates this variation for the 2020 election cycle.

To estimate the effect of primary elections, we run regressions of the following form:

$$\text{ExtremeVote}_{ipcb} = \beta * \text{AfterPrimary}_{ib} + \gamma_{ipc} + \delta_{pb} + \varepsilon_{ipcb},$$

where ExtremeVote_{ipb} is a binary variable indicating whether legislator *i* from party *p* in Congress *c* cast an ideologically extreme vote on bill *b*, AfterPrimary_{ib} is a binary variable indicating whether the primary election date for member *i* has passed by the time bill *b* came up for a vote, γ_{ipc} represents member-Congress fixed effects,[3] and δ_{pb} represents bill-party fixed effects.

Utilizing the method of Fowler and Hall (2012), we classify each roll-call vote as conservative or liberal,[4] and we then classify an extreme vote as a conservative vote taken by a Republican or a

[3] In the rare cases when members change parties in the middle of a Congress, we treat this as a new member, so when we refer to member-Congress fixed effects, this is a shorthand for member-party-Congress fixed effects.

[4] Specifically, we start by making an initial guess about the directionality of the bill using levels of support from each party. For example, if Republicans are more likely than Democrats to vote *yea*, then we classify the yea vote as the conservative vote on that bill. Next, using these codings, we estimate Conservative Vote Probabilities (CVP) for each member in each Congress. Then, we check whether our estimated CVP scores are positively or negatively correlated with the conservative vote on that bill. If the correlation is negative (which, in practice, is extremely rare), we flip the coding for

liberal vote taken by a Democrat. By this measure, 78 percent of all votes in the House and 79 percent in the Senate are classified as extreme. Of course, a member need not be an extremist to cast an extreme vote on a particular bill. For example, if a bill receives unanimous support from one party but not the other, everyone from the unanimous party will likely be classified as casting an extreme vote. However, all of our analyses include bill-party fixed effects, which account for the popularity of extreme voting by party on each bill. When there is relevant variation, our measure should capture the relative ideological extremism of votes cast by members of the same party on the same bill. Furthermore, congressional polarization is typically measured as the average ideological distance between members of the two parties, so if one is interested in polarization per se, our measure of extreme voting is the best way to study member-level roll-call data. More extreme voting, by this measure, will necessarily correspond with greater polarization, and vice versa.

As an alternative outcome, we also study whether members cast a partisan vote, as measured by whether their vote aligns with the majority of their party. By this measure, 92 percent of the votes in the House and 91 percent in the Senate are classified as partisan. As expected, extreme and partisan voting are positively correlated, but the correlation is not especially strong. Specifically, the correlation coefficient is approximately .31 in the House and .36 in the Senate. To understand why, consider a bill in which all Democrats and a majority of (but not all) Republicans vote yea. Most likely, the yea vote will be classified as liberal. Therefore, Republicans voting nay will be classified as taking an extreme but non-partisan vote, while the Republicans voting yea will be classified as taking a partisan but moderate vote.

---

that bill, and re-estimate CVP scores. We repeat this procedure until the CVP scores are positively correlated with conservative voting on every bill.

Our design accounts for the fact that some members are more likely than others to cast an extreme vote, and we allow this to vary idiosyncratically by Congress. We also account for the fact that an extreme vote is more likely on some bills than others, and we allow this to vary idiosyncratically by party. This is implicitly a staggered differences-in-differences design, where the passage of the primary election date occurs for different legislators at different points in time. Specifically, this design leverages within-member changes in voting behavior before and after the primary election date and compares these changes to comparable changes for other members from the same party voting on the same bills but who did not have a primary election during that period. Specifically, our fixed effects mean that we are effectively conducting separate differences-in-differences designs for each Congress by party, and we are pooling together evidence from all of them. Our parallel trends assumption is that differences in roll-call voting around the time of a member's primary election would, in expectation, be the same as differences in the roll-call voting of other members from the same party on the same bills but whose states do not have a primary election at the same time if not for the effect of primary elections. In Figure 1, we present the results of an event study which lend further credibility to the parallel trends assumption. We find no evidence of differential trends in extreme voting before the primary election date.

We modify this baseline specification in several ways to answer different questions and probe mechanisms. We separately analyze cases in which members did and did not face a primary challenger. We also analyze abstention as an outcome of interest to assess the extent to which differential abstention explains our results. We further subset our data to explore heterogeneity over time or across different types of members, bills, or electoral institutions. In the Appendix, we also test whether members change their voting behavior around filing deadlines—the dates by which challengers must announce that they are running in a primary. We find little evidence that members change their voting behavior around filing deadlines, suggesting that they do not appear to cast more extreme or partisan

votes with the goal of deterring potential primary challengers. Also in the Appendix, we combine information on filing deadlines, primary election dates, primary challengers, and primary results to estimate the effect of officially securing the nomination of one's party. In Appendix Table A3, we find that when a candidate knows that she has secured her party's nomination—either because the filing deadline has passed and she has no primary challengers or because she has officially won her primary election—she becomes less likely to cast ideologically extreme votes, but the substantive magnitude of this effect is small.

Caveats

Two-way fixed effects regressions like ours can yield biased estimates under particular types of heterogeneous effects (e.g., Baker, Larcker, and Wang 2022; Goodman-Bacon 2021). For example, if the effect of passing one's primary election date increases or decreases in the time since a member's primary election, this could contaminate our estimates of the bill-party fixed effects, which could, in turn, contaminate our estimates of the effect of interest. One solution is to estimate stacked regressions in which only units that have not yet been treated are used as comparison units (e.g., Cengiz et al. 2019). In Table A5 in the Appendix, we show that this approach yields estimates that are very similar to those arising from our baseline specification, although as expected, the estimates are less precise. Therefore, the potential biases associated with two-way fixed effects seem negligible in this case.

Our research design allows us to estimate the effect of variation in the threat of primary competition on the voting behavior of a member of Congress, holding constant the identity of that member. Therefore, we focus on the incentive effects of primaries rather than any selection effects. Of course, primaries could potentially influence congressional polarization by changing who is nominated and runs in general elections, and these effects have been explored in other studies (e.g.,

Cowburn 2022; Hirano et al. 2010). Our design holds constant the identity of members and focuses on the incentive effects of primaries. Although both effects are potentially important, the incentive effects are of greater interest when thinking about why individual members don't moderate more to perform better in general elections. Furthermore, these incentive effects are what many scholars and commentators argue explain why many elected officials take extreme positions and refuse to compromise (e.g., Anderson, Butler, and Harbridge-Yong 2020; Drutman 2021).

An additional caveat is that our design does not allow us to estimate the effect of *entirely* removing the threat of primary competition. Even after the primary election date has passed, a member of Congress might retain concerns about potential primary challenges in their *next* term. A related concern is that elected officials might consider primary incentives when setting their platform, and they may avoid changing positions because there are reputational costs associated with being seen as a flip-flopper. Consequently, our estimates may understate the overall incentive effects of primary competition. One response to these concerns is that there is evidence of recency bias in electoral politics (e.g., Bartels 2008; Healy and Lenz 2014; Huber, Hill, and Lenz 2012). Whether for rational or irrational reasons, voters appear to put more weight on recent behavior when evaluating incumbents seeking reelection. Therefore, the passage of the primary election date should coincide with a significant decrease in the relevance and salience of primary competition on roll-call voting since members are less likely to be held accountable for votes taken well before their next primary election.

To further address these concerns, in the Appendix, we estimate the effect of the primary election date for members who are unlikely to seek reelection again in the following term. We use a member's age and the competitiveness of their district to predict the probability that they will be seeking reelection in the following term. In cases where a member is very unlikely to seek reelection again, the passage of the primary election date likely means that the member no longer has to worry

about primary competition. For this reason, estimated effects in these cases, while noisier, might better reflect the effect of entirely removing primary election considerations on roll-call voting. In these unusual cases, the estimated effect of the primary election date is slightly larger than but similar to the estimate for the full sample.

Any study of roll-call data must consider the fact that the congressional agenda is selected strategically. One potential concern for our design is that majority party leaders strategically avoid contentious votes during the primary campaign season. We have several responses to this concern. First, the inclusion of bill-party fixed effects accounts for the fact that some bills divide the party more than others, so strategic agenda setting is not a challenge to internal validity. However, there could be concerns about external validity and generalizability since our identification comes from the kinds of bills being voted on during the primary campaign season. Second, as we show in Figure A1, the primary campaign season is so long and variable across states, party leaders may be unable to avoid holding votes on important and divisive bills during this period. Third, we later show analyses that focus only on substantively significant legislation and test for heterogeneity across different types of bills. Lastly, we later show little evidence of heterogeneity between the majority and minority parties. Concerns about the strategic agenda should be greater for the majority party, and yet we find similar results for members of the minority party.

Main Results

Table 1 shows our main results. Columns 1 and 2 show results for the U.S. House, and Columns 3 and 4 show results for the U.S. Senate. Columns 1 and 3 show estimated effects on casting ideologically extreme votes, and the Columns 2 and 4 show estimated effects on casting partisan votes. In the U.S. House, we estimate that members are 0.23 percentage points less likely to cast an ideologically extreme vote and 0.15 percentage points less likely to cast a partisan vote after the primary

Table 1. Effects of Primary Election Dates on Roll-Call Voting, 1995-2022

	(1)	(2)	(3)	(4)
	U.S. House		U.S. Senate	
	Extreme	Partisan	Extreme	Partisan
After Primary	−.0023**	−.0015	−.0069*	.0008
	(.0008)	(.0008)	(.0029)	(.0027)
Member-Congress FEs	X	X	X	X
Bill-Party FEs	X	X	X	X
Observations	7,119,429	7,119,429	844,018	844,018

*Standard errors, corrected for two-way clustering by member and bill, in parentheses; ** p<.01, * p<.05. Columns 1 and 2 show results for the U.S. House, and Columns 3 and 4 show results for the U.S. Senate. Columns 1 and 3 show estimated effects on ideologically extreme voting, while Columns 2 and 3 show estimated effects on partisan voting.*

election. Both estimates are in the expected direction, and the former estimate is statistically distinguishable from zero (p = .006). Therefore, primary elections appear to have a polarizing effect, but the magnitude of the effect is substantively small.

On average, across our period of study, Republicans in the House are 55 percentage points more likely to cast an ideologically conservative vote than Democrats.[5] Therefore, the 0.23 percentage point effect on ideologically extreme voting is approximately half a percent, or one two-hundredth, of the overall partisan polarization that we observe in Congress. Therefore, if we imagine that no member of Congress were worried about primary competition and all of them voted 0.23 percentage points more moderately, then overall polarization in Congress would be approximately 1 percent smaller than it currently is.

Before the primary election date has passed, the average rates of extreme and partisan voting in our sample are 78.02 and 92.53 percent, respectively. Therefore, we estimate that the rate of ideologically moderate voting shifts from 21.98 to 22.21 percent after the primary election date passes,

[5] Using the same data analyzed in Column 1 of Table 1, we regress conservative voting on an indicator for party and bill fixed effects. The resulting coefficient on party is .550.

which is approximately a 1 percent increase. We also estimate that the rate of non-partisan voting increases from 7.47 to 7.62 percent, which is approximately a 2 percent increase.

The substantive magnitude of our estimates is small even relative to the amount of idiosyncratic variation in the roll-call voting of individual members. If we regress extreme voting on member-Congress fixed effects and bill-party fixed effects, we obtain an r-squared statistic of .705, meaning that the fixed characteristics of members and bills appear to explain approximately 70 percent of the variation in extreme voting. When we add the after-primary variable to this regression, the r-squared increases, as expected, but only by a negligibly small amount (from .704868 to .704869). Therefore, the average incentive effect of the primary election date explains very little of the idiosyncratic variation in their roll-call voting.

In the U.S. Senate, we find no evidence that the primary election calendar influences partisan voting, but we do find that members become 0.69 percentage points less likely to cast an ideologically extreme vote after their primary election date has passed, and this estimate is statistically significant ($p = .017$). The magnitude of this estimated effect is approximately three times as large as that in the U.S. House but still substantively small. Because the Senate has fewer members, because only a third of them are up for reelection in a given Congress, and because they serve longer terms, each incumbent who is up for reelection is likely subject to more scrutiny from the public, the media, and potential challengers, which could potentially explain why we detect larger effects in this chamber. However, for the same reasons, our estimates are statistically less precise in the Senate than in the House.

Event Study

Figure 1 shows the results of an event study, assessing the extent to which extreme voting varies with the amount of time until or since the primary election. We compute the number of days between each vote and the primary election, and we generate indicator variables for the number of

Figure 1. Event Study

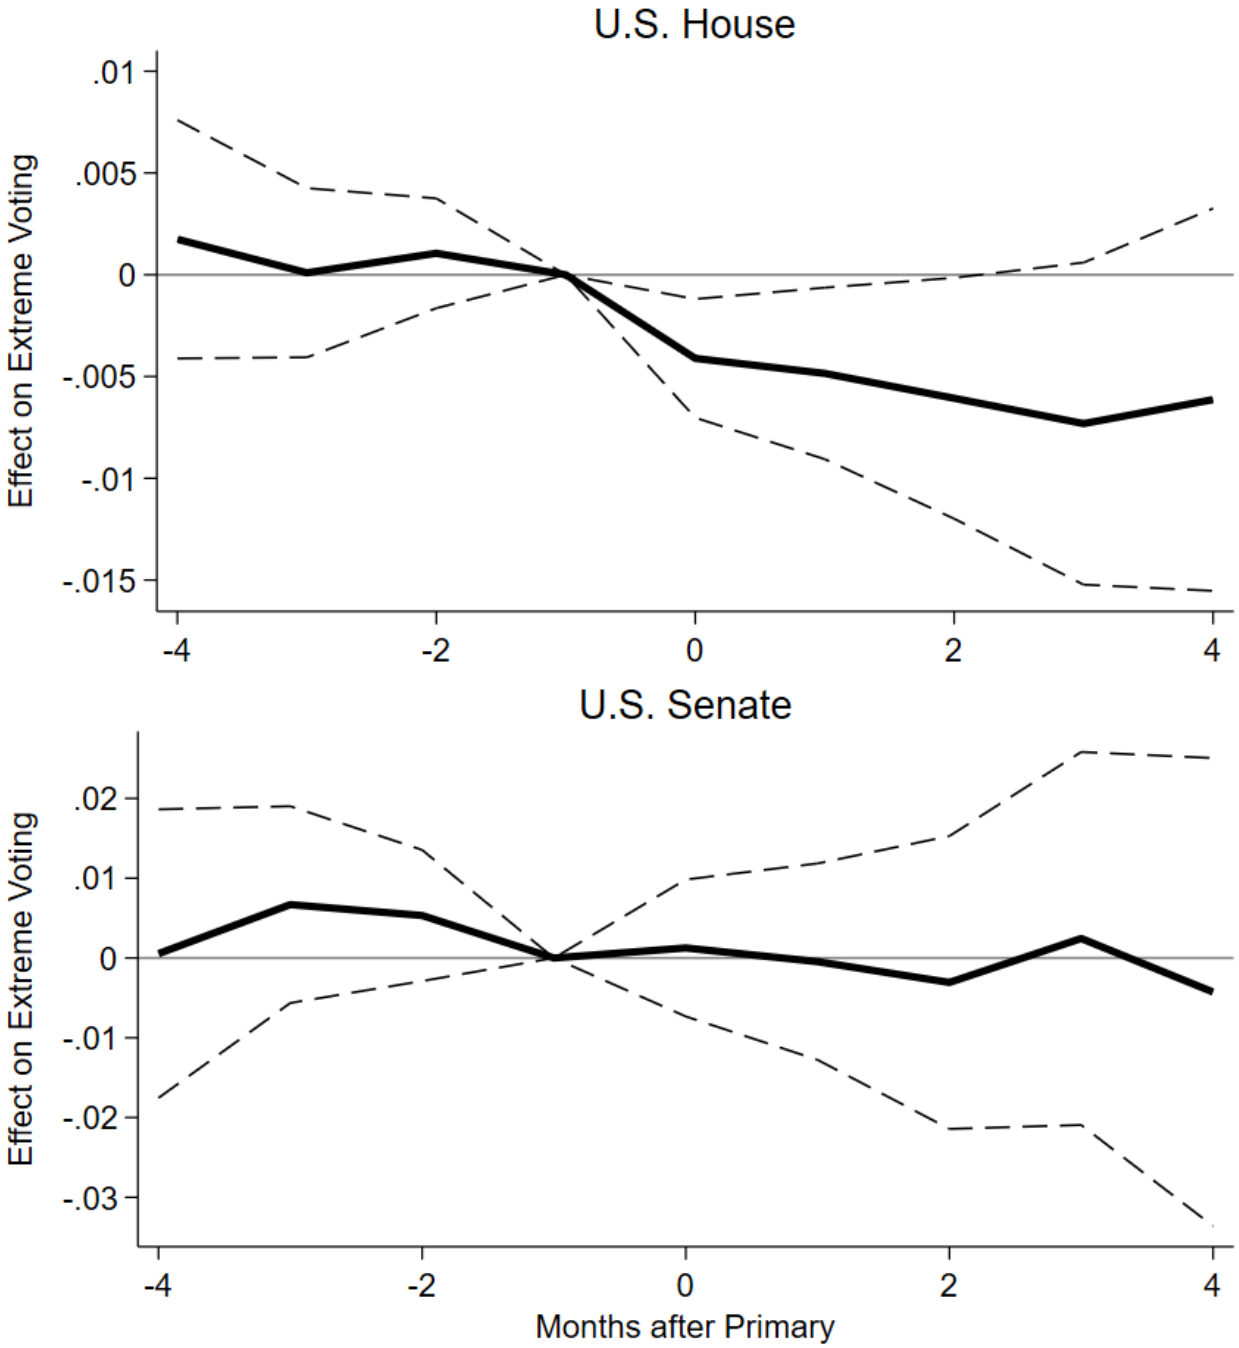

The figure shows the estimated coefficients and upper and lower bounds of the estimated 95% confidence intervals from an event-study regression. Specifically, we utilize the baseline specification but instead of a single indicator for whether a vote took place after the primary election date, we include indicators for the number of months after the primary election date. The baseline category are the votes that took place 1 month (1-30 days) before the primary election.

30-day periods (which we call *months* for convenience) between the primary election day and the roll-call vote. For example, all votes taken 0-29 days after the primary election are classified as 0 months after the primary, votes taken 30-59 days after the primary are classified as 1 month after, and so on. We drop all observations that are classified as more than 4 months before or after the primary election to focus on the most informative period for which we have the most available data.

We replicate our baseline regression specification, except that instead of including a single indicator for whether the vote came after a member's primary election date, we have multiple indicators for 4 months before, 3 months before, and so on through 4 months after. Votes classified as 1 month before are the omitted category. Figure 1 shows the estimated coefficients along with the upper and lower bounds of their estimated 95-percent confidence interval.

Unfortunately, these event-study estimates are imprecise, particularly in the Senate. In the U.S. House, we see that the estimated coefficients for the months before the primary election are all close to zero and statistically insignificant. Although we might expect to see members casting more extreme votes as the primary date draws nearer, we do not find strong evidence of this. Next, we see that the estimated effect of the primary election date emerges soon after the primary election. We also see that the effect in the U.S. House persists in the following months and, if anything, it increases slightly in the months following the primary.

Strategic Abstention

The main analyses exclude cases in which a member abstained from voting. Table 2 assesses the extent to which our results are explained by the possibility that members are more likely to abstain from voting during their primary campaign season. For ease of interpretation, we code a binary variable for each member's primary season, which takes a value of 1 if a roll-call vote takes place after their filing deadline but before their primary election date, and we code a binary variable indicating

Table 2. Does Strategic Abstention Explain the Results?

	DV = Abstain			
	U.S. House		U.S. Senate	
Primary Season	.0062**	.0068**	.0093*	.0144
	(.0012)	(.0017)	(.0045)	(.0084)
Primary Season × Moderate Member		−.0012		−.0095
		(.0022)		(.0095)
Member-Congress FEs	X	X	X	X
Bill-Party FEs	X	X	X	X
Observations	7,417,508		871,585	

*Standard errors, corrected for two-way clustering by member and bill, in parentheses; ** p<.01, * p<.05*. Moderate Member *indicates whether a member is among the more moderate half of members from the same party in the same Congress.*

whether a member abstained on a particular vote. We regress abstention on the indicator for a member's primary season, and as before, we include member-Congress and bill-party fixed effects. As expected, we find that members are more likely to abstain from voting during their primary season. This effect is 0.6 percentage points in the House and 0.9 percentage points in the Senate.

Because members are more likely to abstain during the primary campaign season, strategic abstention could be a mechanism through which primary competition influences rates of extreme and partisan voting. Perhaps members strategically abstain when they are in a competitive primary and they would have otherwise cast a moderate vote on a controversial bill. We cannot directly test for this possibility, but we can test whether ideologically moderate members are particularly likely to abstain during their primary campaign season. To do this, we code a binary variable indicating whether a member is relatively moderate. We measure the Conservative Vote Probability (Fowler and Hall 2012) of each member in each Congress using only votes taken before the first filing deadline in any state. Therefore, member ideology is measured using only votes taken before the primary campaign season began. We classify members as moderate if they are more ideologically moderate than the median member of their party in that Congress. We then interact our indicator for moderate members with our indicator for the primary season.

The interactive coefficients in Table 2 suggest that, if anything, the effect of the primary season on abstention is smaller for moderate members than for extreme members, although the difference is not statistically significant. So the effect of the primary campaign season on abstention is not greater for ideologically moderate members, which casts doubt on the idea that our results are explained by abstention. Of course, moderate and extreme members might both strategically abstain during their primary campaign season, and if they do, we think it would not make our main results any less interesting. But since the abstention effect is comparable for moderate and extreme members, we think the most likely mechanism is that, in rare cases, primary competition causes members to change their votes.

Variation over Time

We explore several different dimensions of potential heterogeneity in our estimates. Because we found consistent effects on ideologically extreme voting but not partisan voting, the remainder of the paper focuses on the effects of the primary election date on ideologically extreme voting. Figure 2 flexibly explores heterogeneity over time. Specifically, we run the specifications from the Columns 1 and 3 in Table 1 separately for different time periods, and we plot coefficients along with the 95% confidence intervals. Specifically, we plot moving average estimates from three contiguous election cycles. For example, the estimates corresponding with 1998 include the 1996, 1998, and 2000 election cycles. We could show separate estimates for each Congress or two-year election cycle, and we do so in Figure A2 in the Appendix, but these election-cycle-specific estimates are less precise.

We see suggestive evidence that these effects have changed nonmonotonically over time. Specifically, the incentive effects of the primary election date appear largest between approximately 2010 and 2014 in the House and between approximately 2006 and 2010 in the Senate. Some of this

Figure 2. Effects of Primary Election Dates on Extreme Voting over Time

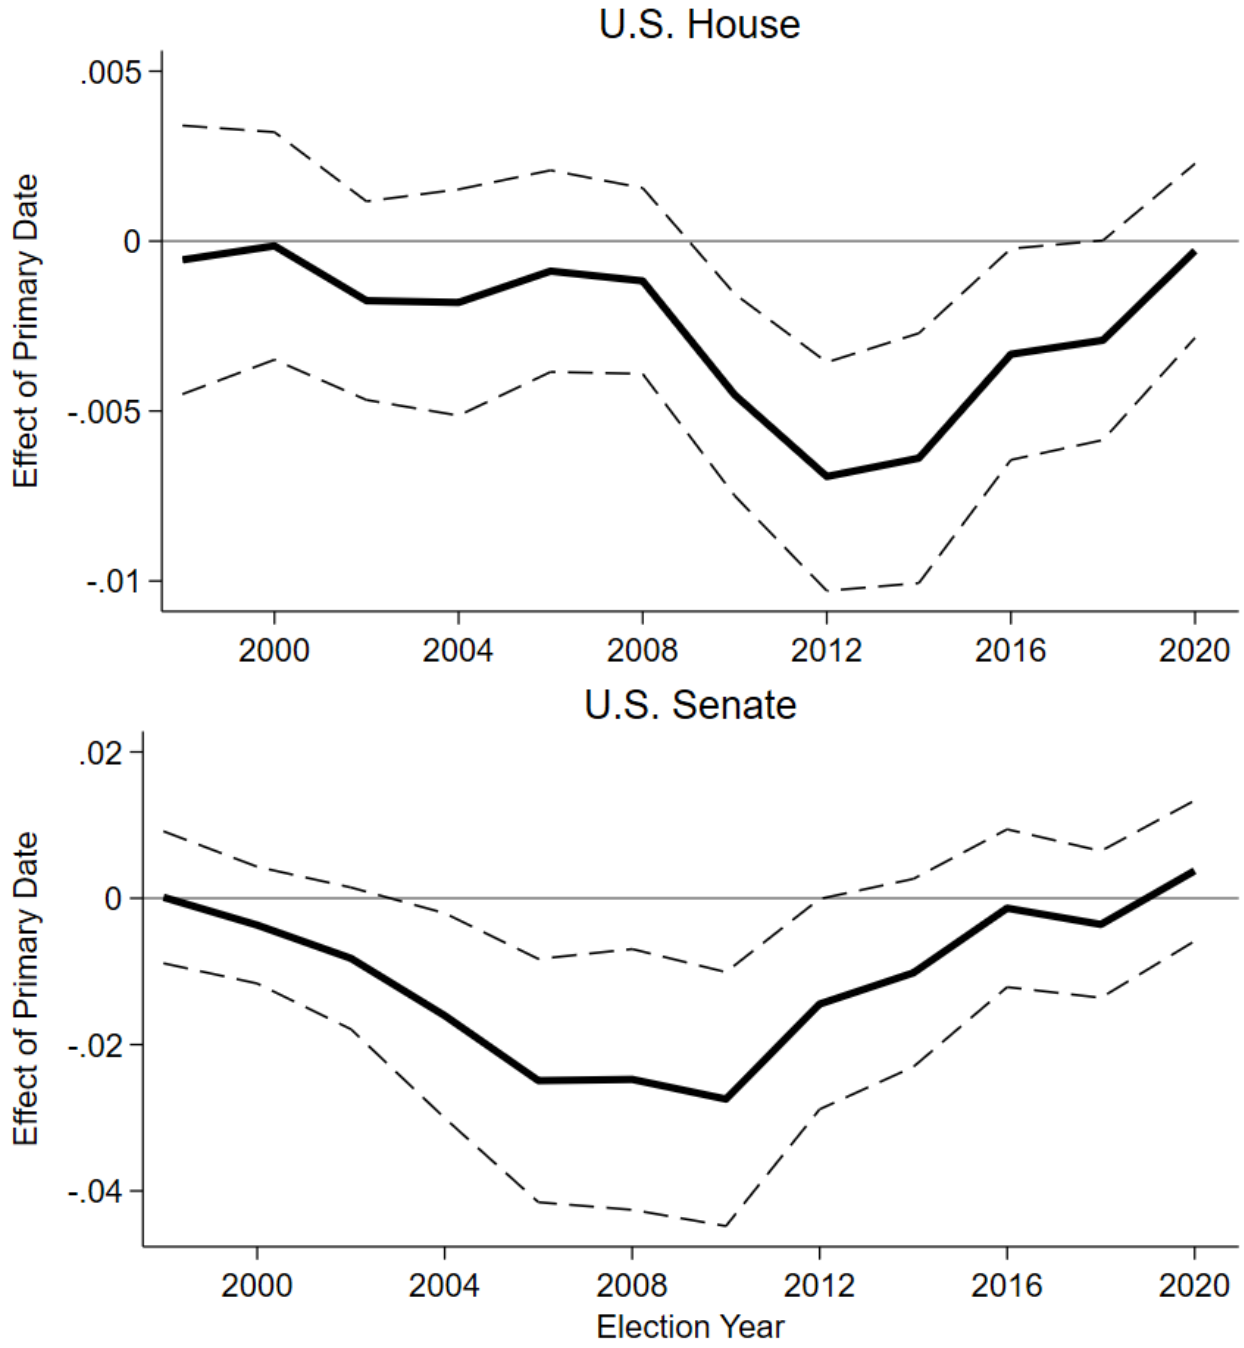

The figure shows how the estimated effects of the primary election date on extreme voting have changed over the period of our analyses. The figure shows the coefficients and 95 percent confidence intervals arising from a replication of the specification from Columns 1 and 3 in Table 1 for different time periods. Each estimate is a moving average that includes the even-year election cycle before and after that year. For example, the estimates corresponding with 1998 show the average effects across the 1996, 1998, and 2000 election cycles.

variability may be due to chance, but interestingly, we do not see evidence that the polarizing effects of primary elections are consistently increasing or decreasing over time.

One of the most prominent instances of a veteran member of Congress losing a primary election was when Eric Cantor, Republican House Minority Leader, lost to the relatively unknown Dave Brat in 2014. Although this was a salient case in the news, it appears to have followed rather than preceded any uptick in the effect of primary election dates on roll-call voting that we observe in Figure 2. Even before Cantor was primaried, members of Congress appear to have moderated their roll-call voting after primary elections, and they did not become more likely to do so after 2014.

Additional Heterogeneity

Table 3 assesses fourteen additional dimensions of heterogeneity in the effect of the primary election date on ideologically extreme roll-call voting in both the House and Senate. We run our baseline specification separately for different subsets of the data, reporting the estimated coefficients and standard errors in the table. For example, the table reports separate estimates for members who were and weren't contested in their primary election, who were retiring or not, and so on. For each potential dimension of heterogeneity, we also test whether the estimated effect of the primary election date differs between these two samples. To do this, we modify our previous specification by also interacting the after-primary indicator with the variable of interest. When we do this, we also interact our bill-party fixed effects and our member-party fixed effects with this variable of interest. The estimated difference and standard error are shown in the "Diff" column.

Unfortunately, the standard errors associated with the estimated differences are typically large relative to even the average effect of the primary election date, so most of our estimated interactive coefficients are statistically insignificant. Below, we summarize these results, and when relevant, we discuss the magnitude of the estimated differences, whether the difference is consistent with

Table 3. Heterogeneous Effects of Primary Dates on Extreme Voting

	U.S. House			U.S. Senate		
	No	Yes	Diff	No	Yes	Diff
1) Contested in primary	−.0015	−.0039**	−.0024	−.0049	−.0156**	−.0107
	(.0010)	(.0013)	(.0015)	(.0045)	(.0058)	(.0072)
2) Retiring	−.0023**	−.0021	.0001	−.0102**	−.0005	.0098
	(.0009)	(.0026)	(.0026)	(.0036)	(.0066)	(.0075)
3) Purple constituency	−.0018	−.0030*	−.0012	−.0016	−.0121**	−.0104
	(.0011)	(.0012)	(.0015)	(.0035)	(.0044)	(.0057)
4) Viable challenger	−.0036*	−.0070*	−.0034	−.0201**	−.0130	.0071
	(.0016)	(.0027)	(.0031)	(.0074)	(.0114)	(.0130)
5) Extreme challenger	−.0020	−.0091*	−.0071	−.0474	−.0121	.0353
	(.0047)	(.0035)	(.0058)	(.0621)	(.0130)	(.0434)
6) Closed primary	−.0032**	−.0001	.0031	−.0055	−.0138*	−.0083
	(.0009)	(.0022)	(.0023)	(.0032)	(.0060)	(.0065)
7) Nonpartisan primary	−.0024**	.0017	.0041	−.0068*	.0041	.0108
	(.0009)	(.0057)	(.0056)	(.0029)	(.0157)	(.0128)
8) Moderate member	−.0013	−.0030*	−.0017	−.0032	−.0115**	−.0083
	(.0009)	(.0012)	(.0015)	(.0036)	(.0042)	(.0057)
9) Republican	−.0012	−.0033*	−.0021	−.0020	−.0110**	−.0090
	(.0011)	(.0013)	(.0018)	(.0038)	(.0042)	(.0057)
10) Majority party	−.0013	−.0031*	−.0018	−.0092*	−.0046	.0046
	(.0012)	(.0013)	(.0019)	(.0046)	(.0032)	(.0054)
11) Experienced	−.0024*	−.0020	.0005	−.0100*	−.0048	.0052
	(.0011)	(.0012)	(.0015)	(.0040)	(.0039)	(.0055)
12) Close vote	−.0018	−.0026*	−.0008	−.0037	−.0118**	−.0081
	(.0010)	(.0011)	(.0013)	(.0035)	(.0038)	(.0046)
13) Priority legislation	−.0023**	−.0019	.0004	−.0117**	.0197	.0314*
	(.0009)	(.0043)	(.0042)	(.0036)	(.0112)	(.0123)
14) Procedural vote	−.0026*	−.0007	.0019	−.0042	−.0094*	−.0052
	(.0010)	(.0009)	(.0012)	(.0032)	(.0037)	(.0039)

*Standard errors, corrected for two-way clustering by member and bill, in parentheses; ** p<.01, * p<.05. The* Yes *and* No *columns show separate estimates of the effect of the primary election date on ideologically extreme voting for different subsamples. The* Diff *columns show the difference in the estimate between two subsamples of interest.*

theoretical expectations, and whether we observe the same difference in both the House and Senate. Since the average effect of the primary election date on extreme roll-call voting has a negative sign, we write that an estimated effect is greater for one subset relative to another if the estimate is more negative.

Table 3 explores heterogeneity across different electoral environments, primary institutions, characteristics of members, and characteristics of bills. The first five rows focus on variation in the

electoral incentives faced by different members. As a validity check, the first row tests whether the effects differ for members with and without a primary challenger. These analyses exclude members who did not seek reelection. Reassuringly for the design and the theory, we do not detect a statistically significant effect of the primary election date for members who did not face a primary challenger, and we detect a greater, statistically significant effect in both chambers for members who did face a primary challenger. In the second row, we separately examine members who are and are not retiring from their current position. Similarly, in both chambers, we detect no significant effect for retiring members and a larger effect that is statistically significant for non-retiring members.

In the third row, we test whether our estimated effect varies according to the partisanship of a member's district. For the House, we compute the average two-party vote share of each district across all presidential elections within the same redistricting cycle. For the Senate, we compute the average two-party vote share of each state across all presidential elections between 1996 and 2020. We classify all districts and states that are above the 25th percentile and below the 75th percentile as purple. All else equal, we would expect bigger effects in districts where members can expect a more competitive general election and a more competitive primary election. Purple districts likely have more competitive general elections but less competitive primary elections, so the differential effect is theoretically ambiguous. We detect a larger effect of the primary election in purple districts and states in both the House and the Senate.

In Row 4 of Table 3, we further test whether the effect of the primary election date is greater for members who faced a particularly viable primary opponent. We focus on incumbents who won their primary, and utilizing data from Bonica (2023), we identify their highest-performing primary opponent. We classify primary challengers who raised more than $50,000 before the primary election date as viable, and we compare the effects of a viable challenger to a non-viable challenger. This data is only available through the 2020 election cycle, so these analyses exclude the 2022 cycle. As expected,

we detect larger effects for members with a viable challenger than for those with a non-viable challenger in the House. We do not find this in the Senate, but these estimates are very imprecise.

Row 5 tests whether, among those members who faced a viable challenger, the effects of primary election dates are greater for those who faced a relatively extreme challenger. We utilize estimates from Bonica (2023) of candidate ideology inferred from the set of donors who contributed to their campaigns. We compare the ideology estimates for the incumbent and her highest-performing primary opponent. If the opponent is more conservative than a Republican incumbent or more liberal than a Democratic incumbent, we classify this as a case of an extreme challenger. We estimate larger effects for members with extreme challengers in the House but not the Senate, but again, the estimates are very imprecise in the Senate.

The next two rows of Table 3 test for variation across different kinds of primary rules or institutions. In Row 6, we test whether the effect of the primary election date is greater in states with a closed primary system, meaning that only those registered with the party can vote in a primary. Norrander and Wendland (2016) find that the voters in open and closed primaries are ideologically similar, so perhaps we wouldn't expect to find heterogeneity on this dimension. Indeed, we find no consistent evidence. Closed primaries appear to have greater effects in the Senate, open primaries appear to have greater effects in the House, and these differences could be attributed to noise.

Row 7 tests whether the effects of primary election dates differ in states utilizing nonpartisan primaries such as the top-two system in California and Washington and the top-four system in Alaska. Grose (2020) finds that nonpartisan primaries have selected more moderate legislators and induced previously elected incumbents to moderate. This could be explained by the fact that incumbents have to worry about primary competition across the ideological spectrum rather than just within their own party. If that's the case, we might expect the effect of the primary election date on extreme voting to be smaller in states with nonpartisan primaries, and that's exactly what we find. In both the House

and Senate, the estimated effect of the primary election date on extreme voting in states with nonpartisan primaries is statistically insignificant and of the opposite sign found in states with partisan primaries.

The next four rows test for heterogeneity across different kinds of members. Using the same measure of member ideology used in Table 2, Row 8 of Table 3 tests whether the effect of the primary election date differs for relatively moderate vs. extreme members. We estimate that the effect of the primary election date is more than twice as large for relatively moderate members in both chambers. Although these moderate members are more likely to represent competitive states and districts where they likely have the most to gain from ideological moderation in the general election, the polarizing effect of primary elections appears to be greater for them. This also means that the incentive effects of primaries likely increase the apparent ideological homogeneity of members within a party.

The next two rows test for differences between parties. The estimated effect of the primary election date is bigger for Republicans than for Democrats in both the House and Senate. We do not find a consistent difference between the majority and minority parties. If anything, the estimated effect is bigger for the majority party in the House and for the minority party in the Senate.

The next row tests for differences between more and less experienced members. We classify a member as experienced if they are in their fifth Congress or later, which is approximately the median level of experience in both chambers. Theoretically, we might expect smaller effects for more experienced members whose policy positions are better known by voters and who have less risk of losing reelection. There is also evidence that elected officials incorporate the positions of their primary challengers even after they have defeated them (Sulkin 2005). Therefore, we might expect experienced candidates who have faced more previous challengers to have already adapted in response to primary competition, meaning they will be less affected by current primary competition. Conversely, as previously discussed, we might also expect bigger effects for older candidates who do not plan to run

for reelection in their next term. On net, we estimate greater effects for less experienced members in both chambers, and in the Senate, the estimated effect is twice as large for less experienced members.

In the final three rows of Table 3, we test for heterogeneity across the types of votes taken by members of Congress. In Row 12, we test whether the polarizing effect of primary elections is greater for close vs. lopsided votes. Following previous studies (e.g., Snyder and Groseclose 2000), we classify a congressional vote as close if more than 35 percent and less than 65 percent of the voting members voted *yea*. In both chambers, we detect larger effects for close votes.

In Row 13, we separately examine bills that are classified by Curry and Lee (2020) as being part of the majority party agenda. In both chambers, we detect greater effects for non-priority legislation, and in the Senate, this estimated difference is statistically significant. However, this is the only heterogeneity test out of 28 in the table that is statistically significant, and we would expect a significant result to arise by chance, so this could be a false positive.

In Row 14, we test for differences between procedural and non-procedural votes. Votes to pass a bill, amend a bill, agree to a conference report, override a veto, confirm a nominee, or issue a verdict are classified as non-procedural, and all other votes are classified as procedural. We would expect more substantive disagreement over non-procedural votes, so we might expect greater effects in these cases, and that's what we find in the House but not the Senate.

Not shown in Table 3, we have also examined heterogeneity across different policy domains as classified by Clausen through the 113th Congress in the U.S. House.[6] Although we do not have clear theoretical predictions, readers may be interested to know the kinds of bills for which we detect the greatest effects. We find the largest effects of the primary election date on bills pertaining to miscellaneous policy (0.53 percentage points) and government management (0.38 percentage points),

[6] See voteview.com/articles/issue_codes.

we detect suggestive effects in the domain of foreign and defense policy (0.28 percentage points), and we do not detect substantively meaningful or statistically significant effects for bills pertaining to social welfare, agriculture, and civil liberties.

Discussion and Conclusion

Why are elected officials so polarized when voters are so moderate? This has been one of the most perplexing puzzles in American politics over the past several decades. A common explanation among pundits, practitioners, and some scholars is that elected officials are worried that if they moderate too much, they will lose their primary election to a more ideologically extreme candidate. Even if primary losses are rare in practice, nevertheless, the fear of "getting primaried" may explain much of the polarization we see among elected officials.

This paper offers the most direct test yet of this hypothesis. We exploit the fact that primary election dates vary by state, and we test whether members change their roll-call voting around these critical dates. Our differences-in-differences designs implicitly account for ideological differences between members in each Congress, and they also account for differences by party across each bill.

We find that members cast more ideologically extreme roll-call votes before their primary election date, and this is especially true of members facing a primary challenger. Therefore, the evidence is consistent with the possibility that members of Congress believe extremism helps them to defeat their primary challengers—or alternatively, moderation puts them at greater risk of losing a contested primary. We detect clearer and more consistent effects on ideologically extreme voting than on partisan voting, so primaries appear to increase ideological extremism more than they increase party cohesion.

Our results suggest that there would likely be less congressional polarization if incumbent members of Congress did not have to worry about primary competition. However, the substantive

magnitude of this effect is small. The polarizing effect of the primary campaign explains approximately one percent of the average difference between Democrats and Republicans in the U.S. House. Even when studying members who are unlikely to seek reelection again and will therefore never have to worry about another primary election, we continue to estimate substantively modest effects. So on the whole, primary elections do not explain a meaningful share of congressional polarization, and they do not resolve the puzzling disconnect between moderate voters and their polarized elected officials.

Lending additional credibility to our results, we find that the effects of primary election dates mostly vary in ways we would theoretically expect. For example, we detect larger effects in cases where a member faces a viable and relatively extreme primary challenger. We detect no significant effect in cases where a member is retiring or has no primary challenger. We detect greater effects for less experienced members who have less of an ideological reputation and potentially more reason to worry about primary challengers. We also find that the polarizing effect of primary elections is greater for close votes and non-procedural votes, suggesting that these effects could have consequences for legislative outcomes of substantive importance.

We find suggestive evidence that the polarizing effect of primary elections is greater for moderate members and those from competitive constituencies, which means primaries likely increase the homogeneity of party members. The irony is that those representing purple constituencies have the greatest risk of losing their general election, so the polarizing effect of primary elections may cause the incumbent's party to lose support in the cases where it matters most.

Even if our results were stronger than they are, this paper should not be interpreted as a recommendation against primary elections. Even if primary elections do exacerbate polarization, they also provide accountability for elected officials in partisan districts, and they likely have many other important effects that are outside the scope of this study. We do find suggestive evidence that the polarizing effect of primaries is greater for partisan primaries than for nonpartisan primaries such as

the top-two system. Therefore, although the average effect of primaries on roll-call voting is small, how we nominate candidates likely has consequences for policy, representation, and polarization.

Acknowledgements

We thank Sarah Anderson, Pam Ban, Chris Berry, Ruth Bloch Rubin, Jason Byers, Nathan Canen, Michael Crespin, Jesse Crosson, Andy Eggers, Jon Eguia, Sandy Gordon, Cathy Hafer, Andy Hall, Leander Heldring, Will Howell, Michael Kistner, Stefan Krasa, Jeff Jenkins, Dimitri Landa, Jie Lu, Cherie Maestas, Mellissa Meisels, Bryant Moy, Rachel Porter, Leah Rosenstiel, Boris Shor, Jim Snyder, Silvia Vannutelli, Adam Zelizer, and seminar and conference participants at New York University, Northwestern, Purdue, Stanford, the University of Chicago, the University of Houston, Renmin University of China, Yale, and the Midwest Political Science Association for helpful feedback.

References

Anderson, Sarah E., Daniel M. Butler, and Laurel Harbridge-Yong. 2020. *Rejecting Compromise: Legislators' Fear of Primary Voters*. Cambridge University Press.

Ansolabehere, Stephen, Jonathan Rodden, and James M. Snyder, Jr. 2006. Purple America. *Journal of Economic Perspectives* 20(2):97-118.

Ashworth, Scott. 2012. Electoral Accountability: Recent Theoretical and Empirical Work. *Annual Review of Political Science* 15:183-201.

Bafumi, Joseph and Michael C. Herron. 2010. Leapfrog Representation and Extremism: A Study of American Voters and Their Members in Congress. *American Political Science Review* 104(3):519-542.

Baker, Andrew C, David F. Larcker, and Charles C.Y. Wang. 2022. *Journal of Financial Economics* 144(2):370-395.

Bartels, Larry M. 2008. *Unequal Democracy: The Political Economy of the New Gilded Age*. Princeton and Oxford: Princeton University Press.

Barton, Richard. 2023. The Primary Threat: How the Surge of Ideological Challengers is Exacerbating Partisan Polarization. *Party Politics* 29(2):248-259.

Boatright, Robert G. 2013. *Getting Primaried: The Changing Politics of Congressional Primary Challenges*. University of Michigan Press.

Bonica, Adam. 2023. Database on Ideology, Money in Politics, and Elections: Public version 3.0. Stanford, CA: Stanford University Libraries.

Burden, Barry C. 2001. The Polarizing Effects of Congressional Primaries. *Congressional Primaries and the Politics of Representation*, eds. Peter F. Galderisi, Mami Ezra, and Michael Lyons. Lanham: Rowman & Littlefield.

Canes-Wrone, Brandice, David W. Brady, and John F. Cogan. 2002. Out of Step, Out of Office: Electoral Accountability and House Members' Voting. *American Political Science Review* 96(1):127-140.

Canes-Wrone, Brandice and Michael R. Kistner. 2022. Out of Step and Still in Office? Incumbent and Challenger Accountability in Congressional Elections across Time. *Quarterly Journal of Political Science* 17(3):389-420.

Cengiz, Doruk, Arindrajit Dube, Attila Linder, and Ben Zipperer. 2019. The Effect of Minimum Wages on Low-Wage Jobs. *Quarterly Journal of Economics* 134(3):1405-1454.

Cowburn, Mike. 2022. Partisan Polarization in Congressional Nominations: How Ideological & Factional Primaries Influence Candidate Positions. Doctoral Dissertation. Berlin: Graduate School of North American Studies.

Curry, James M. and Frances E. Lee. 2020. *The Limits of Party: Congress and Lawmaking in a Polarized Era*. Chicago: University of Chicago Press.

Drutman, Lee. 2021. What We Know about Congressional Primaries and Congressional Primary Reform. *New America* <bit.ly/4fcweSm>.

Fowler, Anthony. 2020. Partisan Intoxication or Policy Voting? *Quarterly Journal of Political Science* 15(2):141-179.

Fowler, Anthony and Andrew B. Hall. 2012. Conservative Vote Probabilities: An Easier Method for Summarizing Roll Call Data. Social Science Research Network <dx.doi.org/10.2139/ssrn.2120720>.

Fowler, Anthony and Andrew B. Hall. 2016. The Elusive Quest for Convergence. *Quarterly Journal of Political Science* 11(1):131-149.

Fowler, Anthony, Seth Hill, Jeff Lewis, Chris Tausanovitch, Lynn Vavreck, and Christopher Warshaw. 2023. Moderates. *American Political Science Review* 117(2):643-660.

Goodman-Bacon, Andrew. 2021. Difference-in-Differences with Variation in Treatment Timing. *Journal of Econometrics* 225(2):254-277.

Grose, Christian R. 2020. Reducing Legislative Polarization; Top-Two and Open Primaries Are Associated with More Moderate Legislators. *Journal of Political Institutions and Political Economy* 1(2):267-287.

Hall, Andrew B. 2015. What Happens When Extremists Win Primaries? *American Political Science Review* 109(1):18-42

Hall, Andrew B. 2019. *Who Wants to Run? How the Devaluing of Political Office Drives Polarization*. Chicago: University of Chicago Press.

Hall, Andrew B. and James M. Snyder, Jr. 2015. Information and Wasted Votes: A Study of U.S. Primary Elections. *Quarterly Journal of Political Science* 10(4):433-459.

Hall, Andrew B. and Daniel M. Thompson. 2018. Who Punishes Extremist Nominees? Candidate Ideology and Turnout Out the Base in U.S. Elections. *American Political Science Review* 112(3):509-524.

Healy, Andrew and Gabriel S. Lenz. 2014. Substituting the End for the Whole: Why Voters Respond Primarily to the Election-Year Economy. *American Journal of Political Science* 58(1):31-47.

Hill, Seth J. 2015. Institution of Nomination and the Policy Ideology of Primary Electorates. *Quarterly Journal of Political Science* 10(4):461-487.

Hirano, Shigeo and James M. Snyder, Jr. 2014. Primary Elections and the Quality of Elected Officials. *Quarterly Journal of Political Science* 9(4):473-500.

Hirano, Shigeo and James M. Snyder, Jr. 2019. *Primary Elections in the United States*. Cambridge: Cambridge University Press.

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.

Huber, Gregory A., Seth J. Hill, and Gabriel S. Lenz. 2012. Sources of Bias in Retrospective Decision Making: Experimental Evidence on Voters' Limitations in Controlling Incumbents. *American Political Science Review* 106(4):720-741.

Jacobson, Gary C. 2013. *The Politics of Congressional Elections*. 8th edition. Pearson.

Jewitt, Caitlin E. and Sarah A. Treul. 2019. Ideological Primary Competition and Congressional Behavior. *Congress & the Presidency* 46(3):471-494.

Lewis, Jeffrey B., Keith Poole, Howard Rosenthal, Adam Boche, Aaron Rudkin, and Luke Sonnet. 2022. Voteview: Congressional Roll-Call Votes Database. https://voteview.com/

Mayhew, David R. 1974. *Congress: The Electoral Connection*. New Haven: Yale University Press.

McCarty, Nolan. 2019. *Polarization: What Everyone Needs to Know*. Oxford: Oxford University Press.

McGhee, Eric, Seth Masket, Boris Shor, Steven Rogers, and Nolan McCarty. 2014. A Primary Cause of Partisanship? Nomination Systems and Legislator Ideology. *American Journal of Political Science* 58(2):337-351.

Meyer, Chase B. 2022. Getting "Primaried" in the Senate: Primary Challengers and the Roll-Call Voting Behavior of Sitting Senators. *Congress & the Presidency* 49(2):230-251.

Norrander, Barbara and Jay Wendland. 2016. Open Versus Closed Primaries and the Ideological Composition of Presidential Primary Electorates. *Electoral Studies* 42:229-236.

Palfrey, Thomas R. 1984. Spatial Equilibrium with Entry. *Review of Economic Studies* 51(1):139-156.

Poole, Keith T. 2005. *Spatial Models of Parliamentary Voting*. Cambridge: Cambridge University Press.

Sides, John, Chris Tausanovitch, Lynn Vavreck, and Christopher Warshaw. 2020. On the Representativeness of Primary Electorates. *British Journal of Political Science* 50(2):677-685.

Snyder, James M., Jr. and Tim Groseclose. 2000. Estimating Party Influence in Congressional Roll-Call Voting. *American Journal of Political Science* 44(2):193-211.

Stone, Walter J. 1980. The Dynamics of Constituency: Electoral Control in the House. *American Politics Quarterly* 8(4):399-424.

Sulkin, Tracy. 2005. *Issue Politics in Congress*. New York: Cambridge University Press.

Theriault, Sean. 2016. Review of Getting Primaried: The Changing Politics of Congressional Primary Challenges. *American Review of Politics* 35(2).

Utych, Stephen M. 2020. Man Bites Blue Dog: Are Moderates Really More Electable than Ideologues? *Journal of Politics* 82(1):392-396.

Biographical Statements

Anthony Fowler is the Sydney A. Stein, Jr. Professor in the Harris School of Public Policy at the University of Chicago, Chicago, IL 60637.

Shu Fu is an Assistant Professor in the School of International and Public Affairs at Shanghai Jiao Tong University, Shanghai, China.

Do Primary Elections Exacerbate Congressional Polarization?

Online Appendix

Variation in Filing Deadlines and Primary Election Dates across States

Figure A1 visually illustrates the variation in filing deadlines and primary election dates across states. Specifically, the figure plots the period between the filing deadline and the primary election date for each state during the 2020 election cycle. We see considerable variation in both filing deadlines and primary election dates. Even excluding Louisiana, there is an approximately 6-month period during which some states have already held their primary elections but other states have not. Furthermore, the states in the figure are sorted according to their primary election date, and no obvious patterns stand out. For example, larger states do not systematically have earlier primaries than smaller states. California and Texas have early primaries, New York is in the middle, Florida is later. Similarly, the early or late states do not clearly overrepresent particular regions, richer or poorer states, or Democratic or Republican states.

Effects of Filing Deadlines

In this section, in addition to estimating the effect of the primary election date, we also estimate the effect of filing deadlines. If primary elections exacerbate polarization, we have ambiguous predictions for the effects of filing deadlines. On one hand, members might cast more extreme or partisan votes before the filing deadline to deter primary challengers. On the other hand, if they indeed experience a primary challenge, they might cast more extreme or partisan votes after the filing deadline. Another complication is that challengers can announce their candidacy before the filing deadline, so

Figure A1. Primary Campaign Season by State in 2020 Election Cycle

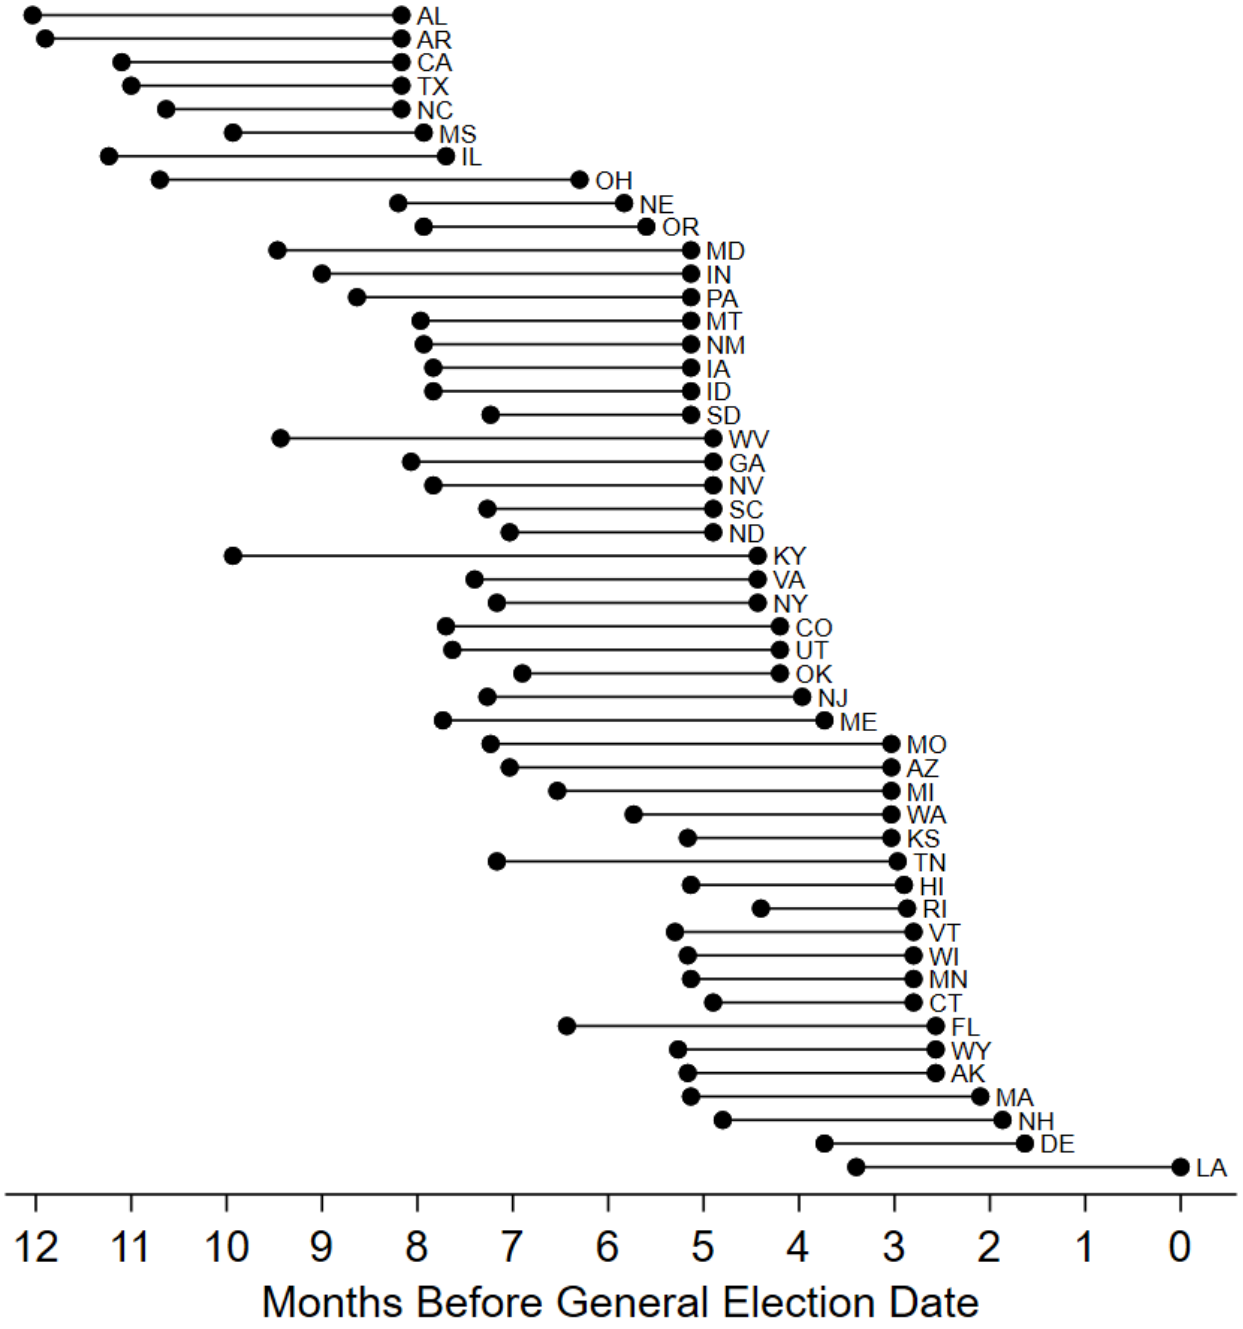

The lines show the period ranging from the filing deadline to the primary election date in each state during the 2020 election cycle. States are sorted according to their primary election date.

Table A1. Effects of Primary and Filing Dates on Roll-Call Voting, 1995-2022

	(1)	(2)	(3)	(4)
	U.S. House		U.S. Senate	
	Extreme	Partisan	Extreme	Partisan
After Primary	−.0027**	−.0018*	−.0062	.0005
	(.0009)	(.0008)	(.0036)	(.0031)
After Filing	.0024**	.0017*	−.0009	.0005
	(.0008)	(.0008)	(.0026)	(.0023)
Member-Congress FEs	X	X	X	X
Bill-Party FEs	X	X	X	X
Observations	7,119,429	7,119,429	844,018	844,018

*Standard errors, corrected for two-way clustering by member and bill, in parentheses; ** p<.01, * p<.05. Columns 1 and 2 show results for the U.S. House, and Columns 3 and 4 show results for the U.S. Senate. Columns 1 and 3 show estimated effects on ideologically extreme voting, while Columns 2 and 4 show estimated effects on partisan voting.*

if there is an effect of learning that one has a primary challenger, it need not perfectly coincide with the filing deadline.

Table A1 replicates the main results from Table 1 but also adds an additional independent variable indicating whether a roll-call vote was taken after the primary challenger filing deadline in a member's state. We find that members of the House become more likely to cast an ideologically extreme or partisan vote after their filing deadline has passed. The magnitudes of these estimated effects in the House are comparable to (but in the opposite direction of) those of the primary election date. In other words, our estimates suggest that members of the House become slightly more ideologically extreme and partisan after the filing deadline has passed, and after the primary election date, they revert back to their pre-filing-deadline levels of extremism or partisan voting. Alternatively, members of the Senate do not appear to significantly change their voting behavior after the filing deadline, but they do become less extreme (but not less partisan) after their primary election date.

As previously discussed, the average effects of the filing deadline are theoretically ambiguous. On one hand, a member might have an incentive to cast more extreme votes after the filing deadline once they learn that they will face a primary challenger. On the other hand, she might also have an incentive to cast more extreme votes before the filing deadline with the hope of deterring challengers.

Table A2. Separate Estimates for Members with and without a Primary Challenger

	U.S. House		U.S. Senate	
	Extreme	Partisan	Extreme	Partisan
After Primary	−.0041**	−.0031*	−.0154*	−.0026
	(.0013)	(.0013)	(.0066)	(.0053)
After Filing	.0014	.0013	−.0003	.0014
	(.0013)	(.0012)	(.0043)	(.0039)
Member-Congress FEs	X	X	X	X
Bill-Party FEs	X	X	X	X
Primary challenger?	Yes			
Observations	2,162,493		314,646	
After Primary	−.0020	−.0011	−.0004	.0019
	(.0011)	(.0010)	(.0053)	(.0058)
After Filing	.0026**	.0021*	−.0063	−.0018
	(.0010)	(.0009)	(.0039)	(.0042)
Member-Congress FEs	X	X	X	X
Bill-Party FEs	X	X	X	X
Primary challenger?	No			
Observations	4,307,539		299,685	

*Standard errors, corrected for two-way clustering by member and bill, in parentheses; ** p<.01, * p<.05. The table replicates the analyses in Columns 2 and 4 of Table 1, but it only includes members who sought reelection, and it separately analyzes those who faced a primary challenger (top panel) and did not face a primary challenger (bottom panel).*

Perhaps the estimates in Table A1 reflect the combination of these two competing mechanisms. In Table A2, we separately analyze these cases. Specifically, we re-run the specifications from Table A1 but focus only on members who sought reelection, and we separately analyze members who did and did not face a primary challenger. The top panel of Table A2 shows results for members who faced a primary challenger, and the bottom panel shows results for those who did not.

Reassuringly for the design, the estimated effects of the primary election date are greater for members who faced a primary challenger, and they are not statistically significant for those who did not face a primary challenger. More interestingly, we see little evidence that members cast extreme or partisan votes with the goal of deterring primary challengers. Members who did not face a primary challenger do not become systematically less likely to cast extreme or partisan votes after the filing deadline passes. We find some suggestive evidence of this in the Senate, although the estimate is not

Table A3. Estimated Effects of Officially Securing Party's Nomination

	U.S. House		U.S. Senate	
	Extreme	Partisan	Extreme	Partisan
Secured Nomination	−.0009	−.0004	−.0091**	.0004
	(.0007)	(.0007)	(.0027)	(.0027)
Member-Congress FEs	X	X	X	X
Bill-Party FEs	X	X	X	X
Observations	6,397,774		608,433	

*Standard errors, corrected for two-way clustering by member and bill, in parentheses; ** p<.01, * p<.05*. Secured Nomination *is a binary variable indicating that the member has secured their party's nomination, meaning they passed their primary election date and won or passed their filing deadline and received no primary challenger. Cases in which the member retired or lost their party's nomination are excluded.*

statistically significant, and surprisingly, we obtain a statistically significant estimate in the opposite direction in the House. This could be interpreted as evidence that some members moderate their roll-call voting in order to deter relatively moderate primary challengers, but since we do not detect a similar pattern in the Senate, it could be a false positive attributable to statistical noise.

To more precisely estimate the overall effect of primary election incentives on roll-call voting, we have conducted additional analyses that estimate the effect of a member officially knowing that she has secured her party's nomination. To conduct this analysis, we drop cases where a member retired or lost their primary election, and we code a binary variable indicating whether the member has officially secured their party's nomination. This variable is equivalent to the after-primary variable in cases where the member faced a primary challenger and the after-filing variable in cases where the member did not. The results of these analyses are shown in Table A3. In the Senate, we estimate that securing the nomination of a member's party decreases a senator's chances of casting an extreme vote by 0.9 percentage points, although we do not detect a statistically significant effect on partisan voting. The estimates are substantively small and statistically insignificant in the House. The explanation for these null results is in Table A2. Although members who faced a primary challenger cast fewer extreme and partisan votes after securing their nomination, members who did not face a challenger surprisingly cast more extreme and partisan votes after the filing deadline, and these comparably sized effects

approximately cancel each other out, meaning that members of the House do not systematically become more or less extreme after they secure their party's nomination.

Separate Estimates for each Two-year Election Cycle

Figure A2 shows separate estimates of the effect of primary election dates on extreme voting for each two-year election cycle in our analysis. Although the cycle-specific estimates are imprecise, this figure tells the same general story as Figure 2. The negative effect of the primary election date on extremism was greatest in the U.S. House around 2012 and greatest in the U.S. Senate around 2008.

Effects of Party Convention Dates in Connecticut

Table A4 estimates the effect of party convention dates in Connecticut in the U.S. House. As discussed in the main text, congressional primaries are rarely contested in Connecticut, and the party nominees are typically selected in party conventions. To assess the effects of these party convention dates, we modify the specifications from Table A1, adding in members from Connecticut. We also code a binary variable indicating whether the party convention date has passed for these members. We only have Connecticut party convention dates from 2000 and onward, so this analysis only includes data from the 106th through the 117th Congresses.

The estimated effects of filing deadlines and primary election dates is virtually unchanged when we include Connecticut in the analysis. More interestingly, we can also estimate the effect of party convention dates for members from Connecticut. Unfortunately, these estimates are statistically imprecise, but there is suggestive but not statistically significant evidence that convention dates have a similar effect in Connecticut as primary election dates in other states. Column 1 of Table A4 estimates that members from the 49 other states become 0.30 percentage points less likely to cast an ideologically extreme vote after the primary election date, and members from Connecticut become

Figure A2. Effects of Primary Dates in Each Election Cycle

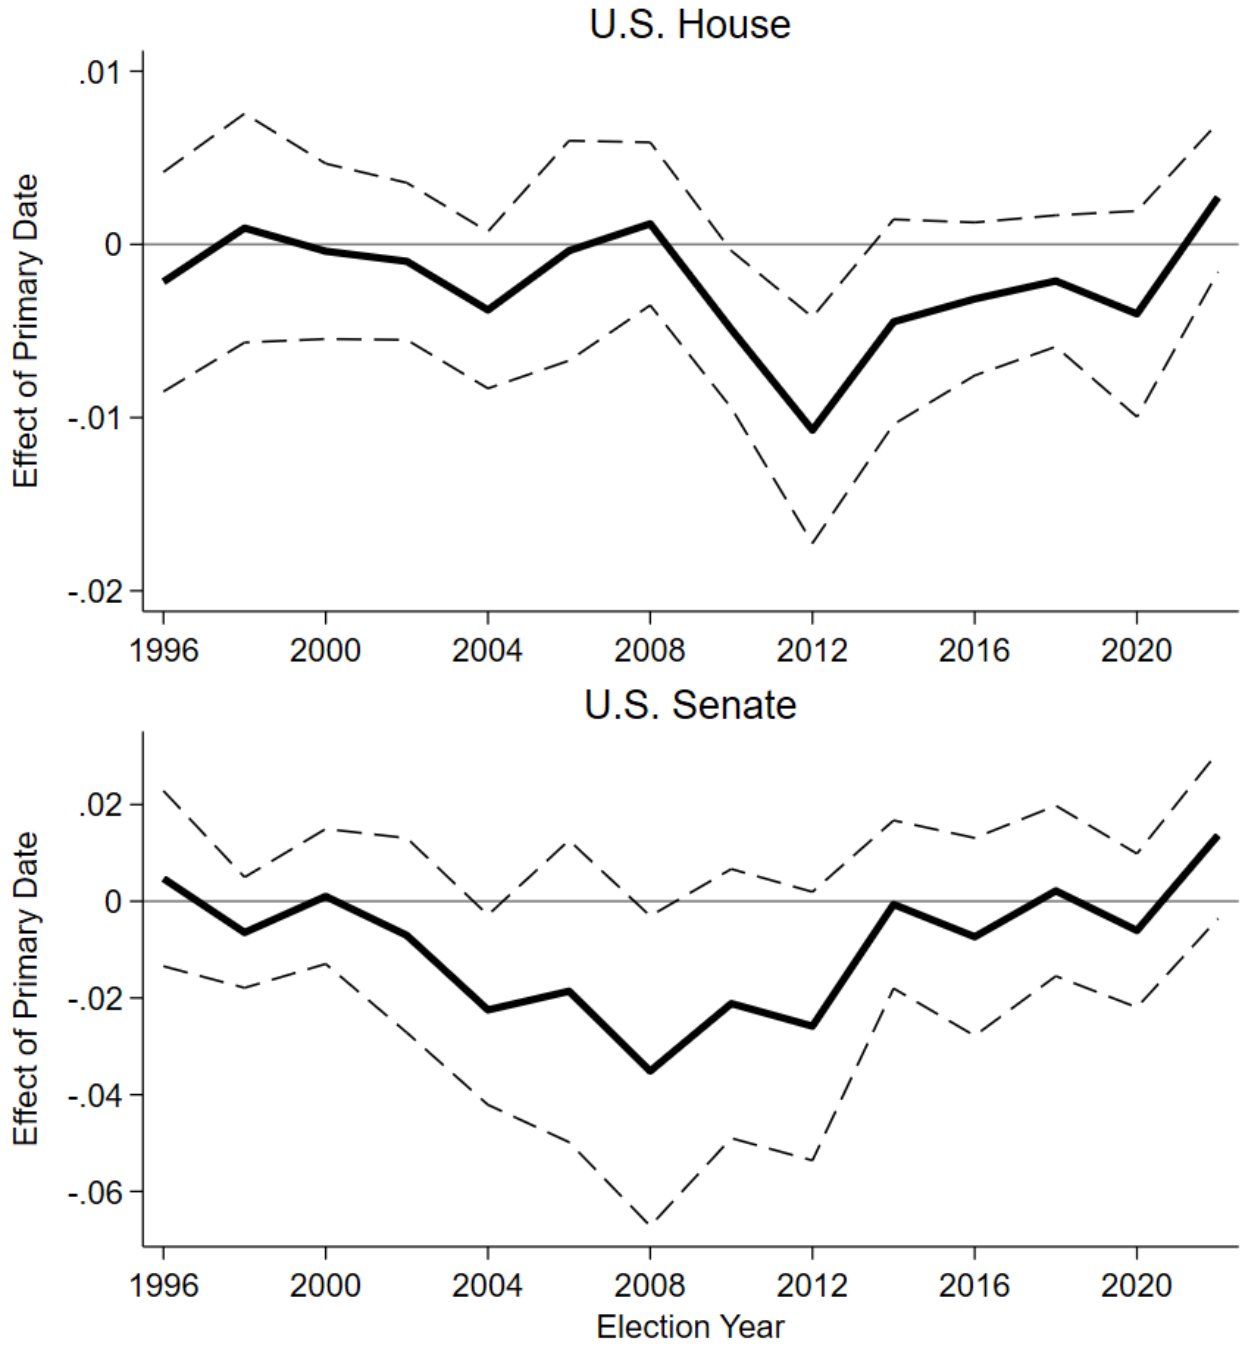

The figure shows how the estimated effects of the primary election date on extreme voting have changed over the period of our analysis. The figure shows the coefficients and 95 percent confidence intervals arising from a replication of the specification from Columns 1 and 3 of Table 1 for each individual election cycle in our analysis.

0.40 percentage points less likely to cast an ideologically extreme vote after their party convention date.

Table A4. Effect of Connecticut Party Convention Dates in U.S. House, 1999-2022

	DV = Extreme Vote	DV = Partisan Vote
After Primary	−.0030**	−.0017*
	(.0009)	(.0008)
After Filing	.0021*	.0011
	(.0009)	(.0008)
After CT Party Convention	−.0040	−.0012
	(.0057)	(.0037)
Member-Congress FEs	X	X
Bill-Party FEs	X	X
Observations	6,178,527	

*Standard errors, corrected for two-way clustering by member and bill, in parentheses; ** p<.01, * p<.05. The variables* After Primary *and* After Filing *take a value of 0 in Connecticut. The variable* After CT Party Convention takes a value of 0 in all other states.

Stacked Regressions to Address Concerns with Two-Way Fixed Effects

Table A5 shows the results of our stacked regressions that account for the potential biases of our two-way fixed effects regressions. In each of these analyses, we create a separate data set for each Congress and primary election date. That data set includes data from the members who had the primary election on this date and a subset of other members that can be thought of as clean comparison units for estimating the effect of the primary election date for the former members. We stack these data sets on top of each other, and we modify our baseline specification by interacting all of our fixed effects with each data set. In other words, we regress extreme or partisan voting on our after-primary indicator, member-Congress-dataset fixed effects, and bill-party-dataset fixed effects.

The table shows four different panels that utilize different samples of roll-call votes and different sets of comparison units. Let us start by discussing the results in the top panel. In each data set, we drop all bills that took place more than 30 days after the primary election date. We include all members that had their primary election at this time (the treated units), and we also include all members that had not yet had their primary election by the end of this period of analysis (the control units). This results in 282 different data sets in each chamber. The resulting estimates are very similar

Table A5. Stacked Regression Estimates

	U.S. House		U.S. Senate	
	Extreme	Partisan	Extreme	Partisan
After Primary	−.0029**	−.0016	−.0059	−.0027
	(.0011)	(.0011)	(.0043)	(.0042)
Member-Congress-Dataset FEs	X	X	X	X
Bill-Party-Dataset FEs	X	X	X	X
Sample	all votes before primary or up to 30 days after			
Comparison Units	haven't yet reached primary date			
Observations	45,374,755		12,252,203	
After Primary	−.0037**	−.0029*	−.0101	−.0014
	(.0012)	(.0012)	(.0055)	(.0055)
Member-Congress-Dataset FEs	X	X	X	X
Bill-Party-Dataset FEs	X	X	X	X
Sample	all votes after primary or up to 30 days before			
Comparison Units	have already passed primary date			
Observations	9,856,493		354,613	
After Primary	−.0059**	−.0039**	−.0007	.0012
	(.0013)	(.0013)	(.0061)	(.0060)
Member-Congress-Dataset FEs	X	X	X	X
Bill-Party-Dataset FEs	X	X	X	X
Sample	votes taken within 30 days of the primary date			
Comparison Units	haven't yet reached primary date			
Observations	5,477,941		1,162,498	
After Primary	−.0031*	−.0026	−.0054	−.0015
	(.0015)	(.0015)	(.0073)	(.0069)
Member-Congress-Dataset FEs	X	X	X	X
Bill-Party-Dataset FEs	X	X	X	X
Sample	votes taken within 30 days of the primary date			
Comparison Units	have already passed primary date			
Observations	5,650,632		205,346	

*Standard errors, corrected for two-way clustering by member and bill, in parentheses; ** p<.01, * p<.05. The table shows results of stacked regressions that account for potential biases associated with two-way fixed effects regressions.*

to our baseline results in Table 1. For example, the estimated effect on extreme voting in the House changed from −.0023 to −.0029, and in the Senate, the estimate changed from −.0069 to −.0059. As expected, the estimates from the stacked regressions are less precise because we use only a subset of roll-call votes, and we use fewer members to construct the counterfactual trend for each treated member. However, the estimated effect of the primary election date on extreme voting remains statistically significant in the House.

The approach in the top panel will be of interest for readers who think of the treatment as being past the primary election date and if we are worried about heterogeneous effects of this treatment. However, we could alternatively think of the treatment of interest as not yet having had a primary election, and we might be worried about biases arising from the heterogeneous effects of primary competition. To address this concern, we replicate the approach from the top panel but we change the set of bills included and the set of clean comparison units. Specifically, for each data set, we drop all votes that occurred more than 30 days before the primary election date, and we only utilize comparison units that had already had their primary election before this period of analysis. The estimates resulting from these analyses are similar but somewhat larger than those in the top panel.

To the extent that the results differ between the first and second panel of Table A5, this could be attributable to biases arising from heterogeneous treatment effects or it could be attributable to the different samples of bills. In Table A4, we found that extreme voting in the House is greatest after the filing deadline but before the primary election date, so we should expect the estimated effect of the primary date to depend upon whether we are comparing votes after the primary to those well before or just before the primary. To assess this possibility, the third and fourth panels of Table A5 replicate the analyses from the first two panels, respectively, but they only include votes taken within 30 days of the primary election date. As expected, these estimates are even less precise because the analyses are limited to a smaller subset of votes. However, the estimates in the third and fourth panel are similar to one another, suggesting that the estimates do not meaningfully depend on the subset of comparison units utilized. Overall, these results suggest that concerns about the biases of two-way fixed effects regressions are substantively negligible in this particular case.

The stacked regressions in Table A5 implicitly put more weight on members for whom we have more relevant variation and more clean comparison units. The benefit of this approach—relative to another that equally weights all members—is precision. However, if there is significant

heterogeneity across members, we might prefer to estimate the average effect across all members. To better estimate the average effect across all treated members, we can re-run the regressions from Table A5 separately for each Congress-party-dataset, record the estimate and the number of treated members, and compute the average of these estimates, weighting by the number of treated members in each Congress-party-dataset. When we do this, the estimates are similar to those in Table A5. For example, utilizing the same comparison units as in the top-panel of Table A5, we estimate that the average effect of the primary date on extreme voting in the Senate is −.0070 in the Senate (as opposed to the estimate of −.0059 that we obtain from the OLS regression in the table). Therefore, the implicit reweighting of OLS does not appear to be consequential for our substantive conclusions. However, because of the imprecision with which many of these dataset-specific ATTs are estimated, we prefer the more precise estimates in the main text and in Table A5.

Better Capturing the Full Incentive Effect

As mentioned in the main text, our aggregate estimates may underestimate the incentive effects of primary elections because even members who have passed their primary election date in the current term may still be worried about primary competition in the next term. This may be particularly relevant for members of the House who serve two-year terms. Even after they have won a primary election, their next primary is only two years away, so perhaps primary competition is always at least in the back of their mind. Members of the Senate, however, serve six-year terms, which may explain why we detect somewhat larger effects in that chamber.

If a member of Congress knows that she is seeking reelection for the last time, then the passage of her primary date means that she will never again have to worry about primary competition. Therefore, if we could identify such cases, we would better capture the full incentive effects of primary election. To this end, we estimate the probability that an incumbent member of Congress seeking

Figure A3. Age and Not Running Again

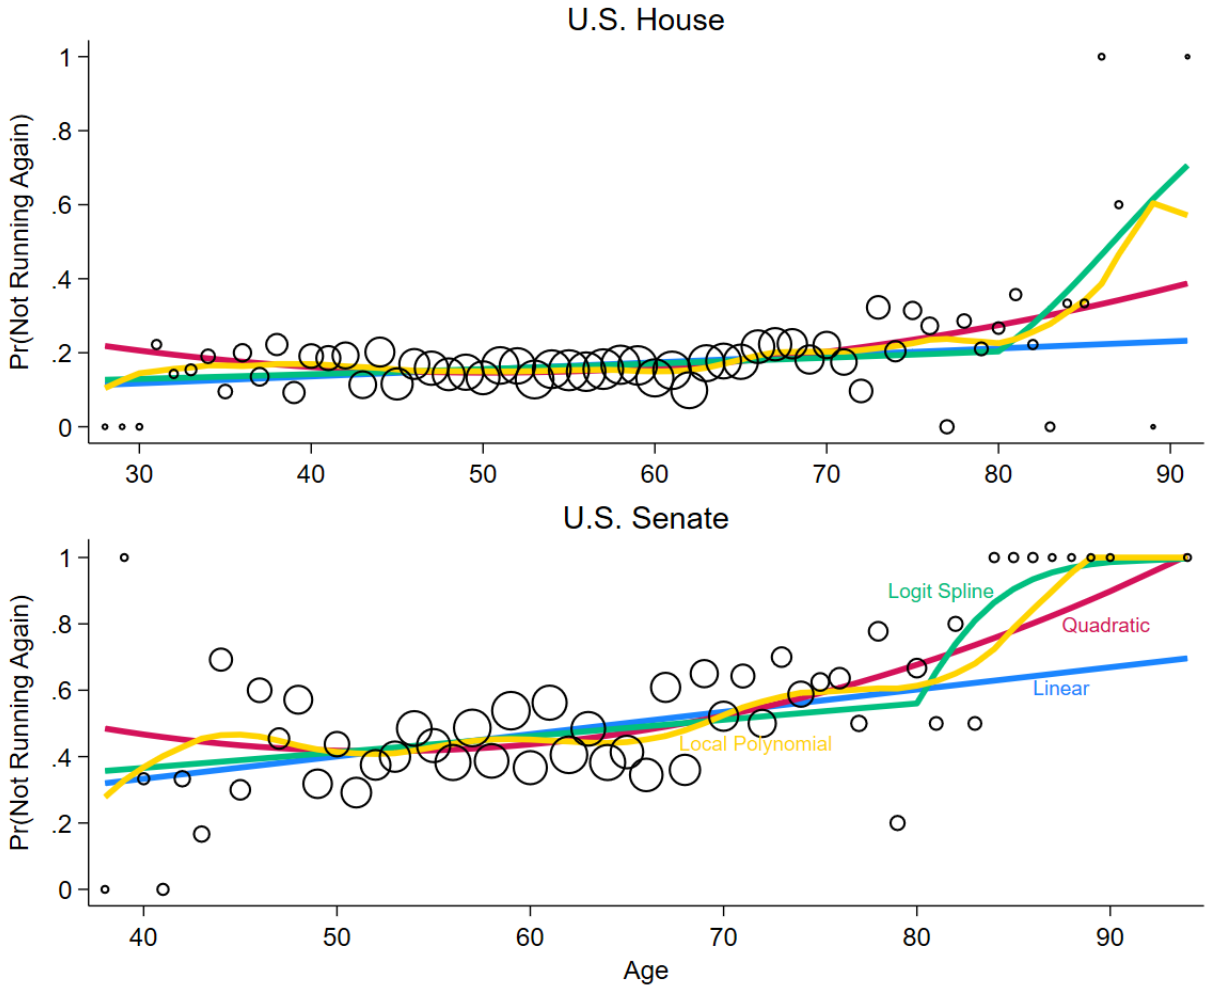

reelection will not be seeking reelection in the following term, and we estimate the effect of the primary election date separately for only those members who are not likely to run again.

Two variables that strongly predict whether a member will seek reelection again are their age and the partisanship of their constituency. Specifically, when a member is older or when a member represents a constituency that is less favorable toward her party (as measured by voting in the most recent presidential election), she is less likely to seek reelection in the next term. Using these two variables, we predict the probability that a member who is currently seeking reelection will not seek reelection again in the following term.

Figure A4. Constituency Partisanship and Not Running Again

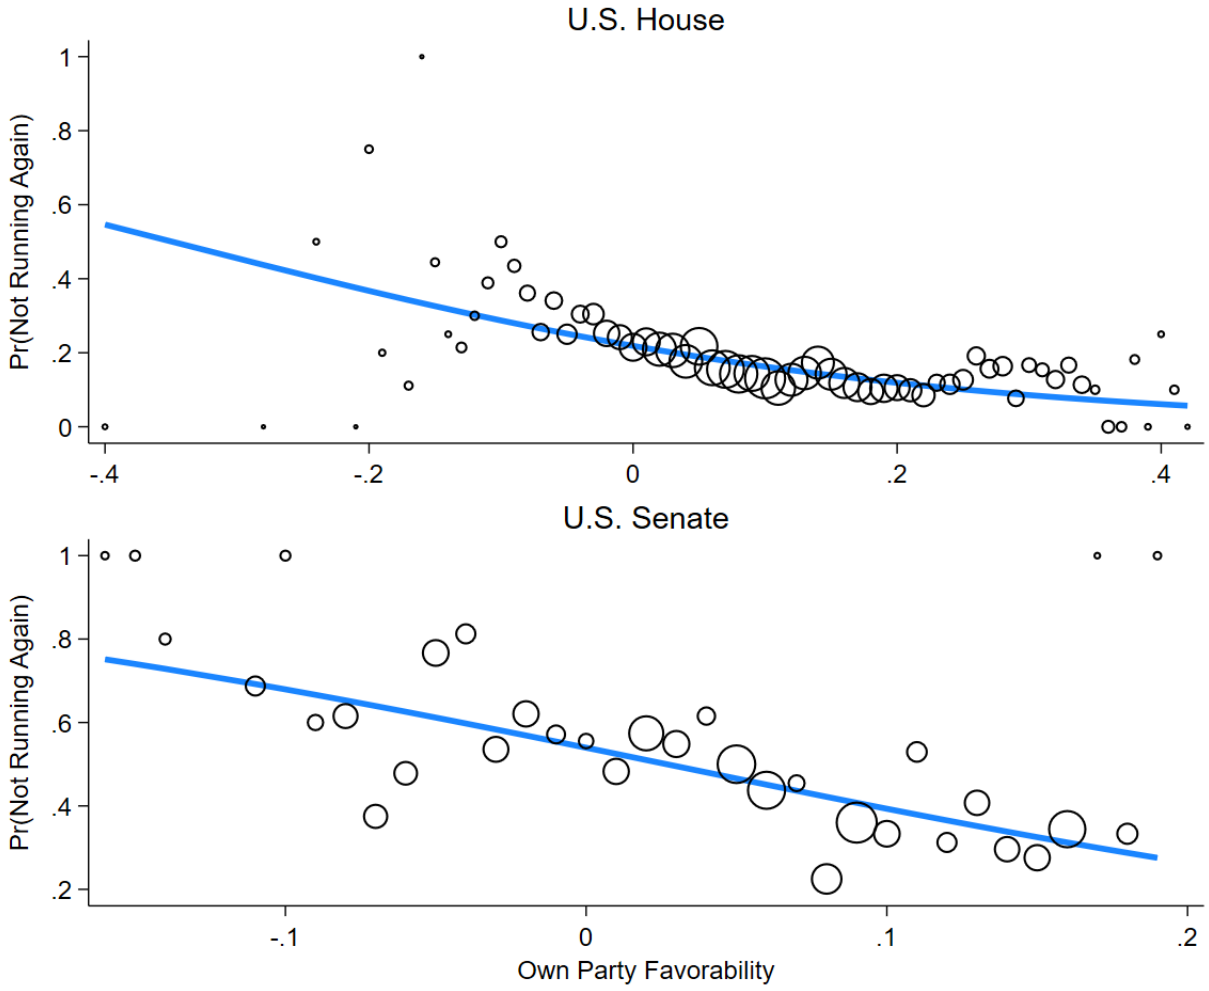

Figure A3 shows the relationship between incumbent age and the probability of not running again in the next term. The top panel shows data from the U.S. House and the bottom panel shows data from the U.S. Senate. Each circle shows the share of incumbents seeking reelection at each year of age who did not run for reelection in the following term. The size of the circles is proportional to the number of cases at that particular year of age. We see that members of the Senate are generally more likely to not run for reelection in the next term than members of the House. We also see that within each chamber, there is a strong, positive correlation between age and not running for reelection

Table A6. Predicting the Probability of Not Running Again in the Next Term

	House	Senate
	DV = Not Running in Next Term	
Age	.018**	.037**
	(.004)	(.008)
$I_{Age \geq 80}$*(Age−80)	.220**	.348
	(.072)	(.194)
Own Party Favorability	−4.031**	−7.265**
	(.391)	(.993)
Constant	−2.246**	−2.026**
	(.223)	(.504)
Observations	5,092	851

*Standard errors in parentheses; ** p < .01, * p < .05.*

in the next term. And lastly, we also see that the relationship between age and not running again is modest and approximately linear for members below the age of approximately 80, but around the age of 80, there is a dramatic kink in the relationship whereby the relationship between age and not running becomes stronger.

Figure A3 also plots four different approaches for predicting the probability of not running again using age. A linear regression line is shown in blue, and it fits the data poorly, particularly at higher ages. Similarly, the predictions from a quadratic regression (red) also fit the data poorly. A local polynomial (yellow) fits the data much better because of its flexibility, but it shows fluctuations that likely reflect overfitting to noise rather than genuine non-monotonicities in the relationship of interest. Lastly, we plot the predictions from a logit spline regression in which we allow for a kink (but not a discontinuity) in the relationship between age and not running at the age of 80. Despite its simplicity and parsimony, this approach yields predictions that fit the data very well, and therefore, this is the approach we later use to generate our ultimate predictions.

Figure A4 shows the relationship between the partisanship of a member's constituency and the probability that they do not seek reelection in the following term. To measure constituency partisanship, we compute the two-party vote share of a member's party in the most recent presidential

Figure A5. Distributions of Predicted Probabilities

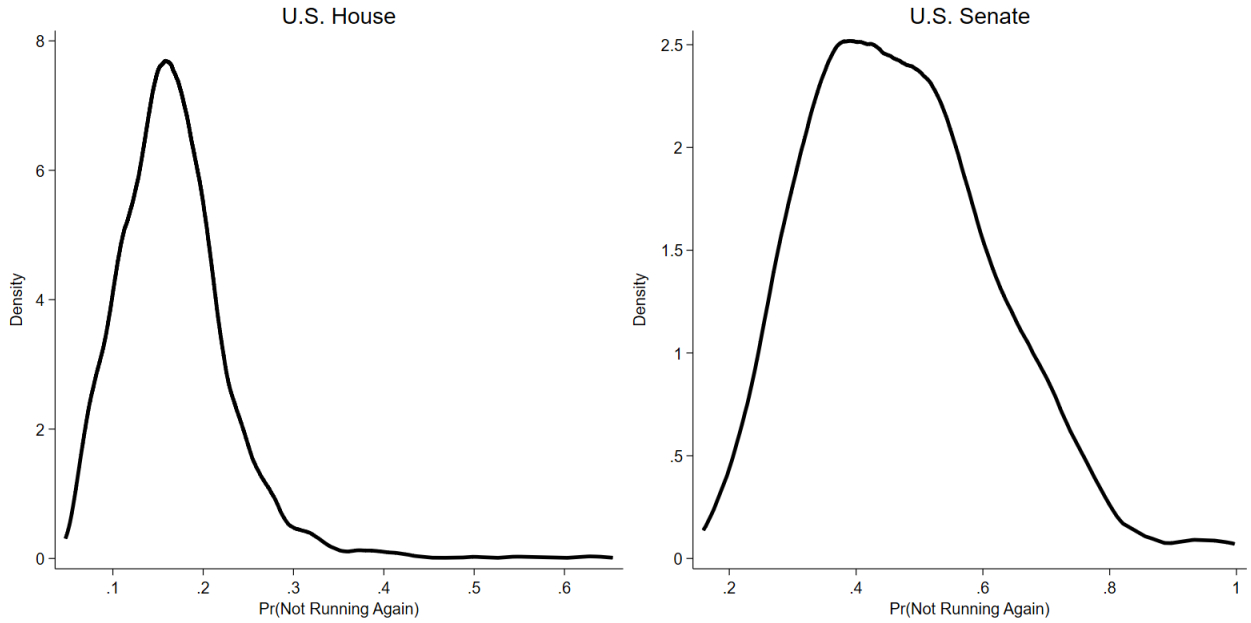

election in her district or state. We subtract the national two-party vote share, leaving us with a measure of how much favorable the constituency is to its member's party relative to the rest of the country. For illustrative purposes in Figure A4, we group constituencies into one-percentage-point bins, and we plot the share of members in each bin that did not run again in the following term. As in Figure A3, the size of the circles corresponds to the numbers of cases in each bin.

As expected, most members of Congress represent constituencies that are more favorable toward her party than the rest of the country, but there is significant variation, with many members representing constituencies that lean against their party. We also see a clear negative association between the favorability of a member's party and the probability that they do not run again. In other words, members are more likely to run again when their constituency is favorable toward their party, and they are more likely to not run again when the constituency is less favorable toward their party. Furthermore, unlike in the case of age, the relationship is approximately linear. The blue curves in

Figure A4 show predicted probabilities from logit regressions, and we see that the logit predictions fit the data quite well.

Combining the insights from Figures A3 and A4, we run a single logit regression in each chamber to predict the probability that a member seeking reelection will not seek reelection in the following term. As independent variables, we include our partisanship variable (measuring the extent to which the constituency leans toward a member's party), the age of the member, and the product of the age of the member minus 80 and an indicator for whether the member is above the age of 80. By including this interaction, we allow for the relationship between age and running again to change at the age of 80, but we do not allow for a discontinuity at 80. This specification matches the logit spline predictions in Figure A3 and the logit predictions in Figure A4, but it simultaneously uses information on both age and constituency partisanship to generate predictions. Results of these logit regressions are shown in Table A6.

Consistent with the results in Figures A3 and A4, we see that older members are more likely to not run again, this relationship is stronger for ages above 80, and members are less likely to not run for reelection when their constituency is more favorable toward their party. Figure A5 shows kernel density plots of the predicted probabilities estimates from the logit regressions in Table A6. We estimate higher predicted probabilities in the Senate than the House. In the Senate, we estimate some predicted probabilities greater than .9, and in the House, we estimate some greater than .6.

To estimate the effect of the primary election date for members who are unlikely to seek reelection again, we code a binary variable indicating whether the predicted probability of not seeking reelection is below a particular cutoff. We then modify our baseline specification by interacting the after-primary variable with this indicator. The coefficient associated with the after-primary variable is an estimate of the effect of the primary election date for members who are not likely to seek reelection. We prefer this interactive specification with the full sample as opposed to analyzing a subset of the

Figure A6. Estimates for Members Unlikely to Seek Reelection Again

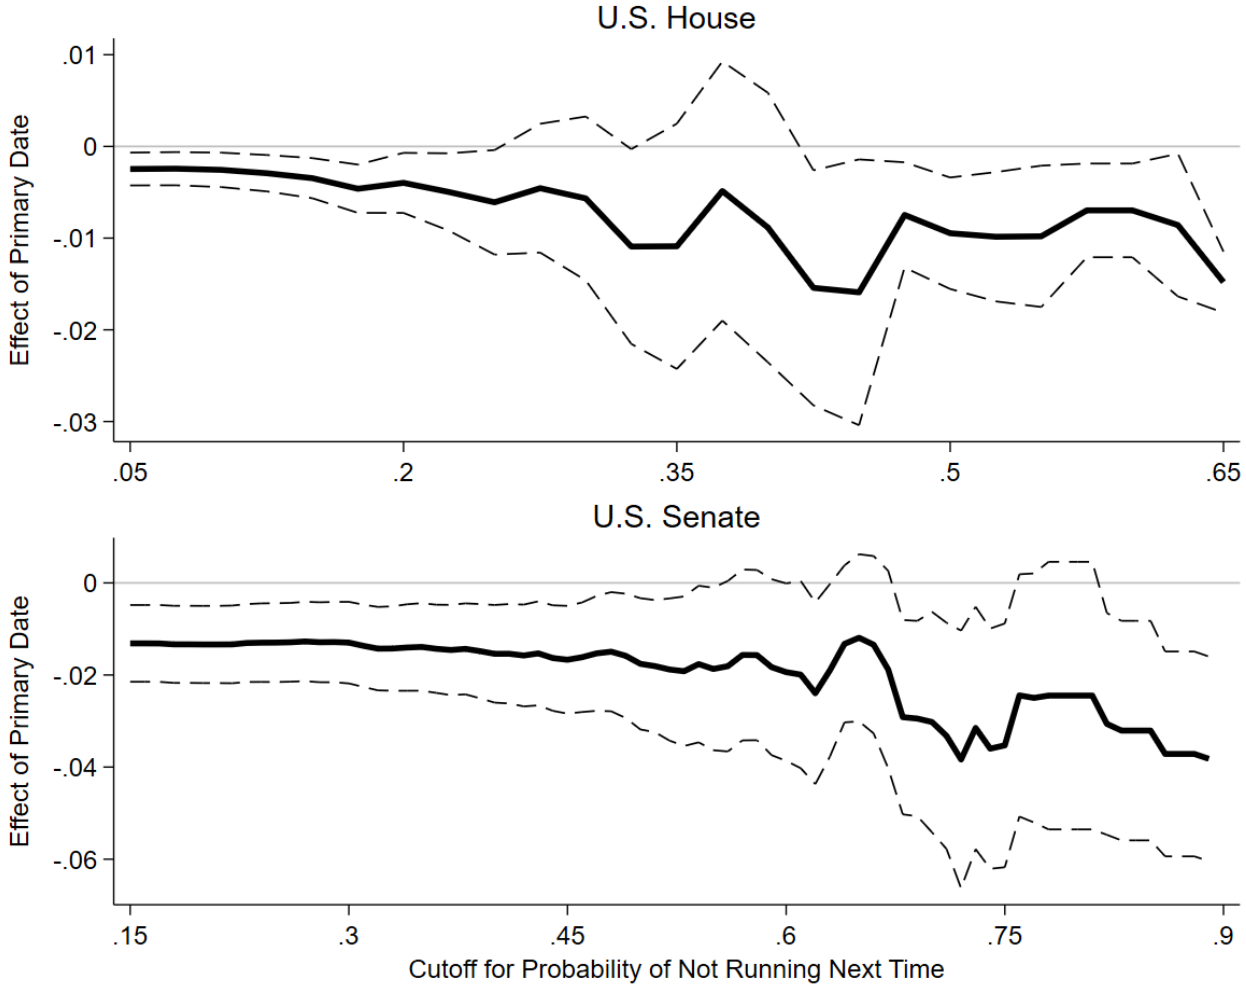

The figure shows the estimated effect of the primary election date on ideologically extreme voting for members who are above a particular predicted probability of not running for reelection in the next term. The figure shows how these estimates (and the corresponding 95 percent confidence intervals) vary as the cutoff varies.

data because the latter approach would generate noisy estimates of the bill-party fixed effects, which would lead to much less precise estimates of the effect of the primary election date.

Results are these analyses are shown in Figure A6. We vary the specific cutoff used for classifying a member as unlikely to run again, and we show estimates for virtually the entire range of feasible cutoffs. The top panel of Figure A6 shows results for the U.S. House, and the bottom panel shows results for the U.S. Senate. As expected, predicted probabilities of not running again are higher in the Senate because, with six-year terms, the rates of re-running in the Senate are generally lower.

The far left-hand side of Figure A6 shows the estimated effect of the primary election date on ideologically extreme voting for the full sample, and as we pan to the right, we can see how this estimate changes as we restrict the sample of interest to members who are less likely to run again. As expected, we detect larger (more negative) effects of the primary election date for members who are less likely to seek reelection again. This is consistent with the idea that the passage of the primary election date has a greater effect on the incentives of members who are unlikely to run again in the next term. However, even in the extreme cases on the far righthand side of Figure A6, the effects are not massive. For example, among members of the Senate who have more than a 90 percent change of not running for reelection again, we detect effects of approximately 4 percentage points, and we can statistically reject estimates of approximately 6 percentage points. Therefore, although the long-run incentive effects of primary elections may contribute to congressional polarization, even in these extreme cases, we can reject the possibility that they explain more than 20 percent of the polarization in Congress.